Locality in Vowel Harmony

Andrew Nevins

Accepted for publication in *Linguistic Inquiry Monographs*,
MIT Press
Final Version of February 20, 2009

Contents

Preface

There is a kind of poetry and elegance in the rules governing the way words are formed. It is a thrill to see how languages across the globe conform to these rules of how the vowels in a word fall into place. In my discovery of the principles and parameters of vowel harmony, I am most appreciative of the guidance and support of my mentors in graduate school, Morris Halle, Michael Kenstowicz, Donca Steriade, and Bert Vaux, for their intellectual guidance in developing a wholly new model of vowel harmony and helping me to find the right case studies to make the arguments. I would like to thank many students and colleagues at Harvard, in particular Marc Hauser, Jay Jasanoff, Peter Jenks, Beste Kamali, John Lechner, Patrick Liu, Tim O'Donnell, Gabriel Poliquin, Bridget Samuels, Kobey Shwayder, and Daniel Wallach, for the opportunity to discuss this model extensively in courses and individually. David Braun, Brian Gainor, Max Guimarães, John McCarthy, Raphael Nevins, and Cilene Rodrigues read through a number of chapters of the manuscript, and provided invaluable feedback. The reports of three anonymous *LI Monographs* reviewers were extremely insightful and enabled me to clarify and rethink many details of the proposals advanced here. I would also like to extend my gratitude to Ada Brunstein and Jay Keyser for their extensive editorial support. And finally, I would like to thank *you*, dear reader, without whom this book would be but a tree falling in a forest.

The following symbols are used throughout this book:

$\nleftarrow$ or $\nrightarrow$: search terminated early
$\Lsh$ or $\Rsh$: indicates an element being copied from

Chapter 1

What's vowel harmony, how does it vary, and why study it?

It is a common occurrence when typing an email that our fingers might hit the wrong key, and that a savvy spell-checker will underline the misspelled word in red. Suppose I am writing about the rodent protagonist of an animated film and intend to type *rat*. The resulting typo might create a real word that is different from the one I intended (e.g. *ray* instead of *rat*) or it might create an "impossible word", such as *rta*. The three-letter sequence *rta* could never be a word of English because of a constraint prohibiting the consonant sequence *rt* at the beginning of a syllable. The study of syntax begins with the observation that the words of a sentence cannot go in any order they like, and the study of phonology with the same observation for the consonants and vowels (the *segments*) of a word. Thus, *tra* is a possible segmental sequence in an English syllable, while *rta* is not. Vowel harmony, in languages that have it, is a set of restrictions that determine the possible and impossible sequences of vowels within a word.

These co-occurence restrictions are largely based on the principle of dividing the vowels of the language into two sets – let's call them "even vowels" and "odd vowels" for now– and ensuring that no mixing-and-matching of vowels from the even set and the odd set can occur in the same word. Thus, in the idealized vowel harmony language, the only permissible words would be those which are composed entirely of even vowels (e.g. "2426") or entirely of odd vowels (e.g. "1153"). The example with "even" and "odd", while formally very close to how vowel harmony works, is analogical. The actual way that a particular vowel-harmony language will divide up its inventory of vowels into two sets is along some phonetic/phonological dimension. For example, one articulatory parameter that divides up vowels neatly is whether they are pronounced with the body of the tongue aimed towards the front or the back of the mouth. In Turkish, the front vowels *i,ü,e,ö* cannot mix with the back vowels *ı,u,a,o* within the same word if it is to be considered "harmonic". Even when suffixes pile up, they keep to this restriction of back vowels after back vowels, as the word formed by adding the following thirteen suffixes to the stem "Europe" illustrates:

(1)

Avrupa-	lı-	laş-	tır-	a-	ma-	yacak-	lar-	ımız-	dan-
Europe-	from-	become-	caus.-	abil.-	neg.-	fut.-	plural-	1pl-	abl.-

mı-	y-	dı-	nız
Q-	cop.-	past-	2pl

"Were you one of those whom we are not going to be able to turn into Europeans?"

The root *Avrupa* has vowels in which the tongue is pulled back, and due to harmonization, all thirteen suffixes will have the tongue body pulled back as well. By contrast, if the root ends with a fronted vowel, all thirteen of the following suffixes will have front vowels:

(2)

Akdeniz-	li-	leş-	tir-	e-	me-	yecek-	ler-
Mediterranean	from-	become-	caus.-	abil.-	neg.-	fut.-	plural-

imiz-	den-	mi-	y-	di-	niz
1pl-	abl.-	Q-	cop.-	past-	2pl

"Were you one of those whom we are not going to be able to turn into Mediterraneans?"

Like the phonotactic restrictions banning certain sequences of consonants, such as *rta* (disallowed in English but allowed in Russian), the restrictions banning certain sequences of vowels are subject to a number of language-specific variations. But the restrictions governing possible orders and combinations of vowels within a word rest on completely different sets of principles. This is because, for one, the articulatory dimensions along which these "harmonic sets" can be divided require a certain degree of symmetry in the set of vowels in a language (witness Turkish, with four back and four front vowels); the instantiation of harmony from one language to the next may differ because not every language has a perfectly balanced vowel inventory. Furthermore, the description of the syntax of vowel sequences is more complicated in large part because vowels are rarely strictly adjacent to one another, and the resulting co-occurrence restrictions become a type of "action at a distance", unlike the strictly local rules which say that *t* cannot immediately follow *r* at the beginning of a word. Vowel harmony is one of the only phonotactic processes across human languages that consistently instantitates a "long-distance relationship".

1.1. The Computation of Locality

Though we will examine the variety of long-distance dependency relationships in vowel-harmonizing languages, this book about much more than vowel harmony. Instead, it is about a rather non-intuitive discovery of a much more general aspect of cognition – the way that humans measure distance – and illustrates this principle with the patterns of vowel harmony found in many languages across the world. The non-intuitive discovery about distance-measurement (or the computation of "locality", more technically put) is easily introduced when considering the subway map of Boston and Cambridge.

Boston and Cambridge are two cities in Eastern Massachusetts that are separated by a river, the Charles. The subway system in the Boston area is one of the oldest in the country, and has a characteristic organization into a network of color-coded subway lines. The subway system was incremented over time as the city and important sites (often organized into "Squares", such as Harvard Square, Central Square, Porter Square, Kenmore Square, etc.) within it grew in various ways. Now, suppose we are sitting in the subway station in Harvard Square (which is on the Cambridge side of the river, and on the "red" line) and you turn to me and ask, "I want to take the subway into Boston. What is the closest Boston stop we can get to from here?"

The answer that I or anyone else would give without even thinking twice would be "Charles Street", which is in the Back Bay. But of course, strictly speaking, this answer is completely false. As the crow flies, the closest stop from Harvard Square (our "source") is not Charles Street at all. As measured by absolute distance, BU Central (where the Boston University is) sits much closer to Harvard Square than the Back Bay. From Harvard Square, it is 1.7 miles to BU Central, as compared to 2.9 miles to

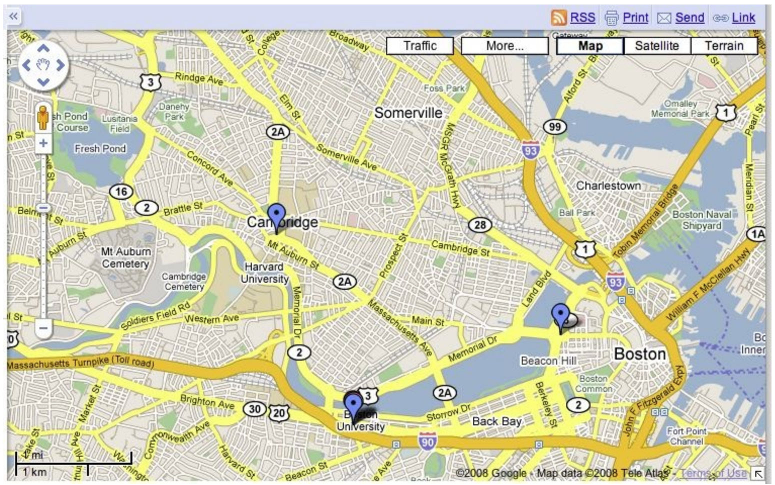

Figure 1.1: Actual Locations and Relative Distances of Subway Station Locations

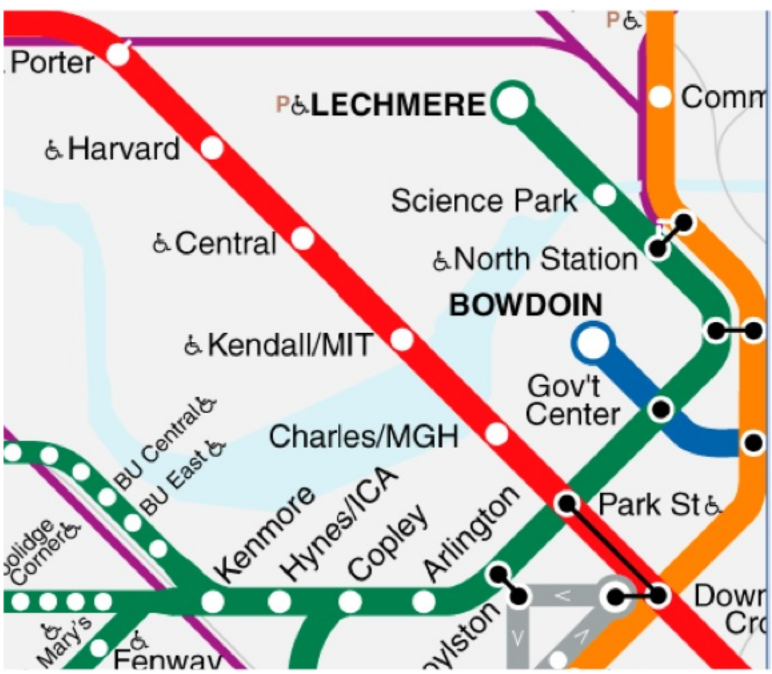

Figure 1.2: Boston Subway Connection Map, **Not Drawn to Scale!**

Charles Street; these three stations form a textbook example of Pythagoras' hypotenuse theorem, as shown in Figure 1. However, when you asked "What is the closest stop in Boston from here", if I had actually answered "BU Central", this would have been an irrelevant and smart-alecky response that took your question far too literally (much like when pedants respond to "Can I have a glass of water" with "I don't know, *can* you?"). The reason is, when you said the word *closest*, you meant "closest as measured by stops on the subway", and not closest in an absolute, as-the-crow-flies measurement of pure linear distance. And on the subway from Harvard Square, it is in fact only 3 stops to Charles Street, but 12 stops to BU Central, as Figure 2 shows (and if I had a nickel for every time I heard a complaint about this, I could buy the subway system). In fact, while Harvard Square and Charles Street are both on the "red" line, getting from Harvard Square to BU Central requires switching trains and transferring to the "green" line, a different subway line altogether. Therefore, when you asked "What is the closest Boston stop", I provided the answer that was most relevant given the structure we were traveling within, namely the subway system. The closest Boston-stop to Harvard Square *that is on the red line* is definitely Charles Street. Identical scenarios can no doubt be constructed for subway systems in a variety of other cities, and perhaps for flight routes and highway layouts as well.

This short subway dialogue illustrates the surprising centrality of *relativized locality* within human cognition. When measuring from point A (the "source") to the closest "goal" (a stop in Boston being our goal above), and there are two possible targets, B and C, it can turn out that C is decided upon as closer to A than B, *even when further away in terms of linear distance*, if C is closest as measured by some other system of measurement based on the properties it has, such as *being a stop on the red line*. Thus, Charles Street is not the closest-Boston-stop to Harvard Square, but it *is* the closest-Boston-stop relativized to the property of being on the red line, and the closest-Boston-stop as measured by the logic of subway lines, rather than absolute distance.

One of the key arguments in this book will be that in languages with vowel harmony, the relationships between vowels are governed by this relativized, rather than absolute form of locality. Thus, somewhat surprisingly, when a vowel A has to decide which of vowels B or C is closer, in certain cases, it will choose C, even when C is further away in terms of linear distance as-the-crow-flies, because C turns out to be the closest *relevant* vowel. In a certain sense, this is a very surprising result, because one might expect that in order to figure out the closest goal from a given source, you just get out your measuring tape, so to speak. But human cognition, particularly in the domain of language, does not seem to work that way; rather, we employ a metric of minimality that says that when we search for the closest Boston-stop in a subway, or the closest harmony-valuer among vowels, that locality must be defined with respect to the *properties* borne by the elements in question. Not all subway systems are laid out identically, and similarly, not all vowel harmony systems look at the same types of properties in order to determine closest-relevant-goal. One of our central objectives in this book will be to delimit a system of parametric options that vowel-harmony languages may choose from. Given the overarching and invariant principle of relativized minimality that holds across all such systems, these limited parameters determine how the variation in "what is relevant" is computed.

As a simple demonstration that vowel harmony may flaunt simple metrics of absolute distance, we may consider the ways that vowels compute distance in Yoruba (a Niger-Congo language spoken by around 18 million people in Nigeria and Benin). In Yoruba, vowel harmony is a pattern that determines where the **advanced** mid-height vowels e,o will show up as opposed to the **retracted** mid-height vowels ɛ,ɔ. The latter two, written with the less familiar symbols, are lax vowels, produced with a retracted tongue root, much like the short vowels in English *bet* and *cot*. While Yoruba and English share these four vowels, there is an important difference: a pattern like English *robot* [ro.bɔt] (or *apex* [e.pɛks]) would be against the rules of vowel harmony in Yoruba's native vocabulary, because the first o is tense and the second ɔ is lax. The rules of vowel harmony for Yoruba words ensure that lax ɛ,ɔ can only occur when another lax vowel follows. Thus, the following words are all well-formed according to the rule that "ɛ,ɔ are allowed when the next vowel is lax". (The diacritics ´ and ` indicate high and low tone, respectively).

(3) *Yoruba vowel harmony in disyllabic words*:

a. èwé 'lip'
b. olè 'thief'
c. èrò 'crowd'
d. òd͡ʒò 'rain'
e. ɛ̀gɛ́ 'cassava'
f. ɔsɛ 'soap'
g. ɛ̀fɔ́ 'vegetable'
h. ɔwɔ́ 'hand'

However, the concept that is used to measure "next vowel" in Yoruba is not, as might appear solely based on the above examples, strictly based on absolute distance. Since what is relevant to vowel harmony is the rule that forces all mid-height vowels e,o,ɛ,ɔ to share the same feature along the tense/lax dimension, other vowels that are not relevant to this distinction could presumably intervene without interfering in the dependency relation among these four types of vowels.

Examining the variety of Yoruba spoken in Ifẹ, and comparing it with standard Yoruba, we find that these two closely-related dialects differ only in how they measure locality. Both dialects have the same inventory of vowel sounds, and both dialects share a large part of the lexical vocabulary. However, the rule of vowel harmony operates rather differently in these two dialects. In Ifẹ Yoruba, the first vowel in a word determines its tense/lax harmony according to the next *relevant* vowel. In other words, if the second vowel in the word is a high vowel like *i* or *u* (which are tense vowels), it is not counted as "next" in Ifẹ Yoruba. This is because the tense/lax distinction *is not important* in the high vowels *i,u* in Ifẹ Yoruba, and so these vowels are not counted in the computation of "adjacent" vowel. (In Chapter 5, we return to the pattern of the low vowel *a* in Yoruba harmony). This relativization of the measurement of "next-to" in such a way that it does not measure absolute distance, but rather distance as determined by the relevance of the tense/lax contrast, exemplifies the core of our study of the locality of vowel harmony.

(4) *Ifẹ Yoruba skips the irrelevant second vowel (Ola Orie, 2001)*:

a. ɔrúkɔ 'name'
b. ɛ̀lùbɔ́ 'yam flour'
c. ɛúrɛ́ 'goat'
d. ɔdídɛ 'parrot'
e. ɔ̀títɔ́ 'truth'

What is most interesting about the pattern of Ifẹ Yoruba is that it counts locality differently from Standard Yoruba, which does measure locality in absolute, as-the-crow-flies terms (i.e. counts the "closest vowel" as indeed the closest vowel), as can be seen by comparing the pronunciation of these same words:

(5) *Standard Yoruba counts the second vowel as the next vowel*:
a. orúkɔ
b. èlùbɔ́
c. ewúrɛ́
d. odídɛ
e. òtítɔ́

In Standard Yoruba, the tense/lax value of *any* vowel matters for counting what's next and therefore ɔ cannot precede u. In Ifẹ Yoruba, harmony only looks at tense/lax among the vowels e,o,ɛ,ɔ, and so the intervening *u* in (4-a) is not even counted in the computation.

This brief comparison highlights the fact that two languages with the same vowel inventory – indeed two dialects of the same language – may conduct their search for vowel harmony features entirely differently, depending on what counts as closest. The locality of vowel harmony in Ifẹ Yoruba is determined by the closest vowel contrastive for the tense/lax distinction, while the locality of vowel harmony in Standard Yoruba is determined by the closest vowel, period. In essence, each language may relativize its concept of closest vowel based on properties of the vowel inventory as a whole.

1.2. Vowel Harmony is distinct from Coarticulation

While discussing Yoruba dialects, a slight twist arises in the harmony patterns of Central Ede dialects, which include Ijẹṣa Yoruba and Akurẹ Yoruba. Recall that in Standard Yoruba, the high vowels *i,u* have no lax counterparts ɪ,ʊ. However, in the Central Ede dialects, there *is* a contrast between tense *i,u* and lax ɪ,ʊ, as seen in Ijẹṣa *rín* 'laugh' vs. *rín* 'walk' and *mún* 'be sharp' vs. *mʊn* 'drink' (Bamgboṣe, 1967; Akinlabi, 2008). Since the high vowels i,u,ɪ,ʊ can vary along the tense/lax dimension, they too participate in vowel harmony:

(6) *Ijẹṣa/Ekiti Yoruba vowel harmony affects the high vowels*:
 a. ɔ̀ʊ́kɔ'he-goat'
 b. eʊ́rɛ́ 'she-goat'
 c. ɛrʊ̀pɛ̀ 'sand'
 d. ɔrʊ́kɔ 'name'
 e. ɛ̀lʊ̀bɔ́ 'yam flour'

What makes Ijẹṣa high vowels undergo harmony, while Standard Yoruba high vowels don't? Some scholars of dialect variation might consider the participation of high vowels in harmony an innovation, and as with all instances of phonological distinctions, then investigate the seeds for these types of change. In a broader sense, the fact that Ijẹṣa high vowels undergo harmony and Standard Yoruba high vowels don't undergo it can be seen as the question of why some languages develop vowel harmony and others don't, writ small.

One idea that has been proposed is that vowel harmony rules in a language are the result of elevating certain phonetic tendencies of the articulators to the status of hard-and-fast grammatical rules. For example, it is well-known from experimental phonetics that in a great deal of languages, the movement of the tongue body involved in articulating vowels, like many muscular articulations, displays anticipatory movements (approaching its target before it actually needs to) and carryover movements (staying at its target longer than it actually should). These types of articulatory "bleedthrough" of an articulation from one sound to the one that follows or precedes happen in all languages, and are known as coarticulation. Coarticulation does not happen all the time for every vowel, and it happens to varying degrees, depending on many factors such as how fast and how carefully one is speaking.

According to the foremost proponent of this view, Ohala (1994, p.491), "Vowel harmony, *phonological* co-occurrence constraints between the features of vowels in polysyllabic words, is a *fossilized* remnant of an *earlier* phonetic process involving vowel-to-vowel assimilation" (italics mine).

Provided that we interpret *fossilized* to mean phonologized, i.e. part of the phonological computation, and *earlier* to include the possibility of a no longer extant phonetic process, this statement is orthogonal to an analysis of harmony in a synchronic description. Indeed, Ohala emphasizes that this fossilization involves a transformative difference between what was once a gradient, variable, coarticulatory, phonetic process and what has become a categorical and symbolic phonological process, stating "whereas the magnitude of the phonetic variation may vary continuously as a function of the strength or proximity of the contextual environment, the phonological variation is typically of an 'all-or-none' sort" (p.491).

The proposal of vowel harmony as the phonologized all-or-none version of a once variable phonetic process of coarticulation, while speculative, has sparked much interest and research. In fact, pursuant to the question of whether the harmony affecting the high vowels in Central Ede dialects of Yoruba is a development arising from coarticulatory trends existing in the language family, Przezdziecki (2005) attempted to study whether dialects *without* harmony of the high vowels nonetheless showed phonetic trends in that direction.

What Przezdziecki (2005) found was that in Standard Yoruba and in the Mọba dialect of Yoruba, while these languages' high vowels do not undergo harmony, there was nonetheless a variable but statistically significant trend towards the vowel *i* being more laxed when it preceded a lax vowel. Przezdziecki (2005) did not find such a trend for *u*, and concluded "this leaves the dilemma of explaining how it is that /i/ and /u/ pattern together as targets for harmony [in Central Ede dialects], since they do not both exhibit harmony-like patterns in these [Mọba and Standard Yoruba] data" (p.225). As a result, it is not clear that the participation of the high vowels *i,u* in harmony in Central Ede dialects such as Ijẹṣa and Akurẹ results from a simple thresholding / bootstrapping of an existing coarticulatory process into a phonologized pattern, as Przezdziecki (2005) had originally hoped to find.[1] Rather, as we will show throughout this book, rules of vowel harmony are based on natural classes of symbolic phonological features, such as [+high]. These rules vary across dialects and languages in how they compute locality: in some cases, they treat *i,u* as completely invisible to vowel harmony (as in Ifẹ Yoruba); in others, as participating donors of vowel harmony (as in Standard Yoruba), and finally in yet others as elements that themselves undergo vowel harmony (as in Ijẹṣa Yoruba). This variation is of a cognitive and symbolic nature that cannot be reduced to coarticulatory trends of varying strengths: even the most careful acoustic studies do not find evidence for these trends where they would be expected if all harmony ultimately corresponded to below-the-radar coarticulation. Microvariation in vowel harmony found between closely related dialects of the same language is the result of how locality is computed, and can be modeled by a small set of parametric options that restrict the possible variation in what may count as relevant.

1.3. Locality is not Measured in Pure Distance in Syntax Either

An additional example in which locality is not measured in terms of distance alone may be drawn from syntactic dependencies. Anaphoric noun phrases such as *himself* and *herself* mean nothing on their own, and have as their defining property that they must find another noun phrase within the same sentence to fix their reference. In a sentence such as *Mary thinks that John is proud of himself*, we understand *himself* to mean *John*, the closest noun phrase that can fix the reference of this anaphor. Thus, a reasonable description of elements like *himself* is that there must be a nearby referential element (with matching gender and number) in order to license the presence of the anaphor.

We can minimally change this sentence by flipping the position of the two noun phrases, yielding *John thinks that Mary is proud of himself*. We then observe that the result is ungrammatical. The intuition that we have is that in this case the word *John* is "too far" from *himself* in order to fix its reference. The measuring-tape theory of locality can get this result: there is another noun phrase closer to *himself*: *Mary* is only 4 words away, while *John* is 7 words away. In terms of absolute distance, *Mary* is closer, but since *himself* is specified as masculine, its closest goal is of the wrong gender. The sentence thus "crashes": the search for a closest goal in terms of absolute distance resulted in failure to find a noun phrase that matched in gender.

But the measuring-tape theory makes a prediction that is immediately wrong. Consider a sentence such as *Mary thinks that a sister of John is proud of himself*. Much like the first sentence we considered, the word *John* is 4 words away from himself, and so should by all rights be the closest element. But this sentence is ungrammatical. The reason is, we don't count distance in terms of number of words. Much like the subway system of

Boston, it is the relevant distance within a structure that matters. The sentence *Mary thinks that a sister of John is proud of himself* has an internal structure in which *a sister of John* is the subject of the predicate *proud*, and it turns out that what's relevant for anaphoric distance is being an argument of the predicate. It's *a sister of John* who is proud, and hence it's *a sister of John* that will be considered as the closest noun phrase to *himself*, and indeed, it's *a sister of John* which fails to match the gender of *himself* and causes the search to fail and the sentence to crash. What's relevant in determining the closest noun phrase for fixing the referent of *himself* is, yet again, not absolute distance, but distance as relativized to being an argument of the same two-place predicate *proud* that the anaphor *himself* is an argument of[2]. This fact about human language is one that never ceases to cause problems for strictly word-by-word models of syntax: the relevant distance to determine anaphoric locality is not computed as the crow flies, but relativized to certain properties borne by the elements being measured. In the chapters that follow, we explore a number of cases that are somewhat more intricate from the domain of vowel harmony to demonstrate that this basic principle is universal across seemingly different levels of linguistic structure.

1.4. Overview of the Major Claims

One source of inspiration for this book began with the suspicion that a casual remark made by Noam Chomsky in the fall of 2000 might be incorrect, namely the remark that there was little indication that the strong minimalist thesis (that language provides optimal solutions to linking interfaces) could hold for phonology. One of the central objectives of this book is to demonstrate that subsegmental phonology may be insightfully modeled according to the core tenets of the Minimalist program, namely (a) an emphasis on interface-driven computations, (b) a principle of minimal search and efficient computation, and (c) an attempt to derive crosslinguistic variation from the structure of the inventory rather than violable principles. The principal empirical focus of syntactic theory is on long-distance dependencies, such as the agreement relation (*Agree* in Chomsky (2000) et. seq) between a verbal predicate (the *Probe*) and an argument noun-phrase (the *Goal*).

Vowel harmony provides one of the best opportunities within phonology to explore the calculus of the *Agree* relation, as vowel harmony typically involves cases of a long-distance dependency between the dependent vowel (typically, an alternating suffixal vowel) and the determiner of its harmonic value, namely the closest eligible leftward vowel within the root. In fact, alternating suffixal vowels, which by definition have no stable underlying value, may be modeled as underlyingly un-valued; however, an interface requirement (namely phonetic convergence) requires full specification for these values.

Vowel harmony emerges as a computational operation whose purpose is to provide a solution for valuing for this "uninterpretable" suffixal vowel through the principle of minimal search for the closest *Goal* to act as a feature-source. Importantly, I argue that the application of the *Agree* model to asymmetric dependencies among subsegmental features in phonology goes far beyond simple resemblance: for example, the perplexing phenomena of "parasitic harmony" within phonology (whereby a Goal providing a value for a feature F bears a prerequisite value for a feature G) are argued to receive a unified analysis under the rubric of *Defective Intervention* developed for the *Agree* system by Chomsky (2000) and Hiraiwa (2001). Finally, the model of vowel harmony, necessarily "target-centered" (i.e. one in which the suffixal vowel initiates the search and in which locality is defined from the perspective of the value-copying, rather than value-providing, element) is argued to be empirically superior to traditional "trigger-centered" accounts of vowel harmony for cases of non-constituent copying in Turkish and bidirectional copying in Woleaian.

In comparing the execution of interface-driven feature-copying at the levels of both syntactic structure and phonological structure, I introduce the notion of *Crossmodular Structural Parallelism*, a hypothesis about the nature of human language that seeks to minimize differences between levels that do not follow from a difference in alphabet. I will argue that the core locality computation driving *Agree* and vowel harmony differs only in the alphabet of data structures to which it applies. This is a unorthodox hypothesis, driven by the goal of a higher-order synthesis between linguistic phenomena, and a similar hypothesis has been pursued to some extent in Anderson (1992). The attempt to track as closely as possible the syntactic phenomena of locality in tandem with the conditions on long-distance vowel harmony is a unique effort, and one whose success will ultimately be measured in terms of its empirical accuracy.

The data structures adopted in this book follow the tradition of generative phonology (see Odden (2003) for a recent introduction) in employing binary features as the basic currency of intersegmental relations. The theoretical framework adopted in this book is the Principles-and-Parameters approach (Chomsky, 1981) to linguistic universals and variation, in which a core set of invariant principles hold across all languages and a limited set of parameters restrict the possible variation. The book is typologically oriented through a wide range of languages from diverse families, drawn upon in order to demonstrate the limits of crosslinguistic variation in locality effects in vowel harmony.

There are a few things that this book will not cover. In particular, while the discussion focuses on the locality of vowel harmony in over thirty different languages, these cases are ones in which the relation between two vowels is clearly one of a phonotactic dependency. However, cases of "morphemic" harmony, in which vowels or consonants undergo "mutations" as the exponent of a morphological process, will not be handled. While phonological harmony is about featural identity *between* elements, morphemic harmony requires that a feature be realized across a certain domain. For example, in Kanembu, the distinction between two types of aspectual interpretation requires that all vowels in the word are [+ATR] or [−ATR], even though there is no visible affix for completive or incompletive aspect:

(7) *Kanembu completive is [+ATR] and incompletive is [−ATR] (Akinlabi, 1996):*

gɔnəkɪ	'I took'	gonʌki	'I am taking'
dalləkɪ	'I got up'	dʌllʌki	'I am getting up'
barɛnəkɪ	'I cultivated'	bʌrenʌki	'I am cultivating'

Kanembu instantiates what is called a case of "featural affixation" (Akinlabi, 1996) in which the relevant affix is only expressed by a single phonological feature, which is then realized on specific elements of the word to which it attaches. Such cases of morphological exponence represent a wholly different type of process than phonological vowel harmony (see Cole (1987) and Finley (2009) for a thorough comparison).

Vowel harmony of the type exemplified in the Altaic, Finno-Ugric, Niger-Congo and Bantu languages, among others studied throughout this book, is a locality computation initiated by a single morpheme, which copies features from a single source. Morphemic harmony, by contrast, appears to be a globally-sensitive process of realizing a single feature, e.g. [+nasal] (sometimes on as many segments in the word as possible, as the exponent of a morphological category that does not have its own segmental content. For example, in Terena (Piggott, 1988), the first person possessor is grammatically realized by [+nasal] features showing up on every segment of the word possible:

(8) *Terena Nasalization*:

ajo	'his brother'	ãjõ	'my brother'
arine	'his sickness'	ãr̃ĩnẽ	'my sickness'
owoku	'his house'	õw̃õᵑgu	'my house

Similarly, in Mixtec (Piggott, 1992), inflection for the 2nd person is achieved by nasalizing as many segments as possible from the final segment leftward. While vowel harmony may fail to occur altogether for a given morpheme if the vowel fails to find a source of harmony within the bounds permitted by locality, morphemic harmony is different, in that it always occurs in order to realize a morphological category; the only thing that varies is how many segments it will get a chance to apply in a given word. In Korean ideophone harmony, used only to express 'light' and 'dark' senses of ideophones, both [± ATR] and [± high] sometimes change the initial vowel of a word only, e.g. *minduŋ*∼*mɛnduŋ* 'bald', while sometimes change only the initial and final vowel, e.g. *umullək*∼*omullak* 'chewing'. These processes are related to exponence of "floating" features via mutation of the edges of a word.

Morphemic harmony can also require that a class of vowels appear only in the presence of a particular morpheme. For example, [−ATR] vowels only occur in the presence of the masculine singular suffix ʊ in Pasiego: "Lax vowels may not occur in a word without this suffix" (McCarthy, 1984, 294). When the masculine singular morpheme is added, every vowel possible becomes [−ATR], but never outside of the highly specific inflectional category of the masculine singular.

Such processes are well-accounted for by the literal notion of a "floating" feature, that associates to one or more edges of the word and attempts to express itself on as many segments as possible within some bounded range (Akinlabi, 1996; Zoll, 1996; Wolf, 2007) . Vowel harmony, however, is not: I will argue that it is crucially needy-vowel-centric, rather than being computed from the point of view of the donor. The omission of morphemic harmony from the present investigation is a principled exclusion.

A second topic, fascinating in its own right, but not within the aims of the current investigation, concerns the functional motivations or "benefits" of vowel harmony. This book focuses on the formal, rather than substantive properties of this phonological process, attempting to characterize what is a possible or impossible vowel harmony system, which is an independent line of inquiry from how useful such a system might be in terms of its auditory-perceptual properties. Illuminating experimental work such as Suomi et al. (1997). Vroomen et al. (1998) and Mintz and Walker (2006) has demonstrated that vowel harmony can be used as a heuristic in word segmentation, providing cues to the listener for where one word stops and the next word begins – in much the same way that a fontsizechangecanbeused as a replacement for whitespace in perceiving separate words in text (Kaye, 1989). The research area concerned with what vowel harmony may be functionally useful for will be left to other ambitious researchers: our focus is on the procedures and constraints responsible for how it is that harmony is executed and computed within a word, and on how the notion of closest vowel is measured, rather than how it may be used in connected speech. This methodological narrowing of focus is parallel to the choices that are made within syntactic theorizing to explore the restrictions on the locality of agreement relations independently from the pursuit of what verbal agreement may be used for.

The narrative of the rest of this book is broadly as follows. I will introduce how Search works for harmony and why the locality of harmony should be computed from the point of view of the needy segment. I will then show how locality is computed in a relativized manner and how languages may differ in terms of what is considered relevant. Elements considered "irrelevant" form a theory of what are traditionally called *transparent* elements. Next, I will provide a number of cases where harmony fails to occur. Elements that interfere with the locality of vowel harmony form a theory of what are traditionally called *blocking* elements. I will argue that the right way to model all of these effects is in terms of a Search procedure whose locality instantiates many of the exact same principles and parameters that determine locality of syntactic dependencies. The algorithm developed for vowel harmony will ultimately provide a bridge between phonological theory and syntactic theory.

Chapter 2 introduces the core mechanism of vowel harmony: the Search Principle, in which a search is conducted from a point-of-origin (such as Harvard Square in our subway example above) to find the closest element that satisfies a certain need (such as setting the value of harmony). The main demonstration will be that locality is best modeled in terms of measuring from the dependent element (the recipient of a feature-value) rather than in terms of measuring from the donor element.[3] Illustrating the application of Search in Turkish, Woleaian, and Barra Gaelic harmony, this chapter situates vowel harmony within rubric of *Search for Closest Element* that has become a key focus of inquiry in the Minimalist approach to linguistic computation (Chomsky, 1995).

Chapter 3 introduces an important source of crosslinguistic variation: the definition of "relevant" that is important in pinning down *Closest Relevant Element* as discussed above. Different settings of this parameter may yield dramatically different surface behavior in harmony, as demonstrated in the contrast between Ifẹ Yoruba and Standard Yoruba. The main demonstration of this chapter is that locality may be relativized to certain types of properties of phonological features, namely the contrastive or marked status of the features in the segment that bears them.[4] This exclusion of irrelevant segments subsumes the locality effects called "transparency" in the traditional literature on vowel harmony. In illustrating the parametric variation in how feature-sensitive locality may be relativized in Finnish, in the Turkic and Tungusic language families (Karaim, Sibe, and Sanjiazi Manchu), and in dialectal variation within Kirghiz, Yoruba, and Finnish, this chapter draws important formal parallels between vowel harmony and the theory of relativized minimality that proscribes a theory of *Closest Relevant Element* within syntactic theory (Rizzi, 1990).

Chapter 4 explores what happens when the search fails due to an intervening segment that stands between the value-seeking recipient and its donor in the search domain. This subsumes the locality effects called "blocking" or "opaqueness" in the traditional literature on vowel harmony. A certain segment, due to an intrinsic property, such as its value for an orthogonal feature to the one being searched for, is a "defective" value-source. Since the search procedure has no lookahead, as soon as a defective element in the domain is encountered, the search terminates in failure, even if there is a potentially eligible value-source further downstream. In illustrating the effects of failed copying due to Minimality in Nawuri, Kisa, Khalkha Mongolian, and Jingulu, this chapter draws important parallels with the effects of defective intervention in syntactic agreement – where an inappropriately-case-marked noun phrase with plural number blocks agreement between a verb and a postverbal subject noun phrase, yielding a defective intervention effect in which the verb can look no further than the intervening element that caused the search to fail.

Chapter 5 introduces a final source of parametric variation: how far a Search can go before it gives up. Extrinsic barriers on search yield a type of locality, sometimes called "islandhood" (Ross, 1967), by which certain extrinsically defined domains limit the extent of dependency relations, even where relativized closeness is kept constant. The chapter demonstrates how crosslinguistic and even within-language variation in harmony can result from different settings of "how far a vowel is able to look" – where the search must end, whether something is found or not.[5] In illustrating the parametric variation in *Domain of Search-Extent* in Hungarian and in Gikuyu, this chapter draws important parallels with the role of closed domains of computation in Barrier-based (Chomsky, 1986) or Impenetrability-based (Chomsky, 2001) formalizations of the limitations of syntactic dependency. The second half of the chapter explores the ways in which high-sonority elements may also close off a domain of search, acting as "hurdles", and establishes an implicational generalization about sonority, based on Wolof, Classical Manchu, Hungarian, and Finnish low vowels.

The broader consequences of this principle-and-parameters approach are that the vowel harmony systems of the world can be essentially reduced to one basic principle, with some parametric variables that can be filled in on a language-specific basis, and that what appears to be a different nuanced harmony system in every language reduces to variations on a core theme. The analysis of diverse harmony patterns through a limited set of parametric options adds to the growing body of evidence (e.g. from stress systems in phonology, or from wh- movement in syntax) that the diversity of human languages is not the result of every imaginable Babelian divergence from one to the next. Crosslinguistic and within-language variation results from easily isolated parametric choices (in this case, what's in the search domain and how far it extends) supplied as input to a fundamental algorithm of human cognition. The implications of a single principle of locality, and its consequences for our understanding of the structural parallels between phonology and syntax within the human language faculty, are discussed in Chapter 6.

1.5. The Need for a New Model of Vowel Harmony

Existing models of vowel harmony can be broadly classified into two types: declarative identity-enforcement and sharing-by-spreading. Examining these in turn, with reference to various patterns of vowel harmony, we will conclude that neither type is sufficient to model the locality of vowel harmony, as identity-enforcement models are too permissive (in that they undesirably *allow* patterns of vowel harmony that do not exist), and sharing-by-spreading models are too restrictive (in that they undesirably disallow patterns of vowel harmony that *do* exist).

The distinction between these two types of models has been very well-characterized by Rose and Walker (2004, p.520), who propose that natural language actually contains both mechanisms for harmony: declarative identity-enforcement for processes that show no locality constraints at all (i.e. any two eligible vowels in a word must be identical for the harmonic feature, regardless of their distance), and sharing-by-spreading for processes that are strictly local, and allow no skipping of elements between harmonizing vowels.

Under declarative identity-enforcement (e.g. relations of surface *correspondence* between vowels (e.g. Bakovic (2000); Krämer (2003); Hansson (2001)), it is straightforward to model for the long-distance nature of vowel harmony. Given constraints on what types of vowels need to be identical for which features, there is no issue of locality: the harmonizing vowels are simply subject to a requirement of identity for roundness or whatever the harmonic feature is. Thus, in Ifẹ Yoruba, the requirement might be ALL MID VOWELS MUST BEAR THE SAME VALUE OF LAXNESS. A word like ɛ̀lùbɔ́ would satisfy this requirement, whereas a hypothetical Ifẹ word like ɛ̀lùbó or èlùbɔ́ would not. The issue of the non-lax *u* in the middle of these two vowels would simply not arise, as the constraint on identity does not mention high vowels.

Under sharing-by-spreading (e.g. autosegmental spreading or Firthian prosodies (Palmer, 1970; Goldsmith, 1976; Gafos, 1996)), vowel harmony involves the spreading of a single featural specification (e.g. [+round]) across an entire word, where the extent of spread of the feature may be halted by vowels that are incompatible with the spreading feature or are already specified for the opposite value. Thus, in Standard Yoruba, the laxness feature in a word like èlùbɔ́ would begin spreading from the rightmost vowel, would be stopped by the tenseness value on the vowel in the medial, which would then spread its own tenseness value to the left.

If we only consider these two types of models, one might say that Ifẹ Yoruba employs declarative identity-enforcement, while Standard Yoruba employs sharing-by-spreading. Different dialects would then exhibit different locality restrictions because they employed different mechanisms for achieving vowel harmony.

The issue becomes more complicated, however, when we consider languages in which some vowels are treated non-locally while others are not. Consider the pattern of [± round] harmony in Khalkha Mongolian which is responsible for alternations in the form of the reflexive suffix as either a or round o,ɔ, depending on the roundness values of vowels to its left (where the further distinction between o/ɔ is due to an additional harmony for tense/lax).

(9) *Khalkha Mongolian (Svantesson et al., 2005, 50)*:

  a. poor-ig-o 'kidney-acc-refl.'
  b. xɔɔlʒ-ig-ɔ 'food-acc.-refl'
  c. mʊʊr-ig-a 'cat-acc.-refl.'
  d. suulʒ-ig-e 'tail-acc-refl'.

As (9-a-b) show, the intervening vowel *i* is not counted in the determination of the suffix's harmony: even though *i* is unround, it is irrelevant for the computation of closest vowel. This might lead one to think that Khalkha vowel harmony is the result of declarative identity-enforcement. However, as (9-c-d) show, the high vowels u,ʊ cannot provide a value for round harmony, even though they are round. Now, one might simply counter that u,ʊ are simply not part of the relevant constraint, which is instead ALL NON-HIGH VOWELS MUST BEAR THE SAME VALUE OF ROUNDNESS. In this case, there is simply no identity-requirement that would cause mʊʊr-ig-a to become mʊʊr-ig-ɔ.

As it turns out, however, u,ʊ *are* involved in the computation of locality, because they "prevent" rounding harmony between non-high vowels. In (10), we see that the perfect suffix harmonizes with round o,ɔ in (10-a-b), but not when the closest vowel is round u,ʊ, in (10-c-d):

(10) *Defective Intervention in Khalkha Rounding Harmony*:

  a. tor-oːd 'be.born-perf.'
  b. ɔr-ɔːd 'enter-perf.'

c. tor-uːl-eːd 'be.born-caus.-perf'
d. ɔr-ʊːl-aːd 'enter-caus.-perf.'

We must conclude that declarative identity-enforcement cannot account for Khalkha rounding harmony, as it is "blocked" by vowels of the wrong height. At the same time, the strictly-local spreading-by-sharing model cannot account for Khalkha either, as the vowel *i* is "transparent" and not part of the locality computation.

The spreading-by-sharing proposal of Rose and Walker (2004) is one in which all feature-sharing is strictly local, and Gafos (1996, p.77-81) demonstrates that the autosegmental formalism actually requires all spreading to be strictly local, given an interpretation of autosegmental association lines as temporal overlap. Indeed, the very intuition of sharing-by-spreading theories is an intuition of articulatory "persistence" of a rounding gesture from one vowel to the next, a notion that is very difficult to square with the data in (10-c-d), in which two round vowels in a row nevertheless do not induce lip-rounding in a third vowel that follows both of them.

The intuition underlying the central proposal in this book is that the autosegmental formalism providing the basis for sharing-by-spreading is not the best way to model the locality of vowel harmony, and that a target-centric, relativized search makes an empirically tighter set of predictions. The autosegmental formalism was developed for tonal phenomena, which display very different locality phenomena from vowel harmony, as tonal phenomena essentially show no instances of "blocking" of long-distance spreading.[6] In addition, tonal phenomena often involve the existence of indepedent "melodies", such as HL in Kukuya (Zoll, 2003), that must be mapped to accomodate monomoraic, bimoraic, and trimoraic stems as $\widehat{\text{HL}}$, HL, and HLL, respectively; there is no obvious analogue to this word-length-invariant "melody" in any subsegmental feature sequence. I find it entirely plausible that tonal interactions and subsegmental feature-copying might employ different grammatical mechanisms, and will remain largely agnostic on the question of whether long-distance tone spreading can be assimilated to the target-centric model proposed in this book, or whether traditional autosegmental accounts for tone should be maintained. (Interestingly, Leben (2006), one of the founders of autosegmental approaches to tone, has recently discussed the possibility that alternative formalisms for tone realization are superior to autosegmental theory). In any event, the phenomena of vowel harmony themselves demand a new model that strikes a balance between being overly restrictive and overly permissive.

The present book offers a new theory of vowel harmony, the Search-and-Copy model, which models crosslinguistic variation in locality as the result of different *parameters* on a search procedure. Under the current model, harmony is a search initiated by a "needy" vowel for the features (e.g. laxness, roundness) that it needs. Variations in locality (e.g. between Ifẹ Yoruba and Standard Yoruba) are determined by parameters that cut certain vowels out of the search domain, rendering them irrelevant in the computation of "closest". Some languages may relativize their search only to vowels that are *contrastive* for the harmonic feature, as in Ifẹ Yoruba, as analyzed in Chapter 3; as a consequence, the high vowels will be pruned from the search domain and not counted in the determination of closest. Other languages may relativize their search only to vowels that are *marked* for the harmonic feature, as in Khalkha Mongolian, as analyzed in Chapter 4; as a consequence, the high front vowel *i* will be irrelevant for the search and not counted in the determination of closest.

Yet other languages may determine that, even though the principle of searching for the closest relativized element is crosslinguistically invariant, once the closest relevant vowel is found, additional parameters prevent it from being copied from. Thus, even though Khalkha ʊ is the closest vowel to the perfect suffix in ɔr-ʊːl-aːd, it is of a different height from the suffix vowel, and hence cannot be copied from. The search, however, is over, and even though an eligible vowel is just further upstream, no harmony happens in this case. In the chapters that follow, I argue that patterns of transparency and opacity of intervening vowels can be predicted from a restricted set of parameters based on properties of the languages' inventory and requirements on boundedness, sonority, and orthogonal identity.

The computation of locality in vowel harmony in the Search-and-Copy model is target-driven, as it involves a procedural search initiated by the needy vowel. In a sharing-by-spreading model of harmony, the harmonic value must extend throughout the domain, and failure of a potential target to undergo harmony is treated as a lexical exception. As a result, random occurrences of harmonizing and non-harmonizing affixes are not predicted. However, an important claim in this book is that it is the morphemes that acquire a harmony value that must be morphologically specified as needing this property; morphemes that do not harmonize simply lack this requirement, and this requirement is what is crucial to the initiation of a search by the needy morpheme.

This asymmetric computation of locality (in which 'closest' is counted from the perspective of the feature-seeker, rather than the feature-provider), is defended in Chapter 2, and marks a significant departure from sharing-by-spreading theories, in which locality is computed from the perspective of the feature-provider, or from static / declarative theories in which vowels must simply be identical for the relevant feature. These approaches to locality are empirically distinguishable, and I argue that the right fit between predictions and data is achieved by the target-initiated model.

It is perhaps most broadly interesting that the algorithm developed for vowel harmony based on the empirical demands of coverage that is neither too restrictive nor too permissive results in a model that very closely mirrors the way that locality is computed for syntactic dependencies such as verbal agreement. Previous models of vowel harmony, either in the spreading-by-sharing or declarative identity-enforcement camps, bear no obvious relation to considerations that have led to the development of the *Agree* procedure in syntax, with its emphasis on efficient, interface-driven, no-lookahead, relativized search as a method for providing feature-values. The present model allows for a new line of comparison of the computational systems of syntax and phonology, as discussed in Chapter 6.

In the chapters that follow, readers will notice that successive presentations of the vowel harmony algorithm will be accompanied by pseudocode representations of the search procedure. This algorithm has been implemented in python, and is available along with case studies, at http://www.people.fas.harvard.edu/∼nevins/HarmonyDemo.zip. It is my hope that interested readers will verify the correctness of the predictions of the principles and parameters of the locality of vowel harmony proposed here through experimentation with this software.

Notes

[1]Moreover, Przezdziecki (2005) found near significant coarticulatory effects for the vowel [a], even though in no dialect of Yoruba does [a] categorically participate in harmony. Thus, it seems that if harmony is a phonologization of existing phonetic coarticulatory tendences, cognitive and symbolic aspects of phonological computation step in to 'filter' and extract from those tendencies. Finally, as Adetugbo (1967) and Bamgboṣe (1967) argue that proto-Yoruba had the tense/lax contrast between i/ɪ,u/ʊ, the Central Ede dialects do not represent an innovation (cf. Adetugbo (1967, p.158)); rather the full range of this contrast seems to have been lost in different ways throughout Yoruba.

[2]The exposition here follows Reinhart and Reuland (1993) in their discussion of co-argumenthood as what makes a noun phrase relevant for binding. For a more detailed and technical exposition of binding theory, see Büring (2005).

[3]Mailhot and Reiss (2007) pursue a number of convergent ideas for vowel harmony, from the perspective of learnability, making points of contact with certain aspects of the current model.

[4]The work of Calabrese (1995, 2005) has set the stage for these type of "visibility parameters" in the current theory.

[5]The "distance parameters" of Chapter 5 are inspired by Odden (1994)'s work on harmony.

[6]For example, Cassimjee and Kisseberth (2001) discuss the process of high tone movement in Nguni, whereby a high tone associated with a prefix moves to the antepenult of a toneless verb stem. Although Nguni has depressor consonants that inhibit the local association of high tones, depressor consonants do not block the long-distance process of high tone movement (Cassimjee and Kisseberth, 2001, p.350ff).

Chapter 2

The Search Principle

Natura semper agit per vias brevissimas
'Nature always acts through the shortest pathways" (Burtt 1932:39)

2.1. Harmony and *Agree* are both a Process of Finding a Value

Assimilation refers to any process of copying the features from one segment to another, and making the recipient segment "more alike" with the donor segment. For example, the casual pronunciation of "in Boston" is *im Boston*, where the expected nasal consonant with an alveolar place of articulation is instead pronounced with a bilabial place of articulation. Within phonological theory, assimilation processes (in this case, the nasal consonant copying the place of articulation of the following segment) are analyzed as the result of copying a feature from one segment to another. Phonological features specify and constitute the articulatory instructions that make up segments. Thus, "N" is really just a shorthand for the cluster of features including [+nasal], [coronal], and so forth.[7]

The assimilatory process of copying the bilabial property of "B" to a nasal, turning an "N" into an "M", is the result of copying [labial] from one segment to another.

Vowel harmony is a special kind of assimilation that differs in two ways from the case above. First, it is specifically operative between vowels, and as a result does not simply occur between strictly adjacent segments. Second, vowel harmony often affects or involves all vowels within a word, rather than simply the two segments at the juncture between, say preposition and noun.

One of the reasons that assimilation of vowel features is called "harmony" is because it instantiates a kind of dependency relation, in which a newly-added suffix is forced to "harmonize" with the stem it attaches to. A musical analogy could be made, perhaps, to an open-mic jazz session: a trumpet player who walks in to an ongoing tune must play in the right key that's already been established. The solo that he or she will play depends on the key that has already been set.

With vowels, there are only four or five features along which this kind of harmonization can be required. Three of them have to do with dimensions of tongue movement: [$\pm$ back] determines whether the tongue body is fronted or backed, [$\pm$ high] determines whether the tongue body is raised or not (while [$\pm$ low] determines whether it is lowered or not), and [$\pm$ advanced-tongue-root] (abbreviated ATR) determines whether the root of the tongue (which anchors it towards the pharynx) is advanced forward or not. In addition, the feature [$\pm$ round] determines whether the lips are rounded during the production of a vowel or consonant (distinguishing, for example, French *i* and *ü*), and the feature [$\pm$ nasal] determines whether airflow through the nose occurs (distinguishing, for example, Portuguese *a* and *ã*). In principle, harmony can be required among any of these features.

A straightforward and well-known example where harmony is required is found in Turkish. Turkish does not employ the features [$\pm$ATR] or [$\pm$nasal] in the phonology of its vowels, and [$\pm$ high, $\pm$ back, $\pm$ round] are sufficient to uniquely determine each of the eight vowels of Turkish. As in the example of our trumpet player, when a new suffix waltzes up to an existing root, it must engage in harmony. A newly-added accusative-case suffix is forced to harmonize with the existing "key" set by the root for the features [$\pm$ back] and [$\pm$ round]. Since there are two possible values for [$\pm$ back] and two possible values for [$\pm$ round], the suffix can show four different forms: i is [$-$back, $-$round], ü is [$-$back, $+$round], ɨ is [$+$back, $-$round], and u is [$+$back, $+$round].[8] The choice of which of these four will surface is dependent on what the values for these features in the root are. While the accusative suffix brings its own height feature [$+$high] to the table, it must agree with the root for the other two features. (This is a specific aspect of the way the morphology-phonology interface works in Turkish: specific morphemes must be learned to participate in harmony, while certain others do not.)

(1) *Turkish Vowel Harmony of Accusative Suffix*

stem	ACC.SG	
ip	ip-i	rope
kɨz	kɨz-ɨ	girl
yüz	yüz-ü	face
pul	pul-u	stamp
el	el-i	hand
sap	sap-ɨ	stalk
köy	köy-ü	villa
son	son-u	end

This type of dependency relation, in which one element's form depends on the features of another's, are quite common across varying levels of linguistic structure. Consider verb agreement in Italian: *Tu sei andato*, "You have gone"; *Gioia è andata* "Gioia has gone"; *I ragazzi sono andati* "The boys have gone", which illustrates cases of the verb showing a masculine singular, feminine singular or masculine plural ending, depending on what the gender feature of the subject noun phrase is. Clearly we would not say that the gender or number of *Gioia* or *I ragazzi* depend on the ending of *andato/andati/andata*: it's rather an asymmetric dependency relation in which the verbal elements depend on the subject for their identity.

(2) *Italian auxiliary and participle agreement*

Dependent element	Copies Person:	Copies Number :	Copies Gender:
sei	2nd	[−plural]	none
è	3rd	[−plural]	none
sono	3rd	[+plural]	none
andato	none	[−plural]	[−fem]
andata	none	[−plural]	[+fem]
andati	none	[+plural]	[−fem]

Within syntactic theory, the agreement found on the auxiliary (*è* vs. *sono* vs. *sei*) and the participle (*andata* vs. *andati* vs. *andato*) are modeled as asymmetric dependency relations. With these verbs of motion in the present perfect tense, the auxiliary copies the person and number features from the subject noun phrase, and the participle copies the number and gender features from the subject noun phrase. Notice that the auxiliary does not agree in gender, and that the participle does not agree in person. The auxiliary "needs" person and number, but not gender, whereas the participle needs gender and number but does not "need" person. Agreement is viewed as a process in which recipient nodes find and copy the features that they need.

The goal of this chapter is to present a model of vowel harmony processes like (1) as the result of a search-and-copy procedure that is formally identical to verbal agreement as illustrated in (2), with the obvious difference that syntactic structure defines closeness in terms of hierarchical c-command whereas phonological structure definies closeness in terms of linear precedence. Irrespective of syntactic or phonological structure, the core principles are identical between the *Agree* operation defined in syntactic theory and the *Harmonize* operation, whose preliminary formulation is:

(3) *Statement of Harmony as a Procedure like Syntactic Agree*
A recipient needs a value for a feature F. Search is always initiated *from* the target/recipient of assimilation. Once the target encounters a donor, it copies the value of a feature F

Crucial to the model will be the notion that the segments of a word are in a total order, so that for any x and y that are segments in the word, either x will precede y or y will precede x.[9] The notion of precedence permits one to derive relation of *closer*:

(4) *Definition of Closer*
Given a, b, c: b is *closer* to a than c if either (i) a precedes b and b precedes c or (ii) if b precedes a and c precedes b,
where a, b, c are segments

The search procedure for vowel harmony begins with a recipient and looks for the closest donor in a certain direction δ, which can be either to the left or to the right (or in both directions simultaneously).

(5) *Harmonic Search-and-Copy, in Two Steps: (τ, δ, F)*
a. Find: x = the closest τ to the recipient y in the direction δ
b. Copy: the value of F on x onto y,
where x, y are segments, F is a feature, τ is a predicate over segments

When there is more than one feature to copy, these are all looked for on

myVals V
myPosition P
myFeatsneeded F

while F is not empty:
· Go in direction δ and update P
· **if** P has a value for any $f, f \in F$:
· · Copy Val(P, f) to V
· · Remove f from F

Figure 2.1: Single-Pass Search with all Features Harmonized.

the same search. With Turkish suffixal high vowels, for example, Back-Harmonize and Round-Harmonize occur on the same Search; that is, the same pointer is tracking positions backward, searching simultaneously for each independently and valuing each as soon as possible. The Search mechanism is formalized in the algorithm in Figure 2.1, where variables named with ”my” represent the values updated during Search.

That’s the core of the search algorithm: a forced march leftwards and/or rightwards that stops at each potential donor segment, sees whether there is a [± back] and/or [± round] value to copy, and if there is, takes it. The search-and-copy mechanism is simple and minimalist, but this particular formulation will turn out to have important consequences in the empirical cases examined in the following sections.

2.2. Back-Harmony and Round-Harmony in Turkish Suffixes

Let’s see observe this process in action by returning to the Turkish accusative suffix. The suffix vowel is merged with the root, and needs a value for [±back] and a value for [± round]. It will conduct a search-and-copy procedure for each of these values simultaneously, looking at possible vowel donors.

(6) *Turkish Accusative Case morpheme suffix must:*
Back-and-Round Harmonize: δ = left, F = [$\pm$ back, $\pm$ round]

The procedure is illustrated for *ip-i* 'rope+accusative': The X-slots above each feature-matrix represent the Root node, over which precedence relations are defined, and the arrow represents the direction of search δ. A turned arrow ($\Lsh$) indicates copying from that source.

(7) Accusative Begins Back-Harmonize in *ip-i*:

$$\begin{array}{ccccc} x_1 & & x_2 & \leftarrow & x_3 \\ \begin{bmatrix}\text{+voc}\\ \text{+high}\\ \text{-back}\\ \text{-round}\end{bmatrix} & & \begin{bmatrix}\text{-voc}\\ \text{labial}\\ \text{-cont}\\ \text{-nasal}\end{bmatrix} & & \begin{bmatrix}\text{+voc}\\ \text{+high}\\ \\ \\ \end{bmatrix} \end{array}$$

(8) Accusative Suffix Finds [-back] on x_1 and finds [-round] on x_1:

$$\begin{array}{ccccc} x_1 & \leftarrow & x_2 & & x_3 \\ \begin{bmatrix}\text{+voc}\\ \text{+high}\\ \textbf{-back}\\ \textbf{-round}\end{bmatrix} & & \begin{bmatrix}\text{-voc}\\ \text{labial}\\ \text{-cont}\\ \text{-nasal}\end{bmatrix} & & \begin{bmatrix}\text{+voc}\\ \text{+high}\\ \\ \\ \end{bmatrix} \end{array}$$

(9) Accusative Suffix Copies [-back] and [-round] to itself:

$$\begin{array}{ccccc} x_1 & \Lsh & x_2 & & x_3 \\ \begin{bmatrix}\text{+voc}\\ \text{+high}\\ \textbf{-back}\\ \textbf{-round}\end{bmatrix} & & \begin{bmatrix}\text{-voc}\\ \text{labial}\\ \text{-cont}\\ \text{-nasal}\end{bmatrix} & & \begin{bmatrix}\text{+voc}\\ \text{+high}\\ \textbf{-back}\\ \textbf{-round}\end{bmatrix} \\ i & & p & & i \end{array}$$

The case illustrated above is extremely straightforward, and this basic principle of Search-and-Copy underlies all vowel harmony, though as we will see in Chapters 3, 4 and 5, the items that are searched for and the extent of the search are subject to varying parametric conditions.

2.3. Harmony can Iterate from Morpheme to Morpheme

Harmony occurs whenever a new morpheme that needs a value is added (as well as when a new vowel is added to a morpheme by epenthesis), and what makes the process iterative is that a morpheme M_j that has found a value for harmony can then in turn provide a value for a linearly subsequent morpheme M_{j+1}. This iterative process is illustrated below in the interaction of harmony in Turkish with both a Genitive Case morpheme and a Plural Number morpheme.

(10) *Turkish Genitive Case morpheme suffix must:*
Back-and-Round Harmonize: δ = left, F = [$\pm$ back,$\pm$ round]

(11) *Turkish Plural morpheme suffix must:*
Back-Harmonize: δ = left, F = [$\pm$ back]

(10) says that the genitive case must find values for both [$\pm$back] and [$\pm$round], while (11) says that the plural must find a value for [$\pm$ back]. The plural suffix is inherently [$-$round]. (In fact, all non-initial [$-$high] vowels in Turkish are [$-$round], unless lexically specified otherwise.) An important consequence of the Plural's failure to undergo harmony for [$\pm$round] means that any harmonically-dependent element that immediately follows the Plural will be [$-$round]:

(12) *Harmony in Turkish Nouns plus Plural plus Genitive*:

stem	+PL+GEN	
ip	ip-ler-in	rope
kɨz	kɨz-lar-ɨn	girl
el	el-in	hand
el	el-ler-in	hand
sap	sap-lar-ɨn	stalk
yüz	yüz-ler-in	face
pul	pul-lar-ɨn	stamp
köy	köy-ler-in	villa
son	son-lar-ɨn	end

Observing (12), one might think that the genitive suffix is limited to alternating between *-in* and -ɨn. However, this is not the case: the form of the genitive singular (i.e. just the genitive suffix) is not equal to just taking out *-lar/ler* above and subtracting it away to get the result. The relationship between the suffix and the root is not direct: it is mediated by the harmony that the intervening plural suffix happened to (not) undergo. The following examples show that the noun plus the genitive alone is distinct from subtracting the plural out of (12):

(13) *Harmony in Turkish nouns plus Genitive alone*:

stem	+GEN	
ip	ip-in	rope
kɨz	kɨz-ɨn	girl
el	el-in	hand
sap	sap-ɨn	stalk
yüz	yüz-ün	face
pul	pul-un	stamp
köy	köy-ün	villa
son	son-un	end

Compare, for example, the form of the genitive suffix in *pul-un* and in *pul-lar-ɨn*. There is a rounded vowel in the genitive when it immediately follows the root, and there is an unrounded vowel in the genitive when it immediately follows the plural. There is a simple explanation for this behavior and it is in fact exactly what we would expect given a principle of closest-source as defined in terms of leftward precedence.

(14) Plural Begins Back-Harmonize in *pul-lar*:

$$\underset{}{\overset{x_1}{\begin{bmatrix}\text{+voc}\\ \text{+high}\\ \text{+round}\\ \text{+back}\end{bmatrix}}}\quad \overset{x_2}{\begin{bmatrix}\text{-voc}\\ \text{+sonorant}\\ \text{-nasal}\\ \text{+lateral}\end{bmatrix}}\quad \overset{x_3}{\begin{bmatrix}\text{-voc}\\ \text{+sonorant}\\ \text{-nasal}\\ \text{+lateral}\end{bmatrix}}\quad \overset{\leftarrow}{}\quad \overset{x_4}{\begin{bmatrix}\text{+voc}\\ \text{-high}\\ \text{-round}\\ \end{bmatrix}}\quad \overset{x_5}{\begin{bmatrix}\text{-voc}\\ \text{+sonorant}\\ \text{-nasal}\\ \text{-lateral}\end{bmatrix}}$$

(15) Plural Finds and Copies [+back] :

$$\underset{\text{u}}{\overset{x_1}{\begin{bmatrix}\text{+voc}\\ \text{+high}\\ \text{+round}\\ \textbf{+back}\end{bmatrix}}}\quad \overset{\Lsh}{}\quad \underset{\text{l}}{\overset{x_2}{\begin{bmatrix}\text{-voc}\\ \text{+sonorant}\\ \text{-nasal}\\ \text{+lateral}\end{bmatrix}}}\quad \underset{\text{l}}{\overset{x_3}{\begin{bmatrix}\text{-voc}\\ \text{+sonorant}\\ \text{-nasal}\\ \text{+lateral}\end{bmatrix}}}\quad \underset{\text{a}}{\overset{x_4}{\begin{bmatrix}\text{+voc}\\ \text{-high}\\ \text{-round}\\ \textbf{+back}\end{bmatrix}}}\quad \underset{\text{r}}{\overset{x_5}{\begin{bmatrix}\text{-voc}\\ \text{+sonorant}\\ \text{-nasal}\\ \text{-lateral}\end{bmatrix}}}$$

Added next is the Genitive Suffix, next in line in terms of the morphological structure.

(16) Genitive Begins Back- and Round-Harmonize in *pul-lar-ɨn*:

$$
\begin{array}{ccccccc}
x_1 & x_2 & x_3 & x_4 & x_5 & \leftarrow & x_6 \\
\begin{bmatrix}\text{+voc}\\ \text{+high}\\ \text{+round}\\ \text{+back}\end{bmatrix} &
\begin{bmatrix}\text{-voc}\\ \text{+sonorant}\\ \text{-nasal}\\ \text{+lateral}\end{bmatrix} &
\begin{bmatrix}\text{-voc}\\ \text{+sonorant}\\ \text{-nasal}\\ \text{+lateral}\end{bmatrix} &
\begin{bmatrix}\text{+voc}\\ \text{-high}\\ \text{-round}\\ \text{+back}\end{bmatrix} &
\begin{bmatrix}\text{-voc}\\ \text{+sonorant}\\ \text{-nasal}\\ \text{-lateral}\end{bmatrix} &
 &
\begin{bmatrix}\text{+voc}\\ \text{+high}\end{bmatrix}
\end{array}
$$

(17) Genitive Finds and Copies [+back] and [−round]:

$$
\begin{array}{ccccccc}
x_1 & x_2 & x_3 & x_4 & \Lsh & x_5 & x_6 \\
\begin{bmatrix}\text{+voc}\\ \text{+high}\\ \text{+round}\\ \text{+back}\end{bmatrix} &
\begin{bmatrix}\text{-voc}\\ \text{+sonorant}\\ \text{-nasal}\\ \text{+lateral}\end{bmatrix} &
\begin{bmatrix}\text{-voc}\\ \text{+sonorant}\\ \text{-nasal}\\ \text{+lateral}\end{bmatrix} &
\begin{bmatrix}\text{+voc}\\ \text{-high}\\ \textbf{-round}\\ \textbf{+back}\end{bmatrix} &
 &
\begin{bmatrix}\text{-voc}\\ \text{+sonorant}\\ \text{-nasal}\\ \text{-lateral}\end{bmatrix} &
\begin{bmatrix}\text{+voc}\\ \text{+high}\\ \textbf{-round}\\ \textbf{+back}\end{bmatrix} \\
\text{u} & \text{l} & \text{l} & \text{a} & & \text{r} & \text{ɨ}
\end{array}
$$

We can observe that this process is iterative. The key is, there is no relation between the genitive suffix and the root in (16), and, more generally, there is never any relation stated as "between genitive suffix and the root". Rather, there is a relationship between the genitive suffix and the immediately preceding vowel, *whatever morpheme it is contained in*.

2.4. Turkish Vowel Harmony Reflects a Phonological Computation

As specified in the δ parameter of its pattern of vowel harmony, Turkish copies vowel features from the leftward direction. One interesting question is whether there is any relationship between the algorithmic procedure of copying feature-values from left-to-right and the articulatory production trends for vowel sequences in the language that do not display harmony. For example, one might expect that if vowel harmony somehow represents the crystallization or discretizing threshold of existing articulatory patterns of tongue body backing or fronting in the language (an existing hypothesis among some researchers, as mentioned in Section 1.2), then even words without this thresholded rule might show low-level, transient effects of left-to-right effects of coarticulation. If, on the other hand, vowel harmony is a symbolic phonological computation, as formalized here, then it should be possibly independent of what happens with the coarticulation of vowels in their phonetic implementation.

In a study of Turkish disharmonic sequences of the form *iCa, iCe, eCa* and *iCa* (where C represents a consonant between the two vowels), Beddor and Yavuz (1995) measured the amounts of anticipatory (i.e. right-to-left) and carryover (i.e. left-to-right) coarticulation affecting the fronting or backing of the vowels. They found that all three vowels exhibited effects of anticipatory coarticulation (e.g., *i,e* were reliably more backed when followed by *a*, and *a* was more fronted when followed by *i,e*) but that the effects of carryover coarticulation were much more restricted – essentially only an effect of fronting of *a* when preceded by *i*. They conclude that 'If the primary left-to-right phonological pattern of vowel harmony in Turkish were to reflect active phonetic patterns of vowel-to-vowel coarticulation, then carryover should exceed anticipatory coarticulation in Turkish disharmonic roots. However, the reverse phonetic pattern holds for the set of roots tested here, in that anticipatory coarticulation was found to be a much more general phenomenon than carryover coarticulation...One interpretation to be considered is that perhaps this outcome should cause us to question our underlying assumption of a cause-and-effect relation between vowel-to-vowel coarticulation and vowel harmony....The alternative interpretation that we offer is that once a phonetic behavior is phonologized, it becomes largely distinct from the behavior which gave rise to it" (p.49).

This is an important finding about the modular division between phonological computation and phonetic exigencies. Perhaps a useful analogy is that of the rules of etiquette for holding a wine glass: all epicurean reference guidelines assert that one must hold a wine glass by the stem. While this clearly has a functional origin in laws of thermodynamics, in that holding a wine glass by the bowl can raise the temperature of chilled white wines, it has been generalized and formalized beyond its original context, to places where it no longer has a functional need at all, such as with red wines, or mulled wines served warm, where the rule of etiquette doggedly proceeds with a logic of its own. While the original cause-and-effect relationship between the gradient, variable laws of heat transfer and the rule of etiquette is clear, once formalized as a hard-and-fast rule, it remains distinct from the rules of physics that gave rise to it. Indeed, having your hand down on the stem still does in fact transmit minor amounts of heat to the wine bowl above it, but this is okay in terms of etiquette as long as the rule is categorically followed.

Returning to the fact that Turkish vowel harmony is a phonological computation with a left-to-right direction, while at the same time its transient coarticulatory patterns are stronger from right-to-left, we note that Modern Turkish stress falls on the final syllable. Phonetic coarticulation effects are often stronger from stressed vowels to unstressed vowels (Fowler, 1981), which may explain the directionality of the right-to-left coarticulatory pattern that Beddor and Yavuz (1995) found. Interestingly, Proto-Turkic is posited to have had stress on the initial syllable (Poppe, 1960), and the location of initial stress is thought to have given rise to left-to-right vowel harmony in Turkish (Barnes, 2006). Once the rules of vowel harmony became an entrenched phonological computation in Turkish, even when the stress pattern of the language changed and the pattern of vowel-to-vowel phonetic coarticulation changed, the harmony pattern remained the same, an autonomous search-and-copy computation of the kind specified in (11).

2.5. Non-Undergoing Morphemes can Still be Donors

Up to this point, we have formulated harmony as essentially a statement about a phonological process that a certain morpheme must undergo upon entry into the phonological computation. Just as it is a listed fact in the speaker's mental representation of lexical items that sometimes a grammatical morpheme may not receive any phonological content (such as the plural of *sheep* in English), it is also a property of some morphemes that they must or must not undergo some type of vowel harmony. We thus might expect another suffix in Turkish with similar overall phonological properties to the plural suffix (i.e. being a CVC syllable with a [−high] and [−round] vowel), but which does not have to undergo harmony. The instruction-label for harmony would have to say:

(18) Turkish Nominalizing morpheme suffix must:
(empty)

In other words, while it is a property of the plural morpheme that it must Back Harmonize, showing versions either *-lar* or *ler* depending on the value of the closest source to its left, this is not a property of the nominalizing morpheme *-gen*. It does not harmonize, is not a phonologically "needy" element, and little more needs to be said in the lexicon than that.

In some way of looking at it,*-gen* is an "exception" to vowel harmony, i.e. a suffix that does not "need" to undergo it. But being exceptional with respect to harmony means only that. It is an automatic consequence of the asymmetrically-dependent recipient-centric nature of vowel harmony that:

(19) *A morpheme might not need to search for harmony, but it still must be a donor*:
While a suffix might be "exceptional" in not-undergoing harmony, no suffix is listed as "exceptional in not-triggering harmony".

Finley (2007) argues for the same point: crosslinguistically, there are a number of exceptionally non-searching morphemes, but there are no exceptionally non-donating morphemes. Thus, of the two logically possible types of exceptions in generative phonology distinguished by Coats (1970), exceptionally non-undergoing a phonological process vs. exceptionally non-providing the features/environment needed for a phonological process, in vowel harmony there are only the former type of exceptions. This is exactly what we would expect if harmony is a process initiated by the searching, "needy" morpheme, in which the donor is merely a passive source of a value to be copied.

Non-participating morphemes in vowel harmony are not like noble gases in molecular chemistry: it is not as though these morphemes completely fail to interact in harmony with the other morphemes around them. It's just that *they* don't need a value. While *-gen* does not need to undergo harmony itself, the exception does not impact on whether other morphemes will have such a need or not. We can actually stack the harmonizing plural suffix right next to its stoic foil, the nominalizer, and the result is that *the plural's harmony is determined as usual*, by the value of [−back] immediately to its left, namely the lexical and unalternating value of that feature borne by *-gen*.

(20) *Turkish Non-Harmonizing Nominalizer Plus Harmonizing Plural Suffix*

stem		stem+nominalizer	
üč	'three'	üč-gen-ler	'triangles'
altɨ	'five'	altɨ -gen-ler	'hexagons'
sek^{j}iz	'eight'	sek^{j}iz-gen-ler	'octagons'
čok	'many'	čok-gen-ler	'polygons'

(21) Turkish Pl Initiates Leftward Search in *čok-gen-ler*:

x_1	x_2	x_3	x_4	x_5	x_6	←	x_7	x_8
+voc -high +round +back	-voc +velar -cont -voi	-voc +velar -cont +voi	+voc -high -round -back	-voc coronal +nasal	-voc +sonorant -nasal +lateral		+voc -high -round	-voc +sonorant -nasal -lateral

(22) Turkish Pl Finds and Copies [−back] from Nominalizer:

$$\underset{\text{o}}{\overset{x_1}{\begin{bmatrix}\text{+voc}\\\text{-high}\\\text{+round}\\\text{+back}\end{bmatrix}}}\;\underset{\text{k}}{\overset{x_2}{\begin{bmatrix}\text{-voc}\\\text{+velar}\\\text{-cont}\\\text{-voi}\end{bmatrix}}}\;\underset{\text{g}}{\overset{x_3}{\begin{bmatrix}\text{-voc}\\\text{+velar}\\\text{-cont}\\\text{+voi}\end{bmatrix}}}\;\underset{\text{e}}{\overset{x_4}{\begin{bmatrix}\text{+voc}\\\text{-high}\\\text{-round}\\\textbf{-back}\end{bmatrix}}}\;\Lsh\;\underset{\text{n}}{\overset{x_5}{\begin{bmatrix}\text{-voc}\\\text{coronal}\\\text{+nasal}\end{bmatrix}}}\;\underset{\text{l}}{\overset{x_6}{\begin{bmatrix}\text{-voc}\\\text{+sonorant}\\\text{-nasal}\\\text{+lateral}\end{bmatrix}}}\;\underset{\text{e}}{\overset{x_7}{\begin{bmatrix}\text{+voc}\\\text{-high}\\\text{-round}\\\textbf{-back}\end{bmatrix}}}\;\underset{\text{r}}{\overset{x_8}{\begin{bmatrix}\text{-voc}\\\text{+sonorant}\\\text{-nasal}\\\text{-lateral}\end{bmatrix}}}$$

As shown above, a needy suffix is a greedy suffix: it will copy from the suffix immediately to its left, regardless of the fact that source from which it copies did not itself need to harmonize.

2.6. "Exceptionally" Non-Seeking Vowels in Turkish

The phenomenon of non-seeking elements in vowel harmony languages might at first blush look like a property of a whole morpheme: that the morpheme as a whole is lexically indexed to avoid harmony. However, it turns out that languages have fine-grained specifications of which vowels within a morpheme must undergo vowel harmony and which ones needn't, and this is best represented as a property of each vowel. Let's examine the phenomenon first with the disyllabic suffix used to express the progressive in Turkish:

(23) *Turkish "half-harmonizing" progressive suffix*:

$g^j el^j$-ijor-um	come-prog-1sg
koʃ-ujor-um	run-prog-1sg
$gül^j$-üjor-um	laugh-prog-1sg
bak-ɨjor-um	look-prog-1sg

The first vowel of the suffix *ijor* harmonizes, but the second one does not. Only the first vowel is required to Back and Round Harmonize, which can yield a surface form of either *u*,*ü*,*i*,ɨ. The second vowel is eternally fixed as *o*. Importantly, the next suffix to its right, whatever it is (e.g. 1sg in the case above), will harmonize with this *o*, the closest segment. Thus it is possible that only a single vowel within a morpheme needs to harmonize.

We have seen that Turkish suffixes may differ from each other in that some require Back-Harmony and Round-Harmony, some require Back-Harmony and Round-Harmony for only one of its vowels, and some requires no harmony at all. In Chapter 6, we return to the question of whether and how learners generalize from the harmonic requirements of one suffix to predict the patterning of classes of suffixes in the language as a whole, concluding that relevant insights may come from both first- and second-language acquisition. Importantly, once the lexically-specified harmony requirements for each affix are set, however this may occur through the learning process, the Search-and-Copy procedure operates as described here.

2.7. Non-Harmonizing Root Vowels in Turkish and Finnish and their Non-Fate in Language Games

This brings us to a discussion of whether non-initial vowels in roots need to harmonize. All of the cases of harmony considered thus far take place only across morpheme boundaries. However, in a computation over the lexicon, Harrison et al. (2004) determined that 73% of Turkish unsuffixed stems are harmonic; in other words, the number of roots that agree in backness and roundness among root vowels alone is higher than what would be expected from chance combinatorial shuffling without determinants of allowable co-occurrence. The question is ripe to ask, then, whether harmony is going on internal to a single morpheme.

This question is often difficult to pose, because the primary evidence for vowel harmony in the affixes we considered above comes from the fact that they alternate: for example, the accusative suffix has four different surface variants, determined by the vowel immediately to its left, and dependent only on what noun is being put in the accusative case. However, since Turkish has no prefixation that would provide a variety of leftward vowels from which the root itself could copy,[10] roots themselves do not show alternations – not necessarily because they are immune from harmony, but rather because the environments in which the effects of harmony could be detected are absent – save for one very revealing process.

Somewhat like the formation of English binomial compounds *dilly-dally, knick-knack, ping-pong, pitter-patter*, Turkish has a language game in which roots are doubled but the second copy has an initial back vowel. This is accomplished by "overwriting" the initial vowel in the second copy with [a], e.g. *moloz-maloz* 'debris'. What's crucially different about forming these compounds in Turkish, in contrast to English, is that the *second vowel* within the second copy undergoes harmony, as can be seen in the second and third examples below: (Harrison and Kaun, 2001):

(24) *Turkish Overwriting Reduplication Game: Replace initial vowel with [a]*:

Input	Second Word of Output	
moloz	maloz	'debris'
kibrit	kabrɨt	'match'
bütün	batɨn	'whole'

While the first vowel of the second copy is always [a], the non-initial vowels in the reduplicant stem completely harmonize for back and round with the newly-encountered [a] to their left. This is shown for *kibrit-kabrɨt* in the first and second copy, respectively:

(25) Second vowel of *kibrit* seeks Round & Back Harmony

$$\begin{array}{cccccc} x_1 & x_2 & x_3 & \leftarrow & x_4 & x_5 \\ \begin{bmatrix}\text{+voc}\\ \text{+high}\\ \text{-round}\\ \text{-back}\end{bmatrix} & \begin{bmatrix}\text{-voc}\\ \text{labial}\\ \text{-cont}\\ \text{+voi}\end{bmatrix} & \begin{bmatrix}\text{-voc}\\ \text{+sonorant}\\ \text{-nasal}\\ \text{-lateral}\end{bmatrix} & & \begin{bmatrix}\text{+voc}\\ \text{+high}\end{bmatrix} & \begin{bmatrix}\text{-voc}\\ \text{coronal}\\ \text{-cont}\\ \text{-voi}\end{bmatrix} \end{array}$$

(26) Second vowel of root *kibrit* Copies from Closest Leftward Vowel

$$\begin{array}{cccccc} x_1 & \Lsh & x_2 & x_3 & x_4 & x_5 \\ \begin{bmatrix}\text{+voc}\\ \text{+high}\\ \textbf{-round}\\ \textbf{-back}\end{bmatrix} & & \begin{bmatrix}\text{-voc}\\ \text{labial}\\ \text{-cont}\\ \text{+voi}\end{bmatrix} & \begin{bmatrix}\text{-voc}\\ \text{+sonorant}\\ \text{-nasal}\\ \text{-lateral}\end{bmatrix} & \begin{bmatrix}\text{+voc}\\ \text{+high}\\ \textbf{-round}\\ \textbf{-back}\end{bmatrix} & \begin{bmatrix}\text{-voc}\\ \text{coronal}\\ \text{-cont}\\ \text{-voi}\end{bmatrix} \\ \text{i} & & \text{b} & \text{r} & \text{i} & \text{t} \end{array}$$

(27) Second vowel of nonce *k**a**br_t* seeks Round & Back Harmony

$$\begin{array}{cccccc} x_1 & x_2 & x_3 & \leftarrow & x_4 & x_5 \\ \begin{bmatrix}\text{+voc}\\ \text{-high}\\ \text{-round}\\ \text{+back}\end{bmatrix} & \begin{bmatrix}\text{-voc}\\ \text{labial}\\ \text{-cont}\\ \text{+voi}\end{bmatrix} & \begin{bmatrix}\text{-voc}\\ \text{+sonorant}\\ \text{-nasal}\\ \text{-lateral}\end{bmatrix} & & \begin{bmatrix}\text{+voc}\\ \text{+high}\end{bmatrix} & \begin{bmatrix}\text{-voc}\\ \text{coronal}\\ \text{-cont}\\ \text{-voi}\end{bmatrix} \end{array}$$

(28) Second vowel of nonce *k**a**br_t* Copies from Closest Leftward Vowel

$$\begin{array}{cccccc} x_1 & \Lsh & x_2 & x_3 & x_4 & x_5 \\ \begin{bmatrix}\text{+voc}\\ \text{-high}\\ \textbf{-round}\\ \textbf{+back}\end{bmatrix} & & \begin{bmatrix}\text{-voc}\\ \text{labial}\\ \text{-cont}\\ \text{+voi}\end{bmatrix} & \begin{bmatrix}\text{-voc}\\ \text{+sonorant}\\ \text{-nasal}\\ \text{-lateral}\end{bmatrix} & \begin{bmatrix}\text{+voc}\\ \text{+high}\\ \textbf{-round}\\ \textbf{+back}\end{bmatrix} & \begin{bmatrix}\text{-voc}\\ \text{coronal}\\ \text{-cont}\\ \text{-voi}\end{bmatrix} \\ \text{a} & & \text{b} & \text{r} & \text{ɨ} & \text{t} \end{array}$$

The second root vowel in the compound word *k_br_t* shows the alternation pattern we were looking for: it is *i* in *kibrit* and ɨ in *kabrɨt*. Harmony is in fact found when one rearranges a string of segments into a novel configuration by overwriting the first vowel of a root.

One response to this characterization of re-harmonization of roots might be that *kabrɨt* has simply created a new sequence of segments, but that this doesn't tell us about root-internal harmony per se – one might think that in real words, there is faithful reproduction of the stored root and no harmony, but whenever you shuffle or replace, you get a wholly new string, where no faithfulness to that lemma must apply.

A striking confirmation that root-internal harmony really is a process in which *each vowel* is specified for whether it needs to harmonize or not comes from what happens in this game with disharmonic roots: they don't re-harmonize (Harrison and Kaun, 2001). For example, consider the disharmonic root *butik*, in which the first and second root vowel are not identical in backness or roundness. If root-internal harmony as an active search-and-copy process really is the regular state of affairs, this word, like the second vowel of the progressive suffix *-iyor* encountered above, simply is marked as exceptionally non-undergoing harmony. As a result, even when the first root vowel is overwritten in binominal compound formation, the second root vowel stays non-needy:

(29) *Turkish Overwriting Reduplication Game: No Harmony with Disharmonic Roots*:

Input	Second Word of Output	
butik	batik	'boutique'
bordür	bardür	'edge ornamentation'
kuvvet	kavvet	'strength'

The formal difference in behavior between the second root vowel in *kibrit* and *butik* is the result of whether there is a requirement or not that this vowel must undergo harmony:[11]

(30) Harmony Requirements of Root *kibrit*:
2nd vowel must: Back-Harmonize and Round-Harmonize

(31) Harmony Requirements of Root *butik*:
2nd vowel must: (empty)

The failure of disharmonic roots such as *butik* to reharmonize (say, as *butik*-

batɨk) is evidence that the game of binomial compound formation preserves the lexical requirement of undergoing harmony or not for the second root vowel. The results allow us to conclude that even in non-game utterances of *kibrit*, this second root vowel really is undergoing a process of search and copy with the vowel to its left. By examining the root in isolation, we could not determine whether harmony was occurring or not. But in the special circumstance of overwriting the initial root vowel and observing harmonic alternation, we derive evidence that 73% of Turkish roots do show harmony between the non-initial vowels and the initial vowels within a morpheme as the result of an active harmony process of the same type that occurs across morpheme boundaries. The remaining 27% of disharmonic roots have a representation like that of the progressive suffix *-iyor*, in which a particular vowel of a morpheme is specified as not needing to undergo harmony.

This pattern, in which disharmonic roots fail to "reharmonize", is not specific to this game nor just to Turkish. At this point I will briefly introduce Finnish harmony, although a full treatment awaits the next chapter. Finnish also has [± back] harmony in which, for example, the [+back] *o* cannot co-occur with the [−back] *ö*. There is a game called "kontti kieli" (Campbell, 1980, 1986) which involves overwriting the first CV of a stem with *ko-*. This causes reharmonization of [−back] vowels, like ö, ü, ä:

(32) *Harmony of Root-Vowels in Finnish Kontti Kieli Game:*

Input	First Word of Output	
nähnüt	kohnut	'seen'
mitä	kota	'what'
nälässänsäkö	kolassansako	'in his hunger'

In disharmonic Finnish words, root-internal re-harmonization does not occur, as shown in (33). This is because by hypothesis any disharmonic or non-harmonic vowel is one that does not need a value. The absence of a requirement to copy a value will hold even when the vowel is put in a different context:

(33) *Lack of harmony in Konti Kieli for "non-needy" vowels:*

Input	First Word of Output	
jonglööri	ko+nglööri	'juggler'

No matter whether the game that provides the evidence for harmony is an overwriting game of binomial compound formation or a game of replacement of the initial consonant+vowel sequence, all root vowels to the right of the newly-introduced vowel will undergo harmony as usual if they are in a harmonic root, and will not undergo harmony if they are in a disharmonic root. The representation of harmonic and disharmonic roots can therefore be captured using the same leftward search-and-copy procedure for needy vowels that is independently found in suffixes. We conclude that vowel harmony occurs in Turkish and Finnish roots using this procedure. The interested reader is also directed to the appendix to this chapter, where it is argued that non-final sibilant consonants in a word – including those within a root – must undergo harmony in Chumash.

2.8. Bidirectional Harmony in Woleaian

Thus far, we have seen one important empirical argument for representing harmony as a recipient-initiated search, rather than a donor-initiated process: the empirical generalization (19), repeated below.

(34) *A morpheme might not need to search for harmony, but it still must be a donor*:
While a suffix might be "exceptional" in not-undergoing harmony, no suffix is listed as "exceptional in not-triggering harmony".

This generalization is natural to derive in the present model: the search-and-copy procedure is either initiated or not by a particular vowel. There is simply no way to designate a particular closest leftward vowel to the recipient as exceptionally unwilling to be copied from: the search algorithm is greedy and will opportunistically copy from the first vowel it finds.

As further support for the model of harmony as a recipient-initiated process, we will examine a second case in which the opposite view, that harmony is donor-initiated, would be needlessly cumbersome in comparison. For an alternative class of models that might view harmony as "an imperative to spread", initiated by the donor, in cases where there is *more than one donor*, it becomes immediately clear that harmony should instead be modeled from the perspective of the recipient.

In the Turkish cases discussed above, the direction of search was specified as leftwards. In the case of a suffix, of course, it is somewhat natural to expect search to proceed left, as in the limiting case there will be no material to the right of the suffix. Similarly, one might expect that when prefixes undergo harmony, they will search for a harmonic-value to the right. This prediction is confirmed by the pattern of [$\pm$ ATR] harmony in Akan prefixes, for example, the 3rd person future tense which varies between [+ATR] o and and [−ATR] ɔ in harmony with the leftmost stem vowel (O'Keefe, 2003):

(35) *Akan [± ATR] Prefixal Harmony Contrast*
  a. ɔ-bɛkʊ '3rd.sg-fight'
  b. o-betu '3rd.sg-dig'

(36) *Akan 3rd person singular morpheme must*:
  ATR-Harmonize: δ = right, F = [$\pm$ ATR]

When a past tense suffix is added, this suffix will search to its left for harmony:

(37) *Akan [± ATR] Prefixal Harmony Contrast*
  a. ɔ-bɛkʊ-ɪ '3rd.sg-fight-past'
  b. o-betu-i '3rd.sg-dig-past'

(38) *Akan past tense morpheme must*:
  ATR-Harmonize: δ = left, F = [$\pm$ ATR]

Thus, when a sole suffix is searching, it will look to the left, and when a sole prefix is searching, it will look to the right. However, as will be discussed in Chapter 4, languages may also contain root vowels that copy harmonic features from suffixes to the right. Direction thus must be a free parameter of search. As mentioned previously, there is the possibility that the search direction could be simultaneously left and right. Woleaian, a language spoken throughout the Woleai atoll (Micronesia) has a process of vowel assimilation (Sohn, 1971; Howard, 1972) applying to a thematic formative affix that occupies a position between the verbal stem and the agreement suffix.

Woleaian has three levels of vowel heights, with both long and short versions of each of the vowels.

(39) *Woleaian Vowel Inventory*:

i	ü	u	[+high, -low]
e	ö	o	[-high, -low]
	a	ɔ	[-high, +low]

The features distinguishing Woleaian vowels are [± back], [± high], [± low], and [± round].[12] The harmonizing suffix in question is always [−high,−round], but it is dependent on the surrounding vowels for its feature of [± low]. In Woleaian, this search is not set to be strictly leftwards or strictly rightwards. As a result, in its search for [±low] valuation, detemined by the closest valuation source, it turns out that *both* flanking vowels are equally close. As a consequence, only when both vowels are [−low] will /a/ be raised to /e/; otherwise, a default value of [+low] will have to be inserted.[13]

(40) *Woleaian theme-vowel must:*
Low-Harmonize: δ = L&R, F = [± low]

The results of this harmony are shown in the following set of examples. A number of additional phonological processes affect the form of nouns in Woleaian, such as deletion of final vowels and changes in palatal consonants; for that reason, both the underlying representation (UR) and the surface representation (SR) are shown here.

(41) Woleaian bidirectional [±] low harmony

UR	SR	
ülüm	üːl	drinking object (indp't form)
ülüm-a-ji	ülümej	1st.sg
ülüm-a-mu	ülümem^w	2nd.sg
ülüm-a-la	ülümal	3rd.sg
ülüm-a-ca	ülümaʃ	1st-incl.pl
ülüm-a-mii	ülümemi	2nd.pl
ülüm-a-jire	ülümeːr	3rd.pl

Importantly, the determinants for this process may be *deleted on the surface*, as may be seen in, for example, the 1st.sg form in (41). The "two-sided" nature of the determining environment cannot be attributed to surface coarticulation. Moreover, as shown in the following examples, a non-low vowel on the right is not sufficient to trigger raising. (There is an orthogonal process of low-vowel dissimilation, by which an *a* before another *a* changes to *e*, in the stem below).

(42) *Woleaian failure of harmony when leftward and rightward vowels have different [±low]*:

UR	SR	
mata-ji	metaj	eye, 1.sg
mata-mu	metam^w	eye, 2.sg
mata-mii	metami	eye, 2.pl
mata-la	metal	eye, 3.sg
mata-ca	metaʃ	eye, 1pl.incl

In a donor-initiated theory, there would need to be two rules for Woleaian low harmony in (41): one spreading [-low] Left, and one spreading [-low] Right, and the additional condition that *each could only apply if the other did*. A recipient-initated *Search-and-Copy-From-Closest*, on the other hand, provides an effective way to model *multiple-source valuation* of this sort without any additional stipulations about resolving the application of two-sided spreading rules.

The illustration of Search-and-Copy for the relevant portions of ülüm-a-mii (resulting in ülümemi) is shown below:[14]

(43) Woleaian Theme-Vowel Begins Low-Harmonize in ülümemi:

$$\begin{array}{ccccccc} x_1 & x_2 & \leftarrow & x_3 & \rightarrow & x_4 & x_5 \\ \begin{bmatrix}\text{+voc}\\ \text{+rd}\\ \text{+high}\\ \text{-low}\\ \text{-back}\end{bmatrix} & \begin{bmatrix}\text{-voc}\\ \text{labial}\\ \text{-cont}\\ \text{+nas}\end{bmatrix} & & \begin{bmatrix}\text{+voc}\\ \text{-rd}\\ \text{-high}\\ \\ \end{bmatrix} & & \begin{bmatrix}\text{-voc}\\ \text{labial}\\ \text{-cont}\\ \text{+nas}\end{bmatrix} & \begin{bmatrix}\text{+voc}\\ \text{-rd}\\ \text{+high}\\ \text{-low}\\ \text{-back}\end{bmatrix} \end{array}$$

(44) Woleaian Theme-Vowel Finds [−low] on x_1 and on x_5:

$$\begin{array}{ccccccc} x_1 & \leftarrow & x_2 & x_3 & x_4 & \rightarrow & x_5 \\ \begin{bmatrix}\text{+voc}\\ \text{+rd}\\ \text{+high}\\ \textbf{-low}\\ \text{-back}\end{bmatrix} & & \begin{bmatrix}\text{-voc}\\ \text{labial}\\ \text{-cont}\\ \text{+nas}\end{bmatrix} & \begin{bmatrix}\text{+voc}\\ \text{-rd}\\ \text{-high}\\ \\ \end{bmatrix} & \begin{bmatrix}\text{-voc}\\ \text{labial}\\ \text{-cont}\\ \text{+nas}\end{bmatrix} & & \begin{bmatrix}\text{+voc}\\ \text{-rd}\\ \text{+high}\\ \textbf{-low}\\ \text{-back}\end{bmatrix} \end{array}$$

(45) Woleaian Theme-Vowel Copies [−low] from x_1 and from x_5:

$$\begin{array}{ccccccc}
x_1 & \Lsh & x_2 & x_3 & x_4 & \Rsh & x_5 \\
\begin{bmatrix}\text{+voc}\\ \text{+rd}\\ \text{+high}\\ \textbf{-low}\\ \text{-back}\end{bmatrix} & & \begin{bmatrix}\text{-voc}\\ \text{labial}\\ \text{-cont}\\ \text{+nas}\end{bmatrix} & \begin{bmatrix}\text{+voc}\\ \text{+rd}\\ \text{-high}\\ \textbf{-low}\\ \\ \end{bmatrix} & \begin{bmatrix}\text{-voc}\\ \text{labial}\\ \text{-cont}\\ \text{+nas}\end{bmatrix} & & \begin{bmatrix}\text{+voc}\\ \text{+rd}\\ \text{+high}\\ \textbf{-low}\\ \text{-back}\end{bmatrix} \\
\text{ü} & & \text{m} & \text{e} & \text{m} & & \text{i}
\end{array}$$

The central phenomenon illustrated by Woleaian here is bidirectional harmony. As with the Turkish and Akan harmony processes discussed above, this type of harmony applies to this morpheme, and thus the set of morphemes that participate in harmony are explicitly listed as lexically-specific aspects of their representation.

(46) *Woleaian thematic formative morpheme must:*
Low-Harmonize: δ = L&R, F = [$\pm$ low]

One aspect of the formalization in (46) is that it effectively states that copying of a [$-$low] feature depends on a "tie" in which the closest vowel in both directions bears [$-$low].[15] Given (46), when both the immediately-closest preceding and the immediately-closest preceded vowel to the thematic formative are [αlow] (where α ranges over either $+$ or $-$, where both sides have the same value for α), the thematic formative will successfully copy [αlow]. However, if one side is [$+$low] and the other is [$-$low], there is no way to *copy the value from the closest donor(s)*. The uniqueness presupposition of "*the* value" in this statement fails, since it's impossible to copy $+$ from one side and $-$ from the other side, so nothing can be copied at all. "The Search Terminates in Failure", so to speak.

(47) Woleaian Theme-Vowel Begins Low-Harmonize:

$$\begin{array}{ccccccc}
x_1 & x_2 & \leftarrow & x_3 & \rightarrow & x_4 & x_5 \\
\begin{bmatrix}\text{+voc}\\ \text{-rd}\\ \text{-high}\\ \text{+low}\end{bmatrix} & \begin{bmatrix}\text{-voc}\\ \text{coronal}\\ \text{-cont}\\ \text{-nas}\\ \text{-voi}\end{bmatrix} & & \begin{bmatrix}\text{+voc}\\ \text{-rd}\\ \text{-high}\\ \\ \end{bmatrix} & & \begin{bmatrix}\text{-voc}\\ \text{coronal}\\ \text{+del.rel}\\ \text{-nas}\\ \text{+voi}\end{bmatrix} & \begin{bmatrix}\text{+voc}\\ \text{-rd}\\ \text{+high}\\ \text{-low}\\ \text{-back}\end{bmatrix}
\end{array}$$

(48) Woleaian Theme-Vowel Finds [$+$low] on x_1 and [$-$low] on x_5:

$$\begin{array}{ccccccc} x_1 & \leftarrow & x_2 & x_3 & x_4 & \rightarrow & x_5 \\ \begin{bmatrix} \text{+voc} \\ \text{-rd} \\ \text{-high} \\ \textbf{+low} \end{bmatrix} & & \begin{bmatrix} \text{-voc} \\ \text{coronal} \\ \text{-cont} \\ \text{-nas} \\ \text{-voi} \end{bmatrix} & \begin{bmatrix} \text{+voc} \\ \text{-rd} \\ \text{-high} \end{bmatrix} & \begin{bmatrix} \text{-voc} \\ \text{coronal} \\ \text{+del.rel} \\ \text{-nas} \\ \text{+voi} \end{bmatrix} & & \begin{bmatrix} \text{+voc} \\ \text{-rd} \\ \text{+high} \\ \textbf{-low} \\ \text{-back} \end{bmatrix} \end{array}$$

(49) Woleaian Theme-Vowel Fails to Copy a Value

$$\begin{array}{ccccccc} x_1 & \leftarrow & x_2 & x_3 & x_4 & \rightarrow & x_5 \\ \begin{bmatrix} \text{+voc} \\ \text{-rd} \\ \text{-high} \\ \text{+low} \end{bmatrix} & & \begin{bmatrix} \text{-voc} \\ \text{coronal} \\ \text{-cont} \\ \text{-nas} \\ \text{-voi} \end{bmatrix} & \begin{bmatrix} \text{+voc} \\ \text{-rd} \\ \text{-high} \end{bmatrix} & \begin{bmatrix} \text{-voc} \\ \text{coronal} \\ \text{+del.rel} \\ \text{-nas} \\ \text{+voi} \end{bmatrix} & & \begin{bmatrix} \text{+voc} \\ \text{-rd} \\ \text{+high} \\ \text{-low} \\ \text{-back} \end{bmatrix} \\ \text{a} & & \text{t} & & \check{\text{j}} & & \text{i} \end{array}$$

No feature can be copied in (49), because there are two different values. This is actually identical to what happens in verbal agreement when the subject noun phrase is a conjunction of two different values. In the Italian sentence with a conjoined subject *Gli uccelli e le farfalle sono andati* "The birds and the butterflies have gone", the participle cannot simultaneously agree with both masculine plural *uccelli* and feminine plural *farfalle*. From a hierarchical perspective, each of the parts of this conjunction are equally close to the participle:

(50) *Mixed-Gender Conjunct in Agreement Configuration:*

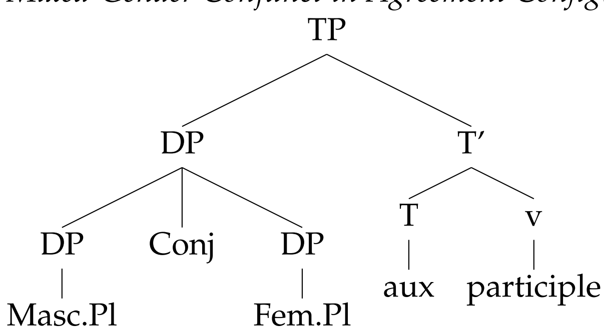

The case of mixed-gender conjunctions and verb agreement in Italian is formally identical to that of Woleaian bidirectional harmony: when there are two equally-close sources of features but they disagree in their feature-value, the otherwise straightfoward procedure of copying fails. All cases of search-based dependencies that seek the closest feature-value to copy must include some "Plan B" when the search fails. Search can fail either because it finds no value to copy or no unique value to copy. We return to syntactic cases in which there is no nominative argument in Chapter 4.

Lest this sound like an esoteric case of feature-purgatory in phonological limbo, consider the fact that verbal agreement, itself modeled as a search-and-copy procedure, may end up with a failed-search and no value copied as a result of the process. Yet this failed search does not result in the option to pronounce an Italian participle as *andat*, without any gender or number encoding whatsoever. There are "default" settings: last-resort feature-values that can be cheaply inserted without ruffling too many feathers. What these values will turn out to be for a given feature in a given language depends in a large part on what is called *markedness theory*, essentially a grammatical statement of which of the values, say $+$ or $-$, for a given phonological or morphological feature, is considered to be the preferred or otherwise-assumed one (we return to this issue in Chapter 3).

A Plan B is needed because a speaker of Woleaian cannot pronounce a vowel that is [$-$round, $-$high] but has no value for [$\pm$ low]: at the interface between phonological representations and phonetic implementation, such a representation is illegible and cannot be properly interpreted. All vowel features that are active in a given language must be fully specified on all vowels as an interface requirement. Similary a Plan B is needed because of speaker of Italian cannot pronounce a participle with no value for gender: at the interface between abstract syntactic features and morphological realization, such a representation is illegible. The Plan B in both cases is insertion of the default value in these contexts: [$+$low] for Woleaian non-round non-high vowels, and [$-$fem] for Italian participles.

(51) *Woleaian thematic formative morpheme must:*
  a. Low-Harmonize: δ = L&R, F = [$\pm$ low]
  b. Failure results in Default Insertion of [$+$low]

While the clause in (51-b) above has not been previously encountered in our model, it is necessary when harmony fails: the "plain vanilla" value of [$\pm$low] in Woleaian, namely [$+$low], is what is chosen for the thematic formative.[16] We will return to a discussion of default value-insertion when harmonic search fails in our discussion of Finnish and Uyghur in the next chapter.

We've covered a lot of ground in these examples of vowel harmony. The important lesson for the theory of locality from Woleaian comes when we consider how naturally the recipient-initated model of search explains what happens in cases of bidirectionally-dependent harmony. Any theory of vowel harmony must deal with the fact that the height of the Woleaian thematic formative is determined by leftward and rightward dependencies. When these two search-based dependencies cannot both resolve to the same value, the result is a default value.

2.9. Directionality in Kalenjin "Dominant-Recessive" Harmony

Directionality is often taken as a property of vowel harmony that is predictable based on the morphological structure of the language: in a wide set of patterns, suffix copy from the left and prefixes copy from the right. However, as we have seen in the discussion above, this is not always the end of the story: vowels within the root also initiate harmonic search under certain conditions, with their direction not predictable by general morphological structure, and in Woleaian, the thematic vowel searches bidirectionally. Finally, as will be discussed in Section 4.7.1, languages such as Jingulu show root vowels copying from suffixes but not vice versa. These latter kinds of examples establish that directionality of search must be a part of roots, and that affixes may search in both roots and suffixes.

One of the apparent challenges to the notion of directionality as an independent parameter of vowel harmony comes from the existence of patterns described as "dominant-recessive", which are often characterized as entirely non-directional, in the sense that a single feature value, say [+ATR] "anywhere within the word" can cause all other vowels to change, regardless of direction or morphological status.

However, upon closer scrutiny, the specification of δ as a free parameter of search specified on individual harmonizing morphemes is perhaps best observed in the interaction of directionality and needy morphemes in the [± ATR] harmony of Kalenjin, a Nilotic language of Kenya claimed to be dominant/recessive. In the discussion below, I present an analysis of Kalenjin's vowel harmony system that demonstrates how dominant-recessive systems may be reanalyzed in terms of directionality specifications, and with positive consequences. The discussion is based on the description of Kalenjin in Hall et al. (1974), which forms the basis for many subsequent analyses in terms of a dominant-recessive pattern. Kalenjin has a symmetric [± ATR] contrast at three heights:

(52) *Kalenjin Vowel Inventory*:

[−back,−round]	[−back,−round]	[+back,+round]	
i		u	[+high,−low,+ATR]
ɪ		ʊ	[+high,−low,−ATR]
e		o	[−high,−low,+ATR]
ɛ		ɔ	[−high,−low,−ATR]
	ɐ		[−high,+low,+ATR]
	a		[−high,+low,−ATR]

One of the properties of Kalenjin vowel harmony that has attracted special attention is the descriptive fact that harmony may affect not only prefixes and suffixes, but also roots. As the following examples show, prefixes and roots copy [± ATR] from their right, while suffixes copy [± ATR] from their left. However, some prefixes, such as *ma-* in (53-d), do not copy, and some suffixes, such as kɛ in (53-e) do not copy.

(53) Kalenjin [± ATR] harmony affects prefixes, roots, and suffixes (Hall et al. (1974, p.247), Lodge (1995, p.32)):

a. ki- ɐ- keːr- in
distant.past- 1sg- see- 2sg.obj
'I saw you'

b. kɪ- a- par- ɪn
distant.past- 1sg- kill- 2sg.obj
'I killed you'

c. ki- ɐ- ker- e
distant.past- 1sg- shut- noncompletive
'I was shutting it'

d. ka- ma- ɐ- keːr- ɐk
recent.past- neg- 1sg- see- 2pl.
'I didn't see you (pl.)

e. ki- ɐ- un- kɛː
distant.past- 1sg- wash- reflexive
'I was washing myself'

f. keːr- un
see- directional
'see it from here'

g. kʊt- ʊn
blow- directional
'blow it here'

As mentioned above, a "dominant-recessive" system is commonly ascribed to Kalenjin, which in intuitive terms means that the value [+ATR] "anywhere in the word" should cause all morphemes in the word to become [+ATR]. Dominant-recessive systems are claimed to be non-directional, featurally asymmetric, and non-local, as summarized in the statement "The presence of a dominant vowel in a word changes the vowels of the non-dominant series" (Aoki, 1968, p.143), and authors such as Bakovic (2000) explicit adopt a dominant-recessive analysis for Kalenjin, whereby [+ATR] values are seen as asymmetrically controlling vowel harmony regardless of directionality.

However, as many of the examples above reveal, exceptionally non-undergoing morphemes are not hard to come by (e.g. the negative, the reflexive), and seem to be as much the rule as the exception. Moreover, prefixes may copy either [+ATR] or [−ATR], depending on what is immediately to their right, regardless of whether an [+ATR] morpheme is found in the root. Summing up, no prefixes cause roots to harmonize, and when prefixes search for a value to copy, they do so based on locality, instead of "dominance".

In fact, once the non-harmonizing morphemes, such as the negative prefix *ma-*, are taken into account in the locality and directionality of vowel harmony, the pattern of Kalenjin harmony reduces to the same directionally-specified Search-and-Copy algorithm developed throughout this chapter. As Lodge (1995) points out, the most straightforward characterization of Kalenjin harmony is not as a "dominant/recessive" system, but rather one in which there is a three-way contrast: morphemes that are inherently [+ATR], morphemes that are inherently [−ATR], and morphemes that need to find a value for [± ATR] through directionally-specified harmony. The specifications for the morphemes in the examples above are thus as follows:

(54) Harmonic specifications of Kalenjin morphemes:

a. *Kalenjin distant past morpheme must*:
ATR-Harmonize: δ = right, F = [± ATR]

b. *Kalenjin recent past morpheme must*:

ATR-Harmonize: δ = right, F = [$\pm$ ATR]

c. *Kalenjin subject agreement morpheme must*:
ATR-Harmonize: δ =right, F = [$\pm$ ATR]

d. *Kalenjin object agreement morpheme must*:
ATR-Harmonize: δ = left, F = [$\pm$ ATR]

e. *Kalenjin directional morpheme must*:
ATR-Harmonize: δ = left, F = [$\pm$ ATR]

f. *Kalenjin negative morpheme must*:
(none; inherently [$-$ATR])

g. *Kalenjin noncompletive morpheme must*:
(none; inherently [+ATR])

h. *Kalenjin reflexive morpheme must*:
(none; inherently [$-$ATR])

i. *Kalenjin roots* BLOW, SHUT, KILL *must*:
ATR-Harmonize: δ = right, F = [$\pm$ ATR]

j. *Kalenjin roots* SEE, WASH *must*:
(none; inherently [+ATR])

In other words, Kalenjin is *not* a "fully bidirectional" system. Roots are specified to copy from the right, and suffixes are specified to copy from the left. Roots do not copy from prefixes, and prefixes never copy from their left (e.g. -ɐ- in ka-ma-ɐ- keːr- ɐk). If no value is found (e.g. in a sequence of a set of entirely unspecified morphemes, as in kʊt-ʊn 'blow-directional', then a default value will be supplied as a last resort.

Harmonization in Kalenjin follows a cyclic order of morphological structure-building whereby root+suffix combinations are concatenated and undergo phonology prior to prefixes. We will consider three types of configurations below: a needy root and a specified suffix, a specified root and a needy suffix, and a needy root and a needy suffix.

The first case is that of a needy root and a specified suffix. The derivation for ki- ɐ- ker- e 'distant.past-1sg-shut-noncompletive' is shown below, with the root copying from the suffix, followed by the prefixes copying from the root:

(55) *Kalenjin Root Begins ATR-Harmonize in ker-e*:

$$
\begin{array}{ccccc}
x_4 & x_5 & \longrightarrow & x_6 & x_7 \\
\begin{bmatrix}\text{-voc}\\ \text{dorsal}\\ \text{-cont}\\ \text{-voice}\end{bmatrix} &
\begin{bmatrix}\text{+voc}\\ \text{-high}\\ \text{-low}\\ \text{-round}\\ \text{+back}\\ \\ \end{bmatrix} & &
\begin{bmatrix}\text{-voc}\\ \text{coronal}\\ \text{+son}\\ \text{-nasal}\\ \text{-lateral}\end{bmatrix} &
\begin{bmatrix}\text{+voc}\\ \text{-high}\\ \text{-low}\\ \text{-round}\\ \text{+back}\\ \text{+ATR}\end{bmatrix}
\end{array}
$$

(56) *Kalenjin Root Finds and Copies [+ATR] in ker-e*:

$$
\begin{array}{ccccc}
x_4 & x_5 & x_6 & \Rsh & x_7 \\
\begin{bmatrix}\text{-voc}\\ \text{dorsal}\\ \text{-cont}\\ \text{-voice}\end{bmatrix} &
\begin{bmatrix}\text{+voc}\\ \text{-high}\\ \text{-low}\\ \text{-round}\\ \text{+back}\\ \textbf{+ATR}\end{bmatrix} &
\begin{bmatrix}\text{-voc}\\ \text{coronal}\\ \text{+son}\\ \text{-nasal}\\ \text{-lateral}\end{bmatrix} & &
\begin{bmatrix}\text{+voc}\\ \text{-high}\\ \text{-low}\\ \text{-round}\\ \text{+back}\\ \textbf{+ATR}\end{bmatrix}
\end{array}
$$

(57) *Kalenjin Prefix Begins [± ATR] Harmonize in ɐ-ker-e*:

$$
\begin{array}{cccccc}
x_3 & \longrightarrow & x_4 & x_5 & x_6 & x_7 \\
\begin{bmatrix}\text{+voc}\\ \text{-high}\\ \text{+low}\\ \text{-round}\\ \text{+back}\\ \\ \end{bmatrix} & &
\begin{bmatrix}\text{-voc}\\ \text{dorsal}\\ \text{-cont}\\ \text{-voice}\end{bmatrix} &
\begin{bmatrix}\text{+voc}\\ \text{-high}\\ \text{-low}\\ \text{-round}\\ \text{+back}\\ \text{+ATR}\end{bmatrix} &
\begin{bmatrix}\text{-voc}\\ \text{coronal}\\ \text{+son}\\ \text{-nasal}\\ \text{-lateral}\end{bmatrix} &
\begin{bmatrix}\text{+voc}\\ \text{-high}\\ \text{-low}\\ \text{-round}\\ \text{+back}\\ \text{+ATR}\end{bmatrix}
\end{array}
$$

(58) *Kalenjin Prefix Finds and Copies [+ATR] in ɐ-ker-e*:

$$
\begin{array}{cccccc}
x_3 & x_4 & \Rsh & x_5 & x_6 & x_7 \\
\begin{bmatrix}\text{+voc}\\ \text{-high}\\ \text{+low}\\ \text{-round}\\ \text{+back}\\ \textbf{+ATR}\end{bmatrix} &
\begin{bmatrix}\text{-voc}\\ \text{dorsal}\\ \text{-cont}\\ \text{-voice}\end{bmatrix} & &
\begin{bmatrix}\text{+voc}\\ \text{-high}\\ \text{-low}\\ \text{-round}\\ \text{+back}\\ \textbf{+ATR}\end{bmatrix} &
\begin{bmatrix}\text{-voc}\\ \text{coronal}\\ \text{+son}\\ \text{-nasal}\\ \text{-lateral}\end{bmatrix} &
\begin{bmatrix}\text{+voc}\\ \text{-high}\\ \text{-low}\\ \text{-round}\\ \text{+back}\\ \text{+ATR}\end{bmatrix} \\
\text{ɐ} & \text{k} & & \text{e} & \text{r} & \text{e}
\end{array}
$$

The next case is that of a needy suffix is concatenated with a root. The suffix it will copy from its left, as in ɐ-keːr-ɐk. On a subsequent cycle, prefixes will copy from their right.

(59) *Kalenjin Suffix Begins ATR-Harmonize in* ka-ma-ɐ-keːr-ɐk:

$$\begin{array}{cccccc} x_6 & x_7 & x_8 & \leftarrow & x_9 & x_{10} \\ \begin{bmatrix}\text{-voc}\\ \text{dorsal}\\ \text{-cont}\\ \text{-voice}\end{bmatrix} & \begin{bmatrix}\text{+voc}\\ \text{-high}\\ \text{-low}\\ \text{-round}\\ \text{+back}\\ \text{+ATR}\end{bmatrix} & \begin{bmatrix}\text{-voc}\\ \text{coronal}\\ \text{+son}\\ \text{-nasal}\\ \text{-lateral}\end{bmatrix} & & \begin{bmatrix}\text{+voc}\\ \text{-high}\\ \text{+low}\\ \text{-round}\\ \text{+back}\\ \\ \end{bmatrix} & \begin{bmatrix}\text{-voc}\\ \text{dorsal}\\ \text{-cont}\\ \text{-voice}\end{bmatrix} \end{array}$$

(60) *Kalenjin Suffix Finds and Copies [+ATR] in* keːr-ɐk:

$$\begin{array}{cccccc} x_6 & x_7 & \Lsh & x_8 & x_9 & x_{10} \\ \begin{bmatrix}\text{-voc}\\ \text{dorsal}\\ \text{-cont}\\ \text{-voice}\end{bmatrix} & \begin{bmatrix}\text{+voc}\\ \text{-high}\\ \text{-low}\\ \text{-round}\\ \text{+back}\\ \textbf{+ATR}\end{bmatrix} & & \begin{bmatrix}\text{-voc}\\ \text{coronal}\\ \text{+son}\\ \text{-nasal}\\ \text{-lateral}\end{bmatrix} & \begin{bmatrix}\text{+voc}\\ \text{-high}\\ \text{+low}\\ \text{-round}\\ \text{+back}\\ \textbf{+ATR}\end{bmatrix} & \begin{bmatrix}\text{-voc}\\ \text{dorsal}\\ \text{-cont}\\ \text{-voice}\end{bmatrix} \\ \text{k} & \text{e} & & \text{r} & \text{ɐ} & \text{k} \end{array}$$

(61) *Kalenjin Prefix Searches for [±ATR] in* ɐ-keːr-ɐk:

$$\begin{array}{ccccc} x_5 & \rightarrow & x_6 & x_7 & x_8 \\ \begin{bmatrix}\text{+voc}\\ \text{-high}\\ \text{+low}\\ \text{-round}\\ \text{+back}\\ \\ \end{bmatrix} & & \begin{bmatrix}\text{-voc}\\ \text{dorsal}\\ \text{-cont}\\ \text{-voice}\end{bmatrix} & \begin{bmatrix}\text{+voc}\\ \text{-high}\\ \text{-low}\\ \text{-round}\\ \text{+back}\\ \text{+ATR}\end{bmatrix} & \begin{bmatrix}\text{-voc}\\ \text{coronal}\\ \text{+son}\\ \text{-nasal}\\ \text{-lateral}\end{bmatrix} \end{array}$$

(62) *Kalenjin Prefix Finds and Copies [+ATR] in* ɐ-keːr-ɐk:

$$\begin{array}{ccccc} x_5 & x_6 & \Rsh & x_7 & x_8 \\ \begin{bmatrix}\text{+voc}\\ \text{-high}\\ \text{+low}\\ \text{-round}\\ \text{+back}\\ \textbf{+ATR}\end{bmatrix} & \begin{bmatrix}\text{-voc}\\ \text{dorsal}\\ \text{-cont}\\ \text{-voice}\end{bmatrix} & & \begin{bmatrix}\text{+voc}\\ \text{-high}\\ \text{-low}\\ \text{-round}\\ \text{+back}\\ \textbf{+ATR}\end{bmatrix} & \begin{bmatrix}\text{-voc}\\ \text{coronal}\\ \text{+son}\\ \text{-nasal}\\ \text{-lateral}\end{bmatrix} \\ \text{ɐ} & \text{k} & & \text{e:} & \text{r} \end{array}$$

In a sequence of more than one prefix, each one copies from its right. If a needy prefix precedes a [−ATR] specified prefix, it will copy from it, even if a [+ATR] affix is further downstream, showing that the "dominant/recessive" system is constrained by closest-source locality, just like any other.

(63) *Kalenjin Prefix Searches for [±ATR] in* ka-ma-ɐ-:

$$\overset{x_1}{\begin{bmatrix}\text{-voc}\\ \text{dorsal}\\ \text{-cont}\\ \text{-voice}\end{bmatrix}} \quad \overset{x_2}{\begin{bmatrix}\text{+voc}\\ \text{-high}\\ \text{+low}\\ \text{-round}\\ \text{+back}\\ \\ \end{bmatrix}} \quad \longrightarrow \quad \overset{x_3}{\begin{bmatrix}\text{-voc}\\ \text{labial}\\ \text{+son}\\ \text{+nasal}\end{bmatrix}} \quad \overset{x_4}{\begin{bmatrix}\text{+voc}\\ \text{-high}\\ \text{+low}\\ \text{-round}\\ \text{+back}\\ \text{-ATR}\end{bmatrix}} \quad \overset{x_5}{\begin{bmatrix}\text{+voc}\\ \text{-high}\\ \text{-low}\\ \text{-round}\\ \text{+back}\\ \text{+ATR}\end{bmatrix}}$$

(64) *Kalenjin Prefix Finds and Copies [−ATR] in* ka-ma-ɐ-:

$$\underset{\text{k}}{\overset{x_1}{\begin{bmatrix}\text{-voc}\\ \text{dorsal}\\ \text{-cont}\\ \text{-voice}\end{bmatrix}}} \quad \underset{\text{a}}{\overset{x_2}{\begin{bmatrix}\text{+voc}\\ \text{-high}\\ \text{+low}\\ \text{-round}\\ \text{+back}\\ \textbf{-ATR}\end{bmatrix}}} \quad \underset{\text{m}}{\overset{x_3}{\begin{bmatrix}\text{-voc}\\ \text{labial}\\ \text{+son}\\ \text{+nasal}\end{bmatrix}}} \quad \Lsh \quad \underset{\text{a}}{\overset{x_4}{\begin{bmatrix}\text{+voc}\\ \text{-high}\\ \text{+low}\\ \text{-round}\\ \text{+back}\\ \textbf{-ATR}\end{bmatrix}}} \quad \underset{\text{ɐ}}{\overset{x_5}{\begin{bmatrix}\text{+voc}\\ \text{-high}\\ \text{-low}\\ \text{-round}\\ \text{+back}\\ \text{+ATR}\end{bmatrix}}}$$

Finally, consider the case of a needy suffix and a needy root, as in kɪ-ɐ-par-ɪn, which begins with the root + suffix cycle. After failing to find any specified value, default insertion will supply the value [−ATR] for the suffix:

(65) *Kalenjin Suffix Begins ATR-Harmonize in par-ɪn*:

$$\overset{x_4}{\begin{bmatrix}\text{-voc}\\ \text{labial}\\ \text{-cont}\\ \text{-voice}\end{bmatrix}} \quad \overset{x_5}{\begin{bmatrix}\text{+voc}\\ \text{-high}\\ \text{+low}\\ \text{-round}\\ \text{+back}\\ \\ \end{bmatrix}} \quad \overset{x_6}{\begin{bmatrix}\text{-voc}\\ \text{coronal}\\ \text{+son}\\ \text{-nasal}\\ \text{-lateral}\end{bmatrix}} \quad \longleftarrow \quad \overset{x_7}{\begin{bmatrix}\text{+voc}\\ \text{+high}\\ \text{-low}\\ \text{-round}\\ \text{+back}\\ \\ \end{bmatrix}} \quad \overset{x_8}{\begin{bmatrix}\text{-voc}\\ \text{coronal}\\ \text{+son}\\ \text{+nasal}\end{bmatrix}}$$

(66) *Kalenjin Suffix Continues Search for ATR in par-ɪn*:

$$\begin{array}{cccccc} x_4 & x_5 & \leftarrow & x_6 & x_7 & x_8 \\ \begin{bmatrix}\text{-voc}\\ \text{labial}\\ \text{-cont}\\ \text{-voice}\end{bmatrix} & \begin{bmatrix}\text{+voc}\\ \text{-high}\\ \text{+low}\\ \text{-round}\\ \text{+back}\\ \\ \end{bmatrix} & & \begin{bmatrix}\text{-voc}\\ \text{coronal}\\ \text{+son}\\ \text{-nasal}\\ \text{-lateral}\end{bmatrix} & \begin{bmatrix}\text{+voc}\\ \text{+high}\\ \text{-low}\\ \text{-round}\\ \text{+back}\\ \\ \end{bmatrix} & \begin{bmatrix}\text{-voc}\\ \text{coronal}\\ \text{+son}\\ \text{+nasal}\end{bmatrix} \end{array}$$

(67) *Kalenjin Suffix Continues Search for ATR in par-ɪn*:

$$\begin{array}{cccccc} x_4 & \leftarrow & x_5 & x_6 & x_7 & x_8 \\ \begin{bmatrix}\text{-voc}\\ \text{labial}\\ \text{-cont}\\ \text{-voice}\end{bmatrix} & & \begin{bmatrix}\text{+voc}\\ \text{-high}\\ \text{+low}\\ \text{-round}\\ \text{+back}\\ \\ \end{bmatrix} & \begin{bmatrix}\text{-voc}\\ \text{coronal}\\ \text{+son}\\ \text{-nasal}\\ \text{-lateral}\end{bmatrix} & \begin{bmatrix}\text{+voc}\\ \text{+high}\\ \text{-low}\\ \text{-round}\\ \text{+back}\\ \\ \end{bmatrix} & \begin{bmatrix}\text{-voc}\\ \text{coronal}\\ \text{+son}\\ \text{+nasal}\end{bmatrix} \end{array}$$

(68) *Kalenjin Suffix Fails ($\nleftarrow$) to ATR-Harmonize in par-ɪn*:

$$\begin{array}{cccccc} \nleftarrow & x_4 & x_5 & x_6 & x_7 & x_8 \\ & \begin{bmatrix}\text{-voc}\\ \text{labial}\\ \text{-cont}\\ \text{-voice}\end{bmatrix} & \begin{bmatrix}\text{+voc}\\ \text{-high}\\ \text{+low}\\ \text{-round}\\ \text{+back}\\ \\ \end{bmatrix} & \begin{bmatrix}\text{-voc}\\ \text{coronal}\\ \text{+son}\\ \text{-nasal}\\ \text{-lateral}\end{bmatrix} & \begin{bmatrix}\text{+voc}\\ \text{+high}\\ \text{-low}\\ \text{-round}\\ \text{+back}\\ \\ \end{bmatrix} & \begin{bmatrix}\text{-voc}\\ \text{coronal}\\ \text{+son}\\ \text{+nasal}\end{bmatrix} \end{array}$$

(69) *Kalenjin Suffix Undergoes Default [−ATR] Insertion in par-ɪn*:

$$\begin{array}{ccccc} x_4 & x_5 & x_6 & x_7 & x_8 \\ \begin{bmatrix}\text{-voc}\\ \text{labial}\\ \text{-cont}\\ \text{-voice}\end{bmatrix} & \begin{bmatrix}\text{+voc}\\ \text{-high}\\ \text{+low}\\ \text{-round}\\ \text{+back}\\ \\ \end{bmatrix} & \begin{bmatrix}\text{-voc}\\ \text{coronal}\\ \text{+son}\\ \text{-nasal}\\ \text{-lateral}\end{bmatrix} & \begin{bmatrix}\text{+voc}\\ \text{+high}\\ \text{-low}\\ \text{-round}\\ \text{+back}\\ \textit{-ATR}\end{bmatrix} & \begin{bmatrix}\text{-voc}\\ \text{coronal}\\ \text{+son}\\ \text{+nasal}\end{bmatrix} \end{array}$$

(70) *Kalenjin Root Begins Rightward [±ATR] Search in par-ɪn*:

$$
\begin{array}{cccccc}
x_4 & x_5 & \rightarrow & x_6 & x_7 & x_8 \\
\begin{bmatrix}\text{-voc}\\ \text{labial}\\ \text{-cont}\\ \text{-voice}\end{bmatrix} &
\begin{bmatrix}\text{+voc}\\ \text{-high}\\ \text{+low}\\ \text{-round}\\ \text{+back}\\ \\ \end{bmatrix} & &
\begin{bmatrix}\text{-voc}\\ \text{coronal}\\ \text{+son}\\ \text{-nasal}\\ \text{-lateral}\end{bmatrix} &
\begin{bmatrix}\text{+voc}\\ \text{+high}\\ \text{-low}\\ \text{-round}\\ \text{+back}\\ \text{-ATR}\end{bmatrix} &
\begin{bmatrix}\text{-voc}\\ \text{coronal}\\ \text{+son}\\ \text{+nasal}\end{bmatrix}
\end{array}
$$

(71) *Kalenjin Root Begins Rightward [±ATR] Search in par-ɪn*:

$$
\begin{array}{cccccc}
x_4 & x_5 & x_6 & \Rsh & x_7 & x_8 \\
\begin{bmatrix}\text{-voc}\\ \text{labial}\\ \text{-cont}\\ \text{-voice}\end{bmatrix} &
\begin{bmatrix}\text{+voc}\\ \text{-high}\\ \text{+low}\\ \text{-round}\\ \text{+back}\\ \textbf{-ATR}\end{bmatrix} &
\begin{bmatrix}\text{-voc}\\ \text{coronal}\\ \text{+son}\\ \text{-nasal}\\ \text{-lateral}\end{bmatrix} & &
\begin{bmatrix}\text{+voc}\\ \text{+high}\\ \text{-low}\\ \text{-round}\\ \text{+back}\\ \textbf{-ATR}\end{bmatrix} &
\begin{bmatrix}\text{-voc}\\ \text{coronal}\\ \text{+son}\\ \text{+nasal}\end{bmatrix} \\
\text{p} & \text{a} & \text{r} & & \text{ɪ} & \text{n}
\end{array}
$$

To summarize, the patterns of Kalenjin [± ATR] harmony based on Hall et al. (1974) show that its computation vowel harmony is determined by directionality and locality, instead of "dominance" as a factor trumping all else. Prefixes copy from the closest local source to their right, regardless of its value, and the interaction of suffixes and roots can be determined based on which is needy, as shown in the three derivations above.

Kalenjin's system of [± ATR] harmony requires no special devices above and beyond the theory of locality already developed. Its departure from the pattern of Akan, in which prefixes copy from roots and suffixes copy from roots, lies only in the fact that some roots may also copy from suffixes. Before concluding, I will mention that one of the challenges in conducting research on putative dominant-recessive systems is the great range of morphological combinations that are needed to confirm all of the potential directionality effects. In particular, for Kalenjin, arbitrating between certain analytic possibilities depends on cases of multiple suffixes are possible, and what value will be computed for a needy suffix flanked by a [+ATR] root to its left and a [−ATR] suffix to its right. Non-final suffixes of this sort may have δ set rightwards, or bidirectionally.[17]

We have seen that a cyclic application of harmony, whereby needy suffixes harmonize first, followed by needy roots, followed by needy prefixes, accounts for all of the patterns in Hall et al. (1974) without any need to appeal to non-local "dominant-recessive" principles. Kalenjin has a variety of needy and non-needy morphemes, and the needy morphemes specify their directionality of search.

2.10. Split-Source Harmony: When A Consonant and A Vowel are Each Copied from in Turkish

Thus far we have devised a relatively simple model: one in which harmony takes place with a given direction of search, starting with the recipient segment "in need"; as soon as the closest-value of the harmonic feature is found, it is copied to the recipient.

Yet if search really is driven in terms of an efficient computation to find the closest source of a harmonic feature, we might expect that in certain configurations, some rather unexpected players could furnish the value for harmony, provided they are the closest to the recipient.

One important case to which we turn, that provides striking confirmation for the recipient-centric model of search, occurs in Turkish vowel harmony already introduced above, but with the additional twist: three pairs of *consonants* may participate in vowel harmony: k/k^j, g/g^j, l/l^j.

In the Turkish consonant inventory, k,g,l have [−back] counterparts (Clements & Sezer 1982: 233); thus, in k∼k^y, g∼g^y, l∼l^y, the feature [± back] is contrastive (and cannot be predicted allophonically):

(72) *Contrastive [± back] in Turkish k,g,l*:

bol	abundant	bol^y	cocktail
kalp	counterfeit	kal^yp	heart
kar	snow	k^yar	profit
gaz	gas	g^yavur	infidel

What is interesting is that these consonants, which show a contrastive and unpredictable [± back] specification, "participate" in vowel harmony. *k,g,l* are contrastively [+back], and hence vowels added after them will find these consonants as the closest source, and copy [+back]. Even in a word whose last vowel is [+back], if the last *segment* is contrastively [−back], then the suffixal alternant will be [−back].

(73) *Effect of [−back] Liquid on Harmony of Following Vowels in Turkish*

usul^j	usul^j-ü	system-acc.sg
petrol^j	petrol^j-ü	petrol-acc.sg
sual^j	sual^j-i	question-acc.sg
okul	okul-u	school-acc.sg
karakol	karakol-u	police station-acc.sg
tʃatal	tʃatal-ɨ	fork-acc.sg
petrol^j	petrol^j-de	petrol-loc.sg
meʃgul^j	meʃgul^j-düm	busy-past-1.sg

Importantly, we can see in the last two examples of (73) that the vowel undergoing valuation and the palatalized liquid need not be strictly adjacent for [−back] valuation to "unexpectedly" occur.

Conversely, in a word whose last vowel is [−back], where we would expect the suffix vowel to be [−back], if an intervening consonant is contrastively [+back], this consonant will determine the harmonic alternant instead (Clements and Sezer, 1982; Levi, 2004), as can be observed with the [+back] velar, that yields [+back] suffix vowels in an otherwise [−back] harmonic sequence.

(74) *Participation of Turkish [k] in [+back] harmony*:

ʃevk	ʃevkɨ	desire
haːl^{j}ik	haːl^{j}ikɨ	creator

These examples clearly show that a consonant may "intercept" [± back] harmony. What is most interesting from the point of view of theoretical models of phonological locality, is when both [± back] and [± round] harmony are operative in high-vowelled suffixes and find their values from different sources. In a form such as meʃgul^j-düm (73), the *ü* of the suffix receives its [+round] value from the preceding root vowel *u*, while it receives its [−back] value from the preceding root consonant l^j.

The intuition to be captured in these Turkish cases is that there are *two sources* of valuation for the needy affix in meʃgul^j-düm. In most cases of vowel harmony, both features, [±back] and [±round], are found in the same source, namely the preceding vowel. However, given that search is opportunistically "greedy", if one feature can be valued immediately, it will.

(75) *Turkish Accusative Case morpheme must:*
Back-Harmonize: δ = left, F = [± back]

Round-Harmonize: δ = left, F = [$\pm$ round]

Not just vowels but any segment that bears [$\pm$ back] can be a potential donor. In fact, parametric variation in what can be a possible donor will occupy the majority of the next chapter. For now our focus will continue to be on the search procedure and the fact that once the set of segments bearing the needed feature is delimited, the search procedure will literally terminate with the closest copy-source.

From the perspective of the needy suffix vowel in meʃgul^j-düm, the l^j is a closer potential source for [$\pm$ back] valuation than the preceding *u*. Hence, in the leftward search for a valuation source, once l^j is encountered, [$-$back] is copied to the suffix, and the search for [$\pm$ back] is terminated. As l^j does not provide a value for [$\pm$ round], the search for [$\pm$ round] continues, until *u* is encountered. This is schematized in (76):

(76) Turkish Acc searches Leftward for Back and Round in meʃgul^j-düm

$$\begin{array}{ccccc} x_1 & x_2 & x_3 & \leftarrow & x_4 \\ \begin{bmatrix} \text{+voc} \\ \text{+high} \\ \text{+back} \\ \text{+round} \end{bmatrix} & \begin{bmatrix} \text{-voc} \\ \text{+sonorant} \\ \text{-nasal} \\ \text{+lateral} \\ \text{-back} \end{bmatrix} & \begin{bmatrix} \text{-voc} \\ \text{coronal} \\ \text{-cont} \\ \text{-nas} \\ \text{+voi} \end{bmatrix} & & \begin{bmatrix} \text{+voc} \\ \text{+high} \\ \\ \end{bmatrix} \end{array}$$

(77) Finds [$-$back]

$$\begin{array}{ccccc} x_1 & x_2 & \leftarrow & x_3 & x_4 \\ \begin{bmatrix} \text{+voc} \\ \text{+high} \\ \text{+back} \\ \text{+round} \end{bmatrix} & \begin{bmatrix} \text{-voc} \\ \text{+sonorant} \\ \text{-nasal} \\ \text{+lateral} \\ \textbf{-back} \end{bmatrix} & & \begin{bmatrix} \text{-voc} \\ \text{coronal} \\ \text{-cont} \\ \text{-nas} \\ \text{+voi} \end{bmatrix} & \begin{bmatrix} \text{+voc} \\ \text{+high} \\ \\ \end{bmatrix} \end{array}$$

(78) Copies [$-$back]

$$\begin{array}{ccccc} x_1 & x_2 & \Lsh & x_3 & x_4 \\ \begin{bmatrix} \text{+voc} \\ \text{+high} \\ \text{+back} \\ \text{+round} \end{bmatrix} & \begin{bmatrix} \text{-voc} \\ \text{+sonorant} \\ \text{-nasal} \\ \text{+lateral} \\ \textbf{-back} \end{bmatrix} & & \begin{bmatrix} \text{-voc} \\ \text{coronal} \\ \text{-cont} \\ \text{-nas} \\ \text{+voi} \end{bmatrix} & \begin{bmatrix} \text{+voc} \\ \text{+high} \\ \textbf{-back} \\ \\ \end{bmatrix} \end{array}$$

(79) Finds [$+$round]

$$
\begin{array}{ccccc}
x_1 & \leftarrow & x_2 & x_3 & x_4 \\
\begin{bmatrix} \text{+voc} \\ \text{+high} \\ \text{+back} \\ \textbf{+round} \end{bmatrix} & &
\begin{bmatrix} \text{-voc} \\ \text{+sonorant} \\ \text{-nasal} \\ \text{+lateral} \\ \text{-back} \end{bmatrix} &
\begin{bmatrix} \text{-voc} \\ \text{coronal} \\ \text{-cont} \\ \text{-nas} \\ \text{+voi} \end{bmatrix} &
\begin{bmatrix} \text{+voc} \\ \text{+high} \\ \text{-back} \\ \end{bmatrix}
\end{array}
$$

(80) Copies [+round]

$$
\begin{array}{ccccc}
x_1 & \Lsh & x_2 & x_3 & x_4 \\
\begin{bmatrix} \text{+voc} \\ \text{+high} \\ \text{+back} \\ \textbf{+round} \end{bmatrix} & &
\begin{bmatrix} \text{-voc} \\ \text{+sonorant} \\ \text{-nasal} \\ \text{+lateral} \\ \text{-back} \end{bmatrix} &
\begin{bmatrix} \text{-voc} \\ \text{coronal} \\ \text{-cont} \\ \text{-nas} \\ \text{+voi} \end{bmatrix} &
\begin{bmatrix} \text{+voc} \\ \text{+high} \\ \text{-back} \\ \textbf{+round} \end{bmatrix} \\
\text{u} & & \text{l}^{\text{j}} & \text{d} & \text{ü}
\end{array}
$$

When inflectional suffixes are added in Turkish, they take their specification for [±back] from, indeed, the **closest** source of valuation, which is, in these cases, a consonant. Again, we see confirmation that the search procedure is greedy: even though both [±back] and [±round] could be copied together from the further-away root vowel, [±back] is copied as soon as it is encountered, removed from the list of needed-values, and the search inexorably proceeds further until all remaining values (in this case [±round]) are furnished.

2.11. Lack of Lookahead in Barra Gaelic Epenthesis Harmony

We have just considered the participation of contrastive [±back] consonants in vowel harmony to illustrate multiple-source valuation. In some configurations in which multiple sources value multiple features, the set of outcomes is even more restricted than we might expect, and provides justification for a stepwise derivational account of the feature-valuing steps as the search proceeds leftward.

In Barra Gaelic (Borgstrom, 1937; Sagey, 1987), a scenario quite similar to that of Turkish obtains: palatalized consonants "intercept" featural valuation that is otherwise entirely determined by a vocalic source.

In Barra Gaelic, there is a process of insertion of an epenthetic vowel to break up consonant clusters composed of sonorant+obstruent. The quality of this epenthetic vowel is often a full copy of *all* vocalic features of the preceding vowel.

(81) *Barra Gaelic Vowel Harmony with Full Copying of the Preceding Vowel*:

	Underlying form	Surface form	
a.	t^jimx^jal	t^jimix^jal	round about
b.	$æms^jir^j$	$æmæs^jir^j$	time

Whenever a consonant specified as [±back] intervenes, what is actually copied from the vowel must be characterized as a "deficient" constituent: all features *except* [± back]. A distinguishing feature of the consonantal inventory in Barra is that *all* consonants in the inventory are contrastive for [± back] except Labials (/m,p,b/) and /n/. The vowel system of Barra Gaelic is as follows:

(82) *Barra Gaelic Vowel system:*

−back	+back
i	ɨ,u
e	ʌ,o
æ	a,ɔ

While Barra Gaelic distinguishes contrastive [± back] extensively among the consonants, it lacks front rounded vowels such as *ü*. Nonetheless, [± round] is contrastive for */u/*, as it distinguishes */u/* from /ɨ/, and hence [± round] will be copied to an epenthetic vowel from */u/*, as we will presently observe.

Barra Gaelic exhibits the phenomenon of leftward *multiple-source valuation* for epenthetic vowels that arise in order to break up sonorant+consonant clusters.

(83) Barra Gaelic Vowel Harmony:

	Underlying form	Surface form	
a.	alpə	alapə	Scotland
b.	s^jærv	s^jærav	bitter
c.	urpel	urupel	tail
d.	ɔrm	ɔrɔm	on me
e.	færk	færak	anger
f.	mʌr^jv	mʌr^jev	the dead
g.	bu$l^j$$k^j$	bul^jik^j	bellows (gen.sg.)
h.	me$r^j$$k^j$	mer^jek^j	rust

In (83-a-d), the intervening consonant has the same [± back] value as the preceding vowel. Hence, the epenthetic vowel is surface-identical to the preceding vowel, even though, on the account developed here, its [± back] value is furnished by the consonant. This divergence in valuation sources is most apparent in (83-e-h), in which the epenthetic vowel and the preceding vowel are identical for all features *except* back. The intervening consonant bears a different value for [± back] than the preceding vowel, and since the consonant is a *closer* valuation source, it is encountered immediately in the leftward search.[18]

The steps involved are the same as those outlined above for Turkish meʃgul^j-düm, except there are more features to be valued. In Barra, a vowel must bear specifications for [±back], [±round], [±high], and [±low], in order to be lexically distinct. On the view that the epenthesis process furnishes only a timing slot with the feature [+vocalic], the vowel must conduct a dynamic search in order to find values for its four vocalic features.

(84) *Barra Epenthetic Vowel Must*:
Back-Harmonize: δ = left, F = [± back]
Round-Harmonize: δ = left, F = [± round]
High-Harmonize: δ = left, F = [± high]
Low-Harmonize: δ = left, F = [± low]

An illustration of one representative case, bul^jik^j, reveals an important aspect of the computation: that the closest source is immediately copied from, *without look-ahead*. Recall that the inventory lacks front rounded vowels, such as /*ü*/. Similar to the Turkish example discussed above, the Barra epenthetic vowel copies [±back] from a preceding consonant, and the remainder of its feature-values from a preceding vowel. Consider bul^{j}ik^j: if there *were* full vowel copy, we would expect an *u* in the second syllable. Let us examine, however, how the directional search proceeds.

Barra vowels require full specification for [±high], [±low], [±round], and [±back]. As mentioned, the feature combination [−back,+round] is banned from the inventory; we may understand this either in terms of an inviolable constraint, or an implicational statement: [−back] → [−round]. For concreteness of illustration, we adopt the latter (though this implementation is a secondary-level assumption). Consider now the steps involved in valuation of the epenthetic vowel in bul^{j}ik^j:

(85) Barra Gaelic Epenthetic Vowel's Search begins

$$\begin{matrix} x_1 & & x_2 & & x_3 \\ \begin{bmatrix} \text{+voc} \\ \text{+back} \\ \text{+round} \\ \text{+high} \\ \text{-low} \end{bmatrix} & & \begin{bmatrix} \text{-voc} \\ \text{+sonorant} \\ \text{-nasal} \\ \text{+lateral} \\ \text{-back} \end{bmatrix} & & \begin{bmatrix} \text{+voc} \end{bmatrix} \end{matrix}$$

(86) Barra Gaelic Epenthetic Vowel Finds and Copies [−back]

$$\begin{matrix} x_1 & & x_2 & \Lsh & x_3 \\ \begin{bmatrix} \text{+voc} \\ \text{+back} \\ \text{+round} \\ \text{+high} \\ \text{-low} \end{bmatrix} & & \begin{bmatrix} \text{-voc} \\ \text{+sonorant} \\ \text{-nasal} \\ \text{+lateral} \\ \textbf{-back} \end{bmatrix} & & \begin{bmatrix} \text{+voc} \\ \textbf{-back} \end{bmatrix} \end{matrix}$$

(87) Barra Gaelic Default Insertion of [−round] when [+voc,−back]

$$\begin{matrix} x_1 & & x_2 & \leftarrow & x_3 \\ \begin{bmatrix} \text{+voc} \\ \text{+back} \\ \text{+round} \\ \text{+high} \\ \text{-low} \end{bmatrix} & & \begin{bmatrix} \text{-voc} \\ \text{+sonorant} \\ \text{-nasal} \\ \text{+lateral} \\ \text{-back} \end{bmatrix} & & \begin{bmatrix} \text{+voc} \\ \text{-back} \\ \textit{-round} \end{bmatrix} \end{matrix}$$

(88) Barra Gaelic Epenthetic Vowel's Search for Other Features Continues

$$
\begin{array}{cccc}
x_1 & \leftarrow & x_2 & x_3 \\
\begin{bmatrix} \text{+voc} \\ \text{+back} \\ \text{+round} \\ \text{+high} \\ \text{-low} \end{bmatrix} & &
\begin{bmatrix} \text{-voc} \\ \text{+sonorant} \\ \text{-nasal} \\ \text{+lateral} \\ \text{-back} \end{bmatrix} &
\begin{bmatrix} \text{+voc} \\ \text{-back} \\ \text{-round} \end{bmatrix}
\end{array}
$$

(89) Barra Gaelic Epenthetic Vowel Finds and Copies Values for Round, High, Low

$$
\begin{array}{cccc}
x_1 & \Lsh & x_2 & x_3 \\
\begin{bmatrix} \text{+voc} \\ \text{+back} \\ \text{+round} \\ \textbf{+high} \\ \textbf{-low} \end{bmatrix} & &
\begin{bmatrix} \text{-voc} \\ \text{+sonorant} \\ \text{-nasal} \\ \text{+lateral} \\ \text{-back} \end{bmatrix} &
\begin{bmatrix} \text{+voc} \\ \text{-back} \\ \text{-round} \\ \textbf{+high} \\ \textbf{-low} \end{bmatrix} \\
\text{u} & & \text{l}^{\text{j}} & \text{i}
\end{array}
$$

Importantly, the resulting epenthetic vowel is *i*, and not *u* or *ü*. This is evident because neither l^{j} nor u alone could yield a high, front, unround vowel. It is crucial here is that there are two potential contrastive sources for valuation of the vowel. If both of them completely get their way, the result will be a front rounded vowel, which is banned by the implicational rule. But consider possible outputs with a front, unrounded vowel, or with a back, rounded vowel: $\text{bul}^{\text{j}}\text{ik}^{\text{j}}$, or the unattested $*\text{bul}^{\text{j}}\text{uk}^{\text{j}}$. The principle of dynamic, directional valuation will always guarantee the former, because [$-$back] is encountered first in the search, in derivational terms.

By contrast, in a donor-centric model, or a theory of global optimization of subsegmental agreement, nothing would guarantee this result. Although principles could be formulated to guarantee that the potential sources of copying are the palatalized liquid and the preceding vowel, such constraints alone will not choose between the attested and unattested output. What is required is a statement to the effect that the first closest contrastive value that is copied is etched in stone, and must be kept, regardless of what might ultimately lead to a better choice from a "global" perspective.[19] Even though in some imaginable grammars, copying all features from the *same* source might lead to a "less polygamous" set of copying relationships, this consideration never seems to matter.[20] Regardless of its ultimate implementation in terms of rules or constraints, all research into multiple-source harmony must reckon with the fact that valuation is a "greedy" operation, and will copy the first value that it can.

In the *search-and-copy* model developed here, this search is initiated *by the recipient*, and upon each match with a contrastive source for the feature in question, valuation occurs. When a contrastively (non)-palatalized consonant is the first segment to the left of the epenthetic vowel, the vowel will copy its value for [$\pm$ back] and commit to that value henceforth, thereafter proceeding to conduct the search for the remaining unvalued features.

2.12. General Conclusion: Search Must be Recipient-Initiated

This chapter has presented the essential search-and-copy algorithm: harmony is a search-and-copy process initiated by a needy vowel in order to find itself one or more feature-values. The search procedure is computed directionally to find the closest value-source and is greedy, in the sense that it operates without lookahead but will copy a value as soon as it finds it. The resulting model bears a striking degree of similarity to the process of *Agree* posited in syntactic theory, in which a value-seeking element (such as a Tense node) engages in a downward search for the closest noun phrase from which to copy person and number features.

We have considered three main arguments for the recipient-initiated process of vowel harmony. The first is that in a recipient-centered model, it is a property of the Turkish plural morpheme *-ler* but not the nominalizing morpheme *-gen* that it needs to copy feature-values from its closest leftward source. The fact that some morphemes may exceptionally non-undergo harmony, but no morpheme may exceptionally not-furnish harmony, requires no modification to the general search procedure. In contrast, in a donor-centric theory, the absence of exceptional non-triggers is surprising.

The second argument is that in a recipient-centered model, when the search is bidirectional, the result of copying-from-closest will require identical feature-values on both sides in case of a literal tie in closeness. In Woleaian theme-vowel harmony, only when both flanking vowels are high will harmony occur, a result which is captured with no existing modification to the general search procedure. In a donor-centric theory, by contrast, this would require serious negotiation between two separate spreading rules.

The third main argument is that when more than one feature-value is needed on the harmonizing vowel, the search will copy from the first source it finds, even if there might be a potentially better global choice downstream. In Barra Gaelic vowel harmony, [−back] is copied from a consonant, even though this prevents subsequent copying of [+round] from the next vowel over due to the fact that the language lacks the front rounded vowel *ü*. This result is natural given the directional leftward march of the search procedure, but would require additional principles to guarantee it in a system that considers multiple candiates for feature-value-sources in parallel.

These empirical considerations strongly establish the need for a recipient-centered model of vowel harmony, and this chapter has established the basic Search principle that accomplishes this goal. The Search principle, that requires copying from the closest element in the domain, has direct syntactic analogues in cases such as the locality of head movement. Consider the C node in an interrogative question, which must attract a contentful head as its sister: it always attracts the closest head (Travis, 1984) and cannot skip the closest element to attract a further one instead (90-c):

(90) a. [$_{CP}$ C [$_{IP}$ they could have left]]
b. [$_{CP}$ Could [$_{IP}$ they t_{could} have left ?]]
c. *[$_{CP}$ Have [$_{IP}$ they could t_{have} left ?]]

The fact that search halts with the closest element of the relevant type finds strong parallels within other modules of linguistic computation. While the search principle for vowel harmony has a broad applicability across languages, however, there are a number of parametric options that determine the nature of the *domain of search* itself, to which we turn in subsequent chapters.

For example, in both the Turkish case of meʃgulʲ-düm and the Barra Gaelic case of bulʲikʲ, harmony was accomplished with a [−back] consonant. These patterns required a revision of the statement that the possible donors of harmony to be changed to "any segment" rather than "vowel". What is traditionally called "back vowel harmony" in fact is harmony with any eligible source of a [−back] value. However, even this statement does not really get at the full picture, since there are some segments that are phonetically [−back] (such as the glide /j/) but that do not participate in harmony at all. The right statement of what is a possible donor in both of these cases is not actually “any segment”, but rather “any segment contrastive for [$\pm$back]”. We begin the next chapter by spelling out exactly what this means.

Notes

[7]I adopt the features of Odden (2003), with two modifications. First, Designated Articulators such as Labial, Coronal, or Dorsal are unary features (cf. Halle (2005)), where vowels are always Dorsal. Second, two features are used to classify vowels, glides, and consonants: [± vocalic] and [± consonantal]. Glides are [−vocalic, −consonantal], vowels are [+vocalic, +consonantal], and consonants are [−vocalic,+consonantal]; see Nevins and Chitoran (2008) for discussion.

[8]I depart from Turkish orthography and represent the back unrounded vowel as ɨ for visual perspicuity.

[9]In the case that partial orders among segments are empirically necessary (perhaps in analyses of metathesis), so that both x and y precede z but there is no lexically specified precedence relation between x and y, the algorithm could be generalized so as to allow traversal of parallel paths, with a mechanism for resolution of contradictory feature values among the closest elements on each path.

[10]The only prefixation in Turkish is reduplicative: e.g. *cip-ciliz*, *dop-doluz*, a process used for intensification of adjectives (Kelepir, 2000). As the prefix is a full copy of the initial root vowel, this process would never induce any alternations in the root.

[11]See Kabak and Vogel (2009) for a similar approach to disharmonic root vowels.

[12]The semantics of these features is that they are binary predicates true or false of a given articulatory configuration, where [−F] means ¬[+F]:

(91) *Semantics of vowel height features*:

a. [+high]: true iff tongue body is raised beyond a certain critical point P_h above the midline
[−high]: true iff tongue body is not raised beyond a certain critical point P_h above the midline

b. [+low]: true iff tongue body is lowered beneath a certain critical point P_h below the midline
[−low]: true iff tongue body is not lowered beneath a certain critical point P_h below the midline

[13]Hyman (1998) analyzes a formally similar process of bidirectional height harmony in the Bantu language Yaka, in which [± high] harmony of the first vowel of the perfective suffix only copies [−high] when both the closest vowel to its left and right are [−high].

[14]Notice that the process of harmony does not determine a value for [± back] for the theme vowel. This particular empirical question awaits decisive evidence from phonological processes about whether the low vowel a of Woleaian is [+back]. The answer to this question does not bear on the illustration of harmony at hand, however, as there is no other [−high, +low] nonround vowel that contrasts with a and no other [−high, −low] nonround vowel that contrasts with e.

[15]It does not matter which vowel is "reached earlier", if for example, there are more irrelevant consonants to traverse on one side than the other.

[16]This can be implemented as the choice between a constraint banning [−low] being what ultimately forces [+low] to surface just in case harmony fails, or as a case of a genuine Plan B process taking place as a rule of insertion. The choice between these two

implementations can't and won't be made here.

[17]Based on similar data from Turkana, e.g. *a-bun-ɛ-rɛ*, (Noske, 2001, p.803-804) concludes that non-final alternating (i.e. needy) suffixes must be specified as harmonizing to their right. Convergent evidence for this conclusion from Nilo-Saharan is found in Levergood (1984, p.278), who determines that Maasai suffixes must copy from their right, based on ɛ-ta-raŋ-iʃ-e '3sg-past-sing-intr.-past', where the non-final suffix copies [+ATR] from its right, rather than [−ATR] from its left.

[18]In the case of epenthesis, however, it looks like the direction of search must be inherently specified. One possibility within the derivational character of the current model is that in the case of sonority-driven epenthesis, given an iterative left-to-right syllabification algorithm, the closest source of visible valuation will be leftwards. A review of the handful of cases in Kawahara (2003) reveals that Kolami copy-epenthesis is from the left, whereas Winnebago and Fula loanwords resulting from word-initial onset-liquid clusters copy from the right. In short, these data are consistent with the possibility that leftward copying is a universal default, except when there is nothing to the left.

[19]A related effect in syntactic theory has been formulated as the Earliness Principle (Pesetsky, 1989; Collins, 2001), requiring that features be valued as soon as possible, even if this will ultimately lead to a sub-optimal outcome. For example, in **John seems that [t_{John} is nice]*, the valuing of the case feature of *John* in the embedded clause cannot be deferred until the matrix clause, leading to inability to satisfy the requirements of the higher tense node.

[20]In a discussion of the Turkish harmony cases where a [−back] liquid intercepts [+back] harmony, Padgett (2002, 94) discusses that when harmony is formalized as the initiative to spread features from donors, an output form like meʃgul^{j}-düm must violate this harmony requirement three times. This seems to be further indication that "spreading" by potential donors is insufficient to model the locality of vowel harmony.

Appendix to Chapter 2: Harmony Necessitated by Deletion

In this appendix I will demonstrate that the Search algorithm developed in Chapter 2 can be extended to cases of harmony that result from deletion processes. While the harmony algorithm is ordinarily used in phonological computations to supply a value for a "needy" segment requiring a feature-value in virtue of lacking one in its memorized representation, the same algorithm can be applied by segments that lack a feature-value due to a deletion rule that is operative later in the phonology. For example, in Chumash, an isolate language of California, certain coronal fricatives require harmony for the feature [± distributed] (affecting tongue tip orientation) not because they lack an underlying value for it, but because a post-cyclic rule deletes their values.

Chumash has a process of leftward [± distributed] harmony affecting its coronal fricatives (henceforth "sibilants") *s* and s̺, as described in Poser (1982) . The contrast between laminal *s* and apical s̺ Chumash is one of [± distributed], following Mithun (1998, p.221) and Hansson (2001, p.58), and thus Poser's *š* will be henceforth transcribed as [−distr] s̺. The apical consonants *d,t,l,n* in Chumash are [−distributed].

What is unique about Chumash's [± distr] harmony is that all non-final sibilants harmonize with the sibilant to their right, regardless of their morphological affiliation, morphological constituency, or underlying value of [± distr]:

(92) Chumash sibilant harmony (Poser 1982: 132)

a.	/k+sunon+us/	[ksunonus]	'I obey him'
b.	/k+sunon+ṣ/	[kṣunotṣ]	'I am obedient'
c.	/su+wayan/	[suwayan]	'cause to hang'
d.	/k+su+ṣoyin/	[kṣuṣoyin]	'I darken it'
e.	/s+ixut/	[sixut]	'it burns'
f.	/s+ilakṣ/	[ṣilakṣ]	'it is soft'
g.	/s+kuti+waṣ/	[ṣkutiwaṣ]	'he saw'
h.	/s+apitṣ^{h}o+it/	[ṣapitṣ^{h}olit]	'I have a stroke of good luck'
i.	/s+apitṣ^{h}o+us+waṣ/	[ṣapitṣ^{h}oluṣwaṣ]	'they had a stroke of good luck'

As this type of harmony affects prefixes, roots, and suffixes alike, without regard to the order of morphological structure, the harmony process is not a cyclic rule of supplying a value for newly-added morphemes. Poser (1982, p.134) shows for (92-i), for example, that no morphological bracketing of [[3SUBJ [[good+luck] 3OBJ]] PAST] would result in morpheme-by-morpheme copying of [−distr] all the way through to every coronal fricative. The distinction between cyclic and post-cyclic processes (see Halle and Vergnaud (1987), among others) is that post-cyclic processes apply in a single pass through the word, without regard to morphological constituency, and that post-cyclic processes apply after all cyclic processes have applied.

Chumash [± distr] harmony occurs after the morphological structure of a word is completely built and all of its feature-values complete. It is therefore somewhat of an unusual process from the point of view of other types of agglutinative-morphology harmony, and has been treated as a post-cyclic process that *changes* rather than supplies feature-values.

In a word such as /s+apitṣ^{h}o+us+waṣ/, let us assume that there is no harmony upon affixation, and that the subject agreement, object agreement, and tense affixes are fully specified for [± distr] in their underlying representations. Prior to the application of the post-cyclic process of harmony, they have the following representations, where only the coronal fricatives are shown, for convenience.

(93) *Representation of Fricatives in /s+apitṣ^{h}o+us+waṣ/ prior to postcyclic block*:

$$
\begin{array}{cccc}
f_1 & f_2 & f_3 & f_4 \\
\begin{bmatrix} \text{+cons} \\ \text{coronal} \\ \text{+cont} \\ \text{-voice} \\ \text{-asp} \\ \text{+distr} \end{bmatrix} &
\begin{bmatrix} \text{+cons} \\ \text{coronal} \\ \text{+cont} \\ \text{-voice} \\ \text{+asp} \\ \text{-distr} \end{bmatrix} &
\begin{bmatrix} \text{+cons} \\ \text{coronal} \\ \text{+cont} \\ \text{-voice} \\ \text{-asp} \\ \text{+distr} \end{bmatrix} &
\begin{bmatrix} \text{+cons} \\ \text{coronal} \\ \text{+cont} \\ \text{-voice} \\ \text{-asp} \\ \text{-distr} \end{bmatrix} \\
\text{s} & \text{ʂ}^{\text{h}} & \text{s} & \text{ʂ}
\end{array}
$$

Following Poser's original proposal, what is particular to Chumash is that all non-final coronal fricatives have their [± distr] specification *deleted* from the representation:[21]

(94) *Representation of Fricatives in /s+apitʂ^{h}o+us+waʂ/ after Deletion of Non-Final [± distr] :*

$$
\begin{array}{cccc}
f_1 & f_2 & f_3 & f_4 \\
\begin{bmatrix} \text{+cons} \\ \text{coronal} \\ \text{+cont} \\ \text{-voice} \\ \text{-asp} \end{bmatrix} &
\begin{bmatrix} \text{+cons} \\ \text{coronal} \\ \text{+cont} \\ \text{-voice} \\ \text{+asp} \end{bmatrix} &
\begin{bmatrix} \text{+cons} \\ \text{coronal} \\ \text{+cont} \\ \text{-voice} \\ \text{-asp} \end{bmatrix} &
\begin{bmatrix} \text{+cons} \\ \text{coronal} \\ \text{+cont} \\ \text{-voice} \\ \text{-asp} \\ \text{-distr} \end{bmatrix}
\end{array}
$$

This post-cyclic deletion rule is all that is unique about Chumash's feature-changing harmony. The subsequent harmony process follows the same algorithm adopted in Chapter 2, with search-and-copy occurring on each element from right to left. All elements contrastive for [± distr] (e.g. the coronal fricatives) within the post-cyclic word must harmonize:

(95) *Chumash post-cyclic word must:*
Distr-Harmonize: δ = right, F = [contrastive: ± distr]

The application of [± distr] harmony to each non-final fricative, from right-to-left, is shown below.

(96) *First Non-Final Fricative Searches for [± distr]*:

$$
\begin{matrix} f_1 & & f_2 & & f_3 & \Rsh & f_4 \\
\begin{bmatrix} \text{+cons} \\ \text{coronal} \\ \text{+cont} \\ \text{-voice} \\ \text{-asp} \end{bmatrix} & &
\begin{bmatrix} \text{+cons} \\ \text{coronal} \\ \text{+cont} \\ \text{-voice} \\ \text{+asp} \end{bmatrix} & &
\begin{bmatrix} \text{+cons} \\ \text{coronal} \\ \text{+cont} \\ \text{-voice} \\ \text{-asp} \\ \textbf{-distr} \end{bmatrix} & &
\begin{bmatrix} \text{+cons} \\ \text{coronal} \\ \text{+cont} \\ \text{-voice} \\ \text{-asp} \\ \text{-distr} \end{bmatrix}
\end{matrix}
$$

(97) *Second Non-Final Fricative Searches for [± distr]*:

$$
\begin{matrix} f_1 & & f_2 & \Rsh & f_3 & & f_4 \\
\begin{bmatrix} \text{+cons} \\ \text{coronal} \\ \text{+cont} \\ \text{-voice} \\ \text{-asp} \end{bmatrix} & &
\begin{bmatrix} \text{+cons} \\ \text{coronal} \\ \text{+cont} \\ \text{-voice} \\ \text{+asp} \\ \textbf{-distr} \end{bmatrix} & &
\begin{bmatrix} \text{+cons} \\ \text{coronal} \\ \text{+cont} \\ \text{-voice} \\ \text{-asp} \\ \text{-distr} \end{bmatrix} & &
\begin{bmatrix} \text{+cons} \\ \text{coronal} \\ \text{+cont} \\ \text{-voice} \\ \text{-asp} \\ \text{-distr} \end{bmatrix}
\end{matrix}
$$

(98) *Third Non-Final Fricative Searches for [± distr]*:

$$
\begin{matrix} f_1 & \Rsh & f_2 & & f_3 & & f_4 \\
\begin{bmatrix} \text{+cons} \\ \text{coronal} \\ \text{+cont} \\ \text{-voice} \\ \text{-asp} \\ \textbf{-distr} \end{bmatrix} & &
\begin{bmatrix} \text{+cons} \\ \text{coronal} \\ \text{+cont} \\ \text{-voice} \\ \text{+asp} \\ \text{-distr} \end{bmatrix} & &
\begin{bmatrix} \text{+cons} \\ \text{coronal} \\ \text{+cont} \\ \text{-voice} \\ \text{-asp} \\ \text{-distr} \end{bmatrix} & &
\begin{bmatrix} \text{+cons} \\ \text{coronal} \\ \text{+cont} \\ \text{-voice} \\ \text{-asp} \\ \text{-distr} \end{bmatrix} \\
\text{ʂ} & & \text{ʂ} & & \text{ʂ} & & \text{ʂ}
\end{matrix}
$$

Thus, Chumash feature-changing harmony can be incorporated into the basic Search-and-Copy algorithm of Chapter 2, provided a postcyclic process of deleting all non-final [± distr] specifications is included. Indeed, other phonological processes of feature-changing copying have been treated with deletion of features on non-final elements; Cho (1990) treats [± voice] assimilation in consonant clusters in Serbo-Croatian (e.g. *rob∼ro**pst**avo* 'slave, slavery'; *svat, sva**db**a* 'wedding guest, wedding') as the result of deletion of voicing on non-final consonants in an adjacent sequence, followed by assimilation. In Chumash, since all representations must exit the phonology with a specification for [± distr], the harmony procedure is enlisted to provide a value for elements that lost them due to deletion.

An additional process of interest in Chumash [± distr] harmony is its interaction with the process of apicalization (called "precoronal palatalization" in Poser (1982)), which supplies a [−distr] value for a sibilant that immediately precedes a [−distr] segment, i.e. *t,l,n*. Since this process only occurs under immediate adjacency, I will call it local apicalization, distinguishing it from the unbounded [± distr] harmony between sibilants contrastive for the feature.

(99) *Chumash local apicalization (Poser 1982:152)*:

a.	/s+nit^{h}oy/	[s̺nit^{h}oy]	'he goes'
b.	/s+tumun/	[s̺tumun]	'its egg'
c.	/s+lok'in/	[s̺lok'in]	'he cuts it'

Following Poser (1993), I assume that the rule of local apicalization precedes any initiation of search-and-copy for harmony.[22] Local apicalization is a process that occurs only under immediate precedence, with additional morphological restrictions, and can be clearly separated from the harmony process. In addition, as *s* and s̺ are contrastive for [± distr] while *t,l,n* are not, local apicalization must clearly be outside of harmony, since neither *n* in *ksunonus* induces apicalization of the sibilant to its left.

Given the greedy and local character of the Search procedure, every sibilant will copy [± distr] from the sibilant source immediately to its right, regardless of whether that value comes from an underlying specification (as in final sibilants), from local apicalization, or from the result of harmony itself. Thus, a sibilant enacting a harmonic search for [± distr] may copy it from another sibilant that has itself acquired it via local apicalization.

(100) *Interaction of locally-apicalized segments in [± distr] harmony (Poser 1982:153)*:

a.	/s+is+tiʔ/	[s̺is̺tiʔi]	'he finds it'
b.	/s+ti+yep+us/	[s̺tiyepus]	'he tells him'
c.	/s+is̺+lu+sisin/	[s̺is̺lusisin]	'they two are gone awry'

In (100-c) there are three non-final sibilants in [s̺is̺lusisin]. The first non-final sibilant copies [± distr] from the closest rightward source, which bears an underlying value. The second non-final sibilant acquired [−distr] via local apicalization. The third non-final sibilant copies [± distr] from its closest rightward source. The derivation proceeds following deletion of [± distr] on all non-final fricatives, where only the fricatives and the *l* are represented, for convenience.

(101) *Representation of fricatives in /s+iʂ+lu+sisin/ after deletion of all non-final [± distr] :*

$$\begin{matrix} f_1 & f_2 & l & f_3 & f_4 \\ \begin{bmatrix}\text{+cons}\\ \text{coronal}\\ \text{+cont}\\ \text{-voice}\end{bmatrix} & \begin{bmatrix}\text{+cons}\\ \text{coronal}\\ \text{+cont}\\ \text{-voice}\end{bmatrix} & \begin{bmatrix}\text{+cons}\\ \text{coronal}\\ \text{+cont}\\ \text{+lateral}\\ \text{-distr}\end{bmatrix} & \begin{bmatrix}\text{+cons}\\ \text{coronal}\\ \text{+cont}\\ \text{-voice}\end{bmatrix} & \begin{bmatrix}\text{+cons}\\ \text{coronal}\\ \text{+cont}\\ \text{-voice}\\ \text{+distr}\end{bmatrix} \end{matrix}$$

(102) *Application of local apicalization*:

$$\begin{matrix} f_1 & f_2 & l & f_3 & f_4 \\ \begin{bmatrix}\text{+cons}\\ \text{coronal}\\ \text{+cont}\\ \text{-voice}\end{bmatrix} & \begin{bmatrix}\text{+cons}\\ \text{coronal}\\ \text{+cont}\\ \text{-voice}\\ \textbf{-distr}\end{bmatrix} & \begin{bmatrix}\text{+cons}\\ \text{coronal}\\ \text{+cont}\\ \text{+lateral}\\ \text{-distr}\end{bmatrix} & \begin{bmatrix}\text{+cons}\\ \text{coronal}\\ \text{+cont}\\ \text{-voice}\end{bmatrix} & \begin{bmatrix}\text{+cons}\\ \text{coronal}\\ \text{+cont}\\ \text{-voice}\\ \text{-dist}\end{bmatrix} \end{matrix}$$

(103) *First Non-Final Fricative Searches for [± distr]::*

$$\begin{matrix} f_1 & f_2 & f_3 & \Rsh & f_4 \\ \begin{bmatrix}\text{+cons}\\ \text{coronal}\\ \text{+cont}\\ \text{-voice}\end{bmatrix} & \begin{bmatrix}\text{+cons}\\ \text{coronal}\\ \text{+cont}\\ \text{-voice}\\ \text{-distr}\end{bmatrix} & \begin{bmatrix}\text{+cons}\\ \text{coronal}\\ \text{+cont}\\ \text{-voice}\\ \textbf{+distr}\end{bmatrix} & & \begin{bmatrix}\text{+cons}\\ \text{coronal}\\ \text{+cont}\\ \text{-voice}\\ \text{+distr}\end{bmatrix} \end{matrix}$$

(104) *Inital Fricative Searches for [± distr]::*

$$\begin{matrix} f_1 & \Rsh & f_2 & f_3 & f_4 \\ \begin{bmatrix}\text{+cons}\\ \text{coronal}\\ \text{+cont}\\ \text{-voice}\\ \textbf{-distr}\end{bmatrix} & & \begin{bmatrix}\text{+cons}\\ \text{coronal}\\ \text{+cont}\\ \text{-voice}\\ \text{-distr}\end{bmatrix} & \begin{bmatrix}\text{+cons}\\ \text{coronal}\\ \text{+cont}\\ \text{-voice}\\ \text{+distr}\end{bmatrix} & \begin{bmatrix}\text{+cons}\\ \text{coronal}\\ \text{+cont}\\ \text{-voice}\\ \text{+distr}\end{bmatrix} \\ \text{ʂ} & & \text{ʂ} & \text{s} & \text{s} \end{matrix}$$

In conclusion, perhaps many cases of "feature-changing" harmony can be accomodated by an initial step of deletion. For example, iterative [± ATR] harmony in dialects of Canadian French (Poliquin, 2006) affects all non-final high vowels within a word; this by be accomplished by deletion feeding harmony. The process of *ikan'e* in Russian changes unstressed /a/ to [i] when preceded by a [−back,+high] consonant. One way to achieve this feature-change is to post-cyclically delete features of all unstressed vowels (a kind of vowel reduction), and copy [±back, ±high] from the closest leftward source. Finally, as mentioned above, feature-changing [± voice] assimilation in languages such as Serbo-Croatian has been handled by Cho (1990) as the result of deletion of [± voice] specifications on non-final consonants in a sequence, followed by feature-copying.

The harmony algorithm developed in Chapter 2 enables a theory of locality for segments that need a feature-value; whether they need this value due to an "innate" or "acquired" property of their representation is irrelevant to the computation employed to supply it.

Chapter 3

Contrastiveness, Markedness and Feature-Based Locality

The search principle introduced in Chapter 2 finds and copies needed feature-values from the closest element in the domain of search. An important source of crosslinguistic variation in delimiting the search domain results from language-specific reference to two paradigmatic properties defined by the inventory of sounds in the language. Once these two properties are defined, differences between languages reduce to use of the same search algorithm over a different set of elements, predictable based on inspection of the featural alphabet of the language.

The chapter will explore a range of crosslinguistic variation in what is excluded from the domain of search, and will conclude with a summary of how the range of "transparency" effects in harmony is quite limited by the grammatical options made available within a relativized-search-domain approach to feature-copying rules.

3.1. Contrastiveness Cuts Potential Donors in Finnish Harmony

We begin by examining Finnish, which has a [± back] harmony system similar to that observed in the previous chapter for Turkish. The essive case suffix alternates between the [+back] vowel *-na* and the [−back] vowel *-nä* depending on the value of the vowel to its left. However, not all vowels count: in particular, the high front vowel *i* (as well as the mid front vowel *e*) acts as if it is not there (Ringen, 1975).

(1) *The high front vowel /i/ is invisible in Finnish Harmony*

a. pöütä-nä 'table-essive'
b. pouta-na 'fine.weather-essive'
c. koti-na 'home-essive'
d. pappi-na 'priest-essive'
e. väkkärä-nä 'pinwheel-essive'
f. makkara-na 'sausage-essive'
g. tühmä-nä 'stupid-essive'
h. tuhma-na 'naughty-essive'

As the examples in (1-a-b) show, the [−back] suffix *-nä* follows a root that is all [−back], and the [+back] suffix *-na* follows a root that is all [+back]. However, (1-c-d) shows that when the closest leftward vowel to the suffix is *i*, it is skipped, and the search continues leftward to the next vowel, which is [+back], eventually yielding the [+back] variant of the suffix.

What's going on here? Shouldn't the search halt immediately with the closest vowel in the domain? To answer this puzzle, we return to one of the central points raised in Chapter 1, that humans often do not compute distance "as the crow flies", but rather within a system. Let's examine the properties of the Finnish vowel inventory. In the table below, I indicate "missing" vowels using the symbol ✗.

(2) *Finnish Vowel Inventory*[23]:

[−back, −round]	[−back, +round]	[+back, +round]	[+back, −round]	
i	ü	u	✗	[+high, −low]
e	ö	o	✗	[−high, −low]
ä			a	[−high, +low]

As (2) shows by the placement of the two ✗ symbols, there are no [+back, −round, −low] vowels in the inventory, and thus, while *ü/u*, *ö,o*, and *ä,a* differ from each other only in [± back], the [−back, −round, −low] vowels *i* and *e* have no harmonic counterpart. The notion of having a harmonic counterpart or not can be formalized once we refer to features[24]:

(3) *Definition of Contrastive:*
A segment S with specification αF in position P is *contrastive* for F if there is another segment S' in the inventory that can occur in P and is featurally identical to S, except that it is -αF

myVals V
myPosition P
myFeatsneeded F

while F is not empty:
· Go in direction δ and update P
· **if** P has a value for any $f, f \in F$:
· · **if** Val(f) is contrastive on P:
· · · Copy Val(P, f) to V
· · · Remove f from F

Figure 3.1: Search Relativized to Contrastive Values.

The definition of contrastive determines whether an element has a "twin" with respect to a certain feature value. We are observing that a paradigmatic property of a vowel, namely whether it is contrastive or not, determines its syntagmatic behavior, i.e. whether it participates in vowel harmony or not. An easy way to remember this pattern might be "If you don't have a twin, you can't be a donor".

Since the vowels *i,e* are non-contrastive for [$\pm$ back], they pattern as invisible in [$\pm$ back] harmony, not acting as donors and not being included in the search. Putting it differently, whether or not a vowel is contrastive for the harmonic feature determines whether or not it is at all included in what the search looks at. The set of items that are potentially included in the search (based on, for example, whether they are contrastive or not) constitute the *domain* of search. The modification of the search algorithm to exclude noncontrastive items as potential donors is shown in Figure 3.1.

Relativization prunes the domain of search: in Finnish, it is not the closest vowel that is copied from, but rather the closest *contrastive* vowel.

(4) *Finnish Essive Case morpheme suffix must:*
Back-Harmonize: δ = left, F = [contrastive: $\pm$ back]

(5) Essive Suffix Begins Back-Harmonize in *koti-na*:

$$\begin{array}{ccccc} x_1 & x_2 & x_3 \quad \longleftarrow & x_4 & x_5 \\ \begin{bmatrix} \text{+voc} \\ \text{-high} \\ \text{-low} \\ \text{+ round} \\ \text{+back} \end{bmatrix} & \begin{bmatrix} \text{-voc} \\ \text{coronal} \\ \text{-continuant} \\ \text{-voice} \end{bmatrix} & \begin{bmatrix} \text{+voc} \\ \text{+high} \\ \text{-low} \\ \text{-round} \\ \text{-back} \end{bmatrix} & \begin{bmatrix} \text{-voc} \\ \text{coronal} \\ \text{+nasal} \end{bmatrix} & \begin{bmatrix} \text{+voc} \\ \text{-high} \\ \text{+low} \\ \text{-round} \end{bmatrix} \end{array}$$

(6) Essive Suffix Finds [+back] on x_1 and copies [+back] to itself:

$$\begin{array}{ccccc} x_1 \quad \Lsh & x_2 & x_3 & x_4 & x_5 \\ \begin{bmatrix} \text{+voc} \\ \text{-high} \\ \text{-low} \\ \text{+ round} \\ \textbf{+back} \end{bmatrix} & \begin{bmatrix} \text{-voc} \\ \text{coronal} \\ \text{-continuant} \\ \text{-voice} \end{bmatrix} & \begin{bmatrix} \text{+voc} \\ \text{+high} \\ \text{-low} \\ \text{-round} \\ \text{-back} \end{bmatrix} & \begin{bmatrix} \text{-voc} \\ \text{coronal} \\ \text{+nasal} \end{bmatrix} & \begin{bmatrix} \text{+voc} \\ \text{-high} \\ \text{+low} \\ \text{-round} \\ \textbf{+back} \end{bmatrix} \\ \text{o} & \text{t} & \text{i} & \text{n} & \text{a} \end{array}$$

A similar pattern of vowel harmony skipping *i* and searching long-distance can be found in Classical Mongolian (Svantesson, 1985), in which the vowel system also leaves *i* noncontrastive for [± back]:

(7) *Classical Mongolian Vowel System*:

[−back, −round]	[−back, +round]	[+back, +round]	[+back, −round]	
i	ü	u		[+high]
e	ö	o	a	[−high]

Classical Mongolian [−high,−round] suffixes copy the feature [± back] from the closest contrastive source, skipping *i* (Svantesson, 1985, 305):

(8) *Classical Mongolian suffixal harmony: copies closest contrastive [± back]*:

ulus	ulus-ača	'nation-abl.'
aman	aman-ača	'mouth-abl.'
üker	üker-eče	'ox-abl.'
mören	mören-eče	'river-abl.'
morin	m**o**rin-ača	'horse-abl.'

In summary, the pattern of transparent *i* in the [± back] contrastive harmony systems of a number of languages exists because those languages lack (some having lost from an earlier stage) [+high,+back,−round] ɨ; this fact, coupled with the definition of contrastiveness, derives the rest.

3.2. Finnish Transparent Vowels are Excluded from Harmony

In concert with a strict interpretation of the autosegmental view of harmony as spreading association lines – so that a single instance of a feature is shared among multiple segments – researchers such as NiChiosain and Padgett (2001) hypothesize that all harmony is strictly local. (See Gafos (1996, p.77-81) for a demonstration that interpreting the autosegmental formalism without underspecification inevitably leads to strict adjacency.[25]). Thus, NiChiosain and Padgett (2001, p.119) conjecture "We assume that locality holds strictly, in two senses of 'strict'. First, spreading respects segmental adjacency. An essential result of this view is that segments are either blockers or participants in spreading: there is no transparency or skipping. Second, segmentally strict locality is inviolable".

The major claim of the current book is that phonological locality is relativized, not strict, and that the model of harmony as spreading a feature from one segment to another is insufficient for capturing the invisibility of segments such as noncontrastive *i*. If strict locality is correct, by contrast, it would mean that the apparent skipping of Finnish *i,e* in harmony is actually inaccurate, and that these vowels are in reality participating in harmony just as fully as any other vowels.

An important discovery of phonetic research is that coarticulation between vowels exists even in languages without vowel harmony (Öhman, 1966; Butcher and Weiher, 1976; Fowler, 1981; Recasens, 1984), such as English, Swedish, and Catalan – and presumably universally. Vowel-to-vowel coarticulation refers to the process through which the phonetic targets of one vowel may be affected by the phonetic targets of a vowel that precedes or follows it. If vowel-to-vowel coarticulation can be found in languages without any vowel harmony to begin with, we should expect to find it in vowel harmonic languages too, even for vowels not undergoing harmony.

If instances of [−back] vowels such as *i* are found to be transiently backed when linearly between [+back]-harmonizing suffix and value-source in words such as *koti-na*, how are we to know whether these effects are the result of a phonetic process of coarticulation, as opposed to the result of the vowel harmony process itself? The answer to this question bears directly on whether harmony is strictly local, or whether backing effects observed on noncontrastive *i* are coarticulatory results independent of the phonology (and equally likely to be found in languages without harmony, as noted above).

Cohn (1990) builds an important synthesis of the distinctions between phonological and phonetic rules, as illustrated in the following table

(9) *Properties of phonological rules vs. phonetic rules (Cohn, 1990, p.26):*

Phonological rules	Phonetic rules
categorial	gradient/quantitative
discrete & timeless segments	continuous in time & space
static effects	segment may vary in quality continuously
full segment affected	part of a segment affected

With these distinctions in mind, we can outline the predictions for Finnish neutral vowels. The acoustic correlate of backness is the second formant, F2 (or $|F2 - F1|$), and so "backness" in what follows refers to these measurable formant values.

(10) *Effects of backing on Finnish i,e in [+back] contexts:*

If vowel harmony	If coarticulation
effect of rule is categorical change in backness	rule may vary with speech style
backness values will cluster into bimodal distribution	vowel backness may show broad unimodal distribution
effects do not depend on properties of preceding vowel	effects may vary in quality continuously
whole vowel affected	only part of the vowel may be affected

Three phonetic studies have looked at effects on Finnish transparent *i,e* in [+back] and [−back] contexts. The results of all three studies squarely place the effects in the righthand column, allowing us to conclude that these vowels are affected by phonetic coarticulation, and not by a phonological process of vowel harmony.

Kim (2005) measured vowels that are both transparent and participating in harmony. For *categorical* harmony, F2 values clustered around discrete values in [+back] and and [−back] environments, and showed a clear separation in their distributions. For gradient coarticulation on *i,e*, on the other hand (of the type observed since Öhman (1966)), there was overlap between F2 values for [+back] and [−back] environments, in other words no categorical separation of distributions. Moreover, Kim found that transparent vowels' F2 values varied directly as function of the preceding vowel's F2, rather than as a function of its binary category. Finally, these lowered F2 values in [+back] contexts did not persist throughout the affected vowel, but rather weakened throughout the duration of the vowel.

Like Kim, Gordon (1999) found that in the case of the transparent vowels, the phonetic difference between the variant occurring after front vowels and the variant occurring after the back vowels is much smaller than the more salient differences found in the contrastive vowels. He conclude that “Vowel harmony for the neutral vowels functions at a low phonetic level, unlike vowel harmony for non-neutral vowels” (p.20).

Välimaa-Blum (1999) compared coarticulatory effects across two speech styles: casual and hyperarticulated (i.e. careful speech). She found a backing effect of neutral vowels in [+back] contexts during casual speech, but no effect at all under the careful speech register. On the other hand, true harmony did not vary with speech style. She concludes for the transparent vowels that “The differences between the formant means disappear in hyperarticulated speech, and this would mean that the assimilation [of *i,e*], if there is one, is style dependent. And if it is style dependent, we propose that it is the result of some late phonetic assimilation rule applying in speech contexts which are less formal or require lesser clarity of presentation”(p.260-261).

One clear formulation of the difference between categorical feature-copying (of the type required in harmony) and phonetic interpolation (resulting from coarticulation) can be found in Keating (1996), who argues that in the phonological component, rules manipulate features and feature-values, and use temporal chunking (parsing into discrete segments), while in the phonetic component, the objects of computation are continuous in time and space, allow overlap, and operate over quantitative values on multiple independent dimensions. Keating (1996, p.264) argues that “Features specify ‘targets’ that the human articulators aim at, moving or ‘interpolating’ from target to target.” Applying the notion of targets to vowel harmony, transparent vowels between harmonizing [+F] vowels do not copy [+F] in the phonology: they are [−F], but may allow for gradient interpolation between [+F] targets when implemented in the phonetics (see also Gick et al. (2006, p.16) for this distinction applied to vowel harmony).

These comparisons can even be made within a single language, as shown by Zsiga (1997), who contrasts Igbo [± ATR] vowel harmony, a categorical and phonological process, with the process of adjacent vowel assimilation (e.g. *nwoke a* → *nwoka a* ‘man DEF’). Zsiga (1997, p.266) argues that “Gradient processes such as Igbo vowel assimilation require reference to the specific temporal information that the gestural approach provides. Categorical alternations such as Igbo vowel harmony neither need nor benefit from detailed timing information. Their similarities are tantalizing, but they cannot be collapsed". Thus, both a module of featurally-based phonological representations and one of inherently quantitative phonetic representation (implemented by Zsiga in terms of articulatory gestural scores) are needed, and neither can be reduced to the other. Igbo vowel harmony is categorical, obligatory, and independent of speech rate and speech style, while Igbo vowel assimilation between adjacent vowels (e.g. across word boundaries) is partial, gradient, variable throughout its time window, highly dependent on the duration of the two vowels, and conditioned by speech rate and speech style.

Thus, as Kim (2005) concludes in her discussion of the truly transparent nature of phonological vowel harmony in Finnish [−round,−low] vowels , "Both vowel harmony and V-to-V coarticulation are independently needed to describe the behavior of vowels. Transparent vowels undergo coarticulation only." Relativized locality is able to compute harmony without involving them.

3.3. What if no Contrastive Source Exists?

When the domain of search for a copying-source for a feature value [±F] includes all potential donors, as was the case in the Turkish examples in Chapter 2, a donor will always be found. However, when the domain of search is relativized to include only contrastive values of [±F] (and to exclude noncontrastive values), then certain configurations will arise in which there is no donor to be copied from. Recall from (2) that the [−back,−round,−low] vowels *i,e* are noncontrastive for the feature [±back]. Since Finnish words can be composed in which noncontrastive vowels are the only vowels, an essive or other needy suffix searching for a value of [± back] will not be able to copy from them.

Nonetheless, a key premise of our model of vowel harmony is that all alternating suffixes come out of the lexicon deficient, and that feature-copying from a sought donor is the way to supply them with a value. When these suffixes are in the unfortunate situation of being concatenated with a root that contains only noncontrastive vowels, they must get a value somehow. As discussed in Section 2.8, a last resort option when harmonic search fails is insertion of the default value. We know already from syntactic theory that in certain configurations in which the verb searches for an argument to agree with, there may be no Nominative argument at all, in which case a last resort is insertion of a set of default feature-values.

(11) *Icelandic Last Resort Feature Insertion when ¬∃ Nominative Argument*:

a. Mig vantar peninga
me-acc lacks-3.masc.sg money-acc (Andrews, 1982)
'I lack money'

Consider cases of a Finnish word with no contrastively [±back] values in the word, such as those in (12).

(12) *Finnish Default Value in case of No Donors is [−back]:*

tie	tie-nä	'road'
veli	velje-nä	'brother'

(13) Essive Suffix Begins Back-Harmonize in *tie-nä*:

$$\begin{matrix} x_1 & x_2 & x_3 & \leftarrow & x_4 & x_5 \\ \begin{bmatrix}\text{-voc}\\ \text{coronal}\\ \text{-cont}\\ \text{-voice}\end{bmatrix} & \begin{bmatrix}\text{+voc}\\ \text{+high}\\ \text{-low}\\ \text{-round}\\ \text{-back}\end{bmatrix} & \begin{bmatrix}\text{+voc}\\ \text{-high}\\ \text{-low}\\ \text{-round}\\ \text{-back}\end{bmatrix} & & \begin{bmatrix}\text{-voc}\\ \text{coronal}\\ \text{+nasal}\end{bmatrix} & \begin{bmatrix}\text{+voc}\\ \text{-high}\\ \text{+low}\\ \text{-round}\\ \\ \end{bmatrix} \end{matrix}$$

(14) Essive Suffix Fails (↚) to Find Contrastive Donor in *tie-nä*:

$$\begin{matrix} \nleftarrow & x_1 & x_2 & x_3 & x_4 & x_5 \\ & \begin{bmatrix}\text{-voc}\\ \text{coronal}\\ \text{-cont}\\ \text{-voice}\end{bmatrix} & \begin{bmatrix}\text{+voc}\\ \text{+high}\\ \text{-low}\\ \text{-round}\\ \text{-back}\end{bmatrix} & \begin{bmatrix}\text{+voc}\\ \text{-high}\\ \text{-low}\\ \text{-round}\\ \text{-back}\end{bmatrix} & \begin{bmatrix}\text{-voc}\\ \text{coronal}\\ \text{+nasal}\end{bmatrix} & \begin{bmatrix}\text{+voc}\\ \text{-high}\\ \text{+low}\\ \text{-round}\\ \\ \end{bmatrix} \end{matrix}$$

(15) Essive Suffix undergoes Last-Resort [−back] Insertion in *tie-nä*:

$$
\begin{array}{ccccc}
x_1 & x_2 & x_3 & x_4 & x_5 \\
\begin{bmatrix} \text{-voc} \\ \text{coronal} \\ \text{-cont} \\ \text{-voice} \end{bmatrix} &
\begin{bmatrix} \text{+voc} \\ \text{+high} \\ \text{-low} \\ \text{-round} \\ \text{-back} \end{bmatrix} &
\begin{bmatrix} \text{+voc} \\ \text{-high} \\ \text{-low} \\ \text{-round} \\ \text{-back} \end{bmatrix} &
\begin{bmatrix} \text{-voc} \\ \text{coronal} \\ \text{+nasal} \end{bmatrix} &
\begin{bmatrix} \text{+voc} \\ \text{-high} \\ \text{+low} \\ \text{-round} \\ \textit{-back} \end{bmatrix} \\
\text{t} & \text{i} & \text{e} & \text{n} & \text{ä}
\end{array}
$$

In the Finnish derivations above, when the search is relativized to exclude noncontrastive values, and the root to the left of the needy suffix *only* contains noncontastive values, in order for the suffix to converge with its requirement of possessing a [± back] value, a last-resort operation inserts the default value in this language, namely [−back].

It is very important to point out, however, that the default value for [± back] in cases in which harmonic search fails need not be identical to what the value of [± back] is in the noncontrastive vowels that failed to act as copy-sources. While in Finnish words such as *tie-nä*, the root is [−back] and the suffix surfaces as [−back], it could have been the other way. Indeed, if *the choice of last-resort value is independent of the vowels that are in the non-copied-from root*, we might expect a language in which the opposite effect from that of Finnish holds.

Uyghur (a Turkic language of Western China) possesses the same vowel inventory as Finnish, as shown in (16), and crucially therefore, the vowels *i,e* are also noncontrastive in Uyghur.

(16) *Uyghur Vowel System*:

[−back, −round]	[−back, +round]	[+back, +round]	[+back, −round]	
i	ü	u		[+high, −low]
e	ö	o		[−high, −low]
ä			a	[−high, +low]

The plural suffix in Uyghur copies from the closest contrastive value, as illustrated in the straightforward cases in (17), which resemble the pattern in Turkish:

(17) *Uyghur Plural Suffix copies from Contrastive [±back]*

yol	yollar	'road'
pul	pullar	'money'
at	atlar	'horse'
köl	köllär	'lake'
yüz	yüzlär	'face'
xät	xätlär	'letter'

However, unlike Turkish, Uyghur harmony is relativized to contrastive feature-values, and as a result, there are cases in Uyghur in which search fails to supply a [±back] value. The last-resort value in Uyghur, however, is [+back], unlike Finnish, *even though both languages possess identical inventories*.

(18) *Uyghur Default Value in case of No Donors is [+back]:*

til	til-lar	'tongue'
deniz	deniz-lar	'sea'

This last-resort value of [+back] in Uyghur is not limited to the Plural suffix; Lindblad (1990) reports this as a general phenomenon and provides examples for the Dative, Locative, and Gerundive suffixes as well.

The conclusion that we must draw from comparing Finnish and Uyghur is that there is indeed a process of last-resort insertion of a default value for the harmonic feature when search fails. The reason search fails in these cases is because it is relativized to only contrastive values for the harmonic feature, and the root to which the alternating affixes have been concatenated happen to be ones which contain no contrastive values for the feature. This scenario is unavoidable as there is no way to prevent an essive or plural suffix from attaching to a root that happens to lack vowels contrastive for [± back], and therefore a "back-up" plan is needed. This back-up plan proceeds independently of what the non-donating root happens to look like.

3.4. Why Contrastiveness?

The importance of contrastiveness in the computation of motor planning and comparison has been established in interesting ways elsewhere in the study of human cognition. Sedivy et al. (1999) made use of the real-time eye-tracking paradigm in an experiment with spoken language and visual

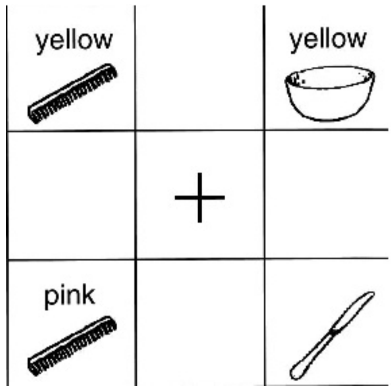

Figure 3.2: Visual setup in Sedivy *et al.* (1999); subjects went for contrastive "yellow" (i.e. looked at the comb).

contexts. Given a scene with a pink comb, a yellow comb, and a yellow bowl (see Figure 3.2), subjects heard instructions such as *Pick up the yellow comb*. Importantly, the use of an eye-tracker allowed Sedivy and her colleagues to measure where subjects looked in real-time, as they heard each incoming word. Sedivy *et al.* found that at the onset of the word *yellow*, subjects looked much faster and more frequently at the yellow **comb**, even before they had heard the head noun.

The only logical explanation is that subjects understood that their interlocutor would be more inclined to use the predicate *yellow* when it was *contrastive* for the object to be manipulated. That is, even though the predicate *yellow* was true of both the comb and the bowl, the subjects preferred to interpret it in a contrastive use within 300 milliseconds, only looking to the object of which it was noncontrastively true much later (at 450 milliseconds). These results are supporting evidence that the preferred contrastive use of a predicate (over a set of more than one object for which it is true) may be a guiding principle in human cognition. Nobody would say that the bowl was "underspecified" for *yellow*, and that is why subjects waited longer to look at it; rather, it seems that when performing a search for something with the relevant feature, contrastive uses of that feature simply get priority.

When a non-contrastive vowel is adjacent to a vowel seeking specification, it is skipped, behaving as invisible. In the present theory, the locality of harmony can largely be predicted based on the structure of the inventory. The most direct source of inspiration for the current work is Calabrese (1995), who argues that "underspecification of feature values becomes an idiosyncratic property of individual rules". Let us paraphrase this as follows, for the case of Finnish, in which vowels not contrastive for [± back] are invisible in a search for harmonic values:

(19) Parametric Visibility: A given harmony rule *R* for the feature [F] may parametrize its search domain to only *contrastive* values for F.

There is a different tradition of treating vowels invisible to harmony (see e.g. Steriade (1987); Dresher (2003)), in which vowels that do not participate in harmony for a feature [±F] (like Finnish *i*) literally lack a feature [−F] in the representation. This solution cannot be pursued for the invisibility of Finnish *i* to [±back] harmony because the feature [−back] is independently needed on *i* in order to explain the phenomenon of Finnish assibilation, whereby the coronal stop /t/ becomes a sibilant [s] when immediately preceding *i*.

(20) Finnish Assibilation: /t/ → [s] before [+high, −back] (Kiparsky, 1973)[26]:
  a. halut-a 'to want (infin.)'
  b. halus-i 'wanted (past)'

Hall and Hamann (2006) characterize the assibilation process as a crucial result of the [−back,+high] features on the aerodynamic realization of obstruents. However, if Finnish /i/ literally lacks [−back], assibilation cannot be characterized in these terms, because the conditioning feature would have to be absent from the representation. While depriving /i/ of [−back] does work in making it invisible for harmony, such a representation leaves us puzzled as to why that same vowel should trigger assibilation. A rule ordering solution, in which */i/* lacks [−back], then harmony happens, then *i* becomes back, then assibilation happens, would be a possible account, but an explanation of why harmony comes first for all such processes would be lacking. In fact, in Section 3.12, it is demonstrated that there are instances of invisibility in harmony that cannot be solved by using the underspecification algorithms based on predictability or redundancy. The proper solution, then, is one in which /i/ is fully specified for [−back] *throughout* the phonology, but certain processes (e.g. harmony) are sensitive to what values of a feature are visible for their search domain.

As the phonology of Finnish requires /i/ to have a [−back] feature in its representation in order to describe assibilation, the right solution to the invisibility of /i/ in harmony must be because harmonic search – but not assibilation – is sensitive to contrastiveness. The "transparency" of *i* in Finnish vowel harmony is because the affixes searching for a donor from which to copy have relativized their search so as to exclude noncontrastive values.

A similar argument that non-participating vowels must still be specified for the feature [−back] comes from Votic, in which *i* is noncontrastive and transparent (e.g. *ko:kkima* 'to dig', *kalli-łła* 'dear-allative'). Blumenfeld and Toivonen (2009) show that while [+back] harmony that skips *i* might suggest that *i* has no [−back] feature, a second process of *l*-fronting before front vowels requires that it does, cf. *łuzikka* 'spoon' vs. *lidna:* 'town'. Like Finnish assibilation, then, a non-harmonic process in Votic requires that *i is* [−back], ruling out models in which it is invisible due to lack of specification.[27] Instead, the harmony process is relativized to *contrastive* values of [−back], thereby excluding *i* from the domain of search, while the adjacent consonant-vowel phonotactics are sensitive to all values of [−back].

There is a second argument that invisibility of noncontrastive vowels to harmony must be due to the nature of the search algorithm itself rather than inherent properties of the vowels, based on a similar phenomenon in Hungarian. Like Finnish, Hungarian has [± back] harmony in which /i/ is not contrastive for [± back]:

(21) *Hungarian vowel inventory (where '(:)' indicates a long version is also possible)*:[28]

[-back,-round]	[-back,+round]	[+back,+round]	[+back,-round]	
i(:)	ü(:)	u(:)		[+high]
e(:)	ö(:)	o(:)	a(:)	[-high]

By the definition of contrastiveness established above, /i/ is also noncontrastive in Hungarian, and due to relativization to contrastive sources only, thereby transparent in [± back] harmony, in a manner parallel to Finnish:

(22) *Hungarian Dative Suffix Harmony: -nak/-nek*:

a. bik**á**-nak 'bull'
b. koszt**ü**m-nek 'costume'

c. k**a**vics-nak 'pebble'
d. **a**kti:v-nak 'active'

The invisibility of *i* in (22-c-d) appears to be evidence that, as in Finnish, this vowel is transparent to harmony. However, Farkas and Beddor (1987) and Ringen and Kontra (1989) discovered that when there is *more than one noncontrastive vowel in a row*, speakers exhibit variation in what they produce. While a one-*i* word such as *akti:v* always takes a [+back] suffix, a two-*i* word such as *aszpirin* shows variation between *-nek* and *-nak*.

(23) *Hungarian "vacillating" stems with more than one noncontrastive vowel*:
a. ✓akti:v-nak, *akti:v-nek
b. ✓aszpirin-nak, ✓aszpirin-nek

The fact that one instance of the vowel *i* does not provide a [−back] value for harmony, but two instances seem to result in a [−back] value, suggests that one cannot simply state that *i* has no value for [−back] throughout the language, because if this were the case we should expect no difference between one and two instances of "nothingness". In Chapter 5, we will return to the reason that words with one noncontrastive vowel differ from words with two noncontrastive vowels, in the context of the role of prosodic constituents in "bounding" the extent of a harmony domain. At this point in the discussion, we can infer that Hungarian *i* must have a representation beyond total invisibility in order for it to rear its head in (23-b). Taking the arguments from Finnish assibilation and Hungarian multiple-noncontrastive-vowel words together, one can see that the correct representation of noncontrastive vowel such as *i* must include the feature [−back], since we can find evidence for its phonological activity outside of the more elementary cases of harmony.

The existence of transparent vowels in harmony systems does not require revision of the fundamental search algorithm, but merely relativization of its domain. Much of the observed crosslinguistic variation in grammatical processes results not from variation in the core computational procedure employed, but rather from the set of elements it operates over. Finnish vowel harmony instantiates a clear case in which the basic algorithm for [± back] harmony, described for Turkish in Chapter 2, remains the same, but whose effects are different once we move to a language with a less symmetric vowel inventory. In particular, Turkish has the [+back,−round, +high] vowel /ɨ/, while Finnish does not, and this makes all the difference.

The discussion above and in Chapter 2 would seem to suggest that Turkish and Finnish differ in that the former has a process of copying [± back] from the closest vowel while the latter has a process of copying [± back] from the closest *contrastive* vowel. Since all vowels in Turkish are contrastive for [± back], the Turkish case is ambiguous between relativization to only-contrastive vowels or not. Based on a number of considerations, we will conclude that the two languages actually both have relativization to only segments contrastive for [± back].

We have already seen evidence that Turkish suffixes sometimes copy their value from a consonant, as in the case of palatalized liquids. This is precisely because these segments are contrastive for [± back] as well. If indeed the relativization of [± back] harmony for Turkish (and throughout the Turkic language family at large) is to contrastive values of [± back], we make two predictions: (1) any consonant in Turkic that is contrastive for [± back] will be included in the search domain, and (2) any vowel in Turkic that is not contrastive for [± back] will not be in the search domain. We now turn to a Turkic language with a very unusual segmental inventory that confirms both predictions and in doing so, provides welcome support for the central role of contrastiveness in delimiting the search domain.

3.5. Contrastiveness Determines Karaim (Consonant) Harmony

Karaim (Kowalski, 1929; Hamp, 1976; Csató and Nathan, 2002) is an endangered Turkic language (spoken in Lithuania) similar to Turkish in its formulation of the harmony process but, due to its consonantal inventory, radically different in terms of the surface effects of harmony. Geographically, Karaim is an outlier language within the Turkic family, and has been embedded in a Baltic- and Slavic-speaking environment and subject to extensive language contact persisting over many centuries. In the late 14th century, a sizeable population of Karaim were relocated from their original homeland on the Crimean peninsula to the town of Trakai (near Vilnius, Lithuania). An important facilitating role for the development of *consonant* harmony and loss of vowel harmony can be attributed to prolonged contact with neighboring Slavic and Baltic languages, where palatalization is contrastive on consonants, while [±back] is not independently contrastive in vowels. The diachronic scenario is discussed further in Nevins and Vaux (2003) and Hansson (2007b).

In Karaim, all consonants (except glides) participate in harmony, because they are contrastive for [± back], and all non-initial vowels don't participate, because they're not contrastive for [± back]. This leads to the interesting phenomeon of "consonant harmony". Karaim suffixal consonants seeking a harmonic value for [± back] can skip right over vowels and copy their value from another consonant.

The transparency of Karaim vowels for [± back] harmony in non-initial syllables may be understood as another instantiation of search-domains that exclude non-contrastive segments. In case of the Karaim, this transparency must be relativized to positional contrast, since it is not the case that the feature [± back] is absolutely non-contrastive in its feature-set on certain vowels, but rather, only non-contrastive on those vowels *in certain positions*. Like Turkish, Karaim features the inventory of eight surface vowels in (24):

(24) *Karaim Vowel Inventory*

	-back		+back	
	-round	+round	-round	+round
+high	i	<ü>	<ɯ>	u
−high	<e>	<ö>	ɑ	o

Unlike Turkish however, the bracketed vowels in (24) are distributionally restricted. In non-initial syllables, there are no contrasts between the non-high non-round vowels ɑ/e or between the round vowels ü,u and ö,o, and there is a limited contrast between the high non-round vowels i/ɯ. Let us turn to the details of these positional restrictions on contrast.

The vowel [e] is only found in initial syllables; elsewhere only [ɑ] occurs. The contrast in initial syllables is observed in k^{hj}el^{j} 'to come' vs. k^{h}ɑl^{ɣ} 'to remain'. Lack of the contrast in non-initial syllables can be seen in forms such as Karaim el^{j}-d^{j}ɑn^{j}'hand.abl' (cf. Turkish *el-den*); only [ɑ] is found, despite the [−back] harmonic value of the root.

The contrasts among the [+round] vowels are confined to absolute word-initial position (i.e. to vowel-initial words): one finds öz^{j} 'self' vs. on 'ten' , as well as üs^{j}t^{j} 'top, upper' vs. us 'reason, intellect' (cf. Turkish *üst, us*). Neutralization of the contrast to [+back] *u,o* is observed in all other environments, e.g. tʰʲuzʲ ‘smooth’ and tʰuz ‘salt’ (cf. Turkish *düz, tuz*).

In Karaim, therefore, vowels with [± back] contrasts among non-high vowels and round vowels occur *only in initial syllables*. The high unrounded pair i/ɯ, as shown in (25) for the genitive -nɯn/-nʲinʲ, does contrast in non-initial syllables (though only [i] is allowed in absolute-initial position, e.g. inʲa ‘needle’ and irlˠa- ‘to sing’ (cf. Turkish iːne, ɨrla)).

In suffixes such as -nɯn/-nʲinʲ, the initial consonant [n/nʲ] of the suffix copies the closest leftward [± back] value, and subsequent segments in turn copy the value of [± back] from it.

(25) *Complete Harmony for [±back] genitive suffixes in Karaim*:

stem	genitive		
tʰɑv	tʰɑv-nɯn	mountain	(cf. Turkish dağ-ɨn)
elʲ	elʲ-nʲinʲ	hand	(cf. Turkish el-nin)

Whereas with suffixes that contain non-high or round vowels, Karaim shares with other Turkic languages the property of [± back] harmony, it differs insofar as the harmonic feature surfaces on consonants rather than vowels. Representative alternations for the plural -lʲɑrʲ/-lɑr and the ablative -tʰʲɑnʲ/-tɑn are shown in (26) and (27).

(26) *Consonantal [± back] harmony for plural and ablative suffix in Karaim*:

stem	ablative		
suv	suv-dɑn	water	(cf. Tk. su-dan)
tʰɑʃ	tʰɑʃ-tʰɑn	stone	(cf. Tk. taş-tan)
mʲenʲ	mʲenʲ-dʲɑnʲ	I	(cf. Tk. ben-den)
kʰʲunʲ	kʰʲunʲ-dʲɑnʲ	day	(cf. Tk. gün-den)
kʰʲunʲ	kʰʲunʲ-lʲɑrʲ-dʲɑnʲ	day-pl.-abl	(cf. Tk. gün-ler-den)
kʰun	kʰun-lɑr-dɑn	servant.pl-abl	(cf. Tk kul-lar-dan)

In addition, Karaim exhibits forms that demonstrate multiple copying of [± back] by one suffix from another, such as that in (27) (Kowalski 1929:69):

(27) *Iteration of Back Harmony through multiple suffixes in Karaim*:
tʰoɾʲɑ-sʲizʲ-lʲigʲ-imʲ-dʲɑnʲ ‘from my injustice’

Notably, the palatal glide /j/ does not group with the palatalized consonants, and can occur in palatalized and non-palatalized contexts alike, as the examples in (28) illustrate.

(28) *Non-Participation of [−back] Glide in Karaim Consonant [+back] Harmony*:

	Karaim		Turkish cognate
a.	jol-daʃ-ɯm	'my fellow traveller'	jol-daʃ-ɨm
b.	koj-maχ	'placing, putting'	koj-mak

Since glides are not contrastive for [± back] (as there is no [−consonantal,−vocalic, +high, −round, +**back**] segment in the Karaim inventory), the fact that the glide does not participate in harmony does not have to be stipulated: glides are noncontrastive in the inventory and are thereby excluded from the search domain.

In the examples above, non-initial [+back] vowels /a,o,u/ are transparent to consonant harmony even though they have [−back] counterparts that are allowed to surface in initial syllables. Supporting evidence for the phonetic transparency of these non-initial vowels is provided in Nevins and Vaux (2003), where spectrographic analysis reveals that the F2 (the primary acoustic correlate of backness) is high for the palatalized consonants and dips low for the intervening back vowels. In other words, the intervening noncontrastive [+back] vowels phonetically "interrupt" the [−back] harmony across them, but phonologically are irrelevant for harmony.

In summary, Karaim exhibits the following positional restrictions:

(29)
a. In Karaim, [−back] is banned from cooccurring with [+round] outside of absolute initial position (i.e. onsetless initial syllables).
b. [−back] is banned from cooccuring with [−high] outside of initial syllables (e.g. initial syllables, onsetless or not).
c. [+back] is banned from cooccuring with [+high,−round] in absolute initial position (i.e. onsetless initial syllables).

There are many languages that exhibit restrictions on the distribution of segments or features in non-initial positions. Karaim is thus not unique in this respect; compare for example Tamil, which systematically bans mid vowels from non-initial syllables (Christdas, 1988). Given the positionally-restricted distribution of these vowels in Karaim's harmony system, the relativization to contrastive status for [±back] is not only a function of the inventory, but also a function of the position within a word:

(30) Positional Contrastiveness: A segment S in position P is **contrastive** for the feature F iff ∃ a segment S in the inventory that is featurally identical to S for all values except F, **and** S′ can occur in position P as well.

(31) *Karaim Genitive, Ablative, Plural suffix must:*
Back-Harmonize: δ = left, F = [positionally contrastive: $\pm$ back]

The effect of positional contrativeness is illustrated for k^{hj}un^{j}-l^{j}ɑr^{j} below; the initial consonant of the plural suffix finds the closest positionally contrastive value ([$-$back] in this case), and the final consonant of the plural suffix subsequently copies [$-$back] from it, across the noncontrastive [$+$back] vowel:

(32) Plural Suffix Begins Positionally-Contrastive Back-Harmonize in k^{hj}un^{j}-l^{j}ɑr^{j}

$$\begin{array}{cccccc} x_1 & x_2 & x_3 \quad \leftarrow & x_4 & x_5 & x_6 \\ \begin{bmatrix} \text{-voc} \\ \text{dorsal} \\ \text{-continuant} \\ \text{-voice} \\ \text{-back} \end{bmatrix} & \begin{bmatrix} \text{+voc} \\ \text{+high} \\ \text{+round} \\ \text{+back} \end{bmatrix} & \begin{bmatrix} \text{-voc} \\ \text{coronal} \\ \text{+nasal} \\ \text{-back} \end{bmatrix} & \begin{bmatrix} \text{-voc} \\ \text{coronal} \\ \text{+lateral} \\ \text{+sonorant} \\ \text{-nasal} \\ \end{bmatrix} & \begin{bmatrix} \text{+voc} \\ \text{-high} \\ \text{-round} \\ \text{+back} \end{bmatrix} & \begin{bmatrix} \text{-voc} \\ \text{coronal} \\ \text{+lateral} \\ \text{-liquid} \end{bmatrix} \end{array}$$

(33) Initial Consonant of Plural Suffix copies [$-$back] in k^{hj}un^{j}-l^{j}ɑr^{j}

$$\begin{array}{cccccc} x_1 & x_2 & x_3 \quad \Lsh & x_4 & x_5 & x_6 \\ \begin{bmatrix} \text{-voc} \\ \text{dorsal} \\ \text{-continuant} \\ \text{-voice} \\ \text{-back} \end{bmatrix} & \begin{bmatrix} \text{+voc} \\ \text{+high} \\ \text{+round} \\ \text{+back} \end{bmatrix} & \begin{bmatrix} \text{-voc} \\ \text{coronal} \\ \text{+nasal} \\ \textbf{-back} \end{bmatrix} & \begin{bmatrix} \text{-voc} \\ \text{coronal} \\ \text{+lateral} \\ \text{+sonorant} \\ \text{-nasal} \\ \textbf{-back} \end{bmatrix} & \begin{bmatrix} \text{+voc} \\ \text{-high} \\ \text{-round} \\ \text{+back} \end{bmatrix} & \begin{bmatrix} \text{-voc} \\ \text{coronal} \\ \text{+lateral} \\ \text{-liquid} \end{bmatrix} \end{array}$$

(34) Final Consonant of Plural Suffix begins Positionally Contrastive Back Harmonize in k^{hj}un^{j}-l^{j}ɑr^{j}

$$
\begin{array}{cccccccc}
x_1 & x_2 & x_3 & x_4 & x_5 & \leftarrow & x_6 \\
\begin{bmatrix} \text{-voc} \\ \text{dorsal} \\ \text{-continuant} \\ \text{-voice} \\ \text{-back} \end{bmatrix} &
\begin{bmatrix} \text{+voc} \\ \text{+high} \\ \text{+round} \\ \text{+back} \end{bmatrix} &
\begin{bmatrix} \text{-voc} \\ \text{coronal} \\ \text{+nasal} \\ \text{-back} \end{bmatrix} &
\begin{bmatrix} \text{-voc} \\ \text{coronal} \\ \text{+lateral} \\ \text{+sonorant} \\ \text{-nasal} \\ \text{-back} \end{bmatrix} &
\begin{bmatrix} \text{+voc} \\ \text{-high} \\ \text{-round} \\ \text{+back} \end{bmatrix} &
 &
\begin{bmatrix} \text{-voc} \\ \text{coronal} \\ \text{+lateral} \\ \text{-liquid} \end{bmatrix}
\end{array}
$$

(35) Final Consonant of Plural Suffix copies [−back]

$$
\begin{array}{cccccccc}
x_1 & x_2 & x_3 & x_4 & \Lsh & x_5 & x_6 \\
\begin{bmatrix} \text{-voc} \\ \text{dorsal} \\ \text{-continuant} \\ \text{-voice} \\ \text{-back} \end{bmatrix} &
\begin{bmatrix} \text{+voc} \\ \text{+high} \\ \text{+round} \\ \text{+back} \end{bmatrix} &
\begin{bmatrix} \text{-voc} \\ \text{coronal} \\ \text{+nasal} \\ \text{-back} \end{bmatrix} &
\begin{bmatrix} \text{-voc} \\ \text{coronal} \\ \text{+lateral} \\ \text{+sonorant} \\ \text{-nasal} \\ \textbf{-back} \end{bmatrix} &
 &
\begin{bmatrix} \text{+voc} \\ \text{-high} \\ \text{-round} \\ \text{+back} \end{bmatrix} &
\begin{bmatrix} \text{-voc} \\ \text{coronal} \\ \text{+lateral} \\ \text{-liquid} \\ \textbf{-back} \end{bmatrix} \\
\text{k}^{\text{hj}} & \text{u} & \text{n}^{\text{j}} & \text{l}^{\text{j}} & & \text{ɑ} & \text{r}^{\text{j}}
\end{array}
$$

Karaim's unusual system of [± back] harmony that skips noncontrastive vowels and includes contrastive consonants enables us to draw an important conclusion. Karaim constitutes a system with an identical harmony rule to that of Finnish (and indeed, Turkish as well), in which [± back] harmony is relativized to only copy from contrastively [± back] donors. However, due to the radically different segmental inventory of Karaim, this single principled search procedure yields surprising effects. As the paradigmatic component of language (in this case, the basic "lexicon" of consonants and vowels) is the locus of idiosyncratic patterning due to arbitrary historical effects (in the case of Karaim, these historical effects included intense language contact between Turkic and Balto-Slavic), this is the component of linguistic structure that we expect to yield variation in syntagmatic patterning in turn. Importantly, however, all three languages, Finnish, Turkish, and Karaim, can be argued to follow the exact same algorithm: suffixal segments in need of a value for [± back] copy it from the closest source which is positionally-contrastive for that value.[29]

We have thus far considered how the search-and-copy algorithm for harmony may exhibit crosslinguistic variation as the result of the search domain itself being pruned of certain segments, specifically the noncontrastive ones. In addition to contrastiveness, the second important inventory-based property of segments bearing a feature [$\pm$F] is whether [α F] is the marked or unmarked value of the feature. Finnish and Karaim have been demonstrated to show "transparency" effects for segments that are not contrastive for [$\pm$F]; we now turn to cases where segments act as transparent because they do not have the marked value for [$\pm$F].

3.6. Markedness in Binary Feature Systems

Trubetzkoy (1931) was one of the first to pioneer the idea that in certain binary oppositions, one member of the opposition is treated by grammatical processes asymmetrically differently from the other. In the binary opposition encoded featurally by [$\pm$ voice], grammar does not treat both values equally, instead seeming to "prefer" [$-$voice] over [$+$voice], in the sense that inventories are more often built out of [$-$voice] segments than out of [$+$voice] segments. In addition, positionally-restricted contrasts neutralize the opposition to [$-$voice], as found in the widespread process of Coda Devoicing of German, Turkish, Russian, among many others. The value of a binary feature that is asymmetrically dispreferred in paradigmatic arrangements and in syntagmatic rules of neutralization is referred to as the *marked* value of the feature, notated here as [mF], as opposed to [uF], the unmarked value.

Jakobson (1941) developed psycholinguistic correlates of the theory of markedness, focusing on the asymmetric treatment of marked features in language acquisition and aphasia, in both of which marked features are produced less frequently than their unmarked counterparts. Greenberg (1963) advanced the notion of markedness in his crosslinguistic work on universals of language inventories, proposing the notion of an implicational relation between the marked and unmarked value of a feature: if a language allows mF in a feature bundle B, it will also allow uF in B. For example, languages that possess voiced fricatives will also possess voiceless fricatives. Both strands of research converged on the concluson that language change favors loss of the marked member of a binary featural opposition rather than the converse.

Following the implicational treatment of markedness by Greenberg and the neutralizational treatment of markedness of Trubetzkoy, markedness enjoyed continued interest in the generative tradition, launched by Chomsky and Halle (1968, ch.9).[30] Phonological research in the last twenty years has continued to focus on markedness values of particular features in predicting implicational relations among vowel inventories, e.g. Calabrese (1988), Archangeli and Pulleyblank (1994), and in predicting patterns of contrast-neutralization to the unmarked feature. de Lacy (2006), in a thorough study of syntagmatic processes of neutralization and assimilation, concludes that phonological processes (encoded as constraints in his formulation) may refer to [±F] or to mF, but never solely to uF. In other words, a phonological process banning the occurrence of a feature in a certain position may refer to both values of a given feature, or to only the marked value of a feature, but never to the unmarked feature alone.

The treatment of included segments in the search domain of a vowel harmony process follows the same logic. A vowel harmony rule may be formulated so as to allow copying from either value of the harmonic feature (e.g. either value of [± low]) in a two-height system) or from the marked value alone (e.g. [+low] alone), but never from the unmarked value alone.

3.7. Sibe Harmony and Context-Free Markedness Value of [+low]

Sibe (pronounced [ʃi-be]) is a Tungusic language of Western China, described by Li (1996), which demonstrates context-free markedness in long-distance harmony. The vowel inventory of Sibe is composed of three binary distinctions: [± back], [± round], and [± low], resulting in eight vowels.

(36) *Sibe Vowel Inventory*:

	[-back, +rd]	[-back, -rd]	[+back, -rd]	+[back, +rd]
[−low]	ü	i	ɨ	u
[+low]	ö	ɛ	a	ɔ

The phenomenon of long-distance harmony in Sibe occurs in a feature-copying process where a suffixal consonant seeks a value from a leftward vowel for [±low]. The search for a valuing-donor can potentially skip intervening vowels, even ones contrastive for this feature. The transparency of all vowels except those that bear the marked feature value for [± low] yields a long-distance harmonic copying between a suffixal velar consonant and a preceding low vowel. Before illustrating this process, we turn to a discussion of the markedness of this feature-value in Sibe.

3.7.1. The Markedness of [+low] in Sibe

I assume that certain feature-values may be subject to language-specific markedness statements, and that while [+ low] may be naturally unmarked, as in Jakobson (1941), it is *logically* marked in Sibe, as determined by the observation and analysis of language-specific phonotactics.[31] We can demonstrate that [+low] is the marked value in Sibe from three sources of distributional evidence, all of which point to the asymmetric grammatical treatment of this feature value as opposed to its unmarked counterpart. The first source of evidence comes from co-occurrence restrictions: Li (1996), based on observations over the Sibe lexicon, states that "Sibe has a restriction on low vowel co-occurrence within the domain of a phonological word." (p. 203). Li's observation is that, by and large, the appearance of multiple low vowels in a word is disfavored. The observation that a marked feature value may occur in a word, but only to a minimal degree, is found also in co-occurrence biases in the lexicon of Japanese, in which more than one instance of marked [+voice] in a word is dispreferred (Ito and Mester, 2003). This often-used diagnostic for markedness points towards a co-occurrence bias in Sibe against multiple instances of the marked value [+low].

A second source of evidence arises from the distribution of vowels in suffixes. Most suffixes in Sibe contain a single high vowel (Li 1996:199-203). The relative distributional bias towards [−low] vowels in suffixes suggests that this is the unmarked value for [± low] in Sibe, especially in light of the recent body of research suggesting that affixes tend to draw from the unmarked pool of segments within a given language. For example, English inflectional affixes only draw from the coronal consonants, Lushootseed has glottalized consonants only in roots and lexical suffixes (Urbanczyk, 1995), and Cuzco Quechua does not have aspirated stops in suffixes (Beckman, 1998). These cases illustrate that affixal inventories are very often reduced in favor of the unmarked values of segmental contrast and support the view that the relative dearth of [+low] vowels in Sibe suffixes is due to their marked status.

In addition, vocalic epenthesis in Sibe is always of a [−low] vowel (either the high back rounded [u] or high back unrounded [ɨ], due to [± round] harmony with the immediately preceding vowel). Examples are shown in (37) and (38).

(37) *Sibe Epenthesis with accusative ending* /-v/:

a. εχ-ɨ-v large.bead-acc.
b. tasχ-ɨ-v tiger-acc.
c. döv-u-v fox-acc.
d. mul-u-v beam-acc.

(38) *Sibe Epenthesis between causative* /-v/ *and present-future tense* /-m/:

a. ömi-v-ɨ-m to drink
b. va-v-ɨ-m to kill
c. ɕöndʑü-v-u-m to elect
d. bu-v-u-m to give

Finally, diachronic evidence also points to the conclusion that low vowels are marked in Sibe. Zhang (1996) provides a comparison of Sibe with its predecessor, Classical Manchu. Classical Manchu had stress on the final syllable, while Sibe has moved stress to the initial syllable, likely under the influence of Mandarin Chinese (Zhang 1996:151). Importantly, Sibe has raised Classical Manchu /a/ to /ɨ/, and has raised Manchu /o/ to /u/. The raising of [+low] vowels under loss of stress can be directly attributed to the markedness of [+low] in vowels. Stem vowels no longer protected by positional prominence lost their marked [+low] feature in Sibe[32]. The diachronic loss of [+low] on what became unstressed vowels therefore supports the claim that [+low] is marked in Sibe.

The totality of evidence to the language learner from distributional asymmetries and from the quality of suffixes and epenthetic vowels that [+low] is the marked value of the feature – coupled with the hypothesis that one way in which vowel harmony may narrow the domain of search is by looking only at marked values of the feature – leads us to expect that in such a language, unmarked values of the [± low] may be treated as "transparent". As a result, search will continue until a marked value-source is found. This may yield the surface appearance of a long-distance harmony process: just like Finnish *k**o**ti-n**a***, in which the suffix copies from a vowel as many as three syllables away because it is the closest contrastive-value source, Sibe has a process in which a suffixal segment may copy from two syllables away because that will be the closest marked-value source.

3.7.2. Velar/Uvular Alternations as [+low] Harmony

The [+low] harmony process in Sibe takes place in a consonant-vowel interaction in which dorsal consonants (i.e., velar/uvular obstruents) engage in a harmonic search for the value of [±low] to their left and copy this value from a vowel.

Sibe opposes four pairs of velar and uvular consonants: /k/ and /q/, /g/ and ɢ, /x/ and /χ/, and /ɣ/ and /ʁ/. These segments are all dorsal consonants, articulated with the tongue body, but the pairs differ in their specification of [± low]. The [−low] segments /k g x ɣ/ are velar, while the [+low] segments /q ɢ χ ʁ/ are uvular.

Languages that oppose velar and uvular consonants may do so featurally with different means, depending on whether the language independently activates the features [± ATR], [± high], and/or [± low] (see Trigo (1991) for a thorough discussion). In Chomsky and Halle (1968), it was suggested that uvulars are [+back, −high]. Trigo (1991) provides arguments that Turkana and Akha uvularization require [−high], and demonstrates that in these languages, both [+ATR] o and [−ATR] ɔ can yield uvularization. Thus, in some languages, vowel height alone is sufficient for uvularization. In addition, it can be shown that [+back] is not always a necessary feature to induce uvularization. Yakut (Krueger, 1962) illustrates such a case.

Yakut has an eight vowel system, as shown in (39), with [±back] and [±round] harmony. [± round] harmony distinguishes two heights, [± high], where [−high] suffixal vowels alternate between *o* and *a*.

(39) *Yakut vowel inventory*:

[-back, +rd]	[-back, -rd]	[+back, -rd]	+[back, +rd]	
ü	i	ɨ	u	[+high]
ö	e	a	o	[−high]

Yakut dorsal consonants require [−high] as a necessary feature for uvularization. Importantly, both /ö/ and /o/ can trigger uvularization, and therefore [+back] is not necessary for uvularization in this language. Consonantal alternations of this sort in suffixes can be illustrated with the 2nd.pl /-gɪt/ (Krueger, 1962, 89). In (40), both [+back] (40-c-d) and [−back] (40-e-f) harmonic vowels may trigger uvularization, as long as they are [−high].

(40) *Yakut velar/uvular alternations*:

a.	tiː-gɨt	your-pl. boat
b.	kel-li-git	you-pl. came
c.	saː-ʁɨt	your-pl. gun
d.	oʁo-ʁut	your-pl. child
e.	kinige-ʁit	your-pl. book
f.	öŋö-ʁüt	your-pl. service

We may conclude that velar∼uvular alternations induced by vowels may be conditioned by [± low], [± high], or [± ATR], depending on the vowel features already present in the language. Returning to Sibe, we adopt the following featural specifications for its dorsal consonants, given that Sibe does not have [± ATR] or [± high] contrasts among the vowels.

(41) *Sibe Dorsal Consonant Inventory*:

[−voi,−cont]	[+voi,−cont]	[−voi,+cont]	[+voi,+cont]	
k	g	x	ɣ	[−low]
q	G	χ	ʁ	[+low]

The diminutive suffix for adjectives has four variants in Sibe (Li 1996:201).[33] The alternation between a velar ([−low]) and uvular ([+low]) consonant in the suffix is determined by whether there is a preceding [+low] vowel anywhere in the word.[34]

(42) *Sibe diminutive suffix shows marked [+low] harmony*:

a.	ɨldɨ(n)-kɨn	bright
b.	ɕümi(n)-kɨn	deep
c.	muxuli(n)-kɨn	round
d.	udʑi(n)-kɨn	heavy
e.	ɕa(n)-qɨn	good
f.	sula-qɨn	loose
g.	ɢɔlmi(n)-qɨn	long
h.	adʑi(g)-qɨn	small
i.	untuxu(n)-kun	empty
j.	ulu-kun	soft
k.	gɨltu(xun)-kun	severe
l.	ɨrsu(n)-kun	ugly
m.	tɔndɔ-qun	honest
n.	χɔdu(n)-qun	quick
o.	dʐalu-qun	full
p.	farχu(n)-qun	dark

In (42-a-d) and (42-i-l), no [+low] vowel precedes the suffix, and as a result, the initial consonant of the suffix is supplied with default [−low] *k*. In (42-e-f) and (42-m), when the stem vowel that is closest to the suffix is [+low], the suffix surfaces with *q*. The most surprising cases are those in (42-g-h) and (42-n-p), in which the suffix surfaces with *q*, even though the determining [+low] vowel is two syllables away, and a [−low] vowel intervenes.

This velar∼uvular alternation may demonstrate long-distance effects even across intervening suffixes. Hence, when the reciprocal suffix *-ndu* precedes the past tense suffix, the uvular variant of the past tense suffix can be copied from a [+low] vowel three syllables away (43-a). When there are no preceding [+low] vowels in the word, the suffix is realized as a velar (43-b).

(43) *Sibe [+low] harmony separated by three syllables*:

a.	qarɨ-ndu-χu	protect-reciprocal-past
b.	nikɨ-ndu-xu	rely.on-reciprocal-past

Velar/uvular alternations are not limited to the [−continuant] series of dorsal consonants. The suffix of the non-self-perceived immediate past tense shows an alternation between the voiceless velar fricative /x/ and the voiceless uvular fricative /χ/ (Li 1996: 202):

(44) *Long-Distance Uvularization Harmony in Sibe Past Tense Suffix*:

a.	dʑi̯-xɨ	to come
b.	tɨ-xɨ	to sit
c.	iɕi-xɨ	to be enough
d.	gɨnɨ-xɨ	to go
e.	tɨsu-xu	to satisfy
f.	türü-xu	to rent
g.	utu-xu	to dress
h.	xinu-xu	to hate
i.	tükɛ-χɨ	to watch
j.	sav-χɨ	to see
k.	fɔndʑi̯-χɨ	to ask
l.	ömi-χɨ	to drink
m.	tɔ-χu	to curse
n.	gö-χu	to hit (the target)
o.	bɔdu-χu	to consider
p.	lavdu-χu	to become more

The Sibe data involving velar/uvular suffixes copying [+low] long-distance are of special interest to theories of the locality of phonological processes. In particular, a [+low] consonant copying across an intervening [−low] vowel exemplifies a case in which even *contrastive* segments are transparent to a harmony process.

The theory of locality that best explains the Sibe data is one that emphasizes the exclusion of certain intervening segments from the domain of search. In an intervener-based theory of locality, a featural relation (such as harmony) may take place between any two segments within the word domain, as long as no segment *of the relevant type* intervenes. The notion of intervention in this approach depends not only on the presence of a feature but its values as well. Recall the main concept of this chapter, that each harmony process may parametrically vary *in its sensitivity to the values* of intervening segments.

Feature-sensitive relativization permits us to formulate the locality conditions for Sibe uvularization: the suffixes in question seek the marked value [+low] anywhere in the word.[35] This process is illustrated for the case of long-distance harmony of the initial consonant of the diminutive suffix in χɔdu-qun 'quick-dim.':

(45) Diminutive Suffix begins Marked [+low] Harmonize in χɔdu-qun

$$
\begin{array}{ccccccc}
x_1 & x_2 & x_3 & \leftarrow & x_4 & x_5 & x_6 \\
\begin{bmatrix} \text{+voc} \\ \text{+round} \\ \text{+back} \\ \text{+low} \end{bmatrix} &
\begin{bmatrix} \text{-voc} \\ \text{coronal} \\ \text{-cont} \\ \text{+voice} \end{bmatrix} &
\begin{bmatrix} \text{+voc} \\ \text{+round} \\ \text{+back} \\ \text{-low} \end{bmatrix} & &
\begin{bmatrix} \text{-voc} \\ \text{dorsal} \\ \text{-cont} \\ \text{-voice} \\ \; \end{bmatrix} &
\begin{bmatrix} \text{+voc} \\ \text{+back} \\ \text{-low} \end{bmatrix} &
\begin{bmatrix} \text{-voc} \\ \text{coronal} \\ \text{+nasal} \end{bmatrix}
\end{array}
$$

(46) Diminutive Suffix Finds and Copies Marked [+low]

$$
\begin{array}{ccccccc}
x_1 & \Lsh & x_2 & x_3 & x_4 & x_5 & x_6 \\
\begin{bmatrix} \text{+voc} \\ \text{+round} \\ \text{+back} \\ \text{+low} \end{bmatrix} & &
\begin{bmatrix} \text{-voc} \\ \text{coronal} \\ \text{-cont} \\ \text{+voice} \end{bmatrix} &
\begin{bmatrix} \text{+voc} \\ \text{+round} \\ \text{+back} \\ \text{-low} \end{bmatrix} &
\begin{bmatrix} \text{-voc} \\ \text{dorsal} \\ \text{-cont} \\ \text{-voice} \\ \textbf{+low} \end{bmatrix} &
\begin{bmatrix} \text{+voc} \\ \text{+back} \\ \text{-low} \end{bmatrix} &
\begin{bmatrix} \text{-voc} \\ \text{coronal} \\ \text{+nasal} \end{bmatrix} \\
\text{ɔ} & & \text{d} & \text{u} & \text{q} & \text{u} & \text{n}
\end{array}
$$

When there is no marked value of [+low] anywhere leftward in the word, the search fails to find a copy-source, and must as a last resort insert the default value of [−low] on the dorsal consonant, as illustrated below for ulu-kun 'soft-dim.'.

(47) Diminutive Suffix Searches for Marked [+low] in ulu-kun

$$
\begin{array}{ccccccc}
x_1 & x_2 & x_3 & \leftarrow & x_4 & x_5 & x_6 \\
\begin{bmatrix} \text{+voc} \\ \text{+round} \\ \text{+back} \\ \text{-low} \end{bmatrix} &
\begin{bmatrix} \text{-voc} \\ \text{coronal} \\ \text{+lateral} \end{bmatrix} &
\begin{bmatrix} \text{+voc} \\ \text{+round} \\ \text{+back} \\ \text{-low} \end{bmatrix} & &
\begin{bmatrix} \text{-voc} \\ \text{dorsal} \\ \text{-cont} \\ \text{-voice} \\ \; \end{bmatrix} &
\begin{bmatrix} \text{+voc} \\ \text{+back} \\ \text{-low} \end{bmatrix} &
\begin{bmatrix} \text{-voc} \\ \text{coronal} \\ \text{+nasal} \end{bmatrix}
\end{array}
$$

(48) Diminutive Suffix Fails to Find Marked [+low]; Last Resort Insertion:

$$
\nleftarrow \quad \begin{array}{cccccc}
x_1 & x_2 & x_3 & x_4 & x_5 & x_6 \\
\begin{bmatrix} \text{+voc} \\ \text{+round} \\ \text{+back} \\ \text{-low} \end{bmatrix} &
\begin{bmatrix} \text{-voc} \\ \text{coronal} \\ \text{+lateral} \end{bmatrix} &
\begin{bmatrix} \text{+voc} \\ \text{+round} \\ \text{+back} \\ \text{-low} \end{bmatrix} &
\begin{bmatrix} \text{-voc} \\ \text{dorsal} \\ \text{-cont} \\ \text{-voice} \\ \textbf{-low} \end{bmatrix} &
\begin{bmatrix} \text{+voc} \\ \text{+back} \\ \text{-low} \end{bmatrix} &
\begin{bmatrix} \text{-voc} \\ \text{coronal} \\ \text{+nasal} \end{bmatrix}
\end{array}
$$

As soon as the Sibe learner has determined that velar$\sim$uvular alternations are based on relativization to the marked value of [+low], the locality properties of the harmony process will follow: the search domain excludes all instances of unmarked values of [$\pm$ low], even if they are contrastive, and proceeds leftward, potentially unboundedly. Like Finnish and Uyghur roots that contain no instances of contrastive vowels and thus require last-resort insertion in their suffixes, Sibe roots that contain no instances of marked vowels require last-resort insertion in their suffixes.

It is worth emphasizing again that crosslinguistic variation in "transparency" – the invisibility of intervening segments to the syntagmatic process of harmony – is not the result of a fundamentally different Search procedure, but rather one small tweak in the feature-values included or excluded in the relativization of the domain to search. These "tweaks" form a parametric space in which the Search procedure may be relativized to *all* values [$\pm$F], contrastive-only values [cF], or marked-only values [mF]. In all of the case studies we have considered, the only ways in which the search procedure itself may vary (in addition to the inventory of segments in the language and the inventory of affixes that require harmony) is in the direction of search δ and the delimitation of the search domain by value-type τ: [$\pm$F, cF, mF].

Under this type of Principles-and-Parameters approach to crosslinguistic variation in harmony systems, in which the fundamental *copy-from-closest* principle of Search is invariant and a small number of parameters restrict the variation in the domain of Search itself, we may well expect to find "microvariation" between closely related languages showing identical harmony processes save for the setting of one of these parameters. A good example is found in the uvularization process of Sanjiazi Manchu, a sister language of Sibe, in which relativization is to [cF] and not to [mF], with ensuing differences in the resulting pattern of harmony.

3.7.3. Contrastive-Value Harmony in Sanjiazi Manchu

This emerging system of parametric differences in value-relativization predicts the possibility of a related language, completely analogous to Sibe, but with access to contrastive values of [± low] determining the locality conditions of V-C uvularization. As another descendent of Classical Manchu, Sanjiazi Manchu is one of the closest languages to Sibe. Li (1996) examines the dialect spoken in Sanjiazi, in the western part of Heilongjiang province. Sanjiazi Manchu has an identical consonant inventory to Sibe and an almost identical vowel inventory:

(49) *Sanjiazi Manchu vowel inventory*:

[−back, −rd]	[−back, +rd]	[+back, −rd]	[+back, +rd]	
i	ü	ɨ	u	[-low]
æ		a	ɔ	[+low]

Like Sibe, Sanjiazi Manchu shows velar∼uvular alternations in the dorsal consonants of the past tense suffix -xa / -xɨ / -xu / -xɔ, with rounding harmony determining the quality of the suffix vowel. However, unlike Sibe, the visibility of intervening feature values in Sanjiazi Manchu is sensitive to all *contrastive* values for [±low].

(50) *Sanjiazi Manchu alternations in the past tense (Li 1996:182)*:

a. qa-χa to obstruct
b. mɨla-χa to roar
c. ʂudʐa-χa to rely on
d. sæ-χa to bite
e. ɔm-χɔ to drink
f. davɨ-xɨ to stride
g. ɨldɨ-xɨ to shine
h. dazi-xɨ to repair
i. tæri-xɨ to plant
j. sü-xu to mix
k. dɔndʐi-xɨ to listen
l. matʃu-xu to grow thinner

As (50-a-e) show, when the contrastive value for [±low] nearest to the dorsal consonant of the suffix is a [+low] vowel, the suffix surfaces with a [+low], uvular [χ], whereas when the closest contrastive value is [−low], the suffix surfaces with [−low] velar [x]. Thus, no long-distance copying of [+low] across an intervening [−low] segment can occur in Sanjiazi Manchu, because the intervening values of [−low] are contrastive[36].

The distinct parametrization of intervener visibility of Sanjiazi Manchu and Sibe is apparent when comparing Sanjiazi Manchu [dɔndʐ̩-xɨ], in which contrastive [−low] is visible and yields a velar alternant in the suffix, with near-minimal Sibe [fɔndʐ̩-χɨ] (*'to ask'*), where only marked [+low] is visible, yielding a uvular alternant. Sanjiazi Manchu uvularization is thus formally identical to Sibe, with the difference of the parametric visibility of the specified features resulting in different harmony behavior.

One might ask whether the differing parametric setting of Sanjiazi Manchu, in having contrastive-only, rather than marked-only harmony for [± low] correlates with any other aspect of the language. It is suggestive to notice that unlike Sibe, Sanjiazi Manchu has a process of contrastive [± low] harmony affecting suffixal vowels as well as the suffix-initial dorsal consonant: compare Sibe past tense -xi/χi with Sanjiazi Manchu -xɨ/-xu/-xa/-xɔ/-χɨ/-χu/-χa/-χɔ. In Sanjiazi Manchu, the suffixal vowel alternates according to [± low] harmony (as well as [± round] harmony). While there is no formal mechanism within the present theory to link the presence of featural harmony for [±F] on the nucleus of an affix to the pattern of featural harmony for [± F] on a separate consonant of that affix (and indeed, the theory should be flexible enough to allow independent harmonic needs on distinct segments of a suffix, as discussed for Turkish progressive *-ijor* in Section 2.6), the existence of [cF] harmony among Sanjiazi Manchu suffixal vowels may bias the learner to adopt a uniform parametric setting for the suffixal consonants.

To conclude, the [± low] harmony values sought by dorsal consonants in two closely related languages, Sanjiazi Manchu and Sibe, differ only in which paradigmatic property they are sensitive to: contrastiveness or markedness. In the following section, we examine other harmony systems that set their relativization to marked values.

3.8. Context-Sensitive Markedness and Sources of [+round]

The distributional evidence for the marked status of [+low] in Sibe in Section 3.7.1 pointed to the conclusion that, along this binary opposition, one value of the feature was clearly dispreferentially treated by grammatical processes. However, it is not always the case that the marked value of a binary feature can be inspected in purely context-free terms. In particular, for many binary features, their marked value is really only marked within the context of other features. For example, even [+voice], one of the most canonically discussed cases of a marked feature-value, is only marked in obstruents, e.g. in the context of [−sonorant]. In languages with neutralization of [± voice] to [−voice] in Coda positions (e.g. Turkish, Russian, German), this neutralization only affects [−sonorant] segments, and indeed, [+sonorant] segments are all [+voice]. Thus even some of the most well-known examples of marked feature-values carry an implicit contextual restriction imposed by other features with which the opposition occurs.

When it comes to features such as [± round], there is no obvious context-free value that is the marked one crosslinguistically or psycholinguistically. However, when paired with certain other feature-values, [+round] is marked (e.g. in combination with [−back]), while with other feature-values, [−round] is marked (e.g. in combination with [+back]). It becomes apparent that while [+round] is not inherently marked, once we fix the context of other features within a bundle to include [−back], the value [+round] is the marked value of the opposition. In other words, [+round] is context-sensitively marked when in the presence of [−back], most likely due to the antagonistic effects that these two features have on the length of the vocal tract "tube" corresponding to the second formant. Similarly, [+round] is marked when in the context of [−high] (the vowel *o* being more marked than *u*), arguably due to the interaction between lip rounding and lower jaw position. In the following section we examine the patterning of [± round] harmony in a number of languages, as harmony for this feature exhibits a number of instructive interactions with context-sensitive markedness.

3.8.1. Context-Sensitive Markedness of [+round]

As discussed above, neither [+round] nor [−back] are marked on their own in vowel systems, but their combination is marked. The context-sensitive markedness of [+round] in [−back] and [−high] vowels is proposed in Chomsky and Halle (1968, p.405), Calabrese (1988, p.22), and Archangeli and Pulleyblank (1994, 78):

(51) *Context-sensitive markedness statements*:

a. [+round] is marked in the context of [−back]
b. [+round] is marked in the context of [−high]

In addition to the implicational relationships between marked and unmarked supporting (51) (e.g. a language with [−back,+round] will also have [−back,−round], and a language with [−high, +round] will also have [−high, −round]), there is a variety of articulatory, acoustic, and perceptual evidence for these context-sensitive markedness values. Based on photographic evidence, Linker (1982) shows that lip rounding activity for [+round] is greater in high vowels than in their non-high counterparts, and that back vowels are more rounded than their non-back counterparts. Within the realm of acoustic phonetics, Stevens (1998, p.293-294) concludes that the acoustic consequences of adding lip rounding to back vowels is greater than would be achieved by adding lip rounding to front vowels; similarly, lip rounding has more dramatic acoustic consequences for high vowels than for non-high vowels. Finally, Kaun (1995, p.121) interprets the results of Terbeek (1977) as demonstrating that rounding is perceptually more robust in *u* than in *o,ü,ö*.

The typology of rounding harmony in the Altaic (e.g. Turkic, Tungusic, and Mongolian) languages is quite well-studied; see Korn (1969), Vaux (1993), and Kaun (1995) for extensive surveys into various types. In the present discussion we will examine the effects of a relativization of the search domain to marked values of [± round] as delimited in (51), and thereby focus on two of the widespread sources of parametric variation.[37]

Altai and Shor (two Turkic languages) have low vowel suffixes that may only copy [+round] from marked sources, namely vowels which are [−high] or [−back] (Korn, 1969, 101). As a somewhat surprising result, these languages do not copy [+round] from *u*, even though this vowel clearly bears the feature [+round]:

(52) *Altai Rounding Harmony: Limited to Marked [+Round] Sources:*

kol	kol-do	'hand-loc.'
kös	kös-tor	'eye-pl'
kün	kün-dö	'day-loc.'
už	už-ar	'fly-aor.'

(53) *Shor Rounding Harmony: Limited to Marked [+Round] Sources:*

kol	kol-don	'hand-abl.'
sös	sös-ton	'word-abl'
külük	külük-tö	'brave-loc.'
ug	ug-ar	'grasp-aor.'

Like the derivations illustrated above for Sibe marked [+low] harmony, Shor (and Altay) marked [+round] harmony involves a relativization of the search domain so as to only include marked values of [± round]. For example, in *sös-ton*, the affix finds a marked [+round] value (context-sensitively marked in the presence of [−back]):

(54) Shor Ablative Suffix Searches for Marked [+round] Source in *sös-ton*:

$$
\begin{array}{ccccccc}
x_1 & x_2 & x_3 & \leftarrow & x_4 & x_5 & x_6 \\
\begin{bmatrix}\text{-voc}\\ \text{coronal}\\ \text{+cont}\\ \text{-voice}\end{bmatrix} &
\begin{bmatrix}\text{+voc}\\ \text{-back}\\ \text{-high}\\ \text{+round}\end{bmatrix} &
\begin{bmatrix}\text{-voc}\\ \text{coronal}\\ \text{+cont}\\ \text{-voice}\end{bmatrix} & &
\begin{bmatrix}\text{-voc}\\ \text{coronal}\\ \text{-cont}\\ \text{-voice}\end{bmatrix} &
\begin{bmatrix}\text{+voc}\\ \text{-back}\\ \text{-high}\\ \\ \end{bmatrix} &
\begin{bmatrix}\text{-voc}\\ \text{coronal}\\ \text{+nasal}\end{bmatrix}
\end{array}
$$

(55) Ablative Suffix Finds and Copies from Marked [+round] Source:

$$
\begin{array}{ccccccc}
x_1 & x_2 & \Lsh & x_3 & x_4 & x_5 & x_6 \\
\begin{bmatrix}\text{-voc}\\ \text{coronal}\\ \text{+cont}\\ \text{-voice}\end{bmatrix} &
\begin{bmatrix}\text{+voc}\\ \text{-back}\\ \text{-high}\\ \text{+round}\end{bmatrix} & &
\begin{bmatrix}\text{-voc}\\ \text{coronal}\\ \text{+cont}\\ \text{-voice}\end{bmatrix} &
\begin{bmatrix}\text{-voc}\\ \text{coronal}\\ \text{-cont}\\ \text{-voice}\end{bmatrix} &
\begin{bmatrix}\text{+voc}\\ \text{-back}\\ \text{-high}\\ \textbf{+round}\end{bmatrix} &
\begin{bmatrix}\text{-voc}\\ \text{coronal}\\ \text{+nasal}\end{bmatrix} \\
\text{s} & \text{ö} & & \text{s} & \text{t} & \text{o} & \text{n}
\end{array}
$$

When a root contains no marked values of [+round], Search will come back empty-handed. For example, in *ug-ar*, the affix's search proceeds leftward, never finding a marked value of [+round], and thereby requiring last-resort insertion of the default value:

(56) Shor Aorist Suffix Searches for Marked [+round] Source in *ug-ar*:

$$
\begin{array}{ccccc}
x_1 & x_2 & \leftarrow & x_3 & x_4 \\
\begin{bmatrix} \text{+voc} \\ \text{+back} \\ \text{+high} \\ \text{+round} \end{bmatrix} &
\begin{bmatrix} \text{-voc} \\ \text{dorsal} \\ \text{-cont} \\ \text{+voice} \end{bmatrix} & &
\begin{bmatrix} \text{+voc} \\ \text{-high} \\ \text{+back} \\ \ \end{bmatrix} &
\begin{bmatrix} \text{-voc} \\ \text{coronal} \\ \text{+sonorant} \\ \text{-nasal} \\ \text{-lateral} \end{bmatrix}
\end{array}
$$

(57) Aorist Suffix Fails to find Marked [+round] Source;
Last Resort Insertion of [−round]:

$$
\begin{array}{ccccc}
\not\leftarrow & x_1 & x_2 & x_3 & x_4 \\
& \begin{bmatrix} \text{+voc} \\ \text{+back} \\ \text{+high} \\ \text{+round} \end{bmatrix} &
\begin{bmatrix} \text{-voc} \\ \text{dorsal} \\ \text{-cont} \\ \text{+voice} \end{bmatrix} &
\begin{bmatrix} \text{+voc} \\ \text{-high} \\ \text{+back} \\ \textbf{-round} \end{bmatrix} &
\begin{bmatrix} \text{-voc} \\ \text{coronal} \\ \text{+sonorant} \\ \text{-nasal} \\ \text{-lateral} \end{bmatrix} \\
& \text{u} & \text{g} & \text{a} & \text{r}
\end{array}
$$

The fact that Shor and Altai low vowels copy [+round] from *ü,ö,o* but not *u*, and that no round harmony system has the opposite pattern (i.e. low vowels copy [+round] from *u* but not from *ü,ö,o*) provides confirmation that the Search algorithm can be relativized to include both values of [± round] or only the marked values of [± round], but can never be relativized so as to only include the unmarked values of [± round].[38]

3.8.2. Microvariation in Kirghiz Dialects

In our earlier discussion of context-free markedness in harmony systems, we compared two closely related languages, Sibe and Sanjiazi Manchu, demonstrating that they differed in a minimal parametric way: one language was set to [search: cF] while the other was set to [search: mF]. A parallel demonstration of minimal parametric variation based on value-relativization can be found with context-sensitive harmony systems, in this case from two different descriptions (arguably two distinct dialects or idiolects) of the same language: Kirghiz. I will refer to these two dialects as Dialect A and Dialect B.

Dialect A of Kirghiz, as described by Comrie (1981), copies round from *u*, thereby instantiating a system of copying from any contrastive [± round] source:

(58) *Kirghiz Rounding Harmony, Dialect A (Comrie, 1981, p.61):*

köl	köl-dön	lake-abl.
üy	üy-dön	house-abl.
tokoy	tokoy-don	forest-abl.
tuz	tuz-don	salt-abl

However, Dialect B of Kirghiz, as described by Hebert and Poppe (1963) (see also Johnson (1980)), copies [± round] only from marked sources (like Altai and Shor), thereby failing to copy from *u*:

(59) *Kirghiz Rounding Harmony, Dialect B (Hebert and Poppe, 1963, p.8):*

köl	köl-dön	lake-abl.
üy	üy-dön	house-abl.
tokoy	tokoy-don	forest-abl.
turmuš	turmus-tan	life-abl

I assume that the differences in vowel harmony patterns described by Hebert and Poppe (1963) on the one hand and Comrie (1981) on the other hand represent coherent dialectal/idiolectal differences, of exactly the type we would expect in a model of harmony where the choice between [search: cF] and [search: mF] represents a restricted means to permit individual grammars to vary while maintaining the same inventories and Search procedure.

(60) Kirghiz A: Ablative Suffix must Search Leftward for Closest Contrastive Value of [± round]
Kirghiz B: Ablative Suffix must Search Leftward for Closest Marked Value of [± round]

The more general conclusion to be drawn is that a system of restricted parametric options for what the domain of harmonic search is relativized to contain can capture the attested pattern of microvariation – without generating unattested patterns – by the mere "flipping of a switch": whether unmarked values of [±F] are excluded from search or not.

3.9. Microvariation in Yoruba dialects

A further example of how dialectal variation can arise from the single difference in a setting of the value-relativization parameter may be illustrated in comparing the Ifẹ dialect of Yoruba with the Standard dialect. As briefly introduced in Chapter 1, Yoruba vowel harmony takes place when all [−high,−low] vowels that are not the final vowel of the word must search and copy a value of [± ATR] from their right. In other words, within a disyllabic root, the initial vowel is dependent on the value of [± ATR] from the final vowel:

(61) *Yoruba [± ATR] Harmony in Disyllabic Roots*:
  a. èwé 'lip'
  b. olè 'thief'
  c. èrò 'crowd'
  d. òd͡ʒò 'rain'
  e. ɛ̀gɛ́ 'cassava'
  f. ɔsɛ 'soap'
  g. ɛ̀fɔ́ 'vegetable'
  h. ɔwɔ́ 'hand'

Only the mid vowels are contrastive for [± ATR]. Under a relativization of vowel harmony to only pay attention to contrastive values of the harmonic feature, the high [+ATR] vowels – which have no contrastive counterpart – will be skipped in harmony, as is found in Ifẹ Yoruba (Ola Orie, 2001):

(62) *Ifẹ Yoruba contrastive ATR harmony skips the irrelevant high vowel*:
  a. ɔrúkɔ 'name'
  b. ɛ̀lùbɔ́ 'yam flour'
  c. ɛúrɛ́ 'goat'
  d. ɔdídɛ 'parrot'
  e. ɔ̀títɔ́ 'truth'

The formalization of Ifẹ Yoruba ATR harmony in terms of the Search principle and its parameters is provided in (63):

(63) Ifẹ Yoruba non-final mid vowels must:
ATR-Harmonize: δ = R, F = [contrastive: ATR]

As a result of the setting in Ifẹ to contrastive values of [± ATR], medial high vowels in (62) exhibit "transparency", being skipped over for [−ATR] copying across them, even though they themselves are [+ATR]. Even more interestingly, the pattern of Ifẹ Yoruba differs from the harmony pattern of the Standard dialect, in which all of the words in (62) enact copying of [+ATR] from the adjacent high vowel:

(64) *Standard Yoruba copies [+ATR] from the adjacent high vowel*:
  a. orúkɔ 'name'
  b. èlùbɔ́ 'yam flour'
  c. ewúrɛ́ 'goat'
  d. odídɛ 'goat'
  e. òtítɔ́ 'truth'

The characterization of [± ATR] harmony in Standard Yoruba, then, is that it is sensitive to *all* values of ATR, whether contrastive or not:

(65) Standard Yoruba non-final mid vowels must:
ATR-Harmonize: δ = R, F = [all: ATR]

Parametrization in (65) does not permit skipping of noncontrastive high vowels in (64); everything is included in the search. (In Chapter 5, we complete the analysis of Yoruba [± ATR] harmony with a discussion of the [+low] vowel *a* in harmony.) Ifẹ and Standard Yoruba are two closely related dialects, with identical inventories and directionality of harmony, but they differ in what values of [± ATR] are "counted as relevant" in the Search.

3.10. Variation in obstruent transparency in [± nasal] harmony

Nasal harmony system (for example, as found throughout Colombia, Venezuela and Brazil), in which the feature [± nasal] is copied by a variety of segments throughout a word, exhibit two distinct patterns of relativization (Piggott, 1992). The first pattern is one in which the inherently [−nasal] consonants such as *p,t,k* are copied from. The second pattern is one in which the [−nasal] consonants such as *p,t,k* are ignored and skipped over, with copying of [+nasal] occurring across them (see Walker (1999) for phonetic evidence of transparency in such cases).

The difference between these two patterns lies in the relativization of the search: if all values of [± nasal] are included in the search, then clearly the [−nasal] voiceless obstruents will be copied from as well. If, on the other hand, only contrastive values of [± nasal] are included, then *p,t,k* will be excluded from the search, as they are not contrastive for [± nasal], there being no other segments in the inventory that are [−continuant, −voiced] but [+nasal].

In the first type of relativization, instantiated by Warao (Osborn, 1966), voiceless obstruents “interrupt” [± nasal] harmony, as illustrated in the following examples. I assume that all non-initial vowels, glides, and /h/ (i.e. the [−consonantal] elements) are needy for the feature [± nasal], and must copy it from the closest source to their left. Like the pattern of [± back] harmony in the Turkic languages, in Warao [+nasal] harmony, the initial vowel in the root is the only non-needy segment. All [+consonantal] segments are inherently [+nasal] or [−nasal].

(66) *Warao segments copies [± nasal] from closest leftward source (Osborn, 1966, p.111–112)*:

a. ĩõ ‘turtle’
b. ũĩ ‘angoleta bird’
c. mõỹõ ‘cormorant’
d. mẽh̃õkohi ‘shadow’
e. nãõte ‘he will come’
f. mõãũpu ‘give them to him’

(67) *Warao [−cons] segments must:*
Nasal-Harmonize: δ = left, F = [all: ± nasal]

The derivation for *nãõte* is shown below; there are two non-initial [−cons] segments, and the first of them copies [+nasal] from the closest element to its left, while the second of them copies [−nasal] from the closest element to its left:

(68) First Non-Initial [−cons] segment searches for [± nasal] in *nãõ-tẽ*:

$$\begin{array}{ccccccc} x_1 & & x_2 & \leftarrow & x_3 & & x_4 \quad\quad x_5 \\ \begin{bmatrix}\text{+cons}\\ \text{+son}\\ \text{coronal}\\ \text{-cont}\\ \text{+nasal}\end{bmatrix} & & \begin{bmatrix}\text{-cons}\\ \text{+son}\\ \text{-high}\\ \text{+low}\\ \text{- round}\\ \text{+back}\\ \text{+nasal}\end{bmatrix} & & \begin{bmatrix}\text{-cons}\\ \text{+son}\\ \text{-high}\\ \text{-low}\\ \text{+round}\\ \text{+back}\\ \end{bmatrix} & & \begin{bmatrix}\text{+cons}\\ \text{-son}\\ \text{coronal}\\ \text{-cont}\\ \text{-nasal}\end{bmatrix} \begin{bmatrix}\text{-cons}\\ \text{+son}\\ \text{-high}\\ \text{-low}\\ \text{-round}\\ \text{-back}\\ \end{bmatrix} \end{array}$$

(69) First Non-Initial [−cons] segment finds [+ nasal]:

$$\begin{array}{ccccccc}
x_1 & & x_2 & \Lsh & x_3 & x_4 & x_5 \\
\begin{bmatrix}\text{+cons}\\ \text{+son}\\ \text{coronal}\\ \text{-cont}\\ \text{+nasal}\end{bmatrix} & & \begin{bmatrix}\text{-cons}\\ \text{+son}\\ \text{-high}\\ \text{+low}\\ \text{- round}\\ \text{+back}\\ \text{+nasal}\end{bmatrix} & & \begin{bmatrix}\text{-cons}\\ \text{+son}\\ \text{-high}\\ \text{-low}\\ \text{+round}\\ \text{+back}\\ \textbf{+nasal}\end{bmatrix} & \begin{bmatrix}\text{+cons}\\ \text{-son}\\ \text{coronal}\\ \text{-cont}\\ \text{-nasal}\end{bmatrix} & \begin{bmatrix}\text{-cons}\\ \text{+son}\\ \text{-high}\\ \text{-low}\\ \text{-round}\\ \text{-back}\\ \end{bmatrix}
\end{array}$$

(70) Final Non-Initial [−cons] segment searches for [± nasal]:

$$\begin{array}{cccccc}
x_1 & x_2 & x_3 & x_4 & \leftarrow & x_5 \\
\begin{bmatrix}\text{+cons}\\ \text{+son}\\ \text{coronal}\\ \text{-cont}\\ \text{+nasal}\end{bmatrix} & \begin{bmatrix}\text{-cons}\\ \text{+son}\\ \text{-high}\\ \text{+low}\\ \text{- round}\\ \text{+back}\\ \text{+nasal}\end{bmatrix} & \begin{bmatrix}\text{-cons}\\ \text{+son}\\ \text{-high}\\ \text{-low}\\ \text{+round}\\ \text{+back}\\ \text{+nasal}\end{bmatrix} & \begin{bmatrix}\text{+cons}\\ \text{-son}\\ \text{coronal}\\ \text{-cont}\\ \text{-nasal}\end{bmatrix} & & \begin{bmatrix}\text{-cons}\\ \text{+son}\\ \text{-high}\\ \text{-low}\\ \text{-round}\\ \text{-back}\\ \end{bmatrix}
\end{array}$$

(71) Final Non-Initial [−cons] segment finds [− nasal]:

$$\begin{array}{cccccc}
x_1 & x_2 & x_3 & x_4 & \Lsh & x_5 \\
\begin{bmatrix}\text{+cons}\\ \text{+son}\\ \text{coronal}\\ \text{-cont}\\ \text{+nasal}\end{bmatrix} & \begin{bmatrix}\text{-cons}\\ \text{+son}\\ \text{-high}\\ \text{+low}\\ \text{- round}\\ \text{+back}\\ \text{+nasal}\end{bmatrix} & \begin{bmatrix}\text{-cons}\\ \text{+son}\\ \text{-high}\\ \text{-low}\\ \text{+round}\\ \text{+back}\\ \text{+nasal}\end{bmatrix} & \begin{bmatrix}\text{+cons}\\ \text{-son}\\ \text{coronal}\\ \text{-cont}\\ \textbf{-nasal}\end{bmatrix} & & \begin{bmatrix}\text{-cons}\\ \text{+son}\\ \text{-high}\\ \text{-low}\\ \text{-round}\\ \text{-back}\\ \textbf{-nasal}\end{bmatrix} \\
\text{n} & \tilde{\text{a}} & \tilde{\text{o}} & \text{t} & & \text{e}
\end{array}$$

In the second type of language, only contrastive values of [± nasal] are visible to the search. This is illustrated for Southern Barasano; a similar pattern is found in Desano (Kaye, 1971) and Guaraní (Walker, 1999). In Barasano, all [± sonorant] segments are contrastive for [± nasal], but the voiceless stops and voiceless fricatives are not.

Interestingly, Barasano has voiced stops that have occasional realization as prenasalized stops (although even when they are prenasalized, they do not induce nasal harmony). I follow the proposal of Piggott (1992) that nasals are [−continuant] (Anderson, 1976) and that all voiced segments, including the prenasalized stops, are in fact sonorants (Piggott, 1992, p.49). Rice (1993) provides evidence from Rotokas, Slavey, and a number of other languages that certain "obstruents" may in fact be [−continuant,+sonorant]. Thus, the voiced stops of Barasano / Tucanoan are [+voice,−cont, +son,−nasal] while the nasal stops are [+voice,−cont, +son,+nasal]. As a result, the voiced stops and the nasal stops are contrastive for [± nasal], and included in the search domain for harmony.

(72) *Southern Barasano Affixes Copy Contrastive [± nasal] from closest Leftward Source (Piggott, 1992, p.47),(Piggott, 2003, p.379)*:

a. mãhã-mã 'go up!'
b. wa-^{m}ba 'come!'
c. ĩã-mĩ 'I saw'
d. wa-^{m}bɪ'I went'
e. mãã-r̃e 'to seize by the handful'
f. baa-re 'to swim'
g. mĩnĩ-ãkã 'small bird'
h. coti-aka 'small pot'

(73) *Barasano [+son] segments must:*
Nasal-Harmonize: δ = left, F = [all: ± nasal]

The derivation for words such as *mĩnĩ-ãkã*, in which the voiceless obstruent is transparent to harmony across it, is shown below for the two affixal vowels.

(74) *First Affixal [+son] segment searches for [± nasal] in mĩnĩ-ãkã*:

$$\begin{array}{ccccc} x_1 & \leftarrow & x_2 & x_3 & x_4 \\ \begin{bmatrix}\text{-cons}\\ \text{+son}\\ \text{+high}\\ \text{-low}\\ \text{-round}\\ \text{-back}\\ \text{+nasal}\end{bmatrix} & & \begin{bmatrix}\text{-cons}\\ \text{+son}\\ \text{-high}\\ \text{+low}\\ \text{-round}\\ \text{+back}\\ \end{bmatrix} & \begin{bmatrix}\text{+cons}\\ \text{-son}\\ \text{dorsal}\\ \text{-cont}\\ \text{-nasal}\end{bmatrix} & \begin{bmatrix}\text{-cons}\\ \text{+son}\\ \text{-high}\\ \text{+low}\\ \text{-round}\\ \text{+back}\\ \end{bmatrix} \end{array}$$

(75) *First Affixal [+son] segment finds [+ nasal]*:

$$
\begin{array}{ccccc}
x_1 & \Lsh & x_2 & x_3 & x_4 \\
\begin{bmatrix} \text{-cons} \\ \text{+son} \\ \text{+high} \\ \text{-low} \\ \text{-round} \\ \text{-back} \\ \text{+nasal} \end{bmatrix} & &
\begin{bmatrix} \text{-cons} \\ \text{+son} \\ \text{-high} \\ \text{+low} \\ \text{-round} \\ \text{+back} \\ \textbf{+nasal} \end{bmatrix} &
\begin{bmatrix} \text{+cons} \\ \text{-son} \\ \text{dorsal} \\ \text{-cont} \\ \text{-nasal} \end{bmatrix} &
\begin{bmatrix} \text{-cons} \\ \text{+son} \\ \text{-high} \\ \text{+low} \\ \text{-round} \\ \text{+back} \\ \\ \end{bmatrix}
\end{array}
$$

(76) *Second Affixal [+son] segment searches for [± nasal]*:

$$
\begin{array}{ccccc}
x_1 & x_2 & x_3 & \leftarrow & x_4 \\
\begin{bmatrix} \text{-cons} \\ \text{+son} \\ \text{+high} \\ \text{-low} \\ \text{-round} \\ \text{-back} \\ \text{+nasal} \end{bmatrix} &
\begin{bmatrix} \text{-cons} \\ \text{+son} \\ \text{-high} \\ \text{+low} \\ \text{-round} \\ \text{+back} \\ \text{+nasal} \end{bmatrix} &
\begin{bmatrix} \text{+cons} \\ \text{-son} \\ \text{dorsal} \\ \text{-cont} \\ \text{-nasal} \end{bmatrix} & &
\begin{bmatrix} \text{-cons} \\ \text{+son} \\ \text{-high} \\ \text{+low} \\ \text{-round} \\ \text{+back} \\ \\ \end{bmatrix}
\end{array}
$$

(77) *Second Affixal [+son] segment finds [+ nasal]*:

$$
\begin{array}{ccccc}
x_1 & x_2 & \Lsh & x_3 & x_4 \\
\begin{bmatrix} \text{-cons} \\ \text{+son} \\ \text{+high} \\ \text{-low} \\ \text{-round} \\ \text{-back} \\ \text{+nasal} \end{bmatrix} &
\begin{bmatrix} \text{-cons} \\ \text{+son} \\ \text{-high} \\ \text{+low} \\ \text{-round} \\ \text{+back} \\ \textbf{+nasal} \end{bmatrix} & &
\begin{bmatrix} \text{+cons} \\ \text{-son} \\ \text{dorsal} \\ \text{-cont} \\ \text{-nasal} \end{bmatrix} &
\begin{bmatrix} \text{-cons} \\ \text{+son} \\ \text{-high} \\ \text{+low} \\ \text{-round} \\ \text{+back} \\ \textbf{+nasal} \end{bmatrix} \\
\text{ĩ} & \text{ã} & & \text{k} & \text{ã}
\end{array}
$$

In sum, the variation between Warao and Barasano as to whether voiceless stops are visible for [± nasal] instantiates yet another case of parametric variation, in which search may be relativized to a certain (sub)set of the values of the harmonic feature.

3.11. Microvariation in Finnish Loanwords

The differences between Kirghiz A and Kirghiz B reduced to whether the value-relativization parameter was set to *contrastive values* or *marked values*, and the differences between Ifẹ Yoruba and Standard Yoruba varied as a function of whether the value-relativization parameter was set to *contrastive values* or *all values*, as did the differences in patterns of nasal harmony in Warao vs. Barasano. In this section we analyze microvariation among speakers of Finnish for a particular class of loanwords that stray from the basic phonotactics of Finnish roots, and are handled by speakers in terms of different parametric visibility as well. Recall the Finnish vowel system:

(78) *Finnish Vowel Inventory*:

[−back, −round]	[−back, +round]	[+back, +round]	[+back, −round]	
i	ü	u		[+high, −low]
e	ö	o		[−high, −low]
ä			a	[−high, +low]

As discussed by Campbell (1980) and Kiparsky (1981), while the pattern of Finnish vowel harmony is straightforward for roots that contain all [+back] vowels, all contrastive [−back] vowels, or noncontrastive [−back] vowels, the pattern of vowel harmony for suffixes concatenated to loanwords that mix contrastive [+back] and contrastive [−back] exhibits idiolectal variability. The words *marttüüri* and *türanni*, for example, contain a [+back] *a* and a [−back] *ü*. Interestingly, while the pattern of vowel harmony varies among speakers for words like *marttüüri*[39], it does not vary among any of them for *türanni*:

(79) Finnish Dialect A:
  a. marttüüri-ä 'martyr-partit.sg'
  b. türanni-a 'tyrant-partit.sg'

(80) Finnish Dialect B:
  a. marttüüri-a 'martyr-partit.sg'
  b. türanni-a 'tyrant-partit.sg'

Putting it differently, these two dialects do not differ in their treatment of contrastively [−back] words followed by [+back] vowels (79-b) vs. (80-b), but they do differ in their treatment of [+back] vowels followed by contrastively [−back] vowels (79-a) vs. (80-a): for Dialect A, the contrastively [−back] vowels are part of the harmonic search, while for Dialect B, they are not.

By now, the explanation of Dialect A versus Dialect B should be familiar: Dialect A's harmonic search is parametrized to include contrastively [−back] vowels, and Dialect B's harmonic search is parametrized to include only [+back] vowels, the marked value of the feature in Finnish (Kiparsky, 1981). In Finnish B, if a word contains all [−back] vowels (i.e. contains no marked values of [+back]), then Last-Resort insertion will insert the default value. This microparametric difference is captured below:[40]

(81) Finnish A: Ablative Suffix must Search Leftward for Closest Contrastive Value of [± back]
Finnish B: Ablative Suffix must Search Leftward for Closest Marked Value of [± back]

One of the interesting results of the formulation of Finnish A vs. Finnish B in (81) is that both dialects will produce identical outputs for words analyzed earlier in this chapter such as *pouta, pöütä, tie, koti*: in all [+back] words, both Dialects will copy [+back], in all [−back] words, both dialects will end up with a [−back] suffix (though through different means), and in [+back] followed by noncontrastive [−back] words, both dialects will copy [+back]. In fact, since the native vocabulary is predominantly composed of these three classes of words, all of which are compatible with both dialect A and dialect B, a clear conclusion to be drawn is that all of these words are *parametrically ambiguous* between [search:cF] and [search:mF], and that the variation emerges as the result of which option is chosen in the face of ambiguous input. Which option was actually chosen will only reveal itself on inputs for which these two parametric settings diverge, i.e. loanwords of the *martüüri*-type. In the words of Andersen (1973, 774), "Innovation in the phonological structure of a language can only be explained on the basis of ambiguities in the corpus of utterances from which the new grammar is inferred". While this statement may not be universally true for all phonological innovations (e.g. epenthesis or syncope may not necessarily involve reanalysis of parametrically-ambiguous input), I would argue that in the domain of microvariation of locality computation for vowel harmony, ambiguity is the wellspring of innovative analyses of value-relativization that lead to divergent results when loanwords introduce new structures into the language.

3.12. Set-Union of Marked and Contrastive Values in Oroch

We have considered a number of cases in which search includes only the marked values of the harmonic feature, and in which search includes only the contrastive values of the harmonic feature. In this section, we consider a fusion of these two visibility parameters.

Before proceeding, we must note that the set of marked values in an inventory and the set of contrastive values are not always in a subset/superset relation. Whenever a language has at least three vowels, {A,B,C} with marked [+F] on B, C, where {A,B} are contrastive for [±F] and C is not contrastive for [±F], then the set of marked values will include {B,C} and the set of contrastive values will include {A,B}. If the language has a fourth vowel, D, which is noncontrastive and unmarked for [± F], then the set of all vowels = {A,B,C,D}, the set of contrastive vowels = {B,C}, and the set of marked vowels = {A,B}. I claim that search-parametrization can include the set-union operator, and hence select the set of vowels that are either contrastive or marked, namely {A,B,C}. Such an operation becomes empirically necessary when considering the [± ATR] harmony of Oroch, a Tungusic language described by Tolskaya (2008).

(82) *Oroch Vowel Inventory*:

[−round,−back]	[−round,+back]	[+round]	
i		u	[+high, +ATR]
		ʊ	[+high, −ATR]
æ	ə		[−high, +ATR]
	a	ɔ	[−high, −ATR]

I assume that *æ* is [+ATR], and that [−ATR] is the marked value of the feature, as Tolskaya (2008) notes that *æ* is in free variation with iə, i or ia for some lexical items, and that it is slightly diphthongal, starting with an ultra-short [i]. Considerations from the Tungusic family suggest that *æ* is related to the [+ATR] *e* of Evenki, cf. Anderson (2004, p.39) "ATR harmony is characteristic of the whole Tungusic family and Evenki is no exception in this regard; the alternation primarily consists of *a* ∼ e/ə". In Oroch, however, the alternation is between *a* and ə, with æ as [+ATR] but unpaired. Most revealingly, Avrorin and Boldyrev (2001, p.30) state that the velars k,g have an allophonic realization as uvular q,ɣ when following the vowels a,ɔ,ʊ, but not when following æ,ə,u, which can be taken as evidence that only the [−ATR] vowels induce uvularization (cf. *baqi* ''lazy' vs. *gæki* 'hawk').

Oroch has two distinct patterns of vowel harmony: [± ATR] harmony and [± round] harmony. Interestingly, while /i,æ/ are invisible for [± ATR] harmony, they participate in the process of [± round] harmony, e.g. ɔtɔŋgo-ɲi-da 'kayak-3sg-foc', analyzed in Section 4.6).

As /u,ʊ/ and /ə,a/ are contrastive for [± ATR], and /ʊ,a,ɔ/ are marked for [± ATR], set-union for the properties of marked and contrastive [± ATR] will include all vowels in the inventory except /i,æ/. The value-relativized sensitivity to contrastiveness or markedness will exclude these vowels from the search and render them invisible for [± ATR] harmony, as shown in (83).

(83) *Invisibility of* i,æ *in Oroch [± ATR] harmony*:
a. ugda-va-da 'boat-acc-foc'
b. xuŋkə-və-də 'canoe-acc-foc'
c. ugda-ɳi-da 'boat-3sg-foc'
d. xuŋkə-ɳi-də 'canoe-3sg-foc'
e. əʒæ lənidə 'just below'
f. ʒaŋgæ-ra 'judge'
g. sɔrɔdæ-da 'greet-foc'

As shown in (83-c-d), *i* is invisible to [± ATR] harmony across it, and and shown in (83-e-g), *æ* is invisible to [± ATR] harmony across it. It is at this point very instructive to compare the theory of value-relativization developed throughout this chapter with one of the other widespread attempts within phonological theory to render certain vowels invisible for harmony, namely underspecification theory. The rationale behind underspecification theory is that there is no need to specify in the featural representation of vowels the values that are predictable on independent grounds. For example, in a language such as Finnish with no [+back,+round,+high] vowel, the [± back] value of a [+high,−round] vowel is predictable, and hence would not included in its representation.

In underspecification theory, vowels that are invisible for harmony literally lack the harmonic feature, until a very late stage of the derivation, if at all. In the theory of Dresher (2005), for example, an algorithm for successive contrast-based division of the inventory assigns feature-values to all vowels in the inventory in a certain order, and certain vowels are never assigned certain features. The assignment of features depends on a set of hierarchical orderings among features that determine the relative scope of a feature over the inventory.

For example, in an inventory consisting of *i,u,a*, and the features [± high, ± back, ± round], there are four possible outcomes of the algorithm, depending on the scope of the features:

(84) *Possible Outcomes of the Successive Division Algorithm for i,u,a*:

a. Choose [± back]: *i* [−back], *u* [−back] *a* [+back]
   Choose [± high] next: *i* [−back], *u* [+back, +high] *a* [+back,−high]
b. Choose [± round]: *i* [−round], *u* [+round] *a* [−round]
   Choose [± high] next: *i* [−round], *u* [+round, +high] *a* [−round,−high]
c. Choose [± high]: *i* [+high], *u* [+high] *a* [−high]
   Choose [± back] next: *i* [+high, −back], *u* [+high, +back] *a* [−high]
d. Choose [± high]: *i* [+high], *u* [+high] *a* [−high]
   Choose [± round] next: *i* [+high, −round], *u* [+high, +round] *a* [−high]

The claim of the Successive Division Algorithm is that cross-linguistic variation in what vowels are active for which processes is the result of the scope of these features. For example, given (84-a), *i* cannot participate in [± high] harmony; given (84-b), *i* cannot participate in [± round] harmony; given (84-c), *a* cannot participate in [± back] harmony; and given (84-d), *a* cannot participate in [± round] harmony.

The Successive Division Algorithm represents a principled way to exhaustively assign just enough feature specifications in order to make every vowel contrastive. Dresher (2003) demonstrates a number of empirical advantages of the SDA over earlier algorithms for assigning feature specification that determine which vowels should be underspecified based on predictability of features, such as the Radical Underspecification Algorithm of Archangeli (1988). One of the celebrated advantages of the SDA is the fact that it can accomodate crosslinguistic variation simply by rearranging the scope of features within the hierarchy.

This stands in stark opposition to the framework adopted in this book, in which all vowels are fully specified for all features, and it is the harmony processes themselves that are relativized to ignore or include certain features. In what follows, I will demonstrate that given the Oroch inventory and the four features [± high, ± ATR, ± round, ± back], any possible scope ordering of the features fails to generate the right specifications. Recall that in Oroch, /i,æ/ are invisible to [± ATR] harmony, and thus must lack a specification for [± ATR] given underspecification theory, while i,æ are visible in [± round] harmony (as shown in Section 4.6), and so must have a specification for [± round].

(85) Desired Outcomes of Underspecification in Oroch, given an underspecification approach to harmonic participation:
  a. Desiderata 1 (D1): i,æ must lack specification for [± ATR], all others must have it
  b. Desiderata 2 (D2): i,æ must bear specifications for [± round]

(86) *Possible Outcomes of the SDA for Oroch given [± high, ± ATR, ± round, ± back]*:
  a. Choose [± ATR] first: FAILS, since it will assign [± ATR] to i,æ, contra D1
  b. Choose [± back] first: FAILS, since *i* will not be later assigned round, contra D2
  c. Choose [± high] first: assigns [+ high] to i,u,ʊ, [−high] to æ,ə,a,ɔ.
    (i) If [± ATR] chosen next, FAILS, since it will assign [± ATR] to i,æ, contra D1
    (ii) If [± back] chosen next: FAILS, since *i* will not be later assigned round, contra D2

    If [± round] chosen next, ɔ will be the only [−high,+round] value, and as ɔ will not be later assigned [± ATR], contra D1, FAILS
  d. Choose [± round] first: assigns [+round] to u,ʊ,ɔ and [−round] to i,æ,ə,a.
    (i) If [± ATR] chosen next, FAILS, since it will assign [± ATR] to i,æ, contra D1
    (ii) If [± high] chosen next: ɔ will be the only [−high,+round] value, and ɔ will not be later assigned [± ATR], contra D1, FAILS
    (iii) If [± back] chosen next: assigns [+back] to u,ʊ,ɔ,ə,a and

[−back] to i,æ

If [± high] chosen next: ɔ will be the only [−high,+back, +round] value, and as ɔ will not be later assigned [± ATR], contra D1, FAILS

If [± ATR] chosen next, FAILS, since it will assign [± ATR] to i,æ, contra D1

As shown in (86), the most successful underspecification algorithm developed to date cannot yield the right pattern of feature specification to guarantee non-participation of Oroch /i,æ/ in [± ATR] harmony and their concomitant participation in [± round] harmony. Rather, all vowels must be specified for all values of these features, regardless of whether one might be predictable from the other.[41] I conclude that the right place to situate feature-visibility is in the harmony processes themselves, rather than in the inventory.

3.13. Conclusion: Transparent Items are Pruned Away due to Irrelevancy

Before concluding this chapter, I would like to reflect on how the model may be viewed as an outgrowth in the steps towards constraints on possible long-distance assimilation rules, from the formal constraints on rules of Howard (1972); Jensen (1974); Battistella (1982); Yamada (1983) to the representationally-based constraints on rules of Clements (1985); Sagey (1986); Steriade (1987); Calabrese (1988); McCarthy (1988); Odden (1994); Archangeli and Pulleyblank (1994). Importantly, as argued in Chapter 2, the current formulation of locality is *target-centric*, with locality defined from the point of view of the needy suffix, and therefore the question of what may "intervene" transparently in a rule is defined in terms of segments that the suffix ignores during its leftward search. I have argued that there are three ways that segments can be ignored in the search: (1) if they do not bear the harmonic feature at all, (2) if they do not bear a contrastive value of the harmonic feature, or (3) if they do not bear a marked value of the harmonic feature. These relativizations serve to prune the search domain of segments that do not bear the relativized values. Once the search domain is set, copying takes place from the closest element, period.

An inspiration for our formulation of locality is the model of relativized minimality developed within syntactic theory, particularly by Rizzi (1990, 2001). Rizzi argued that defining the *most local* syntactic phrase for movement or government relations was not to be measured "as the crow flies", but only once the domain was pruned of irrelevant elements; thus A-bar elements were not relevant for the computation of A-locality, or vice versa. In fact the notion of relativization is implicit in all models of syntactic agreement which take phrases such as Prepositional Phrases to be "transparent" for agreement across them, for example:

(87) There lay(*s) in the grass two frogs

The implicit consensus in syntactic theory is that prepositional phrases do not enter into the search for agreement because they do not even bear ϕ-features to begin with; there is an implicit relativization of the search for an agreement controller to only those phrases that bear agreement features.

One of the closest precursors to the current theory of what is excluded from the domain of search within phonological theory can be found in the work of Yamada (1983). Yamada posits *The Complement Class Condition*, stating that what may benignly intervene between the source of assimilatory copying and the target of copying is anything that bears the complementary set of features from the source.[42] The current set of relativizations can be very naturally stated in these terms:

(88)
a. When [search: ±F], the complement class are those segments which don't bear [± F] at all (e.g. labial consonants do not bear [± back] in Turkish)
b. When [search: cF] the complement class are those segments that do not bear a contrastive value of the feature (e.g. *i* does not bear a contrastive value of [± back] in Finnish)
c. When [search: mF] the complement class are those segments that do not bear a marked value of the feature (e.g. *u* does not bear a marked feature of [± low] in Sibe).

Let us suppose that the notion of complement class can only be defined on the bases of (88) (i.e. the harmonic feature and certain values of it), and not on the basis of other features orthogonal to the one for which harmony is occurring. If so, then the spirit of Yamada's proposal, may be upheld here: elements in the complement of the feature-values for which the search is relativized are allowed to be "transparently" ignored in search.[43]

As far as visibility of interveners is concerned, the three value-relativization parameters constitute a characterization of possible and impossible harmony rules in terms of specification of the domain of search of the rules themselves.[44] Importantly, this parametric space prohibits the harmonic pattern in which both contrastive values of a feature are invisible. This is a strong and falsifiable prediction, not shared by a wide class of alternative theories of possible harmony systems.

This specific case, prohibited by the parameters, would be a harmony pattern in which, say, [$\pm$ back] harmony copied from {o,ö,a,ä}, but skipped {u,ü}; in other words, a system in which *both* members of a contrastive pair would be invisible to harmony, neither valuing nor undergoing the process. Such a system would clearly be prohibited under *all-value* visibility, since all values, including those in {u,ü} would be visible. Such a system would also be prohibited in *marked-value* harmony, since if only [+back] values were visible, then *u* would still remain the visible member of this pair. Finally, such a system would be prohibited under *contrastive-value* harmony, since both {u,ü}, being contrastive for back, would remain visible for harmony.

A second class of cases that would clearly be excluded would be those in which all contrastive values were *transparent* for harmony, so that only Finnish-type noncontrastive *i,e* would be copied from, and all others would result in last-resort insertion of a single feature-value. A third class of cases would be those in which all marked values were transparent, so that a rounding harmony language with the 8-vowel system of Kirghiz, Shor, or Altai would *only* copy round from *u*, and could do so at a distance.

In some previous research on vowel harmony, beginning with Kiparsky (1981), the visibility to search-and-copy of particular classes of feature-values but not others has been implemented in theories of underspecification, which hold that, for example, the noncontrastive or unmarked feature-values are simply not present at the point of representation at which harmony occurs. As discussed earlier in this chapter, underspecification models that treat skipped vowels as missing a value for the harmonic feature altogether face problems when other phonological processes demand the presence of those feature-values in the representation. A concrete case was presented in Section 3.4, in which Finnish *i* must possess [−back] in order to condition the rule of coronal assibilation; a second case mentioned in the same section is the fact that one instance of Hungarian *i* is transparent but two instances are not, which cannot be described if *i* has no specification at all for [$\pm$ back].

τ is either {all values of f_i contrastive for f_i, marked for f_i}
myVals V
myPosition P
myFeatsneeded F

while F is not empty:
· Go in direction δ and update P
· **if** P of type τ for any $f, f \in F$:
· · Copy Val(P, f) to V
· · Remove f from F

Figure 3.3: Parameterized Single-Pass Search until all Features Harmonized.

Further empirical problems with underspecification as a mechanism of determining (non)participation in assimilation have been more generally discussed in Mohanan (1991), McCarthy and Taub (1992), and Steriade (1995). The general spirit of these earlier proposals, that noncontrastive or unmarked feature-values are "invisible" to the harmony process, is captured here by relativization of the search. However, the letter of underspecification proposals, that the feature-values on such segments are literally absent from the representation, cannot be maintained.

Halle et al. (2000) discuss a pattern of Uyghur data in which contrastive *ä*, when unstressed and in a medial open syllable, raises to *i*, thereby becoming a transparent vowel for the purposes of subsequent suffixes added to the right; clearly the change from *ä* to *i* is in height alone, and should not affect the [−back] specification of the vowel, but by virtue of now being a noncontrastive instance of [−back], the vowel is skipped by subsequent harmonic searches. Echoing the conclusion drawn in the present model, they state that the pattern of vowels in harmony systems "derives from their role in the inventory of the language, rather than from their derivational history" (p.399).

The most general statement of what parametrization does is that it cuts out a swath of irrelevant segments, and in so doing, opens up the possibility of long-distance harmony. There is no obvious modification to Optimality Theoretic approaches to harmony, such as correspondence-based identity between vowels or contiguous domains of harmonizing vowels, that can restrict this relation to elements that bear the contrastive value or marked value of the harmonic feature, as the nature of such constraints does not allow reference to inventory properties of this sort. In the present model, the Search algorithm is in a sense, so "dumb" in copying from the closest leftward item in the domain, that if nothing is cut out of the way, all harmony would be strictly adjacent. It is precisely the relativization of adjacency based on these inventory properties that allows the empirical fit with locality patterns of vowel harmony. The Search algorithm we have developed up to this point is schematized in Figure 3.3.

The restrictiveness of the present theory results from the fact that the only way that items can be removed from the search is base on their feature-values, coupled with the inventory-derived properties of contrastiveness and markedness. In the next chapter, we will see that if the needy-segment requires anything of its donor besides contrastiveness or markedness for the harmonic feature, the stop-with-closest nature of minimal search can yield surprising results.

Notes

[21]This rule deletes a [± distr] on a [coronal, +cont, −voice] before another [± distr] on a [coronal, +cont, −voice].

[22]There is some debate about whether local apicalization itself is a cyclic or postcyclic rule. Given that non-final [± distr] deletion is a postcyclic rule, there are two possibilities. The first, as in the text, is that local apicalization follows postcyclic deletion. The second is that local apicalization is a cyclic rule, and that the deletion rule is more sensitive in its foci: it deletes all non-final [± distr] values except those that immediately precede a [−distr] segment. This escape from deletion could be implemented as a condition on deletion that protects segments that have undergone local apicalization (cf. Kiparsky (1993)). Given the complications of the deletion rule under this second possibility, the first possibility is adopted for clarity in the text, where the primary focus of discussion is the interaction of deletion with harmony.

[23]Finnish orthography uses *y* where I use *ü*; I do so to emphasize the parallelism among contrastively [−back] *ü*,*ö*,*ä*.

[24]The definition of contrastiveness here departs from that of Calabrese (1995), who makes crucial use of universal filters and deactivation statements. Bakovic (2003) has a definition of harmonic pairing similar to the one in the text:

(89) Harmonic pairing: A vowel x in a language L with a harmonic feature [± hf] is harmonically paired iff there is another vowel y in L's inventory and y differs from x only in terms of [± hf]

[25]Sagey (1988) demonstrates that the association lines of the autosegmental formalism must be interpreted as temporal overlap, rather than simultaneity (as was originally proposed in Goldsmith (1976)). Gafos (1996) shows that as autosegmental association encodes temporal overlap between feature and segment, then assuming that segments are temporally ordered on their own, locality must be strict. As demonstrated in a number of phonetic studies (Bessell, 1998; Walker, 1999; Shahin, 2002; Nevins and Vaux, 2003), locality is not strict; as a result, autosegmental approaches to locality thus cannot be upheld. See Hansson (2001) for extensive critique of Gafos (1996)'s claims about the non-existence of various harmony types.

[26]This particular rule is subject to certain restrictions in Finnish, as discussed extensively by Anttila (2003), who points out the role of metrical conditioning.

[27]In both Finnish and Votic, the process sensitive to the [−back] nature of *i* is a local consonant-vowel interaction, and the proposal that in both languages these processes are conveniently ordered after harmony would be accidentally fortuitous to an implausible degree; see Steriade (1995) for a similar conclusion.

[28]The short version of /a/ is rounded, and the long version of /eː/ is higher than that of its short counterpart; I return to these facts in Chapter 5.

[29]In fact, even the positionally-relativized aspect of excluding noncontrastive values in Karaim can be argued to be (trivially) at work in Turkish and Finnish as well – those two languages simply lack the positional distributional restrictions of Karaim. That is, all languages with contrastive relativization for harmony may in fact be analyzed as relativization to positionally contrastive values of [± F]. Section 3.11 explores the consequences of parametric ambiguity of this sort.

[30]See Anderson (1985) for an enlightening overview of the historical development of markedness and neutralization in phonological theory.

[31]See Calabrese (2005, p.376) for discussion of this distinction.

[32]The other two [+low] vowels in Sibe, /ɛ/ and /ö/ did not exist in Classical Manchu (CM). They were created in initial syllables in Sibe by an umlaut process, e.g. CM *omi* → Sibe *ömi* 'to drink' and CM *alin* → Sibe *ɛlin* 'mountain'. Since these low vowels were in the initial, stressed syllable in Sibe, they were not subject to markedness reduction.

[33]Whether or not there is a round or unround vowel in the suffix is determined by the roundness of the closest root vowel. Since the suffix vowel is a high vowel, it participates in [± round] harmony, as in Turkish.

[34]Odden (1980) discusses a parallel process of long-distance uvularization in Classical Mongolian, in which a dorsal suffixal consonant may copy marked [+back] (the feature responsible for uvularization in Classical Mongolian) across an intervening vowel: *biči-g* 'write-nmnlz.', *ide-g* 'eat-nmnlz.' vs. *jori-ʁ* 'intend-nmnlz.', *jiru-ʁ* 'draw-nmnlz.'.

[35]The lowering rule does not apply root-internally (cf. ɛdki, *'neighbor'*; üχa, *'to itch'*). This harmony process is thus more specific than a general Left-to-Right spreading from low vowels to dorsal consonants, hence the formulation here in terms of needy suffixes. The analysis differs from Halle et al. (2000) in this regard.

[36]As the Sanjiazi Manchu inventory reveals, the front round high vowel /ü/ is not contrastive for [± low], as it has no [+low] counterpart in the 7-vowel system. However, Li's (1996) book contains no example of suffix alternations for a word in which /ü/ follows a [+low] vowel. Bing Li (personal communication) reports that his field notes contain no such verb roots. The prediction here is that /ü/ should be transparent in such contexts (since it is not contrastive for [±low]) and allow either of [± low] to be copied across it

[37]In Section 4.5, we return to the pattern of "parasitic harmony", whereby the value-source must be of the same height as the value-seeker, which in combination with marked-value harmony, fills out a large portion of the rounding harmony typology in Altaic.

[38]Korn (1969, 102) describes a dialect of Kazakh in which only [−back] vowels are sources for [+round] harmony:

(90) *Kazakh [+round] harmony: only [−back,+round] are marked sources*:

üj-dö	house-loc
köl-dö	lake-loc
som-dan	rubble-abl
kul-da	servant-loc

Apparently, only (51)[a] is active among the context-sensitive marking statements in this language; Vaux (1993), however, notes that Kazakh has a contrast between the non-round non-high vowels e,ä; perhaps (51)[b] may only be active in two-height systems distinguished by [± high].

[39]Other loanwords with this [+back] before contrastively [−back] pattern include *klorofülli, miniatüüri, maniküüri, moleküüli, parfüümi, pseudonüümi, vampürri, amatööri* (Campbell, 1980, 250).

[40]As with many instances of interspeaker variation, loanwords such as *marttüüri* exhibit some *intra*-speaker variation as well. I adopt the general approach to intraspeaker variation taken within generative grammar: such speakers contain "multiple grammars"

(i.e. allow for more than one setting of a given parameter; Kroch (1989), Roeper (1999), Yang (2003)), with the choice among them determined by factors such as lexical item, style, register, and rate of speech.

[41] A further argument that "redundant" specifications are needed for vowels beyond what is required to differentiate the inventory as in the Successive Division Algorithm can be made based on Huave, which has a 5-vowel inventory *i,u,e,o,a*. According to underspecification analyses, it is redundant/predictable to have both [± round] and [± back] in the inventory among the non-low vowels. Nonetheless, the language demonstrates the phonological activity of both [± round] and [± back], allowing us to conclude that full specification for both of these features is needed.

The evidence that [± round] is active in Huave comes from the process of labial dissimilation, whereby /w/ and /f/ delete when preceded by the round vowels *o,u*. The Huave 3rd person plural suffix is /-f/, e.g. *a-rang-af* 'they make it', where the vowel preceding this suffix is a copy of the stem vowel. The process of labial dissimilation causes the /-f/ to dissimilate to the glottal fricative [j] when preceded by a round vowel (Kim, 2008, p.79-80), e.g. *a-xum-uj* 'they find it', *a-nol-oj* 'they have an issue'.

The evidence for [± back] comes from vowel + coda restrictions that result in "vowel breaking", a process of diphthongization (Kim, 2008). Specifically, the [−back] vowels *i,e* cannot surface before non-palatalized consonants, and the [+back] vowels *u,o,a* cannot surface before palatalized consonants (Kim, 2008, p.104-108). In both cases, diphthongization occurs, creating an offglide with the value [± back] of the following coda consonant, e.g /a-nchip/ → [a-nchiop] 'she approaches' and but /a-sap^j/ → [a-saip^j] 'she gives a gift'.

In conclusion, Huave's process of labial dissimilation requires [± round], while its process of diphthongization before coda consonants requires [± back]. The phonology of consonant-vowel interactions in Huave thus presents evidence that even when these features would be predictable one from the other, full specification of both features is still needed.

[42]Yamada's paper followed in the wake of a debate between Jensen (1974) and Odden (1980) regarding the former's attempt to define a "Relevancy Condition" for what may intervene transparently in terms of Major Class features (such as [±vocalic]); the existence, however, as pointed out by Odden, of harmonic features copied from vowels onto consonants across intervening segments clearly renders the appeal to Major Class features irrelevant. See Yamada (1983) for empirical problems with subsequent attempts to revise the relevancy condition, such as Battistella (1982).

[43]However, Yamada's proposal made no appeal to contrastiveness or markedness, and thus had to contain additional subconditions such as "If a complement defined in the Class Complement Constraint contains a neutral vowel, *then no other vowel than this* may intervene between focus and determinant" (p.56, italics added), which clearly run afoul of the spirit of the current proposal, namely to *derive* neutrality from properties of the inventory.

[44]Nevins (2007) provides a demonstration of the application of these same paradigmatically-based parameters of domain relativization to syntactic agreement to capture attested and unattested versions of the Person-Case Constraint, a family of co-occurrence restrictions on clitic clusters.

Chapter 4

Defective Intervention: When Search Comes Back Empty-Handed

"*If two vowels go walking, the first does the talking*."
Theodore Clymer.

4.1. The Grammatical Elements of Minimality

We have seen that the search procedure for vowel harmony can relativize its locality in very limited ways, according to a small parametric space. The establishment of the domain of search leads to an extremely myopic principle of copying from the closest element encountered. We now turn our focus to defining and describing a set of conditions that determine whether an element, once found, can be copied from or not. Elements that cannot be successfully copied from – even though they are included in the domain of search – are called "blockers", especially when there are elements further away in the domain that they block from ever being considered. The resulting effect of blockers, with whom seeking a harmonic value fails as soon as they are encountered, is an inviolable principle of locality. Once the search halts with a defective element in its path, no further-away elements may be considered.

Consider the following configuration: A is a needy *value-seeker* searching for a *source* of a value. The elements z, y, x are in its domain.

(1) [. . . x . . . y . . . z . . .] ← **A**

Ordinarily, the search will terminate with z. However, suppose that z is "defective" in some way; i.e., that, although it is in the domain of search, and the *closest* element to A within that domain, it does not satisfy some additional requirement R. By virtue of not satisfying R, z is defective. What is most interesting about these configurations is that both y and z bear the feature F that A needs, but, by failing to meet R, not only does z become excluded as a source for a value, but z also prevents y (which does meet R) from being a source as well. The concept of *minimality* expresses the idea that search cannot look past a defective element to a further one.

A syntactic example of defective intervention can be schematized with plural agreement in Icelandic:

(2) thadh finnst mörgum stúdentum tölvurnar
There find-sg/*pl some students-dat.pl computers-nom.pl
ljótar
ugly
"Some students find the computers ugly" (Holmberg and Hróarsdóttir, 2004)

In (2), the *value-seeker* is the Tense node, which has the need to value its ϕ-features. The closest element in its domain is the dative 'some students', which indeed has a [+plural] feature that T could potentially copy. However, there is an additional *source condition*: the source must be [+nominative]. The dative fails this condition. Not only does this lead to T failing to agree with the [+plural] Dative, it also leads to failure to agree with the [+plural] nominative, which is the *next* closest element in the domain. In other words, once the closest DP is found, *there is no second chance* (Chomsky, 2001).

When T cannot find a licit value-source in its domain, a last-resort operation supplies a default value for the needed feature(s). The default value for [± plural] in Icelandic is [−plural]. Default values are determinable within a given language based on considerations of morphosyntactic markedness, but to some extent, are language particular. Consider default agreement in Hindi, (3) which is masculine singular, vs. Russian (4) which is neuter singular:

(3) Lakshmi-ne Sumita-ko dekhaa
Lakshmi.fem-erg Sumita.fem-objctv saw-pfctv-**masc**
"Lakshmi saw Sumita" Hindi

(4) Ivanu nužno vrača
Ivan.masc-dat needs-**neut** doctor.masc-acc
"Ivan needs a doctor" Russian

In (3), neither (feminine) noun phrase can be copied from for agreement, because they are non-nominative. The last-resort gender value that is inserted in order for Hindi verbal agreement to interface with the morphology is masculine. In Russian, when neither of the (masculine) noun phrases in (4) can be copied for agreement because they are non-nominative, the value of last-resort default agreement is neuter. We therefore observe that the default-value in case of no licit sources for feature-copying in syntactic agreement must be determined on a language-specific basis. In Section 3.3, it was shown that Finnish and Uyghur, two languages with contrastive [$\pm$ back] harmony and identical inventories, differed in what the last-resort value to be inserted was in case search failed: for Finnish it was [$-$back], while in Uyghur it was [$+$back].

In a target-centric theory of locality, once the *search domain* is delimited, the *value-seeker* (*A* above) will always halt the search with the first element (z) in its domain. If the search has specified an additional *source condition* R on the search, and z fails to meet this condition, search will terminate in failure, **even if there is a potentially farther value-source y that does satisfy R.** This constitutes the essence of *defective intervention*. When search terminates in failure due to defective intervention, however, a last-resort operation of default value insertion provides a value to the value-seeker. (We can diagnose that defective intervention has occurred when a default value surfaces, even though there are non-default values on all elements in the domain).

4.2. No Second Chances after Search Fails

The defective intervention constraint governing the way that searching operates allows for exclusion of a wide class of patterns, and provides a formal foundation for expressing the conditions found in "parasitic harmony" (Steriade, 1981; Cole and Trigo, 1988). Let's begin by briefly exemplifying parasitic harmony based on Yawelmani rounding harmony, which is parasitic on shared height. Yawelmani has two underlying heights, distinguishing [+high] *i,u* from [−high] *a,o*. In order for a suffixal vowel to copy [± round] from a source vowel, the source vowel must bear the same value of [± high] as the value-seeker. The name "parasitic" harmony describes the fact that in order for the suffix to copy [± F] from a source, the two elements must already bear identical values for another feature [± G], the copying being thus "parasitic" on the requirement R of already-shared identity for an orthogonal feature.

(5) Parasitic harmony: Yawelmani [α high] suffix vowels can only copy [± round] from a source that is also [α high].

This is exemplified for the [−high] nondirective gerundial suffix of Yawelmani (Kuroda, 1967, p.14). When the vowel immediately to its left is also [−high], copying for [± round] successfully occurs (6-a-d), but when the vowel immediate to its left is [+high], search fails (6-e-g).

(6) Yawelmani nondir. gerund suffix seeks [± round] & R = [αhigh]:

a	gob-t**o**w	'take care of an infant (nondir. ger.)'
b.	hoyo:-t**o**w	'name (nondir. ger)'
c.	xat-t**a**w	'eat (nondir. ger.)'
d.	pana:-t**a**w	'arrive (nondir. ger)'
e.	giy-t**a**w	'touch (nondir. ger.)'
f.	mut-t**a**w	'swear (nondir. ger.)'
g.	w**o**:wul-t**a**w	'stand up (nondir. ger)'

In (6-e-f), the closest leftward vowel is [+high], and even when it is [+round] as in (6-f), it cannot be copied from, due to bearing "the wrong height". Turning to (6-g), this pattern exemplifies defective intervention: not only does the closest leftward vowel bear the wrong height; it also prevents the next-closest leftward vowel, which does bear the right height, from being copied from. This exemplifies a derivational search procedure in which there is no "lookahead": search immediately terminates in failure with the closest leftward element, even though only one element away there is a licit value-source. Ironically enough – due to the way the search procedure works when orthogonal requirements R such as [αheight]-sharing are imposed on the source – in (6-g), no value of [+round] at all is copied, even though both stem vowels are [+round]! [45]

In cases of parasitic harmony, the conditional requirement R is always the requirement for identity between the value-seeker and the value-source for a feature-value orthogonal to the harmonic value. In Yawelmani, in order for [$\pm$ round] to be successfully copied to the gerundive suffix, the source for that value must meet an additional subcondition of identity for height with the suffix.

Mailhot and Reiss (2007) provide a useful metaphor for search-termination due to the failure to meet a conditional requirement R. Suppose I ask you to find the first man with a parrot, and if it's red, tell me the parrot's name. That's very different from asking you to find the first man with a red parrot, and to tell me the parrot's name. In this second case, the search will be carried out until any red parrot is found. In the first set of instructions, however, the search will be carried out only until a parrot is found, whatever its color, and if that parrot happens not to be red, then the search terminates in failure. Defective intervention occurs when, once a potential value-source is found, if it does not meet the conditional requirement R, the search terminates. The essential premise of the current chapter is that phonology does not allow "search for a red parrot", if you will – in linguistic terms, there are never cases of parasitic harmony (searches for a source for the harmonic feature [F], with an additional source requirement on what licit sources must bear for the value [G]) that are of unbounded distance and blind to defective interveners.

Vowel harmonic searches may be simple "search for a parrot", e.g. [$\pm$ back] harmony in Finnish, with no additional source requirements, or they may be "searches for a parrot, copying *if it's red*", e.g. [$\pm$ round] harmony, parasitic on [αhigh], which fails if there is any closer [$-\alpha$ high] source. However, phonology does not allow vowel harmonic searches that are a "search for a red parrots", in other words, [$\pm$round] harmony parasitic on [αhigh] that is not blocked by any closer [$-\alpha$ high] source and can look past such defective interveners.

In what follows we examine a number of different types of defective intervention, subsuming the representative cases of what have been traditionally called "opaque segments" or "blockers" in the vowel harmony literature. The present model places important constraints on what can be a "blocker": (1) it must be something that is already included within the relativized domain of search (e.g. it must bear the harmonic feature), and (2) it must be a fail to be copied from due to a property orthogonal to the harmonic feature. This orthogonal property may be subsegmental or morphological.

Subsegmental conditional requirements are for pre-existing identity with an orthogonal feature – e.g. shared height as in Yawelmani [± round] harmony above and in Khalkha Mongolian (Section 4.5), shared roundness as in Bantu [± high] harmony (Section 4.4.1), or shared [± consonantal], as in Nawuri (Section 4.3). Importantly, such conditional requirements cannot refer to arbitrary properties of the source without reference to the value-seeker, nor can they demand featural *non*-identity between the needy suffix and the source.

Conditional requirements also may place restrictions on the morphological affiliation of the source, as in Jingulu (Section 4.7.1). Such requirements may demand that the source must be of a different morphological constituent than the value-seeker in order to be copied from. Again, failure to meet such a requirement by the closest-element in the domain leads to immediate defective intervention, even when licit downstream sources for the harmonic feature exist.

Importantly, these two types of conditional requirements on copyable sources cannot change the delimitation of the harmonic search domain: only the inventory-derived properties based on the harmonic feature itself, of Chapter 2, may do so. Once the relativized search domain has been established, the affixes are "stuck with" defective interveners in their search path that fail to meet an orthogonal requirement on being a licit source of value-copying. Such defective interveners are really only "partially defective": crucially, they do bear the value that needs to be copied, but by lacking something else, cause the entire search to halt.

4.3. "Opaque" Blockers of Harmony: Consonants in Nawuri

Let us examine a feature-based conditional requirement that leads to defective intervention in the [+round] harmony of Nawuri, a Kwa language of Ghana (Casali, 1995). Nawuri has the following vowel inventory:

(7) *Nawuri Vowel Inventory*:

[−back,−round]	[+back,+round]	[+back,−round]	
i	u		[+high,−low, +ATR]
ɪ	ʊ		[+high,−low, −ATR]
e	o		[−high,−low, +ATR]
ɛ	ɔ		[−high,−low, −ATR]
		a	[−high,+low, −ATR]

The noun class prefix gi-/gɪ-/gu-/gʊ-, copies [± round] (as well as [± ATR]) from a following vowel:

(8) Nawuri noun class prefix copies [± round] (Casali 1995: 651):

a. gɪ-baː 'hand'
b. gɪ-sɪbɪta 'sandal'
c. gʊ-sʊ 'ear'
d. gʊ-lɔ 'illness'
e. gi-ɲi 'tooth'
f. gi-keːliː 'kapok tree'
g. gu-jo 'yam'
h. gu-kuː 'digging'

In addition to a rightward search for [± round], Nawuri imposes an additional source requirement in Nawuri: the closest source of [+round] must be [−consonantal], as indicated by the '& R' condition below:

(9) Nawuri /gV-/ Noun Class Prefix Must:
ATR-Harmonize and Round Harmonize: δ = right, F = [± ATR; ± round & R = −consonantal]

The last-resort value of [± round] is [−round] in Nawuri. The scenario for defective intervention arises because Nawuri has labial consonants (*p,b,m,f*), which in this language contain the feature [+round]. (The glide *w* is a licit source for [+round] harmony, as it is [−consonantal] (Chomsky and Halle, 1968; Clements and Keyser, 1983).[46])

In Nawuri we find that when the closest rightward source of [+round] to copy from is a [+consonantal] element such as *f,b,p,m*, the search fails, because these consonants fail to satisfy the orthogonal conditional requirement. In fact, these consonants crucially *block* rounding harmony from a downstream [+round] vowel, instantiating a case of defective intervention: once the search for [+round] fails with [+consonantal] element in the domain, there is no second chance to copy from the second-closest element.

(10) *Nawuri [+round] copying fails with [+consonantal,+round] nearest source* (Casali 1995: 652):

a. gi-muː 'heat'
b. gi-fufuli 'white'
c. gi-pula 'burial'
d. gi-boːtoː 'leprosy'
e. gʊ-wɛ 'sympathy'
f. gʊ-wʊrʊ 'hat'

The failure of labial consonants to provide a value for [+round] results from the specific requirement of Nawuri that licit sources be [−consonantal].[47] There is evidence from many languages that labial consonants can induce vowel rounding: in Tulu (Sagey, 1986), both labial consonants and round vowels cause adjacent rounding of a high vowel. Conversely, in Akkadian (von Soden, 1969; Odden, 1994), dissimilation between labial consonants is blocked by round vowels.

A successful pattern of prefixal [+round] copying from a following labial consonant may be found in Southern Igbo (Ihiunu and Kenstowicz, 1994):[48]

(11) Igbo Gerundive vowel copies [+round] from following Labial consonant:

a. sị̀-sè 'stir, draw water'
b. ṭị́-ṭá 'bite'
c. mú-má 'know'
d. wù-wè 'take'
e. fụ́-fạ́ 'stuff'

While [+round] can be copied from labial consonants in Igbo, it cannot be copied from labial consonants in Nawuri.[49] As a result of a conditional requirement that the [± round] source must be [−consonantal], any labial consonant in Nawuri causes immediate failure of harmonic search for the prefix, even from a further source. Nawuri constitutes a case of "blocking" by "inert" feature-bearers which neither allow copying from themselves not from downstream licit sources. Lack of lookahead and immediate failure of the search leads to default insertion of [−round] in such cases.

The halting of search with a [+consonantal] element in gi-mu is illustrated by the following derivation.

(12) *Nawuri Noun Class Prefix begins Search for contrastive [± ATR] and all [± round], R = [−consonantal]*:

$$\underset{x_1}{\begin{bmatrix}\text{+cons}\\ \text{dorsal}\\ \text{-cont}\\ \text{+voice}\end{bmatrix}} \quad \underset{x_2}{\begin{bmatrix}\text{-cons}\\ \text{+high}\\ \text{-back}\\ \\ \end{bmatrix}} \quad \longrightarrow \quad \underset{x_3}{\begin{bmatrix}\text{+cons}\\ \text{labial}\\ \text{nasal}\\ \text{+round}\end{bmatrix}} \quad \underset{x_4}{\begin{bmatrix}\text{-cons}\\ \text{+high}\\ \text{+back}\\ \text{+high}\\ \text{+ATR}\\ \text{+round}\end{bmatrix}}$$

(13) *Nawuri Noun Class Prefix Finds Defective [± round]*:

$$\underset{x_1}{\begin{bmatrix}\text{+cons}\\ \text{dorsal}\\ \text{-cont}\\ \text{+voice}\end{bmatrix}} \quad \underset{x_2}{\begin{bmatrix}\text{-cons}\\ \text{+high}\\ \text{+back}\\ \\ \end{bmatrix}} \quad \not\rightarrow \quad \underset{x_3}{\begin{bmatrix}\textbf{+cons}\\ \text{labial}\\ \text{nasal}\\ \textbf{+round}\end{bmatrix}} \quad \underset{x_4}{\begin{bmatrix}\text{-cons}\\ \text{+high}\\ \text{+back}\\ \text{+high}\\ \text{+ATR}\\ \text{+round}\end{bmatrix}}$$

(14) *Nawuri Noun Class Undergoes Default Insertion for [−round]*:

$$\underset{x_1}{\begin{bmatrix}\text{+cons}\\ \text{dorsal}\\ \text{-cont}\\ \text{+voice}\end{bmatrix}} \quad \underset{x_2}{\begin{bmatrix}\text{-cons}\\ \text{+high}\\ \text{-back}\\ \\ \textit{-round}\end{bmatrix}} \quad \longrightarrow \quad \underset{x_3}{\begin{bmatrix}\text{+cons}\\ \text{labial}\\ \text{nasal}\\ \text{+round}\end{bmatrix}} \quad \underset{x_4}{\begin{bmatrix}\text{-cons}\\ \text{+high}\\ \text{+back}\\ \text{+high}\\ \text{+ATR}\\ \text{+round}\end{bmatrix}}$$

(15) *Nawuri Noun Class Prefix continues Search for contrastive [± ATR]*:

$$\underset{x_1}{\begin{bmatrix}\text{+cons}\\ \text{dorsal}\\ \text{-cont}\\ \text{+voice}\end{bmatrix}} \quad \underset{x_2}{\begin{bmatrix}\text{-cons}\\ \text{+high}\\ \text{-back}\\ \\ \text{-round}\end{bmatrix}} \quad \underset{x_3}{\begin{bmatrix}\text{+cons}\\ \text{labial}\\ \text{nasal}\\ \text{+round}\end{bmatrix}} \quad \longrightarrow \quad \underset{x_4}{\begin{bmatrix}\text{-cons}\\ \text{+high}\\ \text{+back}\\ \text{+high}\\ \text{+ATR}\\ \text{+round}\end{bmatrix}}$$

(16) *Nawuri Noun Class Prefix Finds and Copies [+ ATR]*:

$$
\begin{array}{ccccc}
x_1 & x_2 & x_3 & \Rsh & x_4 \\
\begin{bmatrix}\text{+cons}\\ \text{dorsal}\\ \text{-cont}\\ \text{+voice}\end{bmatrix} &
\begin{bmatrix}\text{-cons}\\ \text{+high}\\ \text{-back}\\ \textbf{+ATR}\\ \text{-round}\end{bmatrix} &
\begin{bmatrix}\text{+cons}\\ \text{labial}\\ \text{nasal}\\ \text{+round}\end{bmatrix} & &
\begin{bmatrix}\text{-cons}\\ \text{+high}\\ \text{+back}\\ \text{+high}\\ \text{+ATR}\\ \text{+round}\end{bmatrix} \\
\text{g} & \text{i} & \text{m} & & \text{u}
\end{array}
$$

Although it has not been suggested before, the case of Nawuri defective intervention by labial consonants falls under the logic of parasitic harmony as well. Parasitic harmony is any case in which copying of feature F between segments X and Y applies only if X and Y share the same value for an orthogonal feature G, G≠F. In this case, the orthogonal feature is [−consonantal], a property of both the value-seeker and any licit goal. Nawuri differs from Southern Igbo in that the latter copies [± round] from any source, whereas Nawuri requires that sources be [−consonantal] in order to be copied from.[50] The minor addition of this conditional requirement to the Search procedure developed in Chapter 2 yields great divergences in the surface patterning of Igbo and Nawuri [+round] copying.

The revised Search algorithm is presented in Figure 4.1, where R is a matrix of conditional requirements on each feature that is needed by the copying affix. R may be empty (i.e. when none of the needed features impose conditional requirements on their source).

The revised algorithm in Figure 4.1 states that as soon as the first element P bearing the relativized feature value is found, a check is performed: does the element P meet the source requirement for that feature? If not, the search for that feature terminates, leaving the affix to last-resort means of default value insertion.

4.4. Parasitic Harmony: Conditional Requirements for [± high], [± round] and [± ATR] copying

The discussion of Nawuri illustrated a case in which labial consonants bearing the right feature caused a crash when they intervened between a prefixal vowel and a round vowel, and we absorbed this case under the rubric of parasitic harmony: requirement for identity with respect to an orthogonal feature. In the next three subsections we consider a number of

τ is either {all values of f_i contrastive for f_i, marked for f_i}
myVals V
myPosition P
myFeatsneeded F
myConditionalRequirements(F) = R

while F is not empty:
· Go in direction δ and update P
· **if** P of type τ for any $f, f \in F$:
· · **if** R is true of P:
· · · Copy Val(P, f) to V
· · · Remove f from F
· · **else:**
· · · **exit**

Figure 4.1: Parameterized Single-Pass Search with Conditional Requirements.

“classic” cases of parasitic harmony, that instantiate the inviolable nature of the Search principle halting with the closest element in the domain, even if a “better” one can be found just a bit further away. Section 4.4.1 demonstrates [± high] harmony dependent on identity of [± round] in Kisa. This cases of conditional harmony is then jointly paired with [± ATR] harmony dependent on [± high] in a single language, Kimatuumbi 4.4.2.

4.4.1. Height Harmony in Bantu

Consider the pattern of height harmony found in a majority of the Bantu languages, in which the affixal vowels *i,u* lower when preceded by *e,o*, as found in five-vowel systems such as Kisa:

(17) *Kisa Vowel Inventory*:

[−back,−round]	[+back,+round]	[+back,−round]	
i	u		[+high,−low]
e	o		[−high,−low]
		a	[−high,+low]

Copying of contrastive [± high] eliminates a from the search domain, because the feature [± high] is not contrastive in the context of [+low]. Failure to find anything in the search domain results in last-resort insertion of [+high], as in (18-e).

(18) *Kisa applicative suffix copies contrastive [± high] from its left (Hyman, 1999, 238)*:

a. tsom-el-a ‘pierce-appl.’
b. rek-el-a ‘set.trap-appl.’
c. βis-il-a ‘hide-appl.’
d. fu:ng-il-a ‘lock-appl.’
e. βa:mb-il-a ‘spread.out-appl.’

Up to now this is a straightforward process of contrastive feature-copying of the type seen throughout Chapter 3. However, Hyman notes that Kisa (as well as many other Bantu languages, including Kinyarwanda, Kirundi, Kinyakore, Luganda, Haya, Jita, Shambaa, Shi, Bemba, Yao, and Shona) exhibits a different pattern of harmony with the reversative suffix *-ul-/-ol*. Unlike the applicative suffix *-il/el-*, which may copy [−high] from either *e* or *o*, the reversative copies high from only *o*. This instantiates a case of parasitic harmony: the reversative suffix can only copy [± high] from a source that shares the same roundness value as it. This difference between the two suffixes is contrasted in (19) and (20).

(19) Kisa Applicative Suffix Must:
Height-Harmonize: δ = left, F = [c: high]

(20) Kisa Reversative Suffix Must:
Height-Harmonize: δ = left, F = [c: high & R = +round]

The requirement in (20) dictates that the reversative suffix will not copy contrastive [−high] from *e* – even though *e* bears contrastive [−high] – due to the failure of *e* to satisfy the conditional requirement of sharing [+ round].

(21) *Kisa Reversative suffix copies contrastive [± high] only when source is [+round]*:

a. tsom-ol-a 'pull out'
b. βis-ul-a 'reveal'
c. fu:ng-ul-a 'unlock'
d. βa:mb-ul-a 'spread apart'
e. rek-ul-a 'spring trap'

The fact that (21-b) fails to copy contrastive [−high] from the preceding *e* instantiates a case of parasitic harmony in which the required orthogonal shared feature is [± round].[51] It is well-known from cases such as Yawelmani and from Khalkha Mongolian (to be discussed in Section 4.5) that [± round] harmony may parasitically require orthogonal featural identity for [± high]. The Kisa pattern demonstrates that this height-roundness relation can work in either direction, as [± high] harmony requires orthogonal featural identity for [± round].[52]

Under the rubric of defective intervention, the prediction is that a contrastively [−high] vowel that is not [+round] will terminate the search in failure immediately, even if there is another contrastively [−high] vowel that *is* [+round] further to the left. The Shona repetitive suffix *-urur/-oror-* has the same requirement as the Kisa reversative suffix in (20). When this suffix attaches to disyllabic noun stems of the form [*o. . .e*][53], we can diagnose defective intervention:

(22) *Shona repetitive suffix -urur/-oror (Fortune, 1981): search for contrastive*

[−high] fails with u:

a. dz**o**k-**o**ror 'replant'
b. ts**e**ts-**u**rur 'regrind'
c. z**oté**ng-**ú**rúdz-a 'resell'

Let us resist an explanation of (22) in terms of a restriction that the repetitive suffix cannot look further than one syllable to the left. Many types of harmony are unbounded and can look more than one syllable away (e.g. Sibe [+low] harmony, Finnish [± back] harmony, both discussed in Chapter 3), and hence the question is, why would a putative restriction to one syllable in (22) coincide with the fact that this is a case of parasitic harmony? In other words, if it were merely an incidental property of (22) that conditionally-parasitic search is restricted to one syllable, we would expect to find another language where conditionally-parastic search is unbounded, e.g. in which [−high] copying could look as far as needed for a [+round] source. However, no such cases exist: whenever an orthogonal featural requirement is placed on copyable sources, search halts at the first element in the domain. The logic of defective intervention provides an answer why: such systems are excluded by the theory of *no-second-chance* locality. There are no "long-distance" cases of parasitic harmony at all within the typology of vowel harmony precisely because the search algorithm allows "search and copy [−high] from the closest [−high] vowel only if that vowel is [+round]" and does not allow "search and copy [−high] from the closest [−high, +round] vowel".

The derivation of (22-c), with search terminating immediately when R is not met, is illustrated in (23):

(23) *Repetitive Suffix begins Search for contrastive [± high], R = [+round]*:

$$
\begin{matrix} x_1 & x_2 & x_3 & x_4 & x_5 & x_6 & \leftarrow & x_7 \\
\begin{bmatrix} \text{-voc} \\ \text{coronal} \\ \text{+cont} \\ \text{+voice} \end{bmatrix} &
\begin{bmatrix} \text{+voc} \\ \text{-low} \\ \text{+back} \\ \text{+round} \\ \text{-high} \end{bmatrix} &
\begin{bmatrix} \text{-voc} \\ \text{coronal} \\ \text{-cont} \\ \text{-voice} \end{bmatrix} &
\begin{bmatrix} \text{+voc} \\ \text{-low} \\ \text{-back} \\ \text{-round} \\ \text{-high} \end{bmatrix} &
\begin{bmatrix} \text{-voc} \\ \text{dorsal} \\ \text{+nasal} \end{bmatrix} &
\begin{bmatrix} \text{-voc} \\ \text{dorsal} \\ \text{-cont} \\ \text{+voice} \end{bmatrix} & &
\begin{bmatrix} \text{+voc} \\ \text{-low} \\ \text{+back} \\ \text{+round} \\ \\ \end{bmatrix}
\end{matrix}
$$

(24) *Repetitive Finds Defective Instance of [± high]*:

x_1	x_2	x_3	x_4	↚	x_5	x_6	x_7
-voc	+voc	-voc	+voc		-voc	-voc	+voc
coronal	-low	coronal	-low		dorsal	dorsal	-low
+cont	+back	-cont	-back		+nasal	-cont	+back
+voice	+round	-voice	**-round**			+voice	+round
	-high		**-high**				

(25) *Repetitive Undergoes Last-Resort Insertion of [+high]*:

x_1	x_2	x_3	x_4	x_5	x_6	x_7
-voc	+voc	-voc	+voc	-voc	-voc	+voc
coronal	-low	coronal	-low	dorsal	dorsal	-low
+cont	+back	-cont	-back	+nasal	-cont	+back
+voice	+round	-voice	-round		+voice	+round
	-high		-high			*+high*
z	o	t	e	ŋ	g	u

The effect of conditional requirements on the source introduced in the present chapter differs significantly from the effect of relativization of the domain introduced in Chapter 3. Relativization of the domain prunes certain elements from ever being searched for in the first place, whereas parasitic requirements on sources cannot remove elements from the domain; instead they merely impose strict requirements on whether copying can actually take place.

4.4.2. Kimatuumbi: Conditional [± ATR] and Conditional [±high] Harmony

Kimatuumbi has harmonic copying for two features, each of which imposes its own conditional subrequirement. Kimatuumbi has 7 vowels, with i,ɪ,u,ʊ contrastive for [± ATR] and ɪ,ʊ,ɛ,ɔ contrastive for [± high]:

(26) *Kimatuumbi Vowel Inventory*:

[−round,−back]	[+round,+back]	[−round,+back]	
i	u		[+high, +ATR, −low]
ɪ	ʊ		[+high, −ATR, −low]
ɛ	ɔ		[−high, −ATR,−low]
		a	[−high, −ATR,+low]

Odden (1991, 282) notes that the harmonic copying of [±high, ± ATR] initiated by the applicative suffix in Kimatuumbi "has a peculiar condition on it, namely that the vowel ɛ does not cause *u* to assimilate": while the other [−ATR,−high,−low] vowel ɔ does allow copying from it, as do the other

[−back,−round] vowels i,ɪ. (The [+low] vowel *a* does not allow copying, but is not contrastive for [± high] or [± ATR].)

(27) *Kimatuumbi Applicative Suffix Can Copy from* i,u,ɪ,ʊ,ɔ:

a. tyam-ul-ya 'to sneeze on'
b. tip-ul-ya 'to break off maize for'
c. gul-ul-ya 'to wash dishes for'
d. yʊng-ʊl-ya 'to answer for'
e. tɪk-ʊl-ya 'to break for'
f. lɔnd-ɔl-ya 'to find a witch for'
g. chɛk-ul-ya 'to laugh for'
h. nɛm-ul-ya 'to dance for'

Given the "rectangular" shape of high front, high back, mid front, and mid back vowels, what is unique about the pattern of defective interveners is that only one corner of the rectangle is a defective source: the mid-front vowels. Why should the high front and mid back vowels participate in harmony, forming a "diagonal" natural class of participants to the exclusion of the mid front vowels? The answer, I propose, is that the Kimatuumbi Applicative suffix combines two distinct types of parasitic harmony; one for each feature being copied.

Copying of contrastive [± high] requires that the source be [+round], and copying of contrastive [± ATR] requires that the source be [+high]. Kimatuumbi thus combines the exact parasitic rule harmony observed in Kisa with the parasitic harmony rule observed in Canadian French, whereby [± ATR] harmony only occurs among vowels that are [+high] (Poliquin, 2006); see also Menominee, in Section 5.3.2 for a height-parasitic [± ATR] harmony.[54] I assume that in cases in which multiple features are copied from the same source there may be an ordering imposed among the copying processes;[55] in this case, copying of [± high] (and default Insertion of [+high] in case of failure) precedes copying of [± ATR], as represented by > below.

(28) Kimatuumbi Applicative Suffix Must:
Height-and-ATR Harmonize: δ = left, F = [c: high & R = +round > c:ATR & R = +high]

The combination of these two distinct conditional requirements on each feature being copied mean that [+round] ʊ,ɔ may be copied from for contrastive [± high], and [+high] i,u,ɪ,ʊ may be copied from for contrastive [± ATR]. Neither of ɛ,a may be copied from for either feature. Default insertion for these values are [+high] and [+ATR], as seen in (27-a,g-h).

Both of these harmony requirements are met when the source is ʊ, as illustrated in the following derivation:

(29) Applicative Finds ʊ:

$$\overset{x_1}{\begin{bmatrix}\text{+voc}\\ \text{-low}\\ \text{+round}\\ \text{+high}\\ \text{-ATR}\end{bmatrix}} \quad \overset{\leftarrow}{} \quad \overset{x_2}{\begin{bmatrix}\text{+voc}\\ \text{-low}\\ \text{+round}\\ \end{bmatrix}}$$

(30) Applicative Copies Contrastive [± high] from [+round] source

$$\overset{x_1}{\begin{bmatrix}\text{+voc}\\ \text{-low}\\ \textbf{+round}\\ \textbf{+high}\\ \text{-ATR}\end{bmatrix}} \quad \overset{\Lsh}{} \quad \overset{x_2}{\begin{bmatrix}\text{+voc}\\ \text{-low}\\ \text{+round}\\ \textbf{+high}\end{bmatrix}}$$

(31) Applicative Copies Contrastive [± ATR] from [+high] source

$$\underset{\text{ʊ}}{\overset{x_1}{\begin{bmatrix}\text{+voc}\\ \text{-low}\\ \text{+round}\\ \textbf{+high}\\ \textbf{-ATR}\end{bmatrix}}} \quad \overset{\Lsh}{} \quad \underset{\text{ʊ}}{\overset{x_2}{\begin{bmatrix}\text{+voc}\\ \text{-low}\\ \text{+round}\\ \textbf{+high}\\ \textbf{-ATR}\end{bmatrix}}}$$

A summary of the derivations for other preceding vowels appear in the following table. (I assume, based on the Kimatuumbi vowel inventory, an implicational rule of the form [−high] → [−ATR]).

(32) Summary of Harmonic Copying from Preceding Vowels:

ɔ	[− high] copying succeeds; [± ATR] copying succeeds
u	[+ high] copying succeeds; [± ATR] copying succeeds
i	[± high] copying fails, [+high] inserted; [± ATR] copying succeeds
ɪ	[± high] copying fails, [+high] inserted; [± ATR] copying succeeds
ɛ	[± high] copying fails, [+high] inserted; [± ATR] copying fails
a	[± high] copying fails, [+high] inserted; [± ATR] copying fails

We now observe that the pattern of non-copying with ɛ is no longer peculiar: the fact that ɪ but not ɛ allows [± ATR] copying is due to parasitic harmony for height, and the fact that ɔ but not ɛ allows [±high] copying is due to parasitic harmony for roundness. The combination of two distinct independently attested harmonic requirements involving the same pairs of features provides an explanation for why the Kimatuumbi Applicative suffix refuses to copy from ɛ on two distinct counts.

4.5. Intervention & Skipping Coexist in Khalkha [± round] Harmony

Khalkha Mongolian illustrates an interesting combination of "transparent" and "opaque" interveners in its pattern of [± round] harmony: the high vowels u,ʊ are defective interveners, forcing failed termination of the search, while the high vowel *i* is excluded from the relativization, and skipped over in the search. The combination of marked-value relativization (Chapter 3) and parasitic harmony (this chapter) yields the first system we have considered in which defective intervention and skipped vowels coexist. The vowel system of Khalkha is in (33):

(33) *Khalkha Mongolian Vowel Inventory (Svantesson, 1985)*:

[−round,−back]	[+round,+back]	
i	u	[+high, +ATR]
	ʊ	[+high, −ATR]
e	o	[−high, +ATR]
a	ɔ	[−high, −ATR]

Khalkha has both [± ATR] harmony and [± round] harmony. The vowel *i* is excluded from the search domain of both of these features, but for two different reasons. The value-relativization for [± ATR] harmony is for contrastive values of [± ATR], and hence all vowels except for *i* are included in the search domain. The value-relativization for [± round] harmony, however, is for marked values, as evident by the inclusion of the vowels u,ʊ in the search domain for [± round] harmony, even though they are not contrastive for this feature.

The patterning of *i* on the one hand and u,ʊ on the other are quite different in [± round] harmony: even though none of them are contrastively [± round], the latter pair blocks further leftward harmony, while the former is skipped completely. In the present model, blocking can only be the result of conditional restrictions on copyable sources, and indeed in Khalkha, as in many Altaic [± round] harmony systems, we are dealing with a case of parasitic harmony based on height. These two factors in the harmony system are summarized as follows:

(34) Khalkha [+round] Harmony:
  a. Unmarked Transparency: *i* does not bear the marked value of [± round], stays invisible
  b. Wrong-Height Defective Intervention: u,ʊ are [+round] but block further search

The phenomenon of defective intervention can be observed in the interaction between the perfect suffix, which demands marked [± round] from a source that is [− high], and the causative suffix, which is [+high], and hence can defectively intervene between the causative and the root, yielding a failure of [+round] to be copied from the root, since leftward search halts immediately.[56]

Note that each suffix has no difficulty in copying contrastive [± ATR] from the immediately closest leftward vowel.

(35) *Defective Intervention in Khalkha Rounding Harmony*:
  a. tor-oːd 'be.born-perf.'
  b. ɔr-ɔːd 'enter-perf.'
  c. tor-uːl-eːd 'be.born-caus.-perf
  d. ɔr-ʊːl-aːd 'enter-caus.-perf.'

The harmonic requirements of each Khalkha suffix are shown in (36) and (37).

(36) *Khalkha Causative suffix must:*
  ATR Harmonize:, δ = left, F = [c: ATR]

(37) *Khalkha Perfect suffix must:*
ATR-Harmonize and Round Harmonize: δ = left, F = [c: ATR; m:round & R = $-$high]

The derivation of ɔr-ʊːl-aːd is provided below[57]:

(38) Causative Begins ATR-Harmonize in ɔr-ʊːl:

$$\underset{x_1}{\begin{bmatrix}\text{+voc}\\ \text{-high}\\ \text{+back}\\ \text{+round}\\ \text{-ATR}\end{bmatrix}} \quad \underset{x_2}{\begin{bmatrix}\text{-voc}\\ \text{+sonorant}\\ \text{-nasal}\\ \text{-lateral}\end{bmatrix}} \quad \longleftarrow \quad \underset{x_3}{\begin{bmatrix}\text{+voc}\\ \text{+high}\\ \text{+back}\\ \text{+round}\\ \\ \end{bmatrix}} \quad \underset{x_4}{\begin{bmatrix}\text{-voc}\\ \text{+sonorant}\\ \text{-nasal}\\ \text{+lateral}\\ \\ \end{bmatrix}}$$

(39) Causative Finds and Copies [$-$ATR] :

$$\underset{\text{ɔ}}{\overset{x_1}{\begin{bmatrix}\text{+voc}\\ \text{-high}\\ \text{+back}\\ \text{+round}\\ \text{-ATR}\end{bmatrix}}} \quad \Lsh \quad \underset{\text{r}}{\overset{x_2}{\begin{bmatrix}\text{-voc}\\ \text{+sonorant}\\ \text{-nasal}\\ \text{-lateral}\end{bmatrix}}} \quad \underset{\text{ʊ}}{\overset{x_3}{\begin{bmatrix}\text{+voc}\\ \text{+high}\\ \text{+back}\\ \text{+round}\\ \textbf{-ATR}\end{bmatrix}}} \quad \underset{\text{l}}{\overset{x_4}{\begin{bmatrix}\text{-voc}\\ \text{+sonorant}\\ \text{-nasal}\\ \text{+lateral}\\ \\ \end{bmatrix}}}$$

The next affix to begin harmony is the Perfect Suffix, which follows the Causative in morphological structure:

(40) Perfect Begins ATR-and-Round Harmonize in ɔr-ʊːl-aːd:

$$\overset{x_3}{\begin{bmatrix}\text{+voc}\\ \text{+high}\\ \text{+back}\\ \text{+round}\\ \text{-ATR}\end{bmatrix}} \quad \overset{x_4}{\begin{bmatrix}\text{-voc}\\ \text{+sonorant}\\ \text{-nasal}\\ \text{+lateral}\end{bmatrix}} \quad \longleftarrow \quad \overset{x_5}{\begin{bmatrix}\text{+voc}\\ \text{-high}\\ \text{+back}\\ \\ \\ \end{bmatrix}} \quad \overset{x_6}{\begin{bmatrix}\text{-voc}\\ \text{coronal}\\ \text{-cont}\\ \text{+voice}\end{bmatrix}}$$

(41) Perfect Finds and Copies contrastive [$-$ATR]:

$$\overset{x_3}{\begin{bmatrix}\text{+voc}\\ \text{+high}\\ \text{+back}\\ \text{+round}\\ \text{-ATR}\end{bmatrix}} \quad \Lsh \quad \overset{x_4}{\begin{bmatrix}\text{-voc}\\ \text{+sonorant}\\ \text{-nasal}\\ \text{+lateral}\end{bmatrix}} \quad \overset{x_5}{\begin{bmatrix}\text{+voc}\\ \text{-high}\\ \text{+back}\\ \\ \textbf{-ATR}\end{bmatrix}} \quad \overset{x_6}{\begin{bmatrix}\text{-voc}\\ \text{coronal}\\ \text{-cont}\\ \text{+voice}\end{bmatrix}}$$

(42) Perfect Finds Defectively [+high] Instance of [+round]:

$$\begin{matrix} x_3 & \not\leftarrow & x_4 & x_5 & x_6 \\ \begin{bmatrix} \text{+voc} \\ \textbf{+high} \\ \text{+back} \\ \textbf{+round} \\ \text{-ATR} \end{bmatrix} & & \begin{bmatrix} \text{-voc} \\ \text{+sonorant} \\ \text{-nasal} \\ \text{+lateral} \end{bmatrix} & \begin{bmatrix} \text{+voc} \\ \text{-high} \\ \text{+back} \\ \\ \text{-ATR} \end{bmatrix} & \begin{bmatrix} \text{-voc} \\ \text{coronal} \\ \text{-cont} \\ \text{+voice} \end{bmatrix} \end{matrix}$$

(43) Perfect undergoes Last-Resort insertion of [−round]:

$$\begin{matrix} x_3 & x_4 & x_5 & x_6 \\ \begin{bmatrix} \text{+voc} \\ \text{+high} \\ \text{+back} \\ \text{+round} \\ \text{-ATR} \end{bmatrix} & \begin{bmatrix} \text{-voc} \\ \text{+sonorant} \\ \text{-nasal} \\ \text{+lateral} \end{bmatrix} & \begin{bmatrix} \text{+voc} \\ \text{-high} \\ \text{+back} \\ \textit{-round} \\ \text{-ATR} \end{bmatrix} & \begin{bmatrix} \text{-voc} \\ \text{coronal} \\ \text{-cont} \\ \text{+voice} \end{bmatrix} \\ \text{ʊ} & \text{l} & \text{a} & \text{d} \end{matrix}$$

The high vowels u,ʊ are defective interveners included in the search because they are marked values of round. On the other hand, *i* is transparent for rounding harmony, excluded from the search in the first place, because it is unmarked for [−round].[58] *i* is also transparent for [± ATR] harmony, because it is noncontrastive for this feature. Both types of skipping over *i* are illustrated in the final example below:

(44) *Khalkha [+round] Harmony for the Comitative Suffix:*[59]

a. nar-tai ‘mountain-comit.’
b. ner-tei ‘name-comit.’
c. ʊːl-tai ‘mountain-comit.’
d. xun-tei ‘person-comit,’
e. ɔd-tɔi ‘star-comit.’
f. bičig-tei ‘letter-comit.’
g. mɔri-tɔi ‘horse-comit.

In (44-g), we observe that contrastive [± ATR] harmony and marked [± round] skip over *i*. In (44-c-d), we notice that with this same suffix, u,ʊ fail to provide a [+round] value, even though they are indeed round. Similar effects of transparency in Khalkha Mongolian may be observed with the accusative suffix *-ig*:

(45) *Transparency of suffixal i in Khalkha Mongolian (Svantesson et al., 2005,*

50):

a. poor-ig-o 'kidney-acc-refl.'
b. xɔɔlʒ-ig-ɔ 'food-acc.-refl'
c. mʊʊr-ig-a 'cat-acc.-refl.'
d. suulʒ-ig-e 'tail-acc-refl'.

The Khalkha system is of interest to the theory of locality because of difference in the two high vowels *i* and *u*. Neither ends up providing a value of [± round], but for different reasons and with different effects: *i* because it is excluded from search in the first place, and *u* because it is included in search but defectively of the wrong height.

4.6. Defective Intervention due to All-Value Relativization in Oroch

The comparison of the high vowels in rounding harmony in Khalkha, in which defective intervention is caused by vowels with *marked* values for [± round] of the wrong height, is distinct from what occurs in the Tungusic language Oroch (Tolskaya, 2008), in which defective intervention may be caused by *all* values of [± round] of the wrong height.[60] In Oroch, as in Khalkha, none of the high vowels i,u,ʊ are contrastive for [± round] (since there are no front rounded or back unrounded high vowels in the language), but in fact, all are included in the search, and defectively so. Unlike Khalkha, in Oroch, the high vowel *i* blocks further [± round] copying as well; as soon as search hits any instance of a high vowel, the procedure halts with that closest item, despite further downstream licit instances of [+round].[61] As introduced in Section 3.12, Oroch has the following vowel inventory:

(46) *Oroch Vowel Inventory*:

[−round,−back]	[−round,+back]	[+round]	
i		u	[+high, +ATR]
		ʊ	[+high, −ATR]
æ	ə		[−high, +ATR]
	a	ɔ	[−high, −ATR]

Like Khalkha, Oroch exhibits [± ATR] harmony in addition to round harmony.

(47) *Oroch Accusative Suffix, Focus Suffix must:*

ATR-Harmonize and Round Harmonize: δ = left, F = [c: ATR; ±round & R = −high]

As a result of a relativization in Oroch in which *all* values of [±round] are visible within the domain of search, any high vowel that is the closest to a [± round]-seeking suffix induces defective intervention (48-b-c):

(48) *Oroch Rounding Harmony blocked by [+high] vowels*:

a. ɔtɔŋgo-vɔ-dɔ 'kayak-acc-foc'
b. ɔtɔŋgo-dʊ-da 'kayak-dat-foc'
c. ɔtɔŋgo-ɲi-da 'kayak-3sg-foc'

In Oroch ɔtɔŋgo-ɲi-da (48-c), the focus suffix searches leftward for any value of [± round], the search domain thereby including *i*. As soon as *i* is encountered, its failure to be [−high] causes an immediate end to the search, which terminates in failure, triggering last-resort insertion of [+round]. If only the algorithm were different, the search might have persisted in looking further for a [± round] source that was [−high]. But the derivational nature of the search procedure is such that as soon as the closest element is hit, search stops, whether the encountered item is a good source or not. Oroch blocking by *i* demonstrates yet another instance in which the no-second-chance property of search within a domain is non-negotiable.

4.7. Derived Environment Effects in Harmony

The majority of cases of vowel harmony that have been discussed in preceding sections are cases where an affix copies harmonic features from another morpheme, or where a root copies from within the root. As we have seen in Chapter 2, affixes may be fairly indiscriminate in who they copy from: either a leftward suffix or a root can provide a harmonic value, and what matters is simply what's closest in terms of immediate precedence within the search domain.

We also find languages in which root vowels themselves exhibit harmonic copying from affixes, thereby exhibiting a left-to-right directionality in the case of suffixation (e.g. Kalenjin (section 2.9) Noske (2001) on Turkana, Mahanta (2007) on Assamese). In these cases as well, such root vowels are also indiscriminate in choosing their value-source, enacting copying from either a root vowel or a suffix vowel, depending on phonological closesness alone once the domain has been featurally-relativized.

The model of conditional requirements on sources developed thus far allows for the possibility of value-seekers being "picky" not only about the orthogonal features a licit source may bear, but also about its morphological affiliation. Such cases are rare in the literature and often assumed under the mantle of derived environment effects. However, "derived environment effects" is largely a phenomenological term for which a variety of heterogeneous mechanisms have been proposed. Even when harmonic copying shows apparent derived-environment requirements (namely, requirements that the value-seeker and value-source come from distinct morphemes, and hence that their co-occurrence within the same word is *derived*), one would not necessarily expect them to show blocking effects, whereby a vowel with the wrong morphological affiliation blocks a vowel with the right morphological affiliation. By contrast, modeling the demand for distinct morphological affiliations for the value-seeker and value-source in terms of conditional requirements on copying leads to strong predictions about defective intervention in such cases.

4.7.1. Jingulu: Morphological Requirements on Source

The [± high] harmony in Jingulu (Pensalfini, 2002), a language of North-Central Australia, imposes a morphological requirement on what can serve as a licit source. Jingulu has the three vowels *i,u,a*:

(49) *Jingulu Vowel System*:

[−round,−back]	[+round, +back]	[−round,+back]	
i	u		[+high]
		a	[−high]

In a three-vowel system of this sort, we can assume an implicational statement of the form [−round,+high] → [−back] (see Section 2.11 for an introduction to such "redundancy" statements for languages lacking *ü*). Jingulu [−round] vowels copy marked [+high], starting from left-to-right:

(50) *Jingulu [−round] vowels harmonize with following affixes:*

a.	ngaja	ng**i**j**i**-mindi-yi	'see / see.1.dual.incl-fut.'
b.	bardarba	b**i**rd**i**rb**i**-rni	'younger brother / younger sister'
c.	kunyarrba	kuny**i**rrb**i**-rni	'dog / dog-female'
d.	ngarrabaja	ng**i**rr**i**b**i**j**i**-wurru-nu	'tell / tell.3pl-did'

Pensalfini (2002) characterizes the class of suffixes that may be copied from in terms of morphological domain: gender suffixes with nouns/adjectives and tense/agreement suffixes with verbs.[62] I assume that root vowels alternating between *a/i* are lexically marked as requiring harmony for the feature [± high], and if they fail to find one, [−high] is assigned as the default value (Pensalfini, 1997, Ch.2). Nonalternating root vowels are lexically [+high].

(51) Jingulu Needy Root Vowels Must:
High-Harmonize: δ = right, F = [m: high & R = Non-Root]

In Jingulu, root vowels requiring harmony initiates their own search iteratively, starting with the leftmost such vowel, as shown in for *ngiji-mindi-yi* in the following derivation:

(52) First Root Vowel Begins High-Harmonize in ŋ*ǎja-mindi*:

$$\begin{array}{cccccc c} x_1 & x_2 & \rightarrow & x_3 & x_4 & x_5 & x_6 \\ \begin{bmatrix}\text{-voc}\\ \text{dorsal}\\ \text{+nasal}\end{bmatrix} & \begin{bmatrix}\text{+voc}\\ -\text{round}\\ \end{bmatrix} & & \begin{bmatrix}\text{-voc}\\ \text{coronal}\\ \text{-cont}\\ \text{+voice}\\ \text{-ant}\\ \text{+distr}\end{bmatrix} & \begin{bmatrix}\text{+voc}\\ \text{-round}\\ \end{bmatrix} & \begin{bmatrix}\text{-voc}\\ \text{labial}\\ \text{+nasal}\end{bmatrix} & \begin{bmatrix}\text{+voc}\\ \text{-round}\\ \text{-back}\\ \text{+high}\end{bmatrix} \end{array}$$

(53) First Root Vowel Finds and Copies [+high] on suffix vowel in ŋ*ǎja-mindi*:

$$\begin{array}{ccccc c c} x_1 & x_2 & x_3 & x_4 & x_5 & \Rsh & x_6 \\ \begin{bmatrix}\text{-voc}\\ \text{dorsal}\\ \text{+nasal}\end{bmatrix} & \begin{bmatrix}\text{+voc}\\ -\text{round}\\ \\ \textbf{+high}\end{bmatrix} & \begin{bmatrix}\text{-voc}\\ \text{coronal}\\ \text{-cont}\\ \text{+voice}\\ \text{-ant}\\ \text{+distr}\end{bmatrix} & \begin{bmatrix}\text{+voc}\\ \text{-round}\\ \end{bmatrix} & \begin{bmatrix}\text{-voc}\\ \text{labial}\\ \text{+nasal}\end{bmatrix} & & \begin{bmatrix}\text{+voc}\\ \text{-round}\\ \text{-back}\\ \textbf{+high}\end{bmatrix} \end{array}$$

(54) Second Root Vowel Begins High-Harmonize in ŋ*iǰa-mindi*:

$$
\begin{array}{cccccccc}
x_1 & x_2 & x_3 & x_4 & \rightarrow & x_5 & x_6 \\
\begin{bmatrix}\text{-voc}\\ \text{dorsal}\\ \text{+nasal}\end{bmatrix} &
\begin{bmatrix}\text{+voc}\\ -\text{round}\\ \\ \text{+high}\end{bmatrix} &
\begin{bmatrix}\text{-voc}\\ \text{coronal}\\ \text{-cont}\\ \text{+voice}\\ \text{-ant}\\ \text{+distr}\end{bmatrix} &
\begin{bmatrix}\text{+voc}\\ \text{-round}\\ \\ \end{bmatrix} & &
\begin{bmatrix}\text{-voc}\\ \text{labial}\\ \text{+nasal}\end{bmatrix} &
\begin{bmatrix}\text{+voc}\\ \text{-round}\\ \text{-back}\\ \text{+high}\end{bmatrix}
\end{array}
$$

(55) Second Root Vowel Finds and Copies [+high]:

$$
\begin{array}{ccccccc}
x_1 & x_2 & x_3 & x_4 & x_5 & \Lsh & x_6 \\
\begin{bmatrix}\text{-voc}\\ \text{dorsal}\\ \text{+nasal}\end{bmatrix} &
\begin{bmatrix}\text{+voc}\\ -\text{round}\\ \\ \text{+high}\end{bmatrix} &
\begin{bmatrix}\text{-voc}\\ \text{coronal}\\ \text{-cont}\\ \text{+voice}\\ \text{-ant}\\ \text{+distr}\end{bmatrix} &
\begin{bmatrix}\text{+voc}\\ \text{-round}\\ \\ \textbf{+high}\end{bmatrix} &
\begin{bmatrix}\text{-voc}\\ \text{labial}\\ \text{+nasal}\end{bmatrix} & &
\begin{bmatrix}\text{+voc}\\ \text{-round}\\ \text{-back}\\ \textbf{+high}\end{bmatrix}
\end{array}
$$

(56) Implicational Statement Supplies [−back] for [−round,+high] vowels:

$$
\begin{array}{cccccc}
x_1 & x_2 & x_3 & x_4 & x_5 & x_6 \\
\begin{bmatrix}\text{-voc}\\ \text{dorsal}\\ \text{+nasal}\end{bmatrix} &
\begin{bmatrix}\text{+voc}\\ -\text{round}\\ \text{-back}\\ \text{+high}\end{bmatrix} &
\begin{bmatrix}\text{-voc}\\ \text{coronal}\\ \text{-cont}\\ \text{+voice}\\ \text{-ant}\\ \text{+distr}\end{bmatrix} &
\begin{bmatrix}\text{+voc}\\ \text{-round}\\ \textbf{-back}\\ \text{+high}\end{bmatrix} &
\begin{bmatrix}\text{-voc}\\ \text{labial}\\ \text{+nasal}\end{bmatrix} &
\begin{bmatrix}\text{+voc}\\ \text{-round}\\ \text{-back}\\ \text{+high}\end{bmatrix} \\
\text{n} & \text{a} & \check{\jmath} & \text{i} & \text{m} & \text{i}
\end{array}
$$

Although root vowels requiring harmony copy marked [+ high] from a rightward source, importantly, a [+high] vowel within the same root cannot provide a value for this harmony process. Naturally, due to the very restricted nature of value-parametrization of Chapter 3, there is no way to exclude [+high] root vowels from the search. Nonetheless, the morphological requirement that only non-root vowels may be copied from can be encoded as a conditional requirement on the copying procedure.

A [+high] vowel in the root will *block* the copying process from a suffixal high vowel. Stated differently, an underlying high vowel in the root will prevent any *preceding* vowels from raising. In Jingulu, the conditional requirement R on harmonic sources is that their morphological affiliation is [−root]. As a result, [+high] *root* vowels, although they have the required value, cannot provide a source of [+high] copying. Moreover, not only do these [+high] root vowels fail to provide a source of [+high] copying, they also block [+high] copying from any further potential determinants, thereby constituting defective interveners:

(57) *Defective Intervention in Jingulu [+high] Harmony*:

a.	warlaku	warlaku-rni	'dog/ dog-fem.'
b.	mamambiyaka	mamambiy**i**k**i**-mi	'soft / soft.veg'
c.	ngamurla	ngamurl**i**-rni	' big / big-fem.'
d.	ankila	ankil**i**-rni	'cross-cousin / cross-cousin-fem.'

In (57-a-d), the first two root vowels cannot copy from the suffixal vowel due to a closer, defective root vowel. In (57-c-d) the initial root vowel cannot copy from the suffixal vowel due to a closer defective root vowel. The derivation for *ankili-ṇi* (57-d) is shown below: the first high vowel's search fails as soon as it hits the [+high] second root vowel. As a consequence of defective intervention, default [−high] is inserted. The final root vowel's search then proceeds and successfully copies from a [−root] vowel. In the following diagrams, although [± root] is a morphological affiliation and not a subsegmental feature, I include it in the same column for convenience of visual inspection.

(58) First Root Vowel Begins High-Harmonize in *ankila-ṇi*:

x_1	→	x_2	x_3	x_4	x_5	x_6	x_7	x_8
+voc −round		-voc dorsal +nasal *+root*	-voc dorsal -cont -voice *+root*	+voc -round -back +high *+root*	-voc coronal +sonorant -nasal +lateral *+root*	+voc −round *+root*	-voc coronal -ant -distr +nasal *-root*	+voc -round -back +high *-root*

(59) First Root Vowel Finds Defective [+high] vowel in *ankila-ṇi*:

x_1	x_2	x_3	↛	x_4	x_5	x_6	x_7	x_8
+voc −round	-voc dorsal +nasal *+root*	-voc dorsal -cont -voice *+root*		+voc -round -back **+high** ***+root***	-voc coronal +sonorant -nasal +lateral *+root*	+voc −round *+root*	-voc coronal -ant -distr +nasal *-root*	+voc -round -back +high *-root*

(60) First Root Vowel Undergoes Default Insertion of [−high]:
Implicational [−high] → [+back]:

x_1	x_2	x_3	x_4	x_5	x_6	x_7	x_8
+voc −round **+back** *-high*	-voc dorsal +nasal *+root*	-voc dorsal -cont -voice *+root*	+voc -round -back +high *+root*	-voc coronal +sonorant -nasal +lateral *+root*	+voc −round *+root*	-voc coronal -ant -distr +nasal *-root*	+voc -round -back +high *-root*

(61) Second Needy Root Vowel Begins High-Harmonize:

x_1	x_2	x_3	x_4	x_5	x_6	→	x_7	x_8
+voc −round +back -high	-voc dorsal +nasal *+root*	-voc dorsal -cont -voice *+root*	+voc -round -back +high *+root*	-voc coronal +sonorant -nasal +lateral *+root*	+voc −round *+root*		-voc coronal -ant -distr +nasal *-root*	+voc -round -back +high *-root*

(62) Second Needy Root Finds and Copies [+high] from [−root] vowel:

x_1	x_2	x_3	x_4	x_5	x_6	x_7	↱	x_8
+voc −round +back -high	-voc dorsal +nasal *+root*	-voc dorsal -cont -voice *+root*	+voc -round -back +high *+root*	-voc coronal +sonorant -nasal +lateral *+root*	+voc −round **+high** *+root*	-voc coronal -ant -distr +nasal *-root*		+voc -round -back **+high** *-root*

(63) Implicational Statement provides [−back]:

x_1	x_2	x_3	x_4	x_5	x_6	x_7	x_8
+voc −round +back -high	-voc dorsal +nasal *+root*	-voc dorsal -cont -voice *+root*	+voc -round -back +high *+root*	-voc coronal +sonorant -nasal +lateral *+root*	+voc −round **-back** +high *+root*	-voc coronal -ant -distr +nasal *-root*	+voc -round -back +high *-root*
a	ŋ	k	i	l	i	ṇ	i

In the above derivation, the first root vowel's search ends in failure, due to immediately encountering a [+root] high vowel among its rightward search path. Subsequently, the third root vowel initiates its own search independently, and with no defective [+root] vowels in its path, successfully copies [+high] from an affixal vowel.

The Jingulu restriction on harmonic copying can be viewed as a derived environment restriction: the harmony rule can only be triggered by "derived contexts", and two root-internal vowels do not constitute a derived context. But I will assimilate this case of morphological defective intervention to the subsegmental featural cases because the logic is the same, Moreover, the mechanism through which search halts as soon as the right vowel with the wrong morphological affiliation is found will be general enough to hold even for cases in which word-formation does not contain distinct derivational cycles for root and non-root material.[63]

To summarize our progress, in this section we examined a conditional requirement on the source that was morphological in nature. In the previous sections, we examined more familiar cases of parasitic harmony, in which the source precondition is based on a subsegmental feature. Importantly, the defective intervention rubric unites both sets of cases: once a parrot is found, if it's not red, there is no second chance to look for another. This is an inviolable principle of locality that supersedes the language-specific settings of what constitute licit sources of feature valuation.

4.8. Conclusions

Defective Intervention is the result of the fact that the Search principle inviolably stops at the first element in the domain. The only ways that items can be excluded from the domain are by the two inventory-derived properties of markedness or contrastiveness for the harmonic feature itself. Crosslinguistic variation between harmony systems based on featural properties unrelated to the harmonic feature cannot serve to exclude a segment from the search domain. When such conditions are parametrically imposed on a given set of potential value-sources, they can lead to failure of copying, but cannot render the defective elements invisible.

I have proposed a restrictive set of conditional requirements on licit sources: they can only refer to two types of properties: featural identity and morphological non-identity.

The first type of restrictions that lead to "blocking" are requirements of Orthogonal Featural Identity between Value-Seeker and Source.[64] We have observed cases of parasitic harmony where the conditionally-identical feature may be [± high], [± round], and [± consonantal]. Crucially, however, we find no cases of conditional-requirement harmony in which *non-identity* of an orthogonal feature is required (cf. also Cole and Trigo (1988, 23)); e.g. there is no system in which round harmony copies *only* from an opposite height, and otherwise yields defective intervention.

The second type of restrictions that lead to a terminal halting of the search procedure are requirements of Morphological Distinctness between Value-Seeker and Source. From a functional perspective, these might be somewhat natural conditions to expect: failure to harmonize with segments that are already in the same morpheme. However, in Chapter 2, we have demonstrated through both Turkish and Finnish root harmonization that some languages indeed engage in an active process of harmonic copying between elements within the same morpheme. More importantly, whether or not a condition that elements can only harmonize with heteromorphemic segments is functionally-motivated or not, we would not necessarily expect this condition to yield blocking, but nonetheless it does. Languages that require heteromorphemic sources for harmony cannot "keep looking" past the first element bearing the harmonic feature within the domain.

In summary, what causes "blocking" is never markedness or contrastiveness of the harmonic feature. Blocking can only be a function of the values of orthogonal features, and crucially without respect to their markedness or contrastiveness. Only the predicate of identity between source and goal for the orthogonal feature can cause success or failure of copying.

Defective intervention is an empirical hypothesis about harmony systems, that can be modeled in any target-centric theory of locality, possibly even within differing representations of harmonic features (Kaye et al. (1985), Archangeli and Pulleyblank (1994), van der Hulst and van de Weijer (1995), Clements (2003)). However, the no-lookahead property of the search principle seems to require a target-initiated, procedural search, which is not obviously replicable in declarative models.[65]The conclusion that search terminates with the closest potential source, regardless of more optimal sources downstream, is a generalization about phonological locality with parallels in other modules of linguistic computation where lack of lookahead and strict minimality lead to a failure to copy a needed feature from either the closest or the second-closest source.

Notes

[45]A related example to Yawelmani [+round] harmony's requirement that the source be the same value for [± high] can be found in Kachin Khakass (Korn, 1969, 103), in which [+high] suffixes only copy [+round] from a source that meets the conditional requirement of being [+high]:

(64) *Kachin Khakass failure to copy [+round] from [−αhigh] vowels:*
   a. kün-nu 'day-acc.'
   b. küs-tüŋ 'bird-gen.'
   c. öd-ir 'kill-inf.'
   d. ok-tɨŋ 'arrow-gen.'

A reviewer cites the pattern of Eastern Khanty (Finno-Ugric), as described by Kiparsky and Pajusalu (2003), as another potentially relevant case of defective intervention conditioned by identical-height: it is claimed that the unpaired vowel *i* blocks [± back] harmony between a [+low] vowel in a suffix and [+ back] *a* in a root. A straightforward application of the current model would be one in which [± back] harmony includes all values of [± back] in the search, but is parasitic on shared values of height, such that [−low] *i* would cause termination of the search and insertion of the default value ([−back], as in Uyghur (Section 3.3). Kiparsky and Pajusalu (2003) contains only diagrammatic patterns of the form *[[a . . . i] . . . ä]*, without any actual examples from Eastern Khanty, and contains the claim that "in all the languages [considered in the paper], the vowel *i* is neutral (unpaired)".

Two exhaustive sources on Eastern Khanty, Abondolo (1998) and Filchenko (1979) contain descriptions of the inventory in which *i* is not at all neutral, but rather paired with [+back] ɨ. According to Filchenko (1979, p.4), ɨ occurs word-initially, medially, and finally, e.g. ɨɣeta 'to hang', wɨj 'craftiness', jɨr 'sacrifice', qəlɨ 'corpse', and forms minimal pairs with *i*, e.g. il 'front' vs. ɨl 'down/bottom' (p.9), as well as alternating in harmony wär-i-tä 'to fish with a dam' vs. at-ɨ-ta 'to fence off' (p.10). The picture of Eastern Khanty vowel harmony that Abondolo (1998) and Filchenko (1979) paint is one of a fully unremarkable system of harmony in which i,ɨ fully participate in [± back] harmony, yielding a formally identical pattern to that of Turkish. Against this background, there is sufficient unclarity surrounding Kiparsky and Pajusalu (2003)'s description of the "*[[a . . . i] . . . ä]*" pattern, and as attractive as it may be to apply a defective-intervention analysis in terms of shared-height, the available Eastern Khanty data do not provide sufficient support for it.

[46]Except where discussing glides, I have employed the feature [± vocalic] in discussions and in diagrams. I adopt the feature system whereby glides and vowels are [−consonantal] and glides and consonants are [−vocalic]; see Nevins and Chitoran (2008) for discussion.

[47]Casali shows that the failure to copy from [+round] consonants is not due to a dissimilation effect (whereby round vowels would become [−round] before labial consonants), as there is no root-internal derounding in words like *kufe, kuba, ku:-pu*.

[48]A related pattern in Igbo is described in Hyman (1975, 53).

[49]Labial consonants also interact with [± round] harmony between vowels in Warlpiri (Nash, 1980, p.87ff). The relationship between [+round] vowels and labial consonants could instead be mediated by assuming a language-specific representations in which a

feature [+labial] underlies both round vowels and labial consonants, and to which either vowels or consonants may be sensitive. Such a move, however, would require harmonizing vowels to be searching for [+labial] (and not [+round]) in languages such as Tulu, Igbo, and Nawuri, and a subsequent implicational statement that [+labial] on a vowel requires [+round] in these languages. To keep matters both simpler (avoiding a language-specific search for [+labial] by vowels) and more parallel with other cases of rounding harmony, I will adopt the proposal that labial consonants are [+round] in Tulu, Akkadian, Nawuri, Igbo, and other languages that display consonant-vowel interactions of this sort, acknowledging that its phonetic implementation may be one of either labial tension or labial constriction. An alternative possibility, based on Hyman (1975, p.53-54,154) is an implicational feature statement on [−cons] segments that [+round] ↔ [+labial]; in this case, the search in Nawuri and Igbo would be for [± labial], and successful copying would entail [+round].

[50]Nawuri contains a rounding contrast on labial (pʷ, bʷ, fʷ,mʷ) and non-labial consonants (kʷ, čʷ, sʷ). We follow Casali in analyzing these segments as bisegmental sequences of consonant followed by glide. Coupled with the bisegmental analysis, the defective intervention approach correctly predicts that /pʷ, fʷ, bʷ, mʷ/ block rounding harmony due to the defective [+round] feature on the [+consonant] segment, while /kʷ, čʷ, sʷ/ bear [+round] only on the [−consonantal] glide and therefore allow for copying to the prefix: cf. sʊ-sʷa 'to grease' in which rounding harmony occurs vs. gɪ-bʷaːrːuː 'water yam' (Casali, 1995, p.656). Note that the failure for rounding harmony to occur when p,b,m,f intervene cannot be accounted for in terms of contrastive [± round] alone (i.e. by positing that p,b,m,f are [−round] while pʷ, bʷ, fʷ,mʷ are [+round]), *pace* Halle et al. (2000), since *non-labial* unrounded consonants do not block harmony (cf. gu-kuː 'digging' and kʊkʷɪ 'to be sufficient' (Casali, 1995, p.656).

[51]Hyman (1999) also discusses the pattern of vowel height harmony in Pende, in which, unlike many other Bantu harmony patterns, *e* copies [−high] from the low vowel *a*. In Pende as well, *o* cannot copy [−high] from *a*, even though both are [+back], confirming that [± round] really is the feature on which parasitic harmony is dependent.

[52]In light of the fact that [± high] harmony can depend on source values for [± round] (in Kisa) and that vice versa can exist as well (in Yawelmani), the idea that parasitic harmony derives from asymmetric dependence of one feature on another (e.g. van der Hulst and Smith (1989)) cannot be maintained.

[53]Beckman (1997) discusses the height harmony within Shona roots, yet does not include suffixal harmony with [+round] suffixes. Additional disyllabic verb roots cited in Beckman's article that could potentially induce defective intervention with the [+round] repetitive suffix include *kobodek* 'become empty', and *bover* 'collapse inwards'.

[54]A similar case may be Maasai glide-induced harmony, which Cole and Trigo (1988) analyze as a special case of [± ATR] harmony parasitic on [+high].

[55]Müller (2008) proposes that certain syntactic cases require ordering among the features that are satisfied on a single node.

[56]We observe a similar type of defective intervention in Bashkir, in which [± round] harmony requires the source to be [−high], e.g. *tön-ö* 'night-poss.', *qol-o* 'slave-poss.' versus. ət-ə 'dog-poss.', at-ʌ 'horse-poss'. While *u,i,ü* do not occur in non-initial syllables and thus cannot be used to test defective intervention, Poppe (1964, 20) reports that the

[+high,+round,−voc] glide *w* yields defective intervention for leftward [+round] harmony, due to its wrong height: töðöwsə 'builder', and qorowlʌ 'loaded'.

[57]Vowel length is omitted from the representations.

[58]Hansson (2007a), exploring the role of segmental similarity in harmony, discusses hypothetical predictions for cases of vowel harmony, and argues that "an opaque segment must always be more similar to the target than a transparent segment within the same system". The current model gives teeth to this notion of similarity: an opaque segment is included in the domain of relativization, while a transparent segment is not.

[59]The comitative suffix appears as -tai,-tei, and -tɔi; according to Svantesson, the diphthong *oi* is not allowed in Khalkha on the surface, and *ei* appears where *oi* is expected. Arguably there is copying of [+ATR,+round] at the point of harmony, as this suffix is further copied from for [+round]: ovs-tei-go: 'grass-abl.-refl.'.

[60]Svantesson et al. (2005, 51) note that the Ulaanbaatar dialect of Mongolian is developing a variable process of *i* blocking harmony, in which case its analysis becomes identical to that of Oroch in the text.

[61]In Kaun (1995, 72), additional Oroch examples in which high vowels are reported to block rounding harmony are *oggiča* 'dried out' and *dokčina* 'to hear' (no morpheme boundaries provided for either). Kaun (1995, 77) also reports similar data of *i* blocking rounding harmony for the Tungusic language Ulcha.

[62]Pensalfini (2002) suggests that the only suffixes that the root may copy from are those which are spelled out in the same morphological cycle as the root.

[63]Hualde (1991) discusses a derived environment condition on raising of Basque suffixal *a* by a preceding [+high] root or suffixal vowel, which does not apply within roots. Polgárdi (1998, p.64) discusses [± ATR] harmony in Korop (a Benue-Congo language of Cameroon), for which prefixes (but not roots) need to copy rightward. As pointed out by Łubowicz (2002), in these cases, it may be that only affixal vowels need to harmonize.

[64]See also Rose and Walker (2004) for a discussion of the role of orthogonal identity as a prerequisite for harmony, although they explicitly exclude blocking effects from this rubric.

[65]I leave open whether derivational versions of Optimality Theory, such as McCarthy (2007), are able to model the procedural nature of search halting in failure at the first encountered element when it is defective.

Chapter 5

Domain-Limitations on Search

Imagine you're on a search. Locality (i.e. when you stop searching) has two components: Minimality (stopping at the closest potential element) and Boundedness (an extrinsically defined limit on how far you can go in your search).

In syntactic theory, the principles of Minimality rule out searching for an element beyond a defective one.[66] In *wh-* movement, for example, minimality demands that matrix C's search halt with *whether*, and cannot look further to move *who*:

(1) a. [$C_{\text{[needs: wh]}}$ [did [you [wonder [$_{CP}$ [$\text{whether}_{\text{[wh, -movable]}}$ [$\text{who}_{\text{[wh, +movable]}}$ [danced]]]]]]]]

b. *Who did you wonder whether t_{who} danced?

Boundedness, on the other hand, rules out searching for an element past certain kinds of nodes. In *wh-* movement, matrix C's search cannot look past *the*, regardless of what lies beyond (see Adger (2003, Ch.10) for an introductory exposition):

(2) a. [$C_{\text{[needs: wh]}}$ [did [you [hear [$\text{the}_{boundary}$ [recitals [of [which $\text{poem}]_{\text{[wh, +movable]}}$]]]]]]]

b. *Which poem did you hear the recitals of $t_{\text{which poem}}$ last night ?

In the previous chapter, we discussed the phonological analogue of cases such as (1) within a theory of relativizations and defective interveners. In this we develop the phonological instantion of (2): when search "hits a brick wall" and ends not because a defective element of the relevant type was found, but because the domain of search was bounded.

We will discuss two types of boundedness that constrain the extent of the search in vowel harmony: distance parameters, that regulate the extremes of how far a search can proceed before giving up, and sonority hurdles, that impose barriers to search beyond them.

5.1. Distance Parameters Limit the Extent of Search

While a great deal of phonological theory centers on the minimality side of locality, the fact that some assimilation processes are strictly local, some are long-distance, and yet a third class are in-between (in a way to be made immediately precise) is a dimension that must be developed in a more inclusive model. Just as syntactic theory does not collapse the distinction between minimality and boundedness,[67] the study of vowel harmony must keep separate the delimitation of the search domain – in terms of relativized closesness – from phenomena in which certain elements simply cannot look past an extrinsically-defined limit of segments or syllables. As Yamada (1983, 63) remarks, "In general, the problem of distance should be clearly distinguished from what may intervene in a phonological rule".

5.1.1. Unbounded vs. Syllable-Bounded Search

Odden (1994) contributed a set of important observations that certain assimilation and dissimilation processes have "locality-bounds" that are not solely determined by relativization to a class of interveners. As Odden points out, these distance parameters impose additional constraints on the maximal distance between interacting segments. Many cases of crosslinguistic variation in the same harmony process can only be understood in these terms. Consider the process of long-distance nasal harmony in Kikongo suffixes (Ao, 1991; Odden, 1994; Piggott, 1996):

(3) *Kikongo long-distance nasalization with the affixes -il/in and -ul/un*:

a. sa.ki.di.la 'to congratulate for'
b. man.ti.**n**a 'to climb for'
c. ku.du.mu.ki.si.**n**a 'to cause to jump for'
d. ma.ki.nu.**n**u 'it was planted'
e. wu.man.tu.**n**u 'it was climbed'

As (3) shows, the applicative suffix in Kikongo can extend the search at an unbounded distance, copying marked [+nasal] from up to three syllables away in (3). If no marked [+nasal] is found, the default [−nasal] will be inserted.[68] The formulation of the long-distance harmony rule in Kikongo is straightforward following the domain-relativization parametrization developed in Chapter 3:

(4) Kikongo applicative and passive suffixes must:
Nasal-Harmonize: δ = left, F = [m: ±nasal]

In Lamba, another Bantu language exhibiting nasal harmony, the process imposes stricter demands on the locality of the search domain. Here, the search cannot look further than one syllable away from the needy suffix (I assume that re-syllabification occurs immediately upon suffixation, before even the harmonic search is inititated). The failure of the Lamba reversative and perfective suffix to copy from a [+nasal] source more than one syllable away is illustrated in (5) (Odden, 1994):

(5) *Lamba syllable-bounded nasalization:*
   a. fi.su.**l**u.ka 'plaster-revers.'
   b. mi.nu.**n**u.ka 'swallow-revers.'
   c. pa.ti.**l**e 'scolded-perf.'
   d. u.mi.**n**e 'dry-perf.'
   e. ma.si.**l**e 'plaster-perf.'

While Lamba and Kikongo both exhibit leftward harmony for marked values of [+nasal] with the same relevant segments in the inventory, the parametric difference between these two languages must be grammatically encoded in terms of a boundedness condition. Specifically, there is a distance parameter that is set to *unbounded* in the case of Kikongo but to just *one syllable* in Lamba. In terms of implementation, the search algorithm must be supplemented with a 'counter', so that if the leftward search exceeds one syllable beyond the Lamba suffix, the search is halted. Every time a syllable frontier is passed, the counter is incremented. With each step along the segmental path, the syllable counter is checked whether it has exceeded 1. This is included as a new parameter β, expressed in (6):

(6) Lamba perfect suffix must:
Nasal-Harmonize: δ = left, β = 1 σ, F = [m: ±nasal]

As soon as one syllable boundary has been crossed (as tracked by an independent pointer on the syllable tier), the search among segments must end. This is illustrated for the distinct patterns of the two Lamba words *u.mi.ne* (where a [+nasal] vowel is encountered within one traversed syllable) and *ma.si.le* (where no [+nasal] vowel is encountered within one traversed syllable); their divergent derivations are illustrated below:

(7) Lamba Applicative Begins Search for marked [+nasal] in *u.mi.ne*:

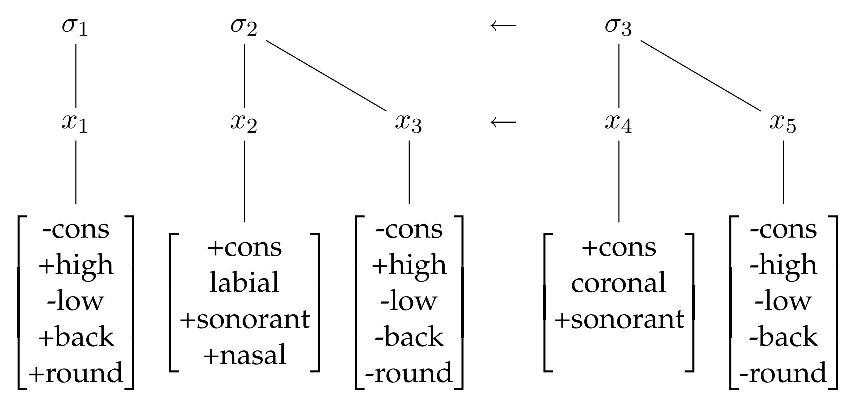

(8) Lamba Applicative Finds Nasal Source:

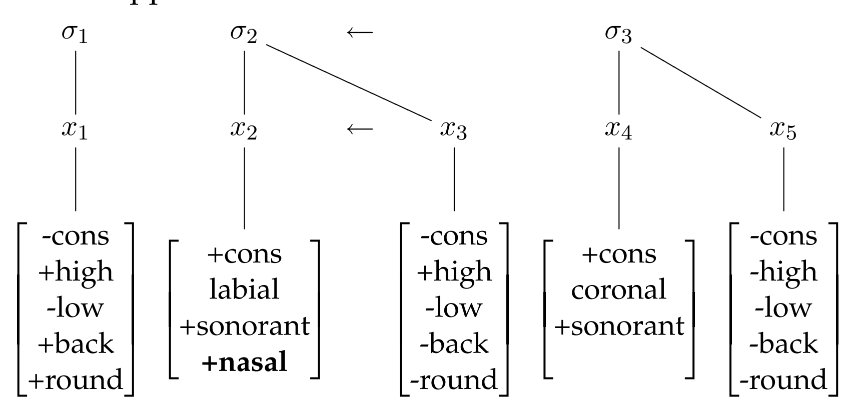

(9) Lamba Applicative Copies Marked [+nasal]:

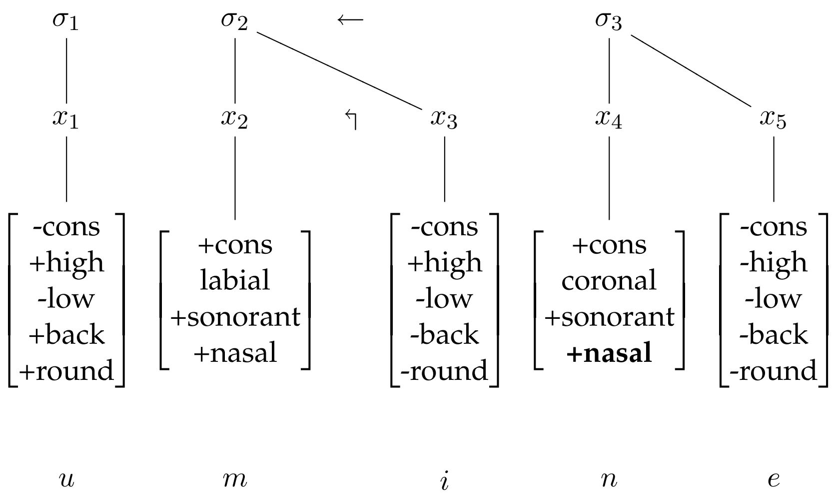

In cases such as *u.mi.ne*, a [+nasal] source is encountered before the syllable-boundary counter has been incremented to 1. However, in *ma.si.le*, the syllable-boundary counter is incremented to 1 before a nasal source has been found, bringing the search to an early end, indicated by ↚ in (11):

(10) Lamba Nasal Search Begins in *ma.si.le*:

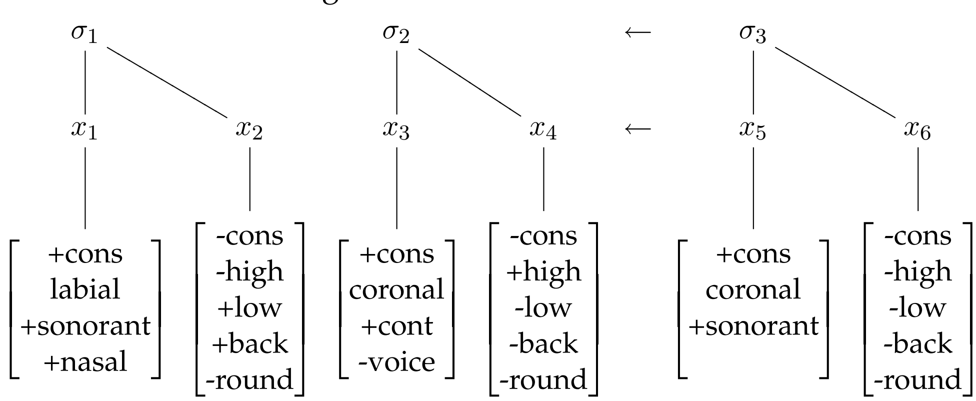

(11) Lamba Nasal Search Terminated when σ_2 Crossed:

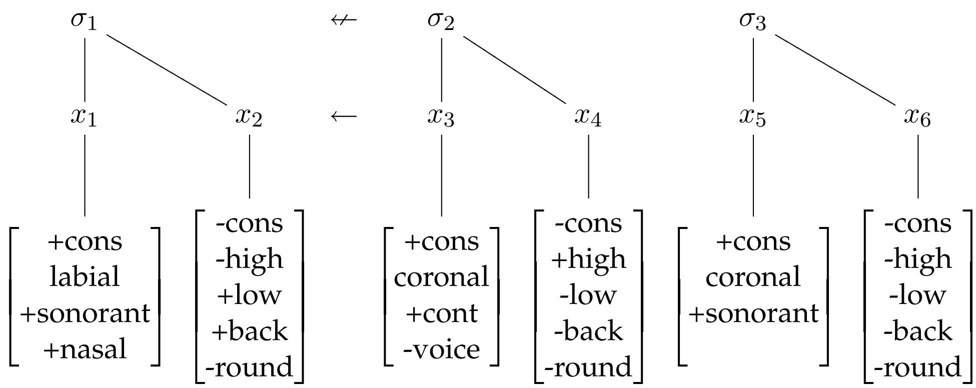

(12) Default Insertion of [−nasal] and default [+son,−nas] → [+lateral]:

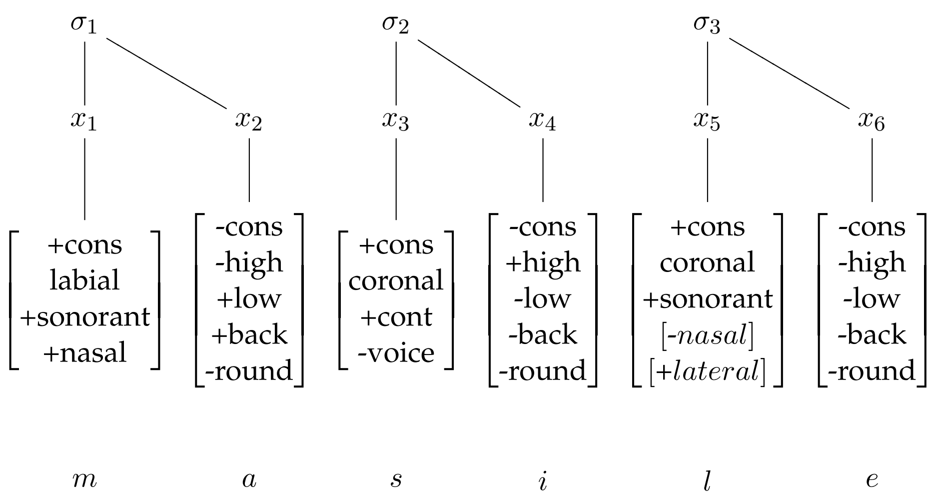

Comparison of the patterns in Kikongo and Lamba reveals that the featural properties of the value-seeker and value-source are the same in both rules. Indeed, both languages have identical inventories in the relevant respects. Nonetheless, Lamba imposes a restriction that the harmonic search may not look beyond one syllable.

Odden (1994) introduced a restricted set of distance parameters that we adopt here, with some modifications to be introduced below.[69] In fact, the rule of nasal assimilation in Kikongo and Lamba may be further compared with that of Chukchi (Odden, 1994, 301), in which a stop assimilates to the nasality of a preceding nasal only when the two are immediately adjacent. While adjacent assimilation rules are not the concern of this book (which focuses on non-adjacent cases of locality), they can be accounted for by the same model as harmony, and clearly demonstrate the need for a distance parameter that allows search to traverse just one segment (i.e. with no intervening segments allowed).

Restrictions on the length of the harmonic search are formally parallel to the role played by Barriers (or phase boundaries) in syntactic theory[70]: even though a harmonic source may not have been found within a relativized domain, and even though an appropriate valuation source may lie just beyond the bound of that domain, the search simply terminates at a given point.

Chomsky (1986), following the work of Rizzi (1982), proposed that one parametric determinant distinguishing the syntax of English and Italian is whether boundedness for *wh-* movement halts after two IP nodes (English) or after two CP nodes (Italian). In English (13) an attempt for the dependency to cross two IP nodes is too far, while Italian parametrically differs in allowing the two IP nodes to be crossed (14)

(13) *Your brother, [$_{IP}$ to whom I wonder [$_{CP}$ what stories [$_{IP}$ they told t]]], was very troubled

(14) Tuo fratello, a cui mi domando che storie abbiano
Your brother, to whom I wonder which stories they.had
raccontato t, era molto preoccupato
told t, was very troubled

In the same way that syntactic searches may parametrically vary in whether they halt with either CP or IP, phonological searches may halt with either the next segment, the next syllable boundary, or the end of the phonological word regardless of its length. In fact, phonological searches, much like the discussion of Barriers above, have the option of allowing the crossing of one extrinsically defined phonological barrier, but not two, as delimited in (15-b,d).

(15) Possible values of β, the distance parameter, in vowel harmony:
a. No intervening segments
b. One intervening segment
c. No intervening syllables
d. One intervening syllable
e. No distance restrictions

The implementation of β in the search algorithm for the parameters in (15) is in Figure 5.1. When no {segment, syllable} is allowed to intervene, the counter that keeps track of traversed segments is incremented if no copying has taken place. In cases of strictly-adjacent assimilation, incrementation of the counter occurs after passing the first segment, and causes shutdown of the search. Similarly, for cases of syllable-adjacent assimilation, incrementation of the counter occurs after passing the first searched-syllable, even though no copying may have occurred. In cases where exactly one {segment, syllable} is allowed to intervene between the value-seeker and value-source, after search has traversed two {segments/syllables} and not found anything, incrementation of the counter will cause the search to terminate as well.[71]

The values of the distance parameter are the result of a combination of three types of quantification (*none, only one, any amount*) with two types of elements (*segments, syllables*), Freely combining these and subtracting one – because allowing any amount of segments is equivalent to allowing any amount of syllables – generates the five distance parameters in (15), with no other possibilities. We have discussed (15-a), (15-c), and (15-e) with respect to nasal harmony in the section above, and now turn to exemplifying the necessity of including (15-b) and (15-d) within the set of universal distance parameters.

5.1.2. Distance Boundary of One-Intervening Segment

Krämer (2001) provides an illustrative example of the existence of the one-intervening segment distance parameter in Yucatec Maya, which allows leftward vowel harmony across one consonant, but not two. Yucatec Maya is a 5-vowel system (*i,u,e,o,a*) distinguished by [± high, ± low, ± back, ± round]. The intransitive imperfective suffix *-Vl* and the subjunctive suffix *-Vk* copy leftward for all four of these vowel features; they may do so across an intervening consonant, as shown in (16)

(16) *Yucatec Maya Total Vowel Harmony across Intervening Consonant*:

a.	ʔah-al ‘wake.up.-impf.’	ʔah-ak‘wake.up.-sbjnct.’
b.	ʔok-ol ‘enter-impf.’	ʔok-ok ‘enter-sbjnct.’
c.	lub’-ul ‘fall-impf.’	lub’-uk ‘fall-sbjnct.’
d.	wen-el ‘sleep-impf.’	wen-ek ‘sleep-sbjnct.’
e.	kiim-il ‘die-impf.’	kiim-ik ‘die-sbjnct.’

τ is either {all values of f_i contrastive for f_i, marked for f_i}
β is either {1,2,∞} and γ is either {$countSylls$ or $countSegs$}
myVals V
myPosition P
mySegsTraversed = 0
mySyllsTraversed = 0
myFeatsneeded F
myConditionalRequirements(F) = R

while F is not empty:
· Go in direction δ and update P
· **if** P is of type τ for any $f, f \in F$:
· · **if** $R(f)$ is true of P:
· · · Copy Val(P, f) to V
· · · Remove f from F
· · **else**:
· · · Remove f from F
· mySegsTraversed = mySegsTraversed + 1
· **if** P is in a new Syllable:
· · mySyllsTraversed = mySyllsTraversed + 1
· · **if** ((mySyllsTraversed $> \beta$ **and** γ = $countSylls$)
· · **or**
· · ((mySegsTraversed $> \beta$) **and** γ = $countSegs$) :
· · · **exit**

Figure 5.1: Parameterized Single- Pass Search with Distance Bounds.

However, Yucatec will not allow such copying across two intervening segments (17); in such cases the default values [−high,+low, +back,−round] are inserted, resulting in *a*.

(17) *Yucatec Maya Leftward Vowel Harmony Quits after one Consonant*:

a. t'oč-b'-al 'harden-impf.'
b. heek'-n-ak 'break-sbjnct'

These affixes' search must give up after two segments are crossed, even without having yet encountered a valid source.

(18) Yucatec Maya Imperfective, Subjunctive Suffixes Must:
Total-Harmonize: δ = left, β = 2 segment, F = [± high, ± low, ± back, ± round]

An identical restriction is encountered in the harmony initiated by root vowels in Assamese, a language of North-East India. Assamese root vowels ordinarily copy marked [+ATR] from a rightward search (Mahanta, 2007); I assume that the rightmost root vowel copies rightward first, and that subsequently the first root vowel copies [+ATR] from the second root vowel. In both cases of copying, the search may cross one intervening segment:

(19) *Assamese rightward [±ATR] allows search past one intervening segment:*

a. k^hɛtɔr, k^hetori 'evil spirit, masc/fem'
b. gɛrɛla, gereli 'fat, masc/fem'
c. lɔg, logori 'company, companion'
d. ʊpɔr, upori 'above, in addition'
e. nɔrɔk, noroki 'hell, sinful'

Vowel harmony is blocked in Assamese when more than one consonant intervenes, and the root vowels undergo last resort insertion of [−ATR] in such cases:

(20) *Assamese rightward [±ATR] harmony quits after two intervening segments:*

a. kɔrmɔ, kɔrmi 'work / active person'
b. xɔbdo, xɔbdit 'sound / resounded'

c. kɛtli 'kettle'
d. sɔkrɔ, sɔkrika 'circle / platelet'
e. tɛz, tɛzɔswi 'strength / powerful'

It is important to point out that in (20-d-e), the two intervening consonants form an onset, and thus that the segmental bounding restriction is in principle independent of syllabic constituency.

Yucatec Maya and Assamese instantiate cases where the distance parameter allows crossing of maximally one intervening segment before the search ends. I have not found any cases in which vowel harmony is specifically bounded by certain subsyllabic constituents, regardless of their size (e.g. vowel harmony that cannot cross a coda, regardless of whether it is one, two, or three segments, yet can cross an onset of any size). We therefore adopt the position that distance parameters are sensitive to segments or syllables, but not to constituents in between these levels of structure.

5.1.3. Distance Boundary of One Intervening Syllable

The distance-parameter that allows one intervening syllable in vowel harmony may be found in many idiolects of Hungarian for trisyllabic stems such as *aszpirin-nek* 'aspirin-dative'. Some background on Hungarian [± back] harmony and on the grammars that generate the *aszpirin-nek* pattern (calling this pattern 'Hungarian ML', for "more local") is in order, before analyzing divergent patterns for these inputs in other idiolects. Once we understand the grammars in which harmonic search is constrained by one-intervening syllable, we will turn to how inter- and intra- speaker variability with these same stems arises.

The pattern of Hungarian vowel harmony is to a certain extent complicated by the fact that the language has long and short vowels with different inventories, and by the fact that long *a:* alternates (in length-alternations) with short ɔ, which is rounded. I assume that this is due to a late rule of rounding, whereby [+back] vowels become [+round] when short (Goldsmith, 1985; Reiss, 2003), but that the grammar of vowel harmony treats short *a* as contrastively [−back] with respect to short ɛ. Similarly, long *e:* alternates with short ɛ, which phonologically patterns as a low vowel (Vago (1975), Hayes and Londe (2006, 62)). Nonetheless, both the long and short versions of these vowels are written in the orthography with the same grapheme (*a* and *e*, respectively). The short and long inventories are in (21) and (22):

(21) Hungarian short vowel inventory

[−back,+round]	[−back,-round]	[+back,+round]	[+back,−round]	
ü	i	u		[+high,−low]
ö		o		[−high,−low]
	ɛ		ɒ	[−high,+low]

(22) Hungarian long vowel inventory

[-back,+round]	[-back,-round]	[+back,+round]	[+back,-round]	
ü:	i:	u:		[+high,−low]
ö:	e:	o:		[−high,−low]
			a:	[−high,+low]

Hungarian suffixes copy contrastive [+back] from the closest leftward vowel,[72] as exemplified below for the [+round] vowels with the dative suffix (an acute accent indicates a long vowel; stress is always initial in Hungarian):

(23) *Hungarian Dative Suffix Copies closest contrastive [±back]*:

a. mókus-nak 'squirrel-dat.'
b. ürü-nek 'sheep-dat.'
c. öröm-nek 'joy-dat.'
d. büro-nak 'bureau-dat.'
e. sofőr-nek 'chauffer-dat.'

Recalling our definition of contrastiveness from Chapter 3, the round vowels and short a,ɛ will be contrastive for [± back]. This relativization to contrastive values leaves the [−round,−low,−back] vowels transparent to vowel harmony[73]. The following dative suffixes copy from the closest leftwards vowel bearing contrastive [±back].

(24) *Hungarian Dative Suffix skips Noncontrastive [−back] Vowels*:

a. kavics-nak 'pebble-dat.'
b. radír-nak 'eraser-dat.'
c. tány'er-nak 'artist-dat.'
d. nüansz-nak 'nuance-dat.'
e. biká-nak 'bull-dat.'
f. bohe:m-nak 'bohemian-dat'
g. mu:ve:sz-nak 'artist-dat.'

While [−back] vowels pattern as transparent to [+back] harmony across them in disyllabic stems such as (24), Farkas and Beddor (1987) and Ringen and Kontra (1989) pointed out that when there is more than one non-contrastively [± back] in a row, a pattern emerges in which suddenly [+back] harmony cannot occur.

Some speakers disallow [+back] harmony in words such as *aszpirin-nek*, even though they allow it in words such as *akti:v-nak*. We will begin by focusing on these speakers/grammars, and subsequently return to the other grammars, which allow [+back] harmony to search further (yielding, for example, *aszpirin-nak*.) Hungarian ML engages in [+back] copying when only one [−back] vowel intervenes between the suffix and a [+back] vowel (e.g. *radír-nak*, in (24)), but disallows such copying from occurring when *two* [−back] vowels intervene:

(25) *Hungarian ML: [+back] harmony fails when two syllables intervene*:

a. a.na.li.zis-nek 'analysis-dat.'
b. a.li.bi-nek 'alibi-dat.'
c. bron.chi.tisz-nek 'bronchitis-dat.'
d. no.vem.ber-nek 'november-dat.'

The grammar of Hungarian ML with respect to vowel harmony represents an intermediate stage between the distance parameter of Lamba and that of Kikongo (Section 5.1.1): while Lamba search fails after traversing one syllable, Hungarian ML's search halts after traversing two syllables. Subsequent failure to find a contrastive value of [± back] within the traversal of 2 syllables results in default insertion of [−back].

(26) Hungarian ML Dative Suffix Must:
Back-Harmonize: δ = left, β = 2 syllable, F = [c: ± back]

For a word like *a.li.bi*, the result of this distance-parameter will terminate the search as soon as the leftward-moving syllable-pointer finishes the syllable *li*, never permitting the search to reach the contrastively [+back] *a*. This is illustrated below:

(27) Hungarian ML [c: back] search begins in *a.li.bi-nek*:

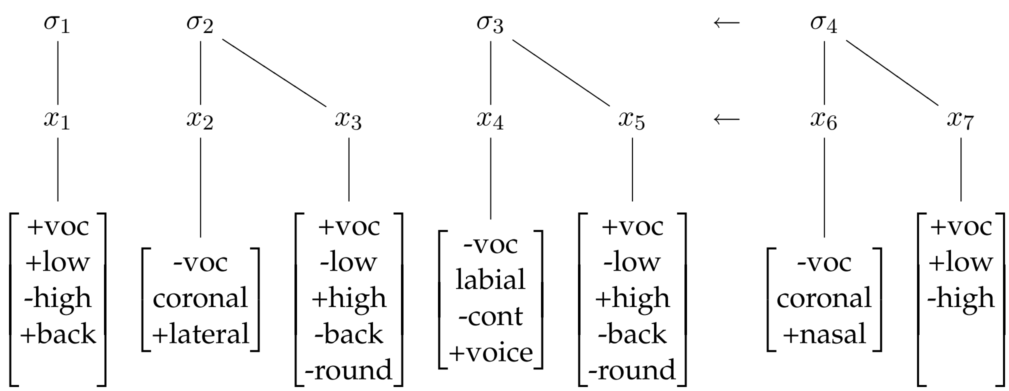

(28) Hungarian ML [c: back] search ends when 2 syllables crossed in *a.li.bi-nek*:

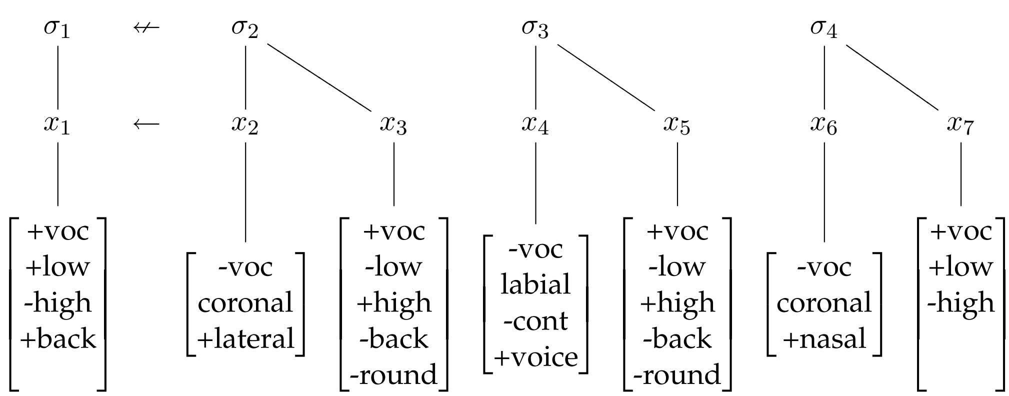

(29) Hungarian ML suffix undergoes Last-Resort Insertion in *a.li.bi-nek*:

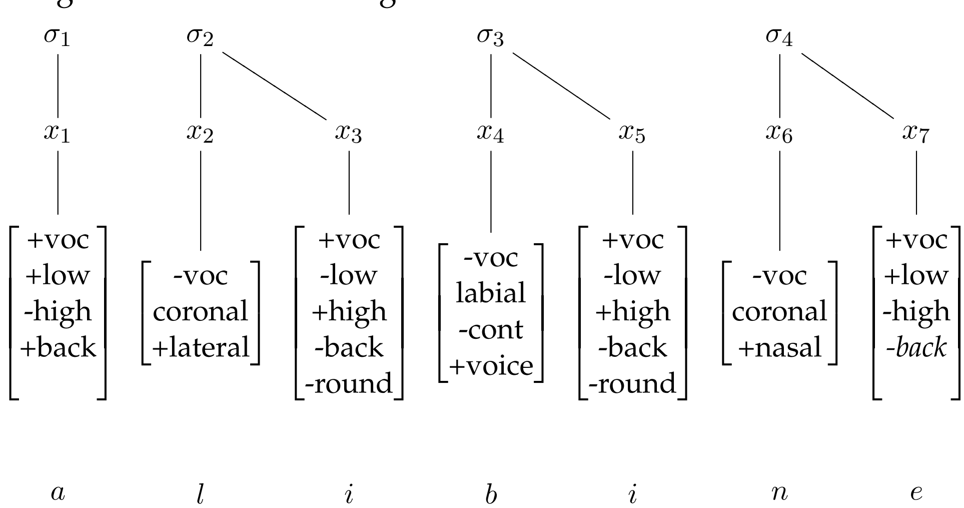

In contrast, the grammar of Hungarian ML generates a successful result in the derivation of a word such as *ak.tiv-nak*: the syllable counter does not surpass 2 syllables before the contrastive [+back] value is found on *a*, and therefore the suffix will successfully copy [+back]. The difference in harmonic [+back] copying for these two patterns of words in Hungarian ML provides evidence for the existence of a one-intervening-syllable distance parameter. While transparent [−back,−low,−round] vowels are not included in the domain of search-and-copy, the additional boundedness on this process tracks two syllables as the maximum that can be traversed before search must end prematurely, leading to default insertion.[74]

Having accounted for Hungarian ML in terms of the intervening-syllable parameter, we turn to a separate grammar of Hungarian that diverges from Hungarian-ML with respect to these trisyllabic Back + Transparent + Transparent words. I will refer to this grammar as Hungarian MD (for "more distant").[75] The pattern of harmony in Hungarian MD differs from that of Hungarian ML only in that the former has an unbounded setting of the distance parameter:

(30) Hungarian MD Dative Suffix Must:
Back-Harmonize: δ = left, F = [c: $\pm$ back]

In Hungarian MD, [+back] copying succeeds both in words such as *akti:v-nak* as well as words such as *alibi-nak*, with the pattern of the latter shown below:

(31) *Hungarian MD: [+back] harmony succeeds when two vowels intervene*:
a. analizis-nak 'analysis-dat.'
b. alibi-nak 'alibi-dat.'
c. bronchitisz-nak 'bronchitis-dat.'
d. november-nak 'november-dat.'

Crucially, both Hungarian ML and Hungarian MD pattern alike with respect to words in which there is only a single intervening syllable, such as *akti:v-nak*. The theory of distance parameters makes a strong prediction: while Hungarian ML and Hungarian MD are the same for *akti:v-nak* and differ for *alibi-nek*∼*alibi-nak*, there could never be idiolectal variation between two grammars in which the opposite held. Specifically, variability between *akti:v-nak*∼*akti:v-nek* but uniformity with *alibi-nak* could never be a possible state of affairs. The reason such a pattern of dialectal microvariation is not possible is because the distance parameters are defined in terms of the value-seeker, and any grammar that allowed finding a [+back] source three syllables away would also allow finding a [+back] source two syllables away. However, as we have seen in Hungarian ML vs. MD, a grammar that allows finding a [+back] source two syllables away does not necessarily allow finding a [+back] source three syllables away.

Our restricted set of distance parameters also predicts that while dialectal variation can arise such that two grammars may be identical with respect to their pattern for copying from two syllables away but differ in ability to copy from three syllables away, there can be no dialectal variation such that two grammars are identical with respect to their pattern for copying from three syllables away but differ in ability to copy from *four* syllables away. The largest distance parameter allowed that is short of unbounded copying is one-intervening-syllable, and the possibility of a grammar that allows copying from a source maximally two-intervening-syllables away is explicitly ruled out.

To summarize so far, the typology of five distance parameters developed in (15) have been exemplified with the following cases:

(32) Example settings of the 5 possible values of β:

a. No intervening segments: Chukchi nasal copying
b. One intervening segment: Yucatec Maya total copy, Assamese [± ATR] copying
c. No intervening syllables: Lamba nasal copying
d. One intervening syllable: Hungarian ML [± back] copying
e. No distance restrictions: Hungarian MD [± back] copying, Kikongo nasal copying

In Chapter 3, we discussed microvariation in Kirghiz, Yoruba, and Finnish with respect to differences only in the value-relativization parameters. In a similar way, distance parameters should be expected to govern not only crosslinguistic variation between typologically unrelated languages, but also microvariation within dialects and idiolects of a single language.

An important question to be raised is: how can microvariation in a single parameter with drastically different surface effects on harmony emerge within a linguistic community? One of the most pervasive sources of within-language variability is parametric ambiguity within the predominant vocabulary and primary linguistic data. Recall that Finnish interspeaker variation arose with loanwords containing mixed contrastive [± back] values, a pattern not encountered within the core vocabulary. Similarly, Hungarian divergence between the ML and MD grammars arises only when noun stems are trisyllabic or longer, items that are loanwords and that, taken on the whole, are less frequent in the input than shorter words, as estimated if token frequency decreases with word length.

The primary 'diet' of harmony patterns on the basis of which the Hungarian learner must set parameters are words that are parametrically ambiguous between an unbounded distance parameter and a one-intervening syllable parameter. Learners attempting to fix a value for the distance parameter may choose one or the other, or may choose based on a default setting for this parameter (if unbounded is the default (Schein and Steriade, 1986, 696)). Alternatively, learners enountering parametrically ambiguous data might end up with a grammar that stochastically chooses between both values. Crucially, stochastic choice between unbounded-distance and one-intervening-syllable will always generate the same result for *akti:v-nak*. Only when longer words are in question will one be able to tell that a stochastic setting of this parameter is present.

The structural property of Hungarian that leads to variation in Back + Transparent + Transparent words is not simply the fact that the input is stochastic to begin with, but rather the fact that Back+Transparent words are ambiguous between different parametric settings of boundedness restrictions, an ambiguity inherent in any parametrized theory of allowable distance between interacting segmental elements.

5.1.4. Dahl's Law Microvariation as an extension

By now we have seen how distance-parameters can regulate the extent to which a search for a subsegmental feature can be restricted by considerations of phonological distance alone. As an interesting confirmation that these parameters circumscribe a space of possible variation, we examine [−voice] dissimilation in the Bantu languages, a process known as Dahl's Law (Davy and Nurse, 1982). In this process, prefixal velar consonants take the value [+voice] when a stem contains a marked [−voice] segment. This is a language-specific instance of featural markedness (perhaps explaining the rarity of [−voice] dissimilation). Bennett (1967) suggests that Dahl's Law results from an earlier stage of Bantu that contrasted aspirated and unaspirated voiceless stops, in which case the [−voice] stops derive from earlier marked aspirated stops. The markedness relationship between the members of this laryngeal opposition has been preserved in contemporary Bantu as marked [−voice].

It turns out that the distance parameters developed above predict exactly the attested within-language variation in Dahl's Law dissimilation, thus offering a possible extension of the model of distance-parameters to other intersegmental dependencies. This process will be exemplified with Gikuyu (the name of which itself, **G**i-**k**uyu, exemplifies an occurrence of Dahl's Law). The application of Dahl's Law can be observed when the gerund prefix becomes [+voice] due to a [−voice] segment in the stem. (A further implicational rule in Gikuyu that [+voice,−son] → [+cont] yields lenition).

(33) *Gikuyu Dissimilation occurs before the voiceless segments s,t,k*:
  a. ko-ruɣa 'to cook'
  b. ko-niina 'to finish'
  c. ɣo-siara 'to give birth'
  d. ɣo-tɛɣa 'to trap'
  e. ɣo-kama 'to milk'

While we have not treated dissimilation at-a-distance in any detail prior to this point, for the purpose of exemplifying the formal parallel with distance-parameters in harmony, we can model dissimilation as a case of search-and-copy in which the *opposite* value is copied from the source. The rightward search seeks a marked value of [± voice], but copies the opposite value to the needy element:

(34) Gikuyu Prefixes Must:
Voice-Dissimilate: δ = right, F = opposite([m: voice])

The question of formal unification of vowel harmony and long-distance consonantal dissimilation is addressed further in Chapter 6. Let us consider for the present discussion the involvement of the distance-parameter for this dissimilation process. When there is only one prefix directly concatenated to the stem, the structural analysis is ambiguous between various limits of boundedness. Dissimilation will be compatible with the distance-parameter set to any of adjacent-syllable, one-intervening-syllable, or unbounded:

(35) *Gikuyu Distance-Parameters for Dissimilation in /ko-ikia/ 'to throw'*:

a. The value-seeker and value-source must be in adjacent syllables: generates [ɣwii.kia]

b. The value-seeker and value-source must have no more than one intervening syllable: generates [ɣwii.kia]

c. The value-seeker and value-source may be at unbounded distance: generates [ɣwii.kia]

The happy convergence encountered among all three parameter settings begins to diverge when there are two instances of *k*-initial prefixes. I assume that copying for [−voice] follows the order of prefix concatenation, so that the object prefix closer to the root attempts dissimilation before the tense prefix does. According to Davy and Nurse (1982), Gikuyu speakers allow two different possibilities for the application of Dahl's Law in /a-kaa-ke-ikia/: either *a.ɣaa.ɣii.kia* or *a.kaa.ɣii.kia*. This microvariation can be modeled as dependent on which distance-parameter was chosen:

(36) *Gikuyu Distance-Parameters for Dissimilation in /a-kaa-ke-ikia/ 'he(1)-will-it(7)-throw'*:

a. The value-seeker and value-source must be in adjacent syllables: generates [a.kaa.ɣii.kia]

b. The value-seeker and value-source must have no more than one intervening syllable: generates [a.ɣaa.ɣii.kia]

c. The value-seeker and value-source may be at unbounded distance: generates [a.ɣaa.ɣii.kia]

In (36-a), the object prefix finds an instance of [−voice] within the adjacent rightward syllable and dissimilates to [+voice]. At the next outward suffix, the tense prefix does not find an instance of [−voice] in the adjacent rightward syllable, and hence surfaces with default [−voice]. In the grammar exemplified by (36-b), the object prefix finds an instance of [−voice] within the adjacent rightward syllable and dissimilates to [+voice]. Subsequently the tense prefix finds an instance of [−voice] within two syllables to the right, and also dissimilates to voice. The same derivation holds for (36-c).

Importantly, although there are two *k*-initial prefixes in (36), and hence 4 logical possibilities for whether dissimilation happens or not, Gikuyu speakers only allow 2 possibilities: dissimilation of the closest prefix, or dissimilation of both prefixes. When we move to three *k*-initial prefixes, although we might logically expect 8 possibilities, only 3 are attested: exactly the three allowed by the set of syllable-based distance-parameters in (15).

Given an input with an object prefix, a tense prefix, and a subject prefix, all of which begin with *k*, Davy and Nurse (1982) report that for inputs such as /ka-kaa-ke-ikia/ 'he(12)-will-it(7)-throw', speakers generate three different options:[76] *ɣa.kaa.ɣii.kia*, *ka.ɣaa.ɣii.kia*, or *ɣa.ɣaa.ɣii.kia*.

(37) *Gikuyu Distance-Parameters for Dissimilation in /ka-kaa-ke-ikia/ 'he(12)-will-it(7)-throw'*:

a. The value-seeker and value-source must be in adjacent syllables: generates [ɣa.kaa.ɣii.kia]
b. The value-seeker and value-source must have no more than one intervening syllable: generates [ka.ɣaa.ɣii.kia]
c. The value-seeker and value-source may be at unbounded distance: generates [ɣa.ɣaa.ɣii.kia]

In (37-a), the object prefix finds [−voice] in the adjacent syllable, and dissimilates to [+voice]; the tense prefix does not find [−voice] in the adjacent rightward syllable, and hence surfaces with default [−voice], and finally the outermost subject prefix finds [−voice] in the adjacent syllable, and hence dissimilates to [+voice]. This yields an "alternating" pattern of dissimilatory voicing, as shown in the derivation below (vowel length not represented):

(38) Innermost prefix begins rightward search in *ɣa.kaa.ɣii.kia*

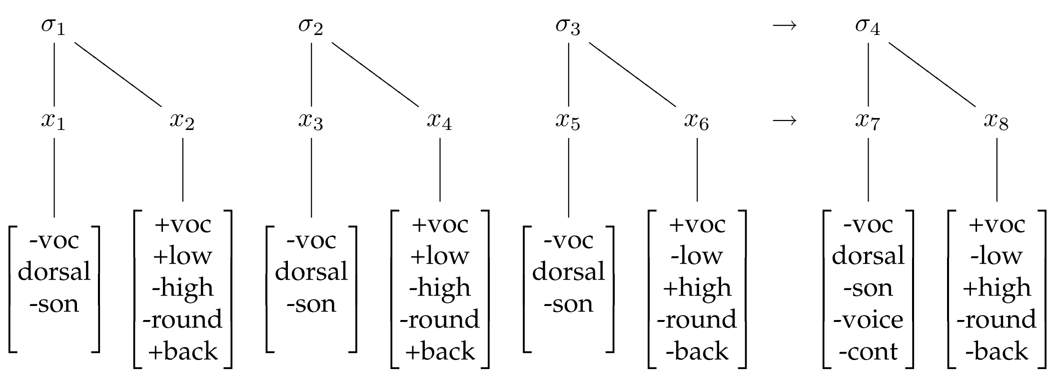

(39) Innermost prefix finds marked [−voice] and copies opposite value:

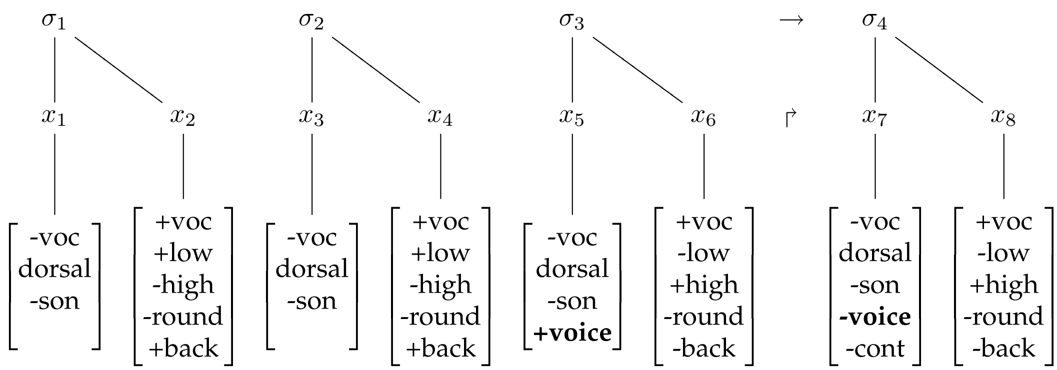

(40) Second prefix begins search for [−voice]:

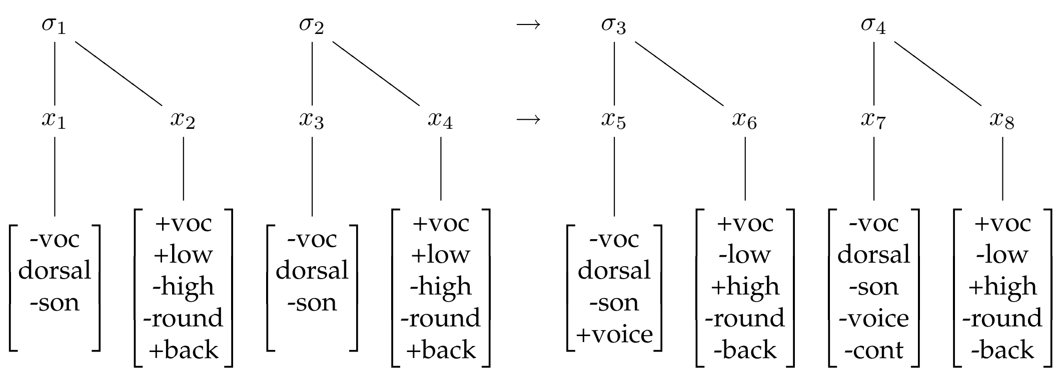

(41) Second prefix passes adjacent syllable and must terminate search for [−voice]:

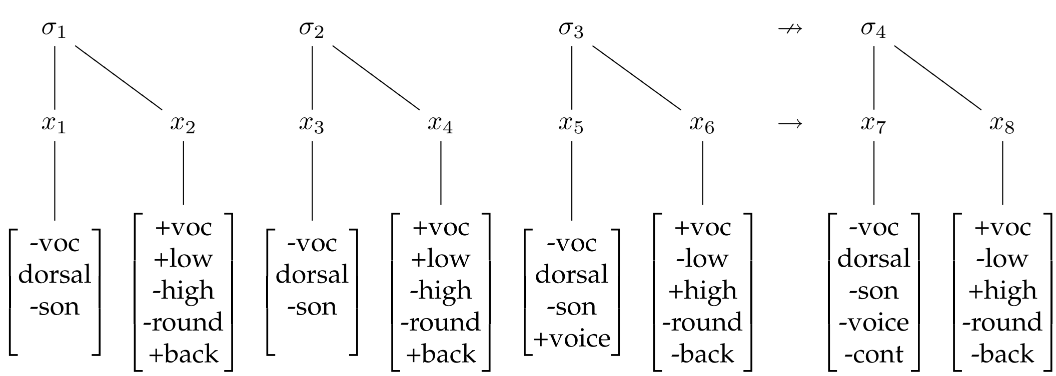

(42) Second prefix undergoes last-resort [−voice] insertion:

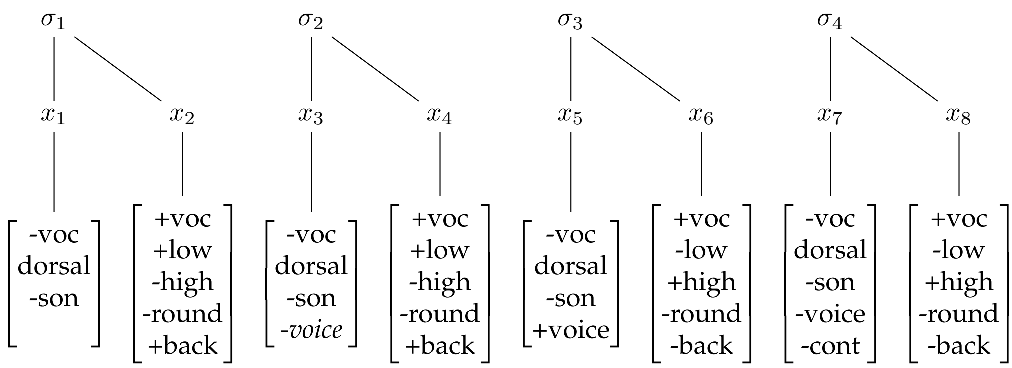

(43) Outermost prefix begins search for [−voice]:

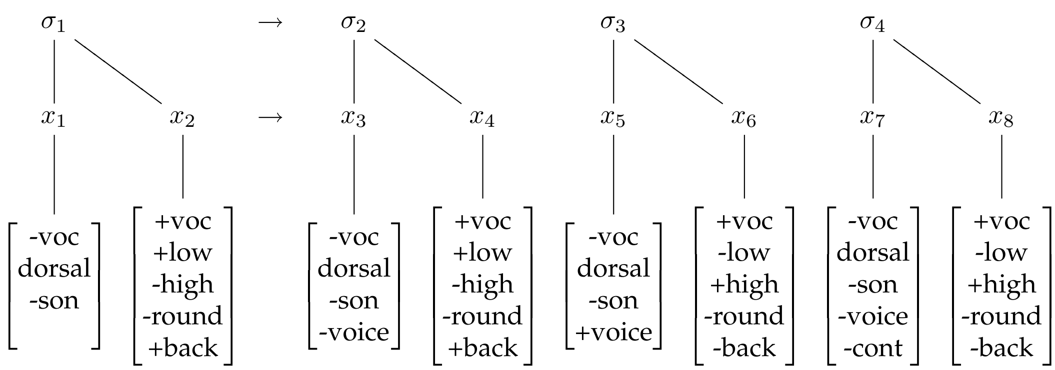

(44) Outermost prefix finds [−voice] and copies opposite value:

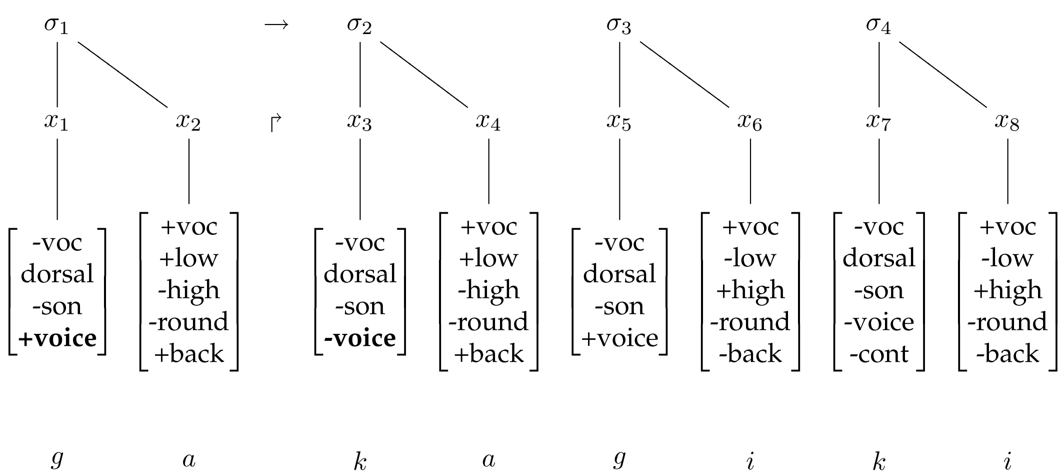

The grammar of Gikuyu under adjacent-syllable distance-boundedness will have an alternating pattern, no matter how many prefixes are attached. The "myopic" nature of adjacent-syllable search by the second prefix does not manage to find the [−voice] segment two syllables to the right. It undergoes last-resort insertion and due to the derivational nature of each affix's search, provides a [−voice] value in turn for the prefix to its left.

In the grammar exemplified by (37)[b] [ka.ɣaa.ɣii.kia], when the distance parameter is set to one-intervening syllable, the object prefix finds [−voice] in the adjacent syllable, and dissimilates to [+voice]. Next, the tense prefix finds [−voice] within one-intervening-syllable away, and hence dissimilates to [+voice]. Finally the subject prefix does not find [−voice] within one-intervening syllable away, surfacing with default [−voice]. This yields a "rightmost-two" pattern of dissimilatory voicing, no matter how many prefixes are attached.

In (37)[c] [ɣa.ɣaa.ɣii.kia], when the distance parameter is set to unbounded, the object prefix finds [−voice] in the adjacent syllable, and dissimilates to [+voice]. Next, the tense prefix finds [−voice] within an unbounded distance, and hence dissimilates to [+voice]. Finally the subject prefix finds [−voice] within an unbounded distance, thereby dissimilating to [+voice]. This yields an "everything dissimilates" pattern of dissimilatory voicing, no matter how many prefixes are attached.

These three options are the only possibilities for Gikuyu speakers, and represent a very small subset of the logically possible options, following exactly the restrictive parameters of distance-sensitivity developed for harmony in (15). While Gikuyu speakers allow multiple possibilities and exhibit intra-speaker variation in cases of three *k*-initial prefixes, this variation is not "random": it is the result of a highly constrained set of options for the boundedness of the search. This variation obtains only with longer sets of potentially dissimilating elements, but shows no optional or stochastic behavior when there is only one dissimilable element. Dahl's Law in Gikuyu supports a model in which parametrically ambiguous inputs (such as the one in (35), compatible with all three distance-settings) give rise to divergent grammars that are only disambiguated in more-distant configurations.

5.2. Sonority Peaks are Barriers

Boundedness in phonological locality is delimited not only by the extent of the search path traversed. In the present section we will see that certain elements can terminate the search as a result of their inherent high-sonority. These sonority-peaks should be excluded from the domain of search by their noncontrastive value, but impose a hurdle past which search cannot proceed.

Ross (1984) discusses a set of syntactic restrictions whereby benefactive and locative prepositional phrases cannot be extracted across negative elements, even though this clearly does not constitute a case of a minimality violation (Frampton, 1991):

(45) a. *For whom didn't you enter the race?
b. *Down the road Sanford doubts we will zoom

Ross (1984) noted that many relative/operator constructions are also blocked by negation:

(46) a. *All these samples don't have to do is contain ytterbium and we're saved
b. *The papers are ready for you not to put relish on

The intuition among syntacticians is that the negative elements create an "island", imposing a boundary across which an A'-relation simply cannot be established, despite the fact that the negative elements need not form a homogenous class in terms of their status as lexical categories (cf. (45-a-b)).

In this section we demonstrate that high-sonority vowels create similar boundaries across which harmonic searches cannot proceed, even though no minimality-based locality is at stake. We thus have two types of extrinsic limits on where search cannot travel: those of *extent* (elements "too far" to cross) and those of *hurdles* (elements "too high" to cross). The syntactic case of crossing two IPs in (13) instantiates a limit of extent, and that of negative elements is a limit of islands / hurdles. The same two limits on the length of search arise in vowel harmony: copying from an element separated by an extent too far, as expressed by distance parameters, and copying past an element that is too high, as I will argue for sonority peaks.

As the result of case studies analyzed below in Classical Manchu, Finnish, Wolof, and Hungarian, I will develop the following implicational generalization:

(47) Given two relativization-excluded (e.g. noncontrastive/nonmarked) vowels ϕ and ψ, where ϕ is of a higher sonority than ψ, it will never be the case that ϕ is transparent while ψ is not.
(In other words: transparency of ϕ implies transparency of ψ)

For example, in Wolof, while *i:,u:,a:* are all noncontrastive for [± ATR], only *a:* acts as a hurdle in harmony. Similarly, while in Classical Manchu both ə,*u* are unmarked for [± ATR], only ə acts as a hurdle in harmony. The claim in (47) is that the opposite cannot occur in natural language: sonority induces a particular asymmetry among the vowels outside of the relativized-value domain. High-sonority vowels may impose hurdles to the exclusion of low sonority vowels, but never vice versa.

Phenomena in which sonority can render otherwise expected transparent elements as blockers in vowel harmony have been discussed by Ultan (1973) and Anderson (1980), with their perceptual basis raised in Harms (1987).[77] The present section represents an attempt to systematically formalize the effects of sonority in a variety of harmony patterns, and develop a strong prediction about the effects of vowel height on vowel harmony, one that is general enough to apply independently of the features involved or the value-relativization.

Before proceeding further, I would like to point out that while high sonority (e.g. lower vowel height) is often correlated with duration, the effect of sonority-hurdles considered in the case studies below cannot be reduced to the effects of duration. For example, Hungarian *long* /i:/ has a greater duration than all short vowels in the inventory (Magdics, 1969), but is nevertheless transparent. Similarly, Hungarian long /e:/ is transparent while short /ε/ is not, even though the former is over 35% greater duration than the latter.[78] Careful examination of the patterns of transparency in languages with long/short vowel distinctions reveals that the relevant factor in the hurdle-effects discussed below is vowel height itself (e.g. high sonority).

We adopt the following scale of sonority levels (based on Parker (2002)). If the distance between each level of the scale is not always evenly-spaced, any reassignment of numerical values will be consistent with the case studies below, as long as the relative ordering of elements is preserved.[79] In the scale below, I assume an inherent ranking of sonority among the segments themselves, where the relevant features are shown to the right. Not every language may phonologically encode a [± ATR] (or even a [± low] contrast), but nonetheless [a] will be more sonorous than [i] in such languages, based on the feature [± high].[80] The accompanying features below are thus provided for a "ten vowel" (e.g. three heights with a [± ATR] contrast) system, and the values provided alongside them will be referred to through this chapter.

(48) *Sonority Levels*:

Segments	Sonority Level	Relevant Features
a,ä,ɒ	8	[+low,−ATR]
ə	7	[+low,+ATR]
ɛ,ɔ	6	[−low,−ATR]
e,o	5	[−low,+ATR]
ɪ,ʊ	4	[+high,−ATR]
i,u	3	[+high,+ATR]
y,w	2	[+high, +ATR, −voc, −cons]
sonorant consonants	1	[+son,+cons]

We adopt the proposal of Gouskova (2004) whereby different languages select different cut-off points in the sonority hierarchy to which their grammar may be sensitive. Under this view, the locus of crosslinguistic variation is the threshold above which certain sonority elements cannot be passed, but the relative ordering in the scale in (48) is itself crosslinguistically invariant.

The ultimate effect of sonority hurdles, like the effect of distance-bounds, is that they halt the search as soon as they are met. This is implemented by a variable maintaining the allowable sonority threshold, as shown in the revised Search Procedure in Figure (48).

While their ultimate effect is to halt the search, there is an important difference between the role of extrinsic bounding imposed by distance-parameters and that of sonority-hurdles. When search is halted by reaching a sonority hurdle, its value for the harmonic feature is copied, *even if not included within the relativized search domain*. The effect of copying a feature-value from a sonority-hurdle is a kind of "penultimate resort": the high-sonority hurdle manages to provide a harmonic value before default/last-resort insertion must be attempted.

τ is either {all values of f_i contrastive for f_i, marked for f_i}
myVals V
myPosition P
myFeatsneeded F
mySonorityTolerance ζ (0 = allows anything)

while F is not empty:
· Go in direction δ and update P
· **if** Sonority(P) $> \zeta$:
· · **if** P has any value for f:
· · · Copy Val (P, F) to V
· · **exit**
· **else-if** P is of type τ for any value for $f, f \in F$:
· · Copy Val (P, F) to V
· · Remove f from F
· ·

Figure 5.2: Single Pass-Search with Sonority Thresholds. (Distance boundaries omitted for conciseness).

Let us examine the effect of sonority-hurdles in two [± ATR] harmony systems and two [± back] harmony systems: Classical Manchu, Finnish, Wolof, and Hungarian. For each of these case studies, an analysis of the language alone could easily include unrelated special statements for the exceptional pattern of ə, *ä*, *a:*, ε, respectively. Within the present analysis, however, the fact that these "exceptions" consistently are the lowest vowels in the set of excluded elements is unified under the generalization about sonority.

As early as Jespersen (1904), the fundamental effects of sonority on phonotactic sequences and intersegmental relations were confirmed. Vowel sonority plays an important role in a number of phonological processes in which vowels are adjacent. Sonority determines the syllablification of vocalic sequences in Spanish (Harris, 1983); Berber, (Dell and Elmedlaoui, 1985; Prince and Smolensky, 1993), and Tahitian (Bickmore, 1995), wherein vowels of lower height are the ones affiliated to the syllabic nucleus. Sonority also may dictate hiatus persistence, namely which vowel stays around in hiatus when one must be deleted (Casali, 1998), wherein the choice of which to keep is not determined by linear order, but by sonority.

Sonority interacts with relative prominence among vowels (i.e. stress) in two directions of causality. Sonority may determine stress placement above and beyond syllable-weight and directionality, as argued for Kobon by Kenstowicz (1997): thus, high-sonority demands stress. Conversely, vowels may undergo lowering when they have already been assigned stress in Chamorro (Crosswhite, 1998), thus, stress demands high-sonority.

The existence of sonority-hurdles within vowel harmony thus falls within a broad class of prominence-based asymmetries in which syntagmatic interactions are sensitive to vowel height.

5.2.1. Classical Manchu

In Classical Manchu, two vowels, namely ə and *u*, are both excluded from the domain of value-relativization. Only one of them manages to wedge its way into the vowel harmonic search, and it is the one of higher sonority among the two. Classical Manchu has an [± ATR] harmony system and the following vowel inventory (Zhang, 1996):

(49) *Classical Manchu Vowel Inventory*

[−back,−round]	[+back,+round]	[+back,−round]	
i	u		[+high,+ATR]
	ʊ		[+high, −ATR]
		ə	[−high, +ATR]
	ɔ	a	[−high, −ATR]

Classical Manchu displays velar/uvular alternations based on the [± ATR] value of immediately preceding vowels. The velars *k,g,x* occur before [+ATR] /i,ə,u/, while the uvulars *q,ɢ,χ* occur before [−ATR] /a,ɔ,ʊ/ (Zhang (1996); Vaux (1999))

(50) *Velar/Uvular Alternations Conditioned by [± ATR] of Preceding Vowel*:

a. ətu-ku ‘ clothing’
b. xərə-ku ‘ladle’
c. kimu-nggə ‘harbouring enmity’
d. urgu-nggə ‘joyous’
e. dərgi-kən ‘somewhat above’
f. taci-qʊ ‘school’
g. ɢɔsin ‘pity’
h. ilχa ‘flower ’
i. nuχa-qan ‘somewhat easy’
j. ɢʊnin ‘thought’
k. fɔχɔlɔ-qan ‘somewhat short’

The consonantal alternations present evidence that Manchu exhibits an [± ATR] contrast with two distinct natural classes of vowels.[81] The marked value is [−ATR], and the vowels contrastive for [± ATR] are ə/a and u/ʊ. Though to a large extent the distribution of u/ʊ is influenced by preceding dorsal consonants, there are minimal pairs demonstrating their phonemic status (Zhang 1996:43):

(51) *Minimal Pairs demonstrating Classical Manchu* u/ʊ *contrast*:

a.	butun	‘hibernation’	butʊn	‘crock, large jar’
b.	mungku	‘a frozen fish’	mʊnggu	‘bird’s nest’
d.	ulən	‘irrigation ditch’	ʊlən	‘house’

Vowel harmony in Classical Manchu is initiated by a variety of suffixes and yields alternations in the suffixal vowels between u and ʊ as well as between ə and a. The vowel *i* is noncontrastive for [± ATR] and is transparent for [−ATR] copying across it. The results of [± ATR] Harmony in the Adjectivalizing, Verbalizing, and Diminutive Suffixes are provided in (52).

(52) *Classical Manchu [−ATR] suffixal variants (Zhang 1996:49)*:

a.	algin-ngga	'famous
b.	malxʊ-ngga	'frugal'
c.	aga-ngga	'rainy'
d.	ilxʊ-ngga	'lying straight'
e.	farxʊ-kan	'somewhat dark'
f.	gʊrgi-la	'to flame'
g.	bakci-la	'to oppose'
h.	bakta-kʊ	'internal organs'
i.	banji-shʊn	'having money'

(53) *Classical Manchu [+ATR] suffixal variants (Zhang 1996:48)*:

a.	xəhə-nggə	'female'
b.	xətu-kən	'somewhat stocky'
c.	dərgi-kən	'somewhat above'
d.	icə-lə	'make new'
e.	juwə-lə	'lean to two sides'
f.	xərə-ku	'ladle'
g.	sidərə-shun	'hobbled/lame'

These suffixal alternations are the result of a leftward search for marked [−ATR]:

(54) Classical Manchu Adjectivalizing, Verbalizing, Diminutive Suffixes Must:
ATR-Harmonize: δ = L; F = [m: ATR]

As can be observed in the following stems, the last-resort value in case nothing is found is [−ATR].

(55) Suffixes undergo default [−ATR] insertion on all-*i* stems (Zhang 1996:56):

a.	fili-kan	'somewhat solid'
b.	ici-ngga	'having direction'
e.	iji-shʊn	'obedient'
f.	sifi-kʊ	'hairpin'

In the inventory of Classical Manchu, both *u* and ə are contrastive for [± ATR] and unmarked for [± ATR]. However, unlike ə, the vowel *u* can co-occur with [−ATR] vowels in roots, and allows [−ATR] copying across it. The transparent behavior of [+ATR] *u* is exemplified by the following cases:

(56) *Classical Manchu transparent u in harmony (Zhang 1996:49)*:
  a. dacu-kan 'somewhat sharp'
  b. gʊsu-la 'tie up with thick rope'
  c. xʊdu-ngga 'speedy'

The transparent patterning of *u* is entirely expected given the relativization to marked values in (54).[82] However, of the two contrastive and unmarked [+ATR] vowels, while *u* may be skipped in search; ə may not. The reason why search halts with ə (but not with *u*) cannot be explained by relativization to contrastive or marked values: there is no way in terms of harmonic feature-values to exclude *u* from the domain without also excluding ə.

The relativization of Classical Manchu [± ATR] harmony to only marked values means that search should continue past [+ATR] values of the feature. However, we have observed search does not pass over ə, despite its unmarked value. Search halting with ə must be due to the effect of its high sonority:

(57) Classical Manchu Adjectivalizing, Verbalizing, Diminutive Suffixes Must:
ATR-Harmonize: δ = L; F = [m: ATR]; ζ = 6

In Classical Manchu, the search for [± ATR] harmony is extrinsically bounded by all vowels of sonority greater than 6, as shown by the threshold parameter ζ in (57). As soon as an element exceeding that cutoff (such as ə) is reached in the search, the search terminates regardless of whether a potential value-source has been encountered or not.

However, given that the purpose of a vowel harmony search in the first place is to provide a value for a suffix for the harmonic feature, the sonority-hurdle is copied from, even though it is outside of the domain of value-relativization. In other words, although the search is relativized to marked [−ATR] values, when a barrier is encountered, the search ends. In examples such as *juwə-lə*, the search reaches an element of sonority $>$ 6, and the hurdle is copied from, as shown below. (The sonority level of each element is included next to its height feature in the diagram for the reader's convenience).

(58) *Verbalizing Suffix Begins Marked [−ATR] Search in juwə-lə*:

$$\begin{matrix} x_1 & & x_2 & \leftarrow & x_3 & & x_4 \\ \begin{bmatrix} \text{-voc} \\ \text{labial} \\ \text{+high} \\ \text{+back} \\ \text{+round} \end{bmatrix} & & \begin{bmatrix} \text{+voc} \\ \text{-high; son = 6} \\ \text{+back} \\ \text{-round} \\ \text{+ATR} \end{bmatrix} & & \begin{bmatrix} \text{-voc} \\ \text{coronal} \\ \text{+sonorant} \\ \text{-nasal} \\ \text{+lateral} \end{bmatrix} & & \begin{bmatrix} \text{+voc} \\ \text{-high} \\ \text{+back} \\ \text{-round} \\ \end{bmatrix} \end{matrix}$$

(59) *Verbalizing Suffix Encounters Sonority hurdle in juwə-lə*:

$$\begin{matrix} x_1 & & x_2 & \nleftarrow & x_3 & & x_4 \\ \begin{bmatrix} \text{-voc} \\ \text{labial} \\ \text{+high} \\ \text{+back} \\ \text{+round} \end{bmatrix} & & \begin{bmatrix} \text{+voc} \\ \textbf{-high; son = 6} \\ \text{+back} \\ \text{-round} \\ \text{+ATR} \end{bmatrix} & & \begin{bmatrix} \text{-voc} \\ \text{coronal} \\ \text{+sonorant} \\ \text{-nasal} \\ \text{+lateral} \end{bmatrix} & & \begin{bmatrix} \text{+voc} \\ \text{-high} \\ \text{+back} \\ \text{-round} \\ \end{bmatrix} \end{matrix}$$

(60) *Verbalizing Suffix Copies from Sonority hurdle in juwə-lə*:

$$\begin{matrix} x_1 & x_2 & x_3 & x_4 \\ \begin{bmatrix} \text{-voc} \\ \text{labial} \\ \text{+high} \\ \text{+back} \\ \text{+round} \end{bmatrix} & \begin{bmatrix} \text{+voc} \\ \textbf{-high; son = 6} \\ \text{+back} \\ \text{-round} \\ \textbf{+ATR} \end{bmatrix} & \begin{bmatrix} \text{-voc} \\ \text{coronal} \\ \text{+sonorant} \\ \text{-nasal} \\ \text{+lateral} \end{bmatrix} & \begin{bmatrix} \text{+voc} \\ \text{-high} \\ \text{+back} \\ \text{-round} \\ \textbf{+ATR} \end{bmatrix} \\ \text{w} & \text{ə} & \text{l} & \text{ə} \end{matrix}$$

By the lights of the inventory and the relativization to marked [−ATR], ə should not be copied from, but its low height and the sonority-threshold of Classical Manchu cause the search to terminate with and copy from it.

5.2.2. Finnish

A formally identical pattern to that of Classical Manchu – in which there is relativization to the marked value of a feature but one member of the unmarked set nonetheless participates in harmony – may be found in Finnish Dialect B (recall Section 3.11). The Finnish vowel inventory is repeated in (61):

(61) *Finnish Vowel Inventory*:

[−back, −round]	[−back, +round]	[+back, +round]	[+back, −round]	
i	ü	u		[+high, −low]
e	ö	o		[−high, −low]
ä			a	[−high, +low]

Recall that in Finnish B, loanwords that have mixed contrastive [± back] vowels yield [+back] harmony in the suffixes. This led us to the conclusion that Finnish Dialect B copies [± back] from marked sources (namely those with [+back]):

(62) Finnish B: Partitive Suffix must:
Back-Harmonize: δ = L; F = [m: back]

The harmony procedure in (62) resulted in transparency of the contrastively [−back,+round] vowels *ü,ö* (in addition to the transparency of *i,e*).

(63) *Finnish B shows Transparency of ü,ö*:
a. marttüüri-a 'martyr-partit.sg'
b. klorofülli-a 'chrolophyll-partit.sg'
c. sutenööri-a 'pimp-partit.sg'
d. amatööri-a 'amateur-partit.sg'

While the relativization to *marked* values of [+back] explains the transparency of *ü,ö* in (63), Campbell (1980, p.251) notes an interesting asymmetry among the contrastive [−back] vowels: "...*sutenööri**ä*** and *amatöör̈i**ä*** are perfectly acceptable [alongside sutenööri**a** and amatöör̈i**a** as options] . However, **hydrosfääri**a*** and the like are impossible." In other words, *ä* does not allow [+back] to be copied across it, in any dialect or register. As all three of *ü,ö,ä* are unmarked for [± back], it is perhaps surprising that Finnish B (the dialect that allows skipping over *ö* in *sutenööri-ä*) *must* copy [−back] from *ä*. The fact that *ä* stubbornly remains a copying-source for [± back] harmony despite its exclusion from value-relativization calls for a parametrization in terms of sonority-hurdles:

(64) Finnish B: Partitive Suffix must:
Back-Harmonize: δ = L; F = [m: back], ζ = 7

Since the sonority-threshold requires all [+low] vowels to halt the harmonic search in (64), items such as *hüdrosfääri-ä* will require suffixal harmony to stop dead in its tracks as soon as *ä* is encountered. The [−low] vowels *ü,ö*, on the other hand, are tolerated by the sonority threshold, and excluded from the domain of value-relativization. The search passes them right by, eventually copying marked [+back] from a leftward vowel that lies past them.

Campbell's observation – that Finnish loanwords with mixed contrastive [± back] vowels allow for copying of [+back], except when the [−back] vowel is [+low] – finds a natural account in terms of sonority-hurdles. All of the [−back] vowels are excluded by marked-value relativization, but the [+low] vowel imposes a boundary that prevents search from looking beyond it. It is striking that a formally identical interaction of marked-value relativization with unmarked low vowels holds in Finnish as in Classical Manchu, even though the harmonic feature is different ([± ATR] in Classical Manchu and [± back] in Finnish)) and the "stubborn" vowel is different as well (ə in Classical Manchu and ä in Finnish). The explanatory power of including sonority-hurdles in the theory is that they allow a unified account of these otherwise disparate phenomena.

5.2.3. Wolof

Sonority-based asymmetry in vowel harmony can be further exemplified in the [± ATR] system of Wolof, a West Atlantic language spoken in Senegal. Wolof has a slightly different inventory of long and short vowels. The short vowel inventory is in (65) and the long vowel inventory is in (66). Notice that long /a:/ has no [+ATR] counterpart, and is thus noncontrastive for the harmonic feature.

(65) *Wolof Short Vowel Inventory*:

[−back,−round]	[−back,−round]	[+back,+round]	
i		u	[+high,−low,+ATR]
e		o	[−high,−low,+ATR]
ɛ		ɔ	[−high,−low,−ATR]
	ə		[−high,+low,+ATR]
	a		[−high,+low,−ATR]

(66) *Wolof Long Vowel Inventory*:

[−back,−round]	[−back,−round]	[+back,+round]	
i:		u:	[+high,−low,+ATR]
e:		o:	[−high,−low,+ATR]
ɛː		ɔː	[−high,−low,−ATR]
	a:		[−high,+low,−ATR]

Suffixes harmonize for [± ATR], including the instrumental/locative *-e/*ɛ, participant *-le/l*ɛ, past-tense -oːn/ɔːn, benefactive -əl/al, possessive -əm/am, and comitative -əndoː/andɔː (Ka, 1994). The following examples illustrate these suffixes copying contrastive [± ATR] from the nearest leftward source:

(67) *Wolof suffixes copy closest contrastive [± ATR]*:

a.	do:r-e	'to hit with'	xɔːl-ɛ	'to look with'
b.	re:r-e	'to be lost in'	dɛm-ɛ	'to go with'
c.	gən-e	'to be better in'	xam-ɛ	'to known in'
d.	do:r-e	'to help hit'	jɔx-lɛ	'to help give'
e.	re:r-le	'to lose one's property'	dɛɛ-lɛ	'to lose a relative'
f.	yəg-le	'to announce'	takk-lɛ	'to help tie'
g.	re:r-o:n	'was lost'	rɛːr-ɔːn	'had dinner'
h.	now-o:n	'came'	jɔx-ɔːn	'gave'
i.	bəgg-o:n	'wanted'	takk-ɔːn	'tied'
j.	le:b-əl	'to tell stories for'	bɛy-al	'to cultivate for'
j.	fo:t-əl	'to launder for'	wɔːr-al	'to fast for'
k.	jënd-əl	'to buy for'	wax-al	'to speak for'
l.	sofoːr-əm	'his driver'	nɛlaw-am	'his sleep'
m.	genn-əndo:	'to go out together'	dɛnd-andɔː	'to be neighbors'
n.	tox-əndo:	'to smoke together'	tɔpp-andɔː	'to imitate'
o.	dəkk-əndo:	'to live together'	wax-andɔː	'to say together'

We find evidence for a sonority-hurdle with the long vowel *a:* which is noncontrastive for [± ATR], but not transparent. In the apt words of Kenstowicz (1994, 354), "[a:] finds its confreres only among the [−ATR] set." The Wolof long [+low] vowel, even though noncontrastive, provides an immediate [−ATR] value source for the vowels to the right:

(68) *Wolof /a:/ imposes a sonority-hurdle*:

a. yobbu-wa:lɛ 'to carry away also'
b. genn-a:lɛ 'to go out with also'
c. do:ra:tɛ 'to hit usually'
d. je:m-əntu-wa:l-ɛɛti 'to try also without conviction once more'
e. indiw-a:lɛ 'to bring in addition'
f. seyta:nɛ 'devil' (from Arabic)
g. kuma:sɛ 'to start' (from French)

The parametrization of Wolof [± ATR] harmony sensitivity to high-sonority peaks is provided in (69). (Note how the cutoff at ζ = 7 will cause the search to not only halt with *a:*, but also halt with short ə/a, though the hurdle effect with these vowels is indistinguishable from contrastive-value copying from them).

(69) Wolof Instrumental, Benenefactive, Past Tense, Possessive, Comitative Suffixes must:
ATR-Harmonize: δ = L; F = [c: ATR], ζ = 7

Importantly, the high vowels are also noncontrastive for [± ATR]. Nonetheless, unlike /a:/, they are transparent to [+ATR] copying across them (Ka 1994: 27-30, Archangeli & Pulleyblank 1994: 231):

(70) *Wolof Noncontrastive High Vowels Skipped in [± ATR] Harmony*:

a. tɛkki-lɛ:n 'untie'
b. moytu-lɛ:n 'avoid'
c. watu-lɛ:n 'have haircut'
d. lettu-le:n 'braid hair'
e. so:bu-le:n 'plunge'
f. gəstu-le:n 'research'
g. to:xi-le:n 'go and smoke'
h. sɔppiwulɛɛn 'you have not changed'
i. tɛɛruwɔɔn 'welcomed'
j. barigɔ 'barrel'
k. kamisɔl 'robe'
l. kɔritɛ 'Muslim holiday'
m. ka:ritɛ 'butter'
n. warugar 'obligation'

We may conclude that the patterning of */a:/* in Wolof harmony is unique among the three noncontrastive vowels, being the only one that is not skipped in the search by contrastive value-relativization. Archangeli and Pulleyblank (1994) propose that [± ATR] harmony must interface with the articulatory-grounded co-occurrence restrictions *[+high,−ATR] and *[+low,+ATR]. While these may represent valid crosslinguistic tendencies, what is missing from such a proposal is an implicational predictive connection between the effect of these two constraints implicationally. The prediction of the sonority-hurdles account is that whenever [+low] vowels are transparent to [−ATR] harmony across them, [+high] vowels must also be transparent to [+ATR] across them. Put differently, whenever the sonority threshold is set very tolerantly and allows search to pass by noncontrastive vowels of higher sonority, it will also allow search to pass across noncontrastive vowels of lower sonority. Wolof long *a:,u:,i:* are all noncontrastive for [± ATR], and the prediction of the sonority-hurdles model is that if only one of them fails to be transparent, it will be the one of highest sonority.

The patterning of [+low] *a:* as a nontransparent sonority hurdle is widespread throughout the [± ATR] vowel harmony systems of West Africa (Ladefoged, 1968, 36-38). This sonority effect is found in all dialects of Yoruba (Bakovic, 2000; Ola Orie, 2001), regardless of whether they parametrize visibility to all or only contrastive values of [± ATR] (cf. Section 3.9).

Wolof becomes an important case study for confirming the role of sonority because the long vowels show the sonority asymmetry among them, thereby reflecting the conclusion that it is not duration but vowel height that imposes early termination of vowel-harmonic searches. As we will see in Section 5.2.4, the effect of vowel height on noncontrastive long vowels creates an asymmetry in Hungarian as well, explaining the differences in transparency between its *ε,e:,a:*. The advantage of the sonority-hurdles model is that it relates the bounding effect of /a:/ in Wolof, an [± ATR] system, to the bounding effect of /a:/ in Hungarian's [± back] harmony, precisely because of role played by sonority.

5.2.4. Hungarian

Many of the case studies in this book have included instances of variability within a single language, in which there are two distinct parametrizations that speakers may employ. For example, the marked vs. contrastive value-relativization in Finnish (Section 3.11), and one-intervening-syllable vs. unbounded distance in Hungarian (Section 5.1.3) generate within-language variability. In this section we examine yet another phenomenon in which the space of parametric options for vowel harmony provides insight into a pattern of "vacillating" behavior: whether or not Hungarian suffixes copy from the low [−back] vowel ɛ.

Recall that one of the intricacies of the Hungarian vowel inventory is the existence of the short vowels *a,e* that alternate according to [± back] harmony, while at the same time *e,é* and *a,á* alternate with each other according to length-modifying processes. In the following examples these vowels are represented in terms of their surface phonetic values, in particular following Vago (2006), the short *a* is phonetically [ɒ].[83]

(71) *Hungarian shortening alternation with the plural or verbalizing suffix (Siptár and Törkenczy, 2000, 53)*:

a.	vi:z	viztɛk	water (sg/pl)
b.	ɒnɒli:zis	ɒnɒliza:l	analysis/analyze
c.	tü:z	tüzɛk	fire (sg/pl)
d.	miniɒtü:r	miniɒtüriza:l	miniature/miniaturize
e.	u:t	utɒk	road (sg/pl)
f.	u:r	uriza:l	gentleman/play the gentleman
g.	pɒsztö:röz	pɒsztörizɒl	pasteurize
h.	lo:	lovɒk	horse (sg/pl)
i.	ɒgo:niɒ	ɒgoniza:l	agony/agonize
j.	ke:z	kɛzɛk	hand (sg/pl)
k.	pre:mium	prɛmiza:l	bonus/ award a bonus
l.	nya:r	nyɒrɒk	summer (sg/pl)
m.	kɒna:lis	kɒnɒlizɒl	canal / canalize

In addition to morphologically-conditioned shortening, Hungarian exhibits a process of vowel lengthening, affecting the last vowel of any morpheme when it precedes a suffix:

(72) *Hungarian morphological ɛ/ɒ lengthening (Siptár & Törkenczy 2000:170)*:

a.	ɒlmɒ	ɒlma:t	apple/ apple-acc.
b.	tɒrtjɒ	tɒrtja:k	he holds it / they hold it
c.	kutyɒ	kutyya:ul	dog / like a dog
d.	órɒ	o:ra:jɒ	watch / his watch
e.	ɛpɛ	ɛpe:s	bilɛ / bilious
f.	vittɛ	vitte:k	he carried it / they carried it
g.	ɛstɛ	ɛste:rɛ	evening / by evening
h.	mɛsɛ	mɛse:jɛ	tale /his tale

The evidence to the learner on the one hand is that ɛ should be a low vowel, since it alternates with ɒ. On the other hand, ɛ should be a mid vowel, since it also alternates with e:. The conundrum is resolved once we realize that what ɛ,ɒ have in common is being [−ATR] (Reiss, 2003), that what ɒ,a: have in common is being [+low], and what ɛ,e: have in common is being [−low]. A revised analysis of the Hungarian Vowel Inventory incorporates these feature assignments (this revision does not affect the analysis of distance-boundaries of Section 5.1.3):

(73) *Revised Features for the Hungarian Vowel Inventory*:

[−back,+round]	[−back,−round]	[+back,+round]	[+back,−round]	
ü,ü:	i,i:	u,u:		[+high,−low,+ATR]
ö,ö:	e:	o,o:		[−high,−low,+ATR]
	ɛ			[−high,−low, −ATR]
		ɒ	a:	[−high,+low, −ATR]

We may observe how the revised feature values of the inventory in (73) provide a means of accomodating the fact that ɛ,ɒ share the same [−ATR] specification and thus are closer in sonority than e and ɒ, while maintaining the conventional wisdom reflected in the orthography and in length alternations that e,ɛ are of the same height in terms of [−low] and that ɒ,a: are of the same height in terms of [+low].

The revised inventory in (73) has significant consequences for the statement of Hungarian vowel harmony in terms of a copying procedure for contrastive values of [± back]. Only the short and long versions of *ü,u* and *ö,o* will be contrastive for [± back]; all other vowels lack a contrastive [± back] counterpart by the definition of contrastiveness in Chapter 3. A relativization of the harmonic search in Hungarian limited to contrastive values would thus exclude all of the following vowels from search:

(74) *Noncontrastive [± back] vowels in Hungarian*:

[−back,+round]	[−back,-round]	[+back,+round]	[+back,−round]	
	i,i:			[+high,−low,+ATR]
	e:			[−high,−low,+ATR]
	ɛ			[−high,−low, −ATR]
		ɒ	a:	[−high,+low, −ATR]

In point of fact, among the vowels in (74), the [−low,+ATR] vowels *i,i:,e:* are excluded from the search domain, and fully transparent to harmony (cf. (24)). The [+low,−ATR] vowels ɒ,aː intrude into the search process by way of a sonority-hurdle, halting the search and providing a [+back] value as soon as they are encountered. The pattern of ɛ, however, is variable, patterning with neither of these groups.

Vago (1975) discusses "vacillating" stems in Hungarian, so-called because speakers vary in whether the suffixes attached to these stems exhibit [+back] or [−back] harmony. Stems with a [+back] vowel in the initial syllable(s) and the [−ATR,−back] vowel ɛ in the final syllable vacillate in which vowel they copy from:

(75) *Hungarian Dative Suffix Variably copies [+back] from Vacillating Stems*

a. ɒgnɛs-nɒk *or* ɒgnɛs-nɛk 'Agnes-dat'
b. Joszɛf-nɒk *or* Joszɛf-nɛk'Joseph-dat.'
c. szɒlɒmɒndɛr-nɒk *or* szɒlɒmɒndɛr-nɛk 'salamander-dat.'
d. hotɛl-nɒk *or* hotɛl-nɛk 'hotel-.dat'
e. pu:dɛr-nɒk *or* pu:dɛr-nɛk 'powder-dat.'

I propose that the vacillation in how suffixes execute search-and-copy with these stems is due to two different sonority thresholds: one grammar has a more Restrictive Sonority (RS) threshold, forcing search to quit when sonority 5 is exceeded (and hence with the [−low,−ATR] vowel ɛ):

(76) Hungarian RS Dative Suffix must:
Back-Harmonize: δ = L; F = [c: back], ζ = 5

The effect of Hungarian RS's low sonority threshold is that ɛ will halt the search immediately (even though it is not contrastive), and thereby impose [−back] as the value to be copied. The alternative pattern is generated by a slightly more Tolerant Sonority (TS) threshold, in which ɛ is breezed by (due to being noncontrastive); and where only the "lowest of the low", namely ɒ,aː, will halt the search and impose [+back].

(77) Hungarian TS Dative Suffix must:

Back-Harmonize: δ = L; F = [c: back], ζ = 6

Speakers who generate ɒgnɛs-nɛk are therefore using the grammar of Hungarian RS when they do so, whereas those who generate ɒgnɛs-nɒk employ the grammar of Hungarian TS when they do so. The inter- and intra-speaker variation is the result of a different parametric value for the sonority threshold, and the choice between which of Hungarian RS or TS to use is conditioned by a variety of lexical and sociolinguistic factors. However, since the sonority hierarchy is strictly ordered, there can be no idiolect or register of Hungarian in which, say *i:* blocks [+back] harmony across it but ɛ does not. A setting of the sonority threshold low enough that *i:* will halt the search automatically requires all vowels of equal or greater sonority to halt the search as well.

5.3. Implicational Sonority Thresholds

This latter portion of this chapter has explored a locality principle that trumps the exclusion of non-relativized elements from the search, by the high sonority of elements along the search path.[84]. The discovery of this principle came about because, despite the extensive empirical coverage afforded by the relativization to only marked or contrastive visibility, there remain a number of cases in which a non-contrastive (or non-marked) vowel unexpectedly intruded into the search, and this vowel is consistently the highest-sonority in its class. Research focused solely on Wolof alone could explain its pattern by assuming a special marking of /a:/ as lexically linked to [−ATR] (as Ka (1994) does). Research solely on Hungarian could explain why noncontrastive /a:/ is opaque, but noncontrastive /i:/ is not, by again assuming that [+back] vowels intrude into the search in Hungarian (though clearly Hungarian harmony is contrastively-relativized, as evidenced by *soför-nek* 'chauffeur-dat.'). Finally, research solely on Classical Manchu could assume that of the two [+ATR] vowels u,ə, the effect of [+ATR] in *u* is somehow weaker given its [+round] specification (e.g. Zhang and Dresher (1996)). While all of these analyses might be perhaps reasonable within the context of a single language, a broad typological survey repeatedly reveals cases of asymmetric behavior between lower and higher vowels in harmony.

In the remaining subsections, we briefly focus on three cases in which low vowels are transparent to harmony across them.[85] Consistent with the predictions of the implicational generalization in (47), repeated below, a higher vowel must be transparent as well.

(78) Given two relativization-excluded (e.g. noncontrastive/nonmarked) vowels ϕ and ψ, where ϕ is of a higher sonority than ψ, it will never be the case that ϕ is transparent while ψ is not.
(In other words: transparency of ϕ implies transparency of ψ)

5.3.1. Kinande

Kinande has ten surface vowels: i,u,ɪ,ʊ,e,o,ɛ,ɔ,ə,a and a system of [± ATR] harmony. Harmony copies the marked value (which is [+ATR]) from a following high vowel. Evidence that [+ATR] is the marked value may be confirmed because, aside from [+high] vowels, [± ATR] contrasts are neutralized in roots, except as the result of harmony. Broadly consistent with this position, Archangeli and Pulleyblank (2002, 149) analyze [−ATR] roots such as ɛrɪ-lɪma as lacking an underlying specification for ATR altogether, with [−ATR] arising by default, and [+ATR] arising in the case of harmony with a following high vowel.[86]

(79) *Kinande Verbal roots copy [+ ATR] from a following [+high] vowel*:

ɔ-mu-lib-i	cover-agentive	ɛ-ri-lib-a	cover-infinitive
ɔ-mu-huk-i	cook-agentive	ɛ-ri-huk-a	cook-infinitive
ɔ-mu-lim-i	cultivate-agentive	ɛ-rɪ-lɪm-a	cultivate-infinitive
ɔ-mu-hum-i	beat-agentive	ɛ-rɪ-hʊm-a	beat-infinitive
ɔ-mu-hek-i	carry-agentive	ɛ-rɪ-hɛk-a	carry-infinitive
ɔ-mu-boh-i	tie-agentive	ɛ-rɪ-bɔh-a	tie-infinitive
ɔ-mu-kər-i	tie-agentive	ɛ-rɪ-kar-a	tie-infinitive

While some research on Kinande had suggested that *a* was transparent to [+ATR] harmony across it (Schlindwein, 1987), Gick et al. (2006) and Kenstowicz (2008) provide instrumental evidence that *a* undergoes [±ATR] harmony, alternating with ə. In short, all vowels in the inventory initiate-and-copy [±ATR] harmony, consistent with the prediction in (47).

5.3.2. Menominee

Menominee is an Algonquian language spoken in Wisconsin and Michigan. Menominee vowel harmony, as described by Bloomfield (1962) and

Milligan (2000), constitutes a case in which *a* is transparent to harmony, while *ä*, another vowel in the inventory, blocks harmony across it. While it is sometimes misunderstood in papers on Menominee harmony, ä is actually a low (æ) rather than mid (ɛ) vowel.[87] This vowel is, by all accounts (Bloomfield, 1962; Miner, 1979; Hockett, 1981; Milligan, 2000), a **low** vowel; however, Bloomfield's transcription with epsilon (ϵ") led some scholars to interpret this vowel as the mid [−ATR] vowel ɛ, with occasional obsfuscatory consequences for subsequent understanding of the process. The features underlying the Menominee vowel inventory are in (80):

(80) *Menominee Vowel Inventory*

[−back, −round]	[+back, +round]	[−back, −round]	
i	u		[−low, +ATR]
ɛ	ɔ		[−low, −ATR]
æ			[+low, +ATR]
		a	[+low, −ATR]

The basic process of vowel harmony in Menominee is illustrated in (81): [−low] long vowels copy [+ATR] from the right, as shown in the second example of each pair:

(81) *Menominee Vowel Harmony: Nonlow Long Root Vowels Copy [+ATR] from the right*:

a. aːtqnɔːhkæw 'he tells a sacred story'
   aːtqnuːhkuwæw 'he tells him a sacred story'
b. nɛːmɔw 'he dances'
   niːmit 'when he dances'

Following Archangeli and Pulleyblank (1994) and Milligan (2000), we analyze the Menominee vowel harmony process as [±ATR] harmony. Milligan (2000) provides an acoustic study of Menominee, stating that the [−low] vowels ɛ,ɔ may have realizations as ɪ,ʊ when short, which consistent with them maintaining a [−ATR] value while undergoing a change in height. The low vowel æ exhibits extensive surface variability in its height; as Doug Pulleyblank has suggested (personal communication), there may be more variability in the realization of vowels with feature-values that are articulatorily antagonistic, as in [+low,+ATR].

As the following pair shows, [+ATR] harmony is blocked by an intervening /æ/, but may skip across an intervening /a/:

(82) *Menominee Vowel Harmony blocked by* /ae/ *but not* /a/ *(Cole and Trigo, 1988)*:

a. kɛːwæːtuaq 'when they go home'
b. muːskamit 'if he emerges'

I propose that Menominee vowel harmony is a parasitic harmony process among [− low] vowels, searching for the marked value of [+ATR].

(83) Menominee Long Root Vowels Must:
ATR-Harmonize: δ = right, F = [m: ATR & R = −low], ζ =1

As a result, [−ATR] *a* will be invisible, since it is excluded by the domain of relativization. By contrast, [+ATR] æ will constitute a defective intervener, being included in the search domain but of the wrong height. Although *a* is transparent to harmony while æ blocks harmony, these vowels are crucially of a different status with respect to the relativization of the harmonic value: [−ATR] *a* is excluded from search altogether, whereas [+ATR] æ is included in search but fails to meet a conditional requirement on its value for [± low]. As a result, the fact that of the two Menominee low vowels, one is transparent to harmony and one is a defective intervener is irrelevant for the predictions of (47). The implicational generalization in (47) is applicable for cases in which both vowels under comparison have the same status with respect to the value-relativization in the language (as in Wolof, Manchu, Hungarian, and Finnish, above). Menominee simply has parasitic harmony.

5.3.3. Londengese

Londengese (Leitch, 1996) is a 7-vowel Bantu language spoken in the Congo that has [± ATR] harmony affecting its [−high,−low] vowels.

(84) *Londengese Vowel Inventory*:

[−round,−back]	[+round,+back]	[−round,+back]	
i	u		[+high, +ATR, −low]
e	o		[−high, +ATR, −low]
ɛ	ɔ		[−high, −ATR,−low]
		a	[−high, −ATR,+low]

Londengese morphemes copy contrastive [± ATR] from the closest leftward source. Under the relativization to contrastive values, *a* is transparent and skipped over by harmony.

(85) *Londengese [a] transparent to [c: ATR] harmony (Leitch 1996:138)*:

a. t-ok-ak-e 'n'écoute pas'
b. t-o-ya-k-e 'ne viens pas'
c. t-ɛnj-ak-ɛ 'ne tire pas'
d. t-ɔngw-ak-ɛ 'ne vole pas'
e. tɛ-lɛ-k-ɛ 'ne mange pas'
f. a-yo-tepy-ak-e 'qu'il aille parler continuellement'
g. a-y-ɔs-ak-ɛ qu'il aille prendre continuellement'

The sonority threshold in Londengese is set to 1, essentially imposing no extrinsic barriers on the search for contrastive [± ATR]. Additional evidence that *a* is transparent comes from a pattern of consonant reduplication formed by a CV- prefix with the fixed vowel *a*. The diminutive plural prefix *to-* copies [± ATR] from the closest contrastive source to the right, skipping right past *a*:

(86) *Londengese reduplicant [a] transparent to diminutive prefix harmonizing across it (Leitch 1996:139):*

a. to-fu-fumbe 'esclave'
c. to-ba-bo 'amende'
d. tɔ-wa-wɔ 'bras'
e. tɔ-ya-yɛ 'feu'
f. tɔ-sa-sɛ 'querelle'
g. tɔ-ta-twɔ 'prix'

The most fortuitous example of transparency across *a* is found with the prefix *ya-*, which means "action at a distance" (Leitch 1996:140), and allows the eponymous process to occur across its high-sonority vowel:

(87) ɛ-sɔmba ɛ-ya-ndɛ njale *'le bateau remonte la rivière là-bas'*

When we apply the implicational generalization in (47), we see that it demands that the noncontrastive vowels must also be transparent to [± ATR] harmony in Londengese. There are not many high-vowelled affixes in Zone C, and as Leitch reports, "there is not a single example of a form with an overt intervening high vowel and retraction after it," (p.141), However, there is indirect evidence that supports the conclusion that high vowels are transparent in Londengese. The causative morpheme (historically *-is-*) is expressed as palatalization of the root-final consonant, and in Babole, and other Zone C languages, this palatalized consonant blocks [−ATR] harmony across it by the following suffix. In other words, in Babole [± ATR] harmony is set to all values. However, in Londengese, these same consonants are transparent to [+ATR] copying across them by the final vowel, as can be seen in the following alternations:

(88) *Londengese palatalized consonants (the reflex of [+high,+ATR]) transparent to [−ATR] harmony (Leitch 1996:141)*:

a.	sin-y-e	'faire écrire'
b.	somb-y-e	'faire acheter'
c.	amb-y-e	'faire raconter'
d.	sɔ k-y-ɛ	'fatiguer'
e.	bots-e (root: bot-)	'faire engendrer'
f.	pits-e (root: pit-)	'faire abimer'
g.	kɛnj-ɛ (root: kɛnd)	'faire aller'

In Babole, [± ATR] copying is for all values, while in Londengese it is for contrastive values of [± ATR]. The transparency of Londengese palatalized consonants (and the high glides in (85),(86),(87)) – in comparison with other Zone C languages in which neither these nor *a* are transparent – constitutes a case in which [± ATR]-noncontrastive *a* is transparent, according to relativization of the search. As predicted by (47), transparent non-contrastive /a/ implicationally requires the transparency of all lower-sonority non-contrastive segments.

5.4. General Conclusion: Extrinsic Bounds on Search

This chapter has introduced two types of extrinsic bounds on search, one of *extent*, i.e. crossing too many nodes on the search path, and one of *hurdles*, in which crossing a node of very high sonority imposes a barrier on further search. I have proposed that distance-parameters that limit the extent of inter-segmental dependencies such as vowel harmony have five possible settings, and that the sonority-threshold imposed on search barriers must follow the universal scale of sonority, with implicational and predictive consequences and implications for possible and impossible harmony patterns.

The empirical terrain of distance boundaries on search is one that, as we have seen, can exhibit a high degree of within-language variability. I have demonstrated that the variable patterning of *aszpirin*-type words in Hungarian and multiple-*k* prefixes in Gikuyu may be understood as a case of microparametric variation of distance-boundaries, whereas the variable patterning of *agnes*-type words in Hungarian may be understood as a case of microparametric variation in sonority-thresholds. In both cases, the *within*-language variability tracked the same restricted possibilities as variation between languages.

The parallels with syntactic search run very deep: Ross' discovery of islands provided one of the confirming demonstrations that a complete theory of long-distance dependencies in syntax would have to eventually incorporate parametrization of its limits. While agreement and *wh*-movement may often create relations between two positions separated by a host of irrelevant intervening nodes, certain boundaries simply cannot be crossed.

The demonstration of harmony failures – cases in which search is unsuccessful because of extrinsic locality boundaries – provide important limits that rein in what may normally be a very long-distance process. While Chapter 3 focused on how far vowel harmony could travel and its apparently non-local aspects, the study of locality would not have been complete without considering its limits.

Our focus on vowel harmony – a long-distance assimilation process par excellence – as the starting point for the model means that simpler cases of strictly adjacent assimilation might be treated as very bounded cases of harmony. In addition, we have seen in Section 5.1.4 that long-distance iterative dissimilation may be tightly modeled using the distance parameters independently developed for harmony, raising the question of how much of the computation serving these two processes might be shared. We turn to these and other potentially far-reaching implications of the parametrized Search model of vowel harmony in the next chapter.

Notes

[66]See Fitzpatrick (2002) for an overview and comparison of different approaches to minimality-based locality within minimalist syntax.

[67]Abels (2003) provides an interesting attempt to assimilate syntactic cases of boundedness as a type of minimality.

[68]An insertion rule supplying [+lateral] by default in the case of [+sonorant, −nasal] segments is additionally necessary.

[69]Suzuki (1998) has developed a theory for distance parameters in dissimilation, concluding that dissimilation may vary in its distance-boundaries in fundamentally similar ways.

[70]See Boeckx and Grohmann (2007) for a comparison of Barriers and Phase Boundaries as extrinsic locality restrictions.

[71]An interesting case arises with respect to what happens when δ is bidirectional (as in the Woleaian, Section 2.8) and β is set to one intervening segment, for cases in which the searching vowel has a single intervening consonant on one side, but two or more intervening consonants on the other side. Does the failure to find a value-source on one side cause the entire bidirectional search to crash, or does the search still copy the value that was successfully found on one side? This question is part of a larger research issue in phonology as to whether structural descriptions may be met universally or existentially in cases of multiple environments (see for example Raimy (2000) for discussion of a uniformity parameter governing the interpretation of structural descriptions in gemination and reduplication), and could be fruitfully answered in artificial grammar experiments of the type discussed in Chapter 6, if no extant languages with this relevant combination of properties are found.

[72]I do not provide a treatment for Hungarian stems such as *hid* 'bridge', that have a [−back] vowel in the stem but require harmonizing suffixes to copy [+back] from them. Within the Agree-based approach to vowel harmony, these stems are akin to 'epicene' nouns such as French *sentinelle* that are semantically masculine but require syntactic agreement with them to feminine (Wechsler and Zlatíc, 2003), necessitating a distinction between inherent features available to the interpretive interface and the feature available for copying in syntagmatic dependencies. See Ringen and Vago (1998) and Krämer (2003) for similar approaches to agreement with such stems.

[73]We return in Section 5.2.4 to a discussion of the [+low] vowels in Hungarian [± back] harmony.

[74]The pattern in which Hungarian ML suffixes surface with [−back] when two but not one noncontrastive [−back] vowels precedes seems crucially different from the "bisyllabic trigger requirement" of Classical Manchu (Zhang and Dresher, 1996; Walker, 2001). In Hungarian ML, harmony with the relativized-value occurs as usual when one [−back] vowel intervenes, but fails when two intervene. In Classical Manchu, suffixes *fail* to harmonize when one [+round] vowel precedes, and only harmonize when two [+round] vowels with the relativized value precede. Classical Manchu may be approached by analyzing its suffixes as required to find two instances of [+round] in their search; a requirement consistent with a target-centric approach.

[75]As mentioned in footnote 33, a single individual may simultaneously exhibit both of these patterns, with lexical and sociolinguistic conditioning, which can be accomo-

dated in a model of intra-speaker variability with multiple parameters such as those of Kroch (1989); Roeper (1999); Yang (2003). See also Hayes and Londe (2006, p.77,fn12) who entertain the possibility of co-existent harmony-processes, one which allows a single intervening syllable and one which is unbounded.

[76]I verified the Gikuyu data with two native speakers, Paul Njoroge and Sam Gakindi of Cambridge, MA.

[77]See also Ringen and Heinämäki (1999) for remarks on sonority in Finnish vowel harmony and Hayes and Londe (2006, p.83,fn19) for remarks on sonority in Hungarian vowel harmony.

[78]In a study of coarticulatory effects in Hungarian [± back] harmony using EMMA magnometry and ultrasound imaging, Benus and Gafos (2007) found that some cases of noncontrastive *i,i:,e:* in Hungarian [+back] harmonic words are produced with a significantly backed articulation, suggesting a susceptibility to coarticulation of the same type that Öhman (1966) and others have found for languages without harmony. Benus and Gafos (2007, p.290) note a puzzling finding with respect to duration. "Long vowels have more time to achieve their target and are thus less prone to contextual coarticulatory influences than their short counterparts.. . .In the transparent vowels used in our experiments, there were two long vowels, [i:] and [e:], and one short vowel [i]. In the EMMA data, [e:] was affected by harmonic type the most. In the ultrasound data, it was the other long vowel [i:] that was affected the most. Therefore, long vowels were affected by harmonic type more than short vowels." While Benus and Gafos (2007) do not provide an explanation for this finding, it seems to provide confirmation that duration is irrelevant to understanding harmony in Hungarian.

[79]The sonority of the vowel ə in this hierarchy is relevant for cases in which it is phonemic, and not for purely allophonic reduced unstressed vowels.

[80]The software implementation of this algorithm includes the option to treat all phonologically inactive values as present for the purposes of computing these sonority values.

[81]A height-based explanation of the u/ʊ and ə/a contrast (e.g. Ard (1984)) cannot account for the velar/uvular alternations: ə, a non-high vowel, takes the velar alternant, while /ʊ/, a high vowel, takes the uvular alternant.

[82]There is some variation in the data for the transparency of *u*; for certain attestations, *u* does participate in harmony. There is thus one pattern of Classical Manchu that allows copying from *u*; this pattern consistenly copies from ə and consistenly skips *i*. Entirely parallel to the analysis of microvariation within Kirghiz in Section 3.8.2, I propose that there are two grammars of Classical Manchu: the one in the text, specified in (57), relativized to marked values of [± ATR], and one, 'Classical Manchu B', with the parameter for value-relativization set to *contrastive* values:

(89) Classical Manchu B Adjectivalizing, Verbalizing, Diminutive Suffixes Must:
ATR-Harmonize: δ = L; F = [c: ATR]

The variant of Classical Manchu B is no less worthy of study than the variant in the text. The one in the text, however, allows us to diagnose the effect of a sonority hurdle, which cannot be detected in Classical Manchu B, since both ə and u *are* included in its domain of search.

[83]Indeed, Hungarian [ɒ], often transcribed as [ɔ], sounds much lower to the unaided ear

– being approximately the height of the Farsi back low round vowel– than, for example, Wolof or Brazilian Portuguese [ɔ].

[84]An interesting question is whether the effects of sonority-hurdles can be found lower down the sonority scale, i.e. with consonantal elements. As discussed in Section 5.1.2, Assamese has [± ATR] harmony. However, there is a systematic set of exceptions to the rightward search-and-copy for [+ATR]: when a nasal consonant intervenes.

(90) Assamese Non-Final Vowels cannot Copy [+ATR] Past a Nasal Consonant:

a. sɛkɔni 'strainer'
b. xɔmɔnia 'colleague'
c. pʊtɔni 'dumping ground'
d. k^hɔmir 'leavening agent'
e. mɔni 'pearl'

This failure of the mid vowels in (90) to copy [+ATR] from the final *i* cannot be a case of defective intervention, as nasals do not bear the feature [±ATR]. Mahanta (2007, 183) assumes that nasals are more sonorous than liquids in Assamese, and follows the proposal in Nevins (2004) (and in the present chapter) that sonority hurdles are responsible for the blocking effect in harmony.

[85]Metaphony processes in Spanish and Italian skip over vowels and target only the stressed syllable. I do not treat such cases of 'morphemic' harmony in this book, and discuss in Section 1.4 that this is principled exclusion from the pursuit of phonological locality.

[86]Orthogonally to the [+ATR] copying yielding a/ə alternations, Kinande has an independent process whereby the [−low] applicative and reversive suffixes harmonize in height and ATR with the vowel to their left, similar to languages described in Section 4.4.2.

[87]Goddard (1987) discusses the inconsistency in the transcription of the low front vowel that Bloomfield originally wrote *ä*, but later, for reasons of typographical preference, changed to "ϵ".

Chapter 6

Minimalist Computation of Vowel Harmony: Implications

We are now ready to review and highlight the broader theoretical relevance of the formal model of vowel harmony developed in previous chapters. As the discussion proceeds, I will attempt to identify areas of potential future collaboration with other branches of linguistics and cognitive science.

I have proposed a Principles-and-Parameters theory of possible vowel harmony systems in human language, with four components. The first component is an invariant *Search Principle*, in which value-seeking elements initiate a search for the feature they need, stop as soon as they find the closest element bearing the relevant feature, and copy the value of that feature to themselves. The second component is the *Relativization Parameters* that determine what is counted as "relevant": following Calabrese (1995), we saw that the three parametric possibilities are *all* values of the harmonic feature, only the *contrastive* instances, or only the *marked* values of that feature. The third component are *Conditional Requirements* on licit sources of feature-copying: in addition to bearing the relativized value of the harmonic feature, an additional subrequirement of identity between source and goal is demanded. The fourth and final parametric component of the search is *Bounding Parameters*, fractionated into a set of extrinsic limits on absolute distance of search and a set of tolerance thresholds determining whether high-sonority elements can be passed or not.

Our inquiry into the locality of vowel harmony has provided empirical arguments for a recipient-initiated Search, based on Split-Source Harmony and Bidirectional Searches. It has also provided an illustration of how microvariation within a language can be insightfully modeled based on different settings of a single parameter. The treatment of defective intervention in Chapter 4 places significant restrictions on what can be a possible "blocker", by restricting orthogonal feature identity and heteromorphemicity as the only possible conditions yielding blocking, with a non-negotiable principle of No Second Chance in Search. Chapter 5 explained how two new adjacency parameters to the ste of those proposed by Odden (1994) expand the typology of boundedness, and established a novel implicational generalization about the role of sonority in blocking harmony.

All of the above contributions to the understanding of vowel harmony might be viewed as advancements in phonological theory. Nonetheless, this book's *raison d'être* is not only to provide extensive empirical coverage and unify a wide set of disparate harmony cases under a single model, but also to explore an important and often neglected question of higher-order synthesis: does phonology instantiate any procedures of "minimalist" computation of the type argued to exist in syntax?

The Minimalist program (Chomsky, 1995) represents a major shift in linguistic theory in the way that language is viewed as an "optimal solution" to design features and interface considerations. Whether or not phonological computations present aspects of optimal solutions to interface conditions has yet to be fully addressed. I hope to convince phonologists that viewing vowel harmony through the lens of Minimalism not only opens the potential to unify different levels of linguistic representation and different domains of empirical inquiry under the same foundational framework, but that moreover the specific implementation of a theory of the locality of dependencies represents a step forward in understanding constraints on possible harmonic languages.

6.1. Expanding the terrain of Minimalist inquiry

Contrary to the title and conclusions of Bromberger and Halle (1989)'s article, phonology may not be as "different" as it seems. One of the goals in this book has been to bring empirical phenomena within phonology closer to the attention of Minimalist theorists, whose training and attention is largely focused on syntactic problems. Due to the surface opacity of phonological theory created by the fact that the content of Optimality-Theoretic constraints and autosegmental representations do not resemble the primitives of syntactic theory, opportunities for unification and cross-fertilization have remained limited simply due to the apparent disparities in representational vocabulary. Howevever, by framing vowel harmony in the context of the mechanism of *Agree*, and demonstrating that the logic of closest-search within a relativized domain can apply in phonology under the relation of precedence in the same way that it applies in syntax under the relation of dominance, one increases the prospects for syntacticians and cognitive scientists more generally to view both subsegmental dependencies and phrase-structural dependencies as both manifestations of the same cognitive architecture.

The minimalist program makes two central claims about the design of human language. The first is that computations are derivational, efficient, and follow principles of least-effort in order to satisfy feature-valuation. We have seen that the Search principle is greedy (cf. Section 2.11), and myopic[88] (cf. Chapter 4). The second claim about language design is that minimalist computations are interface-driven. Vowel harmony lies at the interface between two components of human cognition: the lexicon and phonetic realization. The defining property of harmonizing morphemes is that they lack a value for the harmonic feature. These segments require a value for the harmonic feature in order to be interpreted by the articulatory component of language and perceived by the interlocutor. Vowel harmony seems to be a perfect solution to the problem of supplying a feature-value for a needy vowel: derivationally-initiated feature-copying that stops with the closest possible source.

While many lines of cognitive science question whether syntactic or phonological computations are language-specific (e.g. Langacker (1987), Lieberman (2000); Marcus (2008)), much less work has been devoted to the relationship *between* linguistic modules *within* the faculty of language. Hockett (1960)'s very notion of "duality of patterning" – the existence of combinatorial generation and restrictions at both the phonological and syntactic level – suggests a formal isomorphism may be at work. Rather than claiming that phonology is domain-general or part of generalized cognition, the suggestion I would like to offer is that it bears a closer relation to other linguistic modules than was perhaps previously understood. This view may be traced to older notions predating generative treatments of vowel harmony: Hjemslev (1948) formulated an 'analogie du principe structurel', according to which the units and structural properties that describe one aspect of language (such as syntactic relations within sentences) should be expected to be fundamentally the same as those required by the analysis of other aspects of language (see Bermúdez-Otero and Honeybone (2006) for discussion). While Government Phonology (e.g. Kaye et al. (1985); Charette (1991); Pöchtrager (2006)) embraces the notion that syllable structure follows a binary-branching and headed pattern like that of syntactic constituency, far too little attention has been paid to long-distance dependencies in phonology as a direct structural analogue of syntactic agreement.

One could counterfactually imagine a language faculty in which syntactic agreement was with the most distant element in the domain, while vowel harmony was with the closest element in the domain, or that vice versa were the case. Similarly, one could imagine a phonological component without defective intervention, in which defective elements simply let search go right by them. But I have identified three formal properties of syntactic search-and-copy that are identical with that of vowel harmony: Relativized Minimality, Defective Intervention, and the Minimality/Boundedness distinction. These algorithmic properties of *Harmonize* and *Agree* arguably are due to the same underlying procedure, operating only over different data structures. While syntactic Agree operates over c-command, Harmonize operates over precedence; while Agree copies ϕ-features such as gender and number, Harmonize copies subsegmental features such as [$\pm$ round] and [$\pm$ ATR]. We are confident that the recurrent identical operative principles in both domains are due to the fact that the two literally use the same procedure.

Anderson (1992) discusses the consequences of absence of active cross-modular comparison between syntax and phonology, concluding that excessive modularization can lead to theoretical parochialism and "ill-supported notational alternatives" (p.3). It is my contention that the autosegmental formalism of the No Line-Crossing Constraint, as applied to vowel harmony, is a parochial notation that ultimately obscures the formal identity between relativized locality in syntax and phonology.[89] If the hypothesis of Crossmodular Structural Parallelism is to be pursued, it is clear that the literally geometric statement of the line-crossing constraint is an absurd means for expressing relativized minimality in syntax. The Search-and-Copy algorithm, by comparison, is flexible enough to compute closeness over either dominance (in Agree) or precedence (in vowel harmony), while providing a tight fit to the empirical terrain of locality restrictions.[90]

Once an identical algorithm is in place for syntactic agreement and for phonological vowel harmony – that differs only in the data structures it operates over – the acquisition of language potentially becomes considerably easier. The learnability challenges arise only in acquiring the specifications of each syntactic head and the ϕ features it needs or bears or the specifications of each morpheme and the agreement features it needs or bears, and the parametric specifications of where and how far they can look. The notion of 'micro-parameters' is that linguistic variation is determined only by the specifications of each node entering into computation. This proposal was introduced into syntactic theory by Borer (1984), brought strongly to the fore within the minimalist program by Chomsky (1995), and has been characterized in terms of restricted schema of parameters for a given syntactic node in Longobardi (2005).

The reformulation of parameter within the Minimalist program as a set of featural requirements on individual lexical items dovetails very well with the locus of parametric variation between harmony patterns in the current model. Recall that the only ways that languages differ with respect to their vowel harmony pattern are in (1) what features they search for, (2) whether conditional subrequirements are imposed on licit copying, (3) which direction they look, and (4) how much extrinsic bounding they tolerate. In fact, proposals of a parametrized system of vowel harmony have a rich tradition.

The autosegmental and metrical research tradition for vowel harmony, spanning work including Steriade (1981) to Archangeli and Pulleyblank (1994), often employed a restricted set of parameters along which harmony patterns could differ; more directly the proposals of Odden (1994) and Calabrese (1995) form the basis for the parameters that Chapters 3 and 5 expand upon. However, in these previous proposals, the parameters largely focused on the harmony pattern of the language as a whole, rather than being properties of specific morphemes. As a consequence of situating parametric variation as a property of vowel harmony in the languages as a whole, morphemes that did not undergo harmony required extremely special lexical marking. By contrast, in the present model, we situate parametric specifications as a property of each morpheme, and if a particular morpheme does not undergo harmony, nothing special – in fact, nothing at all, needs to be said about it.

Within all microparametric theories in which the locus of variation is individual morphemes, an important question is the role of inductive biases for uniformity of the parameter settings for broad classes of the items. Although we have seen in Chapter 1 that Turkish suffixes may differ from each other in that one requires Back-Harmony and Round-Harmony, one requires Back-Harmony and Round-Harmony for only one of its vowels, one requires Back-Harmony only, and yet another suffix requires no harmony at all, we have also seen cases where virtually all suffixes in the language harmonize in the same way (e.g. Wolof). Goodman (1955) introduces the notion of an "overhypothesis": if you see that many bags of marbles contain a uniform color (e.g. some bags contain all red marbles, some contain all green marbles), you will eventually try to guess the entire contents of each new bag of marbles just by needing to see one marble from it. While each bag of marbles may be generated by its own independent "rule", an overhypothesis is guess about the uniformity among the rules for each bag. Questions about the uniformity of affixal patterning (into 'classes of morphemes') may be ideally investigated through studies of first- and second-language acquisition.

Microparametric theories of generative grammar have not fully tackled the empirical question of whether language learners quickly generalize from one morpheme's parametric settings to another's. Existing data on language acquisition in Turkish and Finnish indicates that children master vowel harmony very quickly and with few errors (Aksu-Koç and Slobin, 1985; Leiwo et al., 2006). These data indicate that once children learn that a few suffixes are parametrically set to "needy" for feature [$\pm$ F] in their language, they may efficiently generalize to other suffixes, but do not exhibit the incorrect overgeneralization that would be expected if "leftward [$\pm$ back] vowel harmony" were a parameter of the language as a whole. Within the domain of syntactic acquisition, Lightfoot and Westergaard (2007) argue that children do not overgeneralize the featural requirements of complementizers from one clause-type to another, and that they correctly distinguish declaratives and exclamatives with respect to V2. Language acquisition research bearing on the question of "overhypotheses" with respect to microparameter setting is in its early stages empirically and in terms of modelling[91], and questions about the realistic nature of microparametric grammars for phonology and syntax should proceed in tandem, with frequent cross-dialogue.

In conclusion, I would argue that the hypothesis of Crossmodular Structural Parallelism embraces much more than analogical similarities between the syntax of agreement and long-distance vowel harmony, and can be taken literally as a claim that the same Search routine, with its properties of no-lookahead and halting with the first element in the relativized domain, is being "re-used" across different levels of linguistic structure. Once the basic algorithm for Search-and-Copy is in place within the faculty of language, it can be applied to any alphabet of data structures situated at the interface between two components that require an efficient means of feature valuation.

6.2. Are Assimilation and Dissimilation Subroutines of Harmony?

Our focus on vowel harmony in this book was anchored on two pillars (a) vowel harmony is our best chance among phonological phenomena to observe long-distance intersegmental relations that are comparable to syntactic long-distance relations, and (b) vowel harmony displays a wide enough range of crosslinguistic variation to present a non-trivial learning problem, therefore a testing ground for a Principles-and-Parameters approach. In developing a formally explicit parametrized algorithm for vowel harmony, one consequence is that certain routines involved in the process can also serve to carry out the computations for local assimilation or for dissimilation. Let us review the Search algorithm of Figure 5.1.

Minor modifications to this basic algorithm can yield assimilation and dissimilation. For example, as soon as the distance parameter is set to $\beta = 1$, we have segmental adjacency. As soon as the value to be copied for feature f is multiplied by *-1, we have dissimilation. These parametric modifications to the algorithm reveal that once long-distance harmony is accurately handled within a formally explicit theory, assimilation and dissimilation may become special subcases of intersegmental feature relations.

The important questions, however, are empirical: although we can model dissimilation as "opposite-value harmony" and local assimilation as "very bounded harmony", do we still expect the full range of parametric options of long-distance vowel harmony, developed in Chapters 3-5, to reveal themselves in these processes?

Consider the notion of parasitic harmony as applied to local assimilation. What is the typology of cases of assimilation for a feature [$\pm$ F] that occurs only if both segments already bear identity for an orthogonal feature [$\pm$ G]? There are certainly clear cases of assimilation that do require identity along an existing separate feature in order for feature-copying to occur. An example from English is the [$\pm$ anterior] assimilation that only occurs among adjacent segments that are both [+continuant] (e.g. *this shoe* → [ðɪʃ ʃu]). This pattern can be treated within the logic of parasitic harmony as cases of assimilation dependent on a conditional requirement of shared identity.

The possibility of treating orthogonal feature requirements as a prerequisite for local assimilation has not been thoroughly explored to date, and may prove revealing: if on the other hand, parasitic assimilation is not robustly found, then perhaps the extension of the harmonic search-and-copy mechanism to local assimilation is the wrong move. Work such as Cho (1990) and Jun (2004) has extensively formalized implicational conditions on the manner, place, and directionality of assimilation undergoers, but has not unearthed a typology of cases in which the value-source and value-seeker must share an orthogonal feature.

In addition to the extent to which parasitic assimilation exists, we can ask questions about other conditions on vowel harmony that apparently have not yet been documented within local assimilation. One such question would be, does context-sensitive markedness enter into assimilation in the same way that it does in, e.g. Kirghiz B [$\pm$ round] harmony (cf. Section 3.8.2)? If so we would expect patterns of assimilation where only the marked source of a feature may be copied from for local assimilation. A schematic example would be if the marked [+round] feature of *ü* but not *u* could be copied from by a preceding consonant, so that */nü/ → [mü]* but */nu/ ↛ [mu]*. While such predictions certainly seem unlikely, the effects of context-sensitive markedness in Kirghiz vowel harmony were also surprising when first encountered. Only once there is a hypothesis on the table that such phenomena might exist does one begins to look for them actively.

Should it turn out, upon empirical comparison, that long-distance vowel harmony and assimilation might be accomplished by different mechanisms because the latter lacks the full range of effects of parasitic requirements and relativization to marked values, a parallel question should be raised in syntactic theory. The *Agree* formalism was developed for cases of verbal agreement at a distance with noun phrases, and it is appropriate to ask whether the same mechanism underlies the very local agreement between adjectives and noun phrases, known as *Concord*.

Concord seems to be insensitive to hierarchical considerations, and is initated by all elements within a noun phrase constituent; there are no known cases of defective intervention or feature-relativized locality for DP-internal concord.[92] Moreover, Concord never copies person features, but only gender and number; is this lacuna fundamental to the nature of how Concord works, or merely incidental? Further research should focus on the twin questions of whether verbal agreement and DP-internal concord should be formally dissociated and whether long-distance harmony and local assimilation should be formally dissociated, ideally with cross-dialogue between these enterprises. While we pursue answers and await confirmation of these questions, I will suggest that although the subroutine of vowel harmony *can* be used in a limiting case to execute local assimilation, the two phenomena may not actually *be* accomplished through the same procedure. Indeed, they may serve different interfaces: vowel harmony occurs upon the concatenation of two morphemes, while much of local assimilation may be a consequence of post-cyclic articulatorily-driven processes.

The case for dissimilation being the inverse of vowel harmony at first blush fares better: similar to long-distance vowel harmony, there are cases of dissimilation that are (a) clearly long distance, (b) clearly restricted by contrastiveness and (c) clearly "myopic". The extended analysis of Dahl's law in Section 5.1.4 illustrated marked-feature relativization. A well-known dissimilation involving contrastive-value relativization occurs the Latin Adjectival Suffix, which searches for the closest contrastive value for [± lateral] and copying the opposite value:

(1) *Latin Liquid Dissimilation (Gildersleeve and Lodge, 1895)*:

a.	nav-alis	naval
b.	mort-alis	mortal
c.	ven-alis	venal
d.	caud-alis	caudal
e.	milit-aris	military
f.	lun-aris	lunar
g.	consul-aris	consular
h.	vulg-aris	common
i.	flor-alis	floral
j.	sepulchr-alis	funereal
k.	litor-alis	of the shore

The alternation in the suffix between *r/l*, i.e. [± lateral], is asymmetrically dependent on the value of [± lateral] for the closest leftward [+cons, +sonorant, −nasal] element in the stem: when *l* is the closest, the suffix is *-aris* and when *r* is the closest, the suffix is *-alis*. We can model this process as copying of the *opposite* value of the closest contrastive [± lateral]:

(2) Latin Adjectival Suffix Must:
Lateral-Dissimilate: δ = L, F = opposite([c: lateral])

The search is indeed for the *closest* value of [± lateral], as shown by (1-i-k), where both an *l* and *r* occur in the stem, and suffix copies [+lateral] based on the closer [−lateral] element.

The Latin lateral dissimilation process illustrates the potential for long-distance dissimilation to be based on the same algorithm of Search. Nonetheless, one is struck by how the expanded repertoire of vowel harmony parameters have a much more limited extent in dissimilation.

Certainly the value-relativization parameters of Chapter 3 are applicable in delimiting the typology of possible dissimilation patterns. The long-distance [± voice] dissimilation found in Dahl's Law in Gikuyu and Lyman's Law in Japanese (Ito and Mester, 2003) present robust cases of long-distance dissimilation for the marked value of a feature. Moreover, Suzuki (1998) has demonstrated, many of the distance parameters relevant for harmony of Odden (1994) are attested within the typology of dissimilation.

On the one hand the breadth of dissimilation phenomena involving value-relativization and distance-relativization look promising for a unified theory of harmony and dissimilation, as I attempted in Nevins (2004). Analyses that employ the metrics of locality defined by the No Line-Crossing Constraint have no clear no way of presenting a unified theory of assimilation and dissimilation, since dissimilation doesn't involve any "lines" to begin with. As a consequence, the formal unity of harmony and dissimilation was barely raised while phonological theory adopted the autosegmental framework. However, once the locality of vowel harmony is treated as relativized locality along the dimension of the precedence relation and as a copying process, the formal expression of dissimilation suddenly becomes possible to express with the same mechanism. Indeed, Pycha et al. (2003), in an experimental study of adults' abilities to learn artificial grammatical patterns, found that subjects could learn vowel harmony (requiring the same value for [± F]) equally well as disharmony (requiring the opposite value for [± F]).

Despite the promise of treating dissimilation as a kind of "opposite-value" harmony process, however, the paucity of robust cases of iterative cases of dissimilation is somewhat striking. One of the characteristic properties of vowel harmony languages, as exemplified by with the thirteen harmonizing Turkish suffixes in a row in Chapter 1, is its iterative nature: one morpheme M_j harmonizes with the morpheme to its left, M_{j-1}, and M_j subsequently becomes a source of copying for the next morpheme M_{j+1} in turn. Within the realm of dissimilation processes, Dahl's Law appears to be the only case of iterative dissimilation. In particular, the dialect that generates *ɣa.kaa.ɣii.kia* embodies a case in which each morpheme dissimilates in [± voice] from the one to its immediate right, iteratively.

If dissimilation were just the mere flipping of a switch (whether to copy the identical value or not) from the basic harmony procedure, the question would be why we don't observe more cases of iterative dissimilation. Similar questions can be raised for other possible dissimilation phenomena that we might expect if dissimilation were simply a special subcase of harmony. We are led to ask, for example, whether are there any cases of "parasitic dissimilation" – dissimilation for [± F] only if there is already (non)-identity for an orthogonal feature [± G]. A concrete case of parasitic dissimilation would be one where dissimilation of aspiration (the feature [+spread glottis]) would only occur between obstruents that were the same value of [+voice], e.g $/\mathrm{b^hu\text{-}g^ha}/ \rightarrow$ [b^{h}u-ga], but $/\mathrm{b^hu\text{-}k^ha}/ \rightarrow$ [b^{h}u-k^{h}a]. Parasitic dissimilation of this sort would be expected to even yield defective intervention, so that $/\mathrm{b^hu\text{-}k^he\text{-}g^ha}/ \rightarrow$ [b^{h}u-k^{h}e-g^{h}a]. I know of no cases of defective intervention in dissimilation, and one wonders if this is simply due to the limits of small sample size, or reflects a principled lacuna. Moreover, if dissimilation were simply the flip-side of vowel harmony, we might expect that sonority-hurdles could potentially affect dissimilation, blocking long-distance dissimilation by a high-sonority element.

An important consideration in the study of whether dissimilation really implicates the same mechanism as harmonic Search-and-Copy could be drawn from the fact that while syntax is replete with cases of long-distance *Agree*, there are no cases of syntactic *Dis-Agree*, i.e. syntactic dissimilation in which the opposite value of a feature is copied.[93]

Perhaps more disturbing than any of the formal absences of parasitic dissimilation, sonority-thresholded dissimilation, defective-intervention dissimilation, and the rarity of iterative dissimilation is a more pressing question: why is there so little vowel dissimilation? There are a few instances of [+low] dissimilation among vowels (e.g. Woleaian, Kera; (Suzuki, 1998)), but other than these, there are are no cases of dissimilation for other vowel features. I know of no cases of [± back], [±round] , [±ATR], or [±nasal] dissimilation among vowels.

Nespor et al. (2003) suggest that the relative rarity of vowel dissimilation is a phenomenon that must be linked to the relative rarity of consonant harmony. While of course there are a significant number of cases of consonant harmony (e.g. Karaim, Chapter 3; Kikongo, Chapter 5; sibilant harmony (Poser, 1982)), especially documented within Hansson (2001), it is generally agreed upon that consonant harmony is more limited than vowel harmony. Nespor et al. (2003) offer the broad-brush observation that, in terms of long-distance interactions, vowels like to assimilate and consonants like to dissimilate. These observations form part of a far-reaching hypothesis about the nature of vowels and consonants, and their distinctive functional roles.

The observations that Nespor et al. (2003) synthesize draw on diverse sources of evidence in support of the general claim that consonants must remain distinctive and contrastive whereas vowels need not bear any burden of lexical contrast. For example, almost all languages have more consonants than vowels; or, put differently, few, if any languages, have more vowels in their inventory (excluding diphthongs) than they have consonants in their inventory. Compare the consonant to vowel ratios of the following broadly chosen languages: Malay is 20 C: 5 V, Italian is 24 C: 7 V, Hausa is 32 C: 5 V, Arabic is 29 C: 3 V, Igbo is 27 C: 8 V, Sindi is 46 C: 1 0V, Hawaiian is 8 C: 5 V. Nespor et al. (2003) argue that this asymmetry reflects a principled division of labor in which consonants are responsible for filling out the distinctive combinatorics of the lexicon.

The Semitic languages have a morphological structure in which the lexical content of a word is composed by the consonants, that form a lexical skeleton, while the vowels contain inflectional and grammatical information (McCarthy, 1985; Arad, 2005). Importantly, no "Inverse Semitic" – in which the vowels are lexical skeleta but the consonants are the variable grammatical morphemes – is known to exist. Moreover, the Semitic languages are replete with root co-occurrence restrictions among the consonants that seek to avoid co-occurrence of the same feature multiple times within a single consonantal root (Greenberg, 1950; Frisch et al., 2004). Nespor et al. (2003) point out that similarity-avoidance effects are almost always focused on consonants; tongue-twisters, for example, are crosslinguistically based on the difficulty of repetition of certain consonants, and largely do not manipulate vowels.

It seems that vowels are "cheaper" to modify than consonants; notice, for example that vowel reduction is a widespread phenomenon, but that nothing nearly as systematic exists among consonants (for example, no languages have processes that neutralize all coda consonants to a single segment, say *t*). Experimental evidence also reveals that modifying vowels is apparently easier than modifying consonants. Cutler et al. (2000) provided subjects with a task in which they were allowed to change one phoneme to make a word from a nonword. When Spanish and Dutch subjects were provided with a nonword such as *kebra*, they chose to alter the vowel more frequently than the consonant, producing *cobra* more often than *zebra*. Even Voltaire is quoted as saying "Etymology is a science in which the vowels count for nothing, and the consonants for very little" (Müller, 1994, 238). Whether Voltaire was right about the absolute quantities, at least he correctly identified the direction of the asymmetry: consonants are much more reliable backbones of lexical structure than vowels. Surendran and Niyogi (2003) found that consonants carry a functional load (i.e. carry informational content relevant for lexical distinctiveness, even relativized to frequency) of three times more than that of vowels in Mandarin, English, German, and Dutch.

Nespor et al. (2003) develop the idea that vowels and consonants have fundamentally different functional roles: consonants exist in order to generate the lexicon (i.e. the paradigmatic component of syntax), whereas vowels exist in order to facilitate parsing of the speech stream into distinct morphosyntactic constituents. On this view, vowel harmony and consonantal dissimilation may serve fundamentally different interfaces: consonantal dissimilation may be a computation with the purpose of facilitating lexical access of open-class items (e.g. *homonym avoidance*), whereas vowel harmony may serve the purpose of parsing the morphosyntactic words in phrases (e.g. *oronym avoidance*). Other views in the literature adopt a "carrier"/modulation view of vowels vs. consonants (Traunmüller, 1994; Harris and Urua, 2001), and experimental studies such as Owren and Cardillo (2006) find that vowels are used much more for talker recognition and indexical cues than consonants, while consonants are used much more for word recognition and meaning recovery than vowels.

The diachronic sources of vowel harmony and consonantal dissimilation may be fundamentally different as well. Suomi (1983) proposes that vowel harmony arises from perceptual re-analysis of words with weak non-initial (more broadly, non-root) vowels as being asymmetrically dependent on initial syllables for their backness specification. By contrast, Ohala (1981) proposes that the diachronic source for dissimilation is fundamentally one of hypercorrection: a listener encounters a form like Latin *quinque* $[\mathrm{k^w i^{(w)} n k^w e}]$ with rounding on the two consonants that abut a vowel, notices rounding on the vowel, and attributes the rounding to coarticulation on an underlying form with only one source of rounding: $/\mathrm{kink^w e}/$, later becoming Italian *cinque*. If these two distinct perception-based accounts for the emergence of vowel harmony and consonantal dissimilation, respectively, are correct, they point to these two phenomena having crucially different origins in the (mis)learning procedure itself.

The observations and explanations I have offered above regarding the potential differences between harmony and dissimilation are preliminary and require a great deal of promising research ahead. My purpose in the foregoing remarks has been to raise more questions than provide answers, in the hope of inspiring continued comparison of these two phenomena. On the one hand, long-distance dissimilation and vowel harmony share a number of formal properties, such as relativization to contrastive or marked values and distance-parameters. On the other hand, dissimilation is a syntagmatic process with essentially no analogue in syntax, with no formal instantiation of defective intervention, and in terms of its substantive properties, it virtually never operates on vowel features.

The relative weight of the evidence in these two directions merits serious continued attention in terms of both empirical research and theoretical development, in particular because the proposed Search-and-Copy formalism allows – for the first time in generative phonology – the potential to formally unite harmony and dissimilation under the roof of one algorithmic procedure.

6.3. Vowel Harmony and Impossible Languages

> *"Truth is more likely to emerge from error than from vagueness"*
> T.H. Huxley.

Can there be a set of conditions that determine whether a language would be impossible to learn? The aim of generative grammar and the hypothesis within the principles and parameter context posits that certain aspects of human language are universal and inviolable (the principles), while other aspects may vary in very limited ways (the parameters).[94] The way that the Search procedure works is a Principle, while value-relativization, conditional requirements on sources, distance-bounds, and sonority hurdles are parameters, all of which have very restricted means in which they vary. The resulting set of harmony patterns represents a very small fraction of the space of logically possible languages. A wide range of harmony patterns that we could imaginably construct would nevertheless not be allowed by the system of principles and parameters I have proposed. The following is a sample list of harmony languages that are excluded by the current model.

(3) *Impossible Harmony Systems*:

a. Copying from the furthest vowel away from the value-seeker
b. Copying only from vowels noncontrastive for the harmonic feature; exclusion of others
c. Copying only from vowels unmarked for the harmonic feature; exclusion of others
d. Long-distance parasitic harmony: absence of defective intervention when there is a conditional requirement on sources
e. Search that can look past two intervening syllables, but not three intervening syllables
f. Noncontrastive & transparent *a* in the same language as noncontrastive but nontransparent *i*

The impossible languages in (3) are excluded by either the inviolable nature of halting with the closest element, or by the small space of allowable value-relativization parameters, allowable distance-bounds, and allowable sonority hurdles. The claims that harmony patterns in (3) are impossible constitute a strong and falsifiable set of predictions: discovery of any language that possesses one of these patterns will require revision to the theory.

Any contribution to the study of Universal Grammar should be a delimitation of the set of unlearnable languages. In my view, the function of what is called Universal Grammar, is not really to *provide* a grammar, but rather to provide a set of constraints on what can and can't be a possible grammar. Put differently, the function of Universal Grammar in aiding the learner of harmony is not to produce a magic formula for the answer – in fact, as we may observe from cases of optionality and intra-speaker variation, the learner may *never* learn a uniquely correct answer – but rather to *prevent* the learner from ever considering a host of irrelevant answers. And I do mean prevent. On this view, the closest analog to Universal Grammar in another species would be in the "templates" of songbird learning.

W.H. Thorpe's study of trying to teach chaffinches the songs of the tree pipit and observing their failure to learn these songs (although they were sensorily perceptible and motorally executable) led him to conclude that "the chaffinch has the inborn blueprint conferring on it a tendency to learn to pay attention to certain kinds of sounds and certain types of phrase only" Thorpe (1958, p.84). Thorpe's quote, and indeed his views, emphasize that the defining feature of songbird learning is what they cannot and do not try to learn, because they simply exclude attention to certain kinds of sound patterns. This constitutes a strong view of the chaffinch's "blueprint", and Universal Grammar: the languages that it excludes are simply unlearnable; they are grammatical hypotheses that are excluded from the set of possibilities that the learner considers.

With this perspective in mind, what is potentially more realistic than discovering that some newly-encountered language has a heretofore unknown pattern violating the predictions in (3) is the possibility of testing these predictions in an experimental paradigm. The Artificial Grammar Learning (AGL) paradigm is an experimental technique whereby participants (which can be infants, adults, or animals) are taught a "miniature language" that the experimenters construct with specific properties (see Gómez and Gerken (2000) for an overview).[95] After a brief implicit exposure to words and or sentences of this artificial language, participants then engage in a task where they are asked to discriminate and/or produce utterances that would be allowable according to the pattern of this language.

Should an excluded harmony system of the type in (2) arise, either through artificial presentation in the laboratory, or through an intermediate stage of radical language contact (e.g. pidginization), we would predict that learners would reshape the system in order to be consistent with principles given by Universal Grammar. In other words, in a real-life situation of this sort or an "iterated learning" AGL experiment (Kalish et al., 2007), in which one generation of subjects provides input to the next, we would expect a scenario similar to the regularization of non-UG-compatible grammars attested in the Nicaraguan Sign Language scenario (Senghas, 1995). Alternatively, we might observe an all-out failure to learn an impossible grammar (e.g. with non-contrastive-only relativization or only low-sonority hurdles), an outcome parallel to that observed by Smith and Tsimpli (1995), who attempted to teach Christopher, a linguistic idiot savant (i.e a person with lower-than-average general intelligence, but phenomenal ability for learning languages) a non-UG-compatible language ("Epun", in which the sentential-focus-marker always appeared on the third word), and observed his failure to generalize.

Pycha et al. (2003), Koo and Cole (2006), and Bonatti et al. (2005) have carried out revealing experiments within the AGL paradigm to test some questions of possible and impossible patterns of subsegmental dependencies such as vowel harmony and dissimilation. Pycha et al. (2003) studied whether experimental participants could learn harmony (identical-feature copying), disharmony (opposite-feature copying), or a random pattern of suffix allomorphy with equal ease, and found that harmony and disharmony were equally learnable by the participants, none of whom had such patterns in their native languages, but that random patterns were not. Koo and Cole (2006) studied whether experimental participants could learn consonantal harmony, consonantal disharmony, vowel harmony, and vowel disharmony with equal ease, and found that vowel disharmony was hardest for participants to learn. Bonatti et al. (2005) tested whether French-speaking participants could learn triconsonantal 'frames' of the Semitic type with equal ease as learning trivocalic frames, and found that consonantal skeleta were easier to extract from continuous data streams. In summary, experiments of this sort have already established the potential for revealing generalizations about participants' ability to learn typologically 'natural' and 'unnatural' patterns, and demonstrated that the AGL methodology is sensitive enough to allow experimental comparison of harmony vs. disharmony and consonantal vs. vocalic dependencies.

The AGL methodology could be used to test any of the languages in (3). The prediction is that human participants' performance in accurately learning the patterns in (3) would be severely degraded in comparison with their ability to learn the patterns enumerated in (4).

(4) *Possible Harmony Systems*:

a. Copying from the closest vowel to the value-seeker
b. Copying only from vowels contrastive for the harmonic feature; exclusion of others
c. Copying only from vowels marked for the harmonic feature; exclusion of others
d. No second-chance parasitic harmony: presence of defective intervention when there is a conditional requirement on sources
e. Search that can look past one intervening syllable, but not two intervening syllables
f. Noncontrastive & transparent *i* in the same language as noncontrastive but nontransparent *a*

All of the patterns in (4), in contrast to those in (3), are allowed options within the vowel harmony space I have proposed. Vowel harmony is a search-and-copy process that is greedy and economical, vowel harmony can relativize what matters based on salient inventory properties such as contrastiveness and markedness, vowel harmony can count to 1 but not to 2, and vowel harmony can be blocked by high-sonority elements. We have repeatedly considered the possibility that these properties may be the inevitable result of a minimalist computational architecture.

An exciting research direction would be for AGL experiments testing impossible harmony patterns to be carried out in tandem with brain-imaging techniques. As discussed by Moro (2008), even if subjects manage to learn a non-UG-compliant "impossible" language, we want to know if they are doing so using the same neural mechanisms used to process ordinary human langauge, as opposed to employing clever non-linguistic memorization or auditory encoding as an alternative means for dealing with input that the system of linguistic computation simply cannot handle. Moro (2008) contains a very illuminating discussion of how the convergence of formally explicit theoretical models and experimental methodology can allow the possibility of testing "impossible languages", by observing the neural circuitry that humans use in order to deal with patterns that are not permitted by the highly specialized computations of human language. Even more revealing as a confirmation of the hypothesis of Cross-modular Structural Parallelism would be a set of findings that violations of possible harmony systems in (3) generated brain responses identical to those generated by violations of syntactic locality.

Conclusive brain-imaging results from AGL studies of harmony may be difficult to obtain, due to the extreme care needed in designing AGL experiments, and in particular controlling for potential bias from the speakers' first language[96] A second concern in the design of AGL experiments relates to the "ecological validity" of the AGL paradigm for teaching vowel harmony: how does one train subjects to learn that something is an affix asymmetrically dependent on the root, or indeed to perform morphological segmentation at all? How does one encourage subjects to generalize vowel harmony patterns to new roots and treat it as a productive process, rather than a static fact about the words they have heard? Peperkamp et al. (2007) have developed creative methodologies of implicit pattern learning through a distractor or ecologically valid task, and such approaches will continue to provide solutions towards these concerns.

In summary, our lengthy discourse has traversed the locality conditions on harmony patterns where the harmonic feature is [± high], [± low], [± back], [± round], [± ATR], or [± nasal], from over thirty different languages, spoken in the Americas, Africa, Europe, Australia, India, and Asia. The conclusions that emerge overshadow which feature the harmony is for, or where the languages are spoken. Facing locality in linguistic structure, humans do not compute distance as the crow flies, nor do they perform a global optimization. The search for the closest donor vowel may exclude a swath of irrelevant interveners, or may give up with the first defective vowel it finds, even though a better one lies just beyond. These generalizations are reaffirmed repeatedly across a diverse set of alphabets and sequences.

Typological research on understudied languages must continue, and as our knowledge of the world's languages increases, we may find patterns of vowel harmony that call for a revision of the theory. Importantly, even if such languages are never found in natural linguistic communities, the predictions of a formally explicit principles-and-parameters model of vowel harmony are clear enough to be tested experimentally. These predictions will ultimately bear on the nature of what computations of locality are biologically possible.

Notes

[88]Myopia appears to be the right characterization of vowel harmony, independently of defective intervention: Wilson (2003) shows that a globally-optimizing model of harmony wreaks predictive havoc.

[89]See Coleman and Local (1991) for a critique of the Line Crossing Constraint on formal grounds.

[90]A reviewer brings up the interesting possibility that the *Agree* algorithm is limited to PF (the phonological branch of the grammatical computation), citing the proposal of Bobaljik (2008) that verbal agreement is computed post-syntactically, within PF. While much further work is needed to verify Bobaljik's hypothesis, this possibility has the intriguing potential to localize all uses of the Agree algorithm (whether computing over segments or phrases) to a single module.

[91]Kemp et al. (2007) propose a computational model for overhypothesis-learning of items with feature-values.

[92]Carstens (2000) raises a number of important formal questions that arise in the implementation of Concord through the mechanism of Agree.

[93]The closest case of dissimilation for ϕ-features that I have found is the phenomenon by which Hebrew numerals take the opposite gender from their adjacent head noun (Halle, 1994); this is a concord phenomenon, and no long-distance cases exist.

[94]Optimality Theory obliterates the distinction between Principles and Parameters, relegating the universal and inviolable properties of human language to constraints on GEN, the engine responsible for producing candidate output forms.

[95]Yip (2006) has raised the possibility that artificial grammar learning paradigms should be extended to animal experiments, in order to determine whether some of the defining properties of human language phonology are indeed unique to humans. In my proposal, relativized and myopic locality is a shared mechanism between the modules of phonology and syntax, and a potential line of inquiry that could be explored is whether the properties of relativized and myopic locality are unique to the *human* faculty of language.

[96]For this reason, AGL experiments with neonates (e.g. Gervain and Mehler (2008)) can provide potentially some of the most exciting arguments for innateness.

Bibliography

Abels, K. (2003). *The syntax of adposition stranding*. PhD thesis, University of Connecticut.

Abondolo, D. (1998). Khanty. In Daniel Abondolo, editor, *The Uralic Languages*, pages 358–386. London and New York: Routledge.

Adetugbo, A. (1967). *The Yoruba Language in Western Nigeria: Its Major Dialect Areas*. PhD thesis, Columbia University.

Adger, D. (2003). *Core Syntax: A Minimalist Approach*. Oxford University Press.

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

Akinlabi, A. (2008). *A Phonological Grammar of Yoruba*. (in preparation).

Aksu-Koç, A. and Slobin, D. (1985). The acquisition of Turkish. In D. I. Slobin, editor, *Cross-linguistic studies of Language Acquisition, Vol. 1*. New Jersey: Lawrence Earlbaum Associates.

Andersen, H. (1973). Abductive and Deductive Change. *Language*, 49.4:765–793.

Anderson, G. D. (2004). The Languages of Cetnral Siberia: Introduction and Overview. In Vajda, E., editor, *Proceedings of the 13th International Congress of Phonetic Sciences*, pages 1–119. Amsterdam: John Benjamins.

Anderson, J. M. (1992). *Linguistic Representation: Structural Analogy and Stratification*. Berlin: Mouton de Gruyter.

Anderson, L. (1980). Using asymmetrical and gradient data in the study of vowel harmony. In Vago, R., editor, *Issues in Vowel Harmony*, pages 271–340. Amsterdam: John Benjamins.

Anderson, S. R. (1976). Nasal Consonants and the Internal Structure of Segments. *Language*, 52.2:326–344.

Anderson, S. R. (1985). *Phonology in the Twentieth Century*. Chicago and London: The University of Chicago Press.

Andrews, A. (1982). The Representation of Case in Modern Icelandic. In Joan Bresnan, editor, *The Mental Representation of Grammatical Relations*, pages 427–503. Cambridge, MA: MIT Press.

Anttila, A. (2003). Finnish Assibilation. In M. Kadowaki and S. Kawahara, editor, *The Proceedings of NELS 33*. Amherst, MA: GLSA.

Ao, B. (1991). Kikongo nasal harmony and context-sensitive underspecification. *Linguistic Inquiry*, 22:193–196.

Aoki, H. (1968). Toward a Typology of Vowel Harmony. *International Journal of American Linguistics*, 34.2:142–145.

Arad, M. (2005). *Roots and Patterns: Hebrew Morphosyntax*. Dordrecht: Springer.

Archangeli, D. (1988). Aspects of underspecification theory. *Phonology*, 5:183–207.

Archangeli, D. and Pulleyblank, D. (1994). *Grounded Phonology*. Cambridge, MA: MIT Press.

Archangeli, D. and Pulleyblank, D. (2002). Kinande vowel harmony: domains, grounded conditions, and one-sided alignment. *Phonology*, 19:139–188.

Ard, J. (1984). Vowel Harmony in Manchu: A Critical Overview. *Journal of Linguistics*, 20:57–80.

Avrorin, V. and Boldyrev, B. (2001). *Grammatika orochskogo jazyka*. Novosibirsk.

Bakovic, E. (2000). *Harmony, Dominance, and Control*. PhD thesis, Rutgers University.

Bakovic, E. (2003). Vowel harmony and stem identity. In *San Diego Linguistic Papers 1*, pages 1–42. San Diego: Department of Linguistics UCSD.

Bamgboṣe, A. (1967). Vowel Harmony in Yoruba. *Journal of African Languages*, 6.3:268–273.

Barnes, J. (2006). *Strength and weakness at the interface: positional neutralization in phonetics and phonology*. Berlin & New York: Mouton de Gruyter.

Battistella, E. (1982). More on Hungarian Vowel Harmony. *Linguistic Analysis*, 9.2:95–108.

Beckman, J. (1997). Positional faithfulness, positional neutralisation and Shona vowel harmony. *Phonology*, 14:1–46.

Beckman, J. (1998). *Positional Faithfulness*. PhD thesis, UMass, Amherst.

Beddor, P. and Yavuz, H. K. (1995). The Relation between Vowel-to-Vowel Coarticulation and Vowel Harmony in Turkish. In Elenius, K. and Branderud, P., editors, *Proceedings of the 13th International Congress of Phonetic Sciences, vol. 2*, pages 44–51. KTH and Stockholm University, Sweden.

Bennett, P. R. (1967). Dahl's Law and Thagicũ. *African Language Studies*, pages 127–159.

Benus, S. and Gafos, A. (2007). Articulatory characteristics of Hungarian 'transparent' vowels. *Journal of Phonetics*, 35:271–300.

Bermúdez-Otero, R. and Honeybone, P. (2006). Phonology and syntax: a shifting relationship. *Lingua*, 116:543–561.

Bessell, N. (1998). Local and non-local consonant-vowel interaction in Interior Salish. *Phonology*, 15:1–40.

Bickmore, L. (1995). Refining and formalizing the Tahitian stress placement algorithm. *Oceanic Linguistics*, 34:2.

Bloomfield, L. (1962). *The Menomini Language*. New Haven: Yale University Press.

Blumenfeld, L. and Toivonen, I. (2009). A featural paradox in Votic harmony. Paper presented at the Linguistic Society of America meeting.

Bobaljik, J. (2008). Where's Phi? Agreement as a Post-Syntactic Operation. In Harbour, D., Adger, D., and Béjar, S., editors, *Phi-Theory: Phi featuires across interfaces and modules*, pages 295–328. Oxford University Press.

Boeckx, C. and Grohmann, K. (2007). Putting Phases in Perspective. *Syntax*, 10.2:204–222.

Bonatti, L., Peña, M., Nespor, M., and Mehler, J. (2005). Linguistic constraints on statistical computations. *Psychological Science*, 16.6:451–459.

Borer, H. (1984). *Parametric Syntax: Case Studies in Semitic and Romance Languages*. Dordrecht: Foris.

Borgstrom, C. (1937). *The dialect of Barra in the Outer Hebrides*. Oslo: Norsk Tidsskrift for Sprogvidenskap.

Bromberger, S. and Halle, M. (1989). Why Phonology is Different. *Linguistic Inquiry*, 20:51–70.

Büring, D. (2005). *Binding Theory*. Cambridge University Press.

Burtt, E. A. (1932). *The Metaphysical Foundations of Modern Physical Science*. London: Routledge & Kegan Paul.

Butcher, A. and Weiher, E. (1976). An electropalatographic investigation of coarticulation in VCV sequences. *Journal of Phonetics*, 4:59–74.

Calabrese, A. (1988). *Towards A Theory of Phonological Alphabets*. PhD thesis, MIT.

Calabrese, A. (1995). A constraint-based theory of phonological markedness and simplification procedures. *Linguistic Inquiry*, 26:373–463.

Calabrese, A. (2005). *Markedness and Economy in a Derivational Model of Phonology*. Mouton de Gruyter.

Campbell, L. (1980). The Psychological and Sociological Reality of Finnish Vowel Harmony. In Vago, R., editor, *Issues in Vowel Harmony*, pages 245–270. Amsterdam: John Benjamins.

Campbell, L. (1986). Testing Phonology in the Field. In John Ohala and Jeri Jaeger, editor, *Experimental Phonology*, pages 163–173. Academic Press.

Carstens, V. (2000). Concord in Minimalist Theory. *Linguistic Inquiry*, 31.2:319–355.

Casali, R. (1995). Labial Opacity and Roundness Harmony in Nawuri. *Natural Language and Linguistic Theory*, 13:649–663.

Casali, R. (1998). *Resolving Hiatus*. Garland, New York.

Cassimjee, F. and Kisseberth, C. (2001). Zulu tonology and its relation to other Nguni languages. In Kaji, S., editor, *Cross-linguistic studies of Tonal Phenomena*, pages 327–359. Tokyo University of Foreign Studies.

Charette, M. (1991). *Conditions on Phonological Government*. Cambridge University Press.

Cho, Y.-M. Y. (1990). *Parameters of Consonantal Assimilation*. PhD thesis, Stanford University.

Chomsky, N. (1981). *Lectures on Government and Binding*. Dordrecht: Foris.

Chomsky, N. (1986). *Barriers*. Cambridge, MA: MIT Press.

Chomsky, N. (1995). *The Minimalist Program*. Cambridge, MA: MIT Press.

Chomsky, N. (2000). Minimalist Inquiries: The Framework. In Martin, R., Michaels, D., and Uriagereka, J., editors, *Step by Step: Essays on Minimalist Syntax in Honor of Howard Lasnik*, pages 89–155. Cambridge, MA: MIT Press.

Chomsky, N. (2001). Derivation by Phase. In Michael Kenstowicz, editor, *Ken Hale: A Life in Language*, pages 1–52. Cambridge, MA: MIT Press.

Chomsky, N. and Halle, M. (1968). *The Sound Pattern of English*. Harper and Row.

Christdas, P. (1988). *The phonology and morphology of Tamil*. PhD thesis, Cornell University.

Clements, G. (1985). The geometry of phonological features. *Phonology Yearbook*, 2:225–252.

Clements, G. (2003). Feature economy in sound systems. *Phonology*, 20.3:287–333.

Clements, G. and Keyser, S. (1983). *CV Phonology: A generative theory of the syllable*. Cambridge, MA: MIT Press.

Clements, G. N. and Sezer, E. (1982). Vowel and consonant disharmony in Turkish. In van der Hulst, H. and Smith, N., editors, *The Structure of Phonological Representations, Part II*, pages 213–255. Dordrecht: Foris.

Coats, H. (1970). Rule Enviroment Features in Phonology. *Papers in Linguistics*, 2:110–140.

Cohn, A. (1990). *Phonetic and Phonological Rules of Nasalization*. PhD thesis, UCLA.

Cole, J. and Trigo, L. (1988). Parasitic Harmony. In van der Hulst, H. and Smith, N., editors, *Features, segmental structure, and harmony processes*, pages 19–38. Dordrecht:Foris.

Cole, J. S. (1987). *Planar Phonology and Morphology*. PhD thesis, MIT.

Coleman, J. and Local, J. (1991). The "No crossing constraint" in autosegmental phonology. *Linguistics and Philosophy*, 14:295–338.

Collins, C. (2001). Economy Conditions in Syntax. In *The Handbook of Syntactic Theory*, pages 45–61. Blackwell.

Comrie, B. (1981). *The Languages of the Soviet Union*. Cambridge University Press.

Crosswhite, K. (1998). Segmental vs. Prosodic Correspondence in Chamorro. *Phonology*, 15:3:281–316.

Csató, E. and Nathan, D. (2002). *Spoken Karaim*. CD-ROM. Uppsala University and Tokyo University of Foreign Studies.

Cutler, A., Sebastián-Gallés, N., Soler-Vilageliu, O., and van Ooijen, B. (2000). Constraints of Vowels and Consonants on Lexical Selection: Cross-linguistic Comparisons. *Memory and Cognition*, 28.5:746–755.

Davy, J. and Nurse, D. (1982). Synchronic Versions of Dahl's Law: The Multiple Applications of a Phonological Dissimilation Rule. *Journal of African Languages and Linguistics*, 4:157–195.

de Lacy, P. (2006). *Markedness: Reduction and preservation in phonology*. Cambridge University Press.

Dell, F. and Elmedlaoui, M. (1985). Syllabic Consonants and Syllabification in Imdlawn Tashlhiyt Berber. *Journal of African Languages and Linguistics*, 7:105–130.

Dresher, B. E. (2003). Contrast and asymmetry in inventories. In Anna Maria di Sciullo, editor, *Asymmetry in Grammar, Volume 2*, pages 239–257. John Benjamins.

Dresher, B. E. (2005). On the acquistion of phonological contrasts. In van Kampen, J. and Bauuw, S., editors, *Proceedings of GALA 2003, Vol. 1*, pages 27–46. LOT: Utrecht.

Farkas, D. and Beddor, P. (1987). Privative and Equipollent Backness in Hungarian. In Bosch, A., Need, B., and Schiller, E., editors, *Papers from the Parasession on Autosegmental and Metrical Phonology*, pages 91–105. Chicago: Chicago Linguistic Society.

Filchenko, A. Y. (1979). *A Grammar of Eastern Khanty*. PhD thesis, MIT.

Finley, S. (2007). Exceptions in Vowel Harmony are Local. Ms., Johns Hopkins University.

Finley, S. (2009). Morphemic harmony as featural correspondence. *Lingua*, doi:10/1016:to appear.

Fitzpatrick, J. (2002). On Minimalist Approaches to the Locality of Movement. *Linguistic Inquiry*, 33.3:443–463.

Fortune, G. (1981). *Shona Grammatical Constructions*. Dept. of African Languages, University of Rhodesia.

Fowler, C. (1981). Production and Perception of coarticulation among stressed and unstressed vowels. *Journal of Speech and Hearing Research*, 46:127–139.

Frampton, J. (1991). Relativized Minimality, A Review. *The Linguistic Review*, 8:1–46.

Frisch, S., Pierrehumbert, J., and Broe, M. (2004). Similarity Avoidance and the OCP. *Natural Language and Linguistic Theory*, 22:179–228.

Gafos, A. (1996). *The Articulatory Basis of Locality in Phonology*. PhD thesis, Johns Hopkins University.

Gervain, J. and Mehler, J. (2008). The neonate brain's ability to detect linguistic structure. Ms., SISSA, Trieste.

Gick, B., Pulleyblank, D., Campbell, F., and Mutaka, N. (2006). Low vowels and transparency in Kinande vowel harmony. *Phonology*, 23:1–20.

Gildersleeve, B. and Lodge, G. (1895). *Gildersleeve's Latin Grammar*. London: Macmillan.

Goddard, I. (1987). Leonard Bloomfield's descriptive and comparative studies of Algonquian. *Historiographia Linguistica*, 14:179–217.

Goldsmith, J. (1976). An overview of autosegmental phonology. *Linguistic Analysis*, 13:72–105.

Goldsmith, J. (1985). Vowel Harmony Khalkha Mongolian, Yaka, Finnish, and Hungarian. *Phonology*, 2:253–275.

Gómez, R. and Gerken, L. (2000). Infant artificial language learning and language acquisition. *Trends in Cognitive Sciences*, 4.5:178–186.

Goodman, N. (1955). *Fact, fiction, and forecast*. Harvard University Press.

Gordon, M. (1999). The "neutral" vowels of Finnish: How neutral are they? *Linguistica Uralica*, 35:17–21.

Gouskova, M. (2004). Relational hierarchies in Optimality Theory: The Case of Syllable Contact. *Phonology*, 21(2):201–250.

Greenberg, J. (1950). The Patterning of Root Morphemes in Semitic. *Word*, 5:162–181.

Greenberg, J. (1963). Some universals of grammar with particular reference to the meaning of elements. In Joseph Greenberg, editor, *Universals of Language*, pages 73–113. Cambridge, MA: MIT Press.

Hall, B. L., Hall, R., Pam, M. D., Myers, A., Antell, S. A., and Cherono, G. K. (1974). Lenition degrades information: consonant allophony in Ibibio. *Afrika und Übersee*, 57/4:241–267.

Hall, T. and Hamann, S. (2006). Towards a typology of stop assibilation. *Linguistics*, 44.6:1195–1236.

Halle, M. (1994). The morphology of numeral phrases. In Avrutin, S., Franks, S., and Progovac, L., editors, *Annual Workshop on Formal Approaches to Slavic Linguistics: The MIT Meeting, 1993*, pages 178–215. Ann Arbor: Michigan Slavic Publications.

Halle, M. (2005). Palatalization/Velar Softening: What It Is and What It Tells Us about the Nature of Language. *Linguistic Inquiry*, 36:23–41.

Halle, M., Vaux, B., and Wolfe, A. (2000). On Feature Spreading and the Representation of Place of Articulation. *Linguistic Inquiry*, 31.3:387–444.

Halle, M. and Vergnaud, J.-R. (1987). *An Essay on Stress*. Cambridge, MA: MIT Press.

Hamp, E. (1976). Palatalization and harmony in Gagauz and Karaite. In Walther Heissig and John R. Krueger and Felix J. Oinas and Edmond Schütz, editor, *Tractata altaica: Denis Sinor, sexagenario optime de rebus altaicis merito dedicata.*, pages 211–213. Wiesbaden:Harrassowitz.

Hansson, G. (2001). *Theoretical and Typological Issues in Consonant Harmony*. PhD thesis, UC Berkeley.

Hansson, G. (2007a). Blocking Effects in Agreement by Correspondence. *Linguistic Inquiry*, 38.2:395–409.

Hansson, G. (2007b). On the evolution of consonant harmony: the case of secondary articulation agreement. *Phonology*, 24:77–120.

Harms, R. (1987). What Helmholtz Knew About Neutral Vowels. In *A festschrift for Ilse Lehiste*, pages 381–399. Columbus: Ohio State University Working Papers in Linguistics 35.

Harris, J. and Urua, E.-A. (2001). Lenition degrades information: consonant allophony in Ibibio. *Speech, Hearing and Language: Work in Progress*, 2:23–68.

Harris, J. W. (1983). *Syllable structure and stress in Spanish: A nonlinear analysis*. Cambridge, MA: MIT Press.

Harrison, K. D. and Kaun, A. (2001). Patterns, pervasive patterns and featurespecification. In Hall, T. A., editor, *Distinctive Feature Theory*, pages 211–236. Berlin: Mouton de Gruyter.

Harrison, K. D., Thomforde, E., and O'Keefe, M. (2004). The vowel harmony calculator. Available at http://www.swarthmore.edu/SocSci/harmony/public_html/.

Hayes, B. and Londe, Z. C. (2006). Stochastic phonological knowledge: the case of Hungarian vowel harmony. *Phonology*, 23:59–104.

Hebert, R. and Poppe, N. (1963). *Kirghiz Manual*. Indiana University Publications in Uralic and Altaic Series.

Hiraiwa, K. (2001). Multiple Agree and the Defective Intervention Constraint in Japanese. In O. Matushansky, editor, *The Proceedings of the MIT-Harvard Joint Conference (HUMIT 2000), MITWPL 40*, pages 67–80. Cambridge, MA: MITWPL.

Hjemslev, L. (1948). Le verbe et la phrase nominale. In *Mélanges de philologie, de littérature et d'histoire ancienne offerts à J. Marouzeau*, pages 235–281. Reprinted in Essais Linguistiques, 165191 (Travaux du circle linguistique de Copenhague 12).

Hockett, C. F. (1960). The origin of speech. *Scientific American*, 203.3:88–96.

Hockett, C. F. (1981). The phonological history of Menominee. *Anthropological Linguistics*, 23.2:51–87.

Holmberg, A. and Hróarsdóttir, T. (2004). Agreement and movement in Icelandic raising constructions. *Lingua*, 114:651–673.

Howard, I. (1972). *A Directional Theory of Rule Application in Phonology*. PhD thesis, MIT.

Hualde, J. (1991). The strict cycle condition and noncyclic rules. *Linguistic Inquiry*, 20:675–680.

Hyman, L. (1975). *Phonology: Theory and Analysis*. New York: Holt, Rinehart and Winston.

Hyman, L. (1998). Positional prominence and the 'prosodic trough' in Yaka. *Phonology*, 15:41–75.

Hyman, L. (1999). The historical interpretation of vowel harmony in Bantu. In Jean-Marie Hombert and Larry M. Hyman, editor, *Bantu Historical Linguistics: Theorical and empirical perspectives*, pages 235–295. Stanford: CSLI.

Ihiunu, P. and Kenstowicz, M. (1994). Two Notes on Igbo Vowels. Ms., MIT.

Ito, J. and Mester, A. (2003). *Japanese Morphophonemics: Markedness and Word Structure*. Cambridge, MA: MIT Press.

Jakobson, R. (1941). *Kindersprache, Aphasie und allgemeine Lautgesetze*. Uppsala: Almqvist & Wiksell.

Jensen, J. T. (1974). A constraint on variables in phonology. *Language*, 50:675–686.

Jespersen, O. (1904). *Lehrbuch der Phonetik*. Leipzig and Berlin: B.G. Teubner.

Johnson, C. D. (1980). Regular Disharmony in Kirghiz. In Vago, R., editor, *Issues in Vowel Harmony*, pages 89–99. Amsterdam: John Benjamins.

Jun, J. (2004). Place assimilation. In Hayes, B., Kirchner, R., and Steriade, D., editors, *Phonetically Based Phonology*, pages 58–86. Cambridge University Press.

Ka, O. (1994). *Wolof Phonology and Morphology*. Lanham, MD: University Press of America.

Kabak, B. and Vogel, I. (2009). Exceptions to stress and harmony in Turkish: cophonologies or prespecification? In H. Simon and H. Wiese, editor, *Expecting the unexpected: Exceptions in grammar*. Berlin: Mouton de Gruyter.

Kalish, M. L., Griffiths, T. L., and Lewandowsky, S. (2007). Iterated learning: Intergenerational knowledge transmission reveals inductive biases. *Psychonomic Bulletin & Review*, 14.2:288–294.

Kaun, A. (1995). *The Typology of Rounding Harmony: An Optimality Theoretic Approach*. PhD thesis, UCLA.

Kawahara, S. (2003). Locality in Echo Epenthesis. In *The Proceedings of NELS 34*, pages 295–307. Amherst, MA: GLSA.

Kaye, J. (1971). Nasal harmony in Desano. *Linguistic Inquiry*, 2:37–58.

Kaye, J. (1989). *Phonology: A Cognitive View*. Lawrence Erlbaum, Hillsdale, NJ.

Kaye, J., Lowenstamm, J., and Vergnaud, J.-R. (1985). The internal structure of phonological representations. *Phonology*, 2:305–328.

Keating, P. (1996). The Phonology-Phonetics Interface. In Ursula Kleinhenz, editor, *Interfaces in Phonology*, pages 262–278. Berlin: Akademie Verlag.

Kelepir, M. (2000). To be or not to be faithful. In Aslı Göksel and Celia Kerslake, editor, *Studies on Turkish and Turkic languages: Proceedings of the Ninth International Conference on Turkish Linguistics*. Wiesbaden: Harrassowitz Verlag.

Kemp, C., Perfors, A., and Tenenbaum, J. B. (2007). Learning overhypotheses with hierarchical Bayesian models. *Developmental Science*, 10.3:307–321.

Kenstowicz, M. (1994). *Phonology in Generative Grammar*. Blackwell.

Kenstowicz, M. (1997). Quality-Sensitive Stress. *Rivista di Linguistica*, 9.1:157–187.

Kenstowicz, M. (2008). Two Notes on Kinande Vowel Harmony. Paper presented the Linguistic Society of America meeting, Chicago.

Kim, Y. (2005). Finnish Neutral Vowels: subcontrastive harmony or V-to-V coarticulation? Presented at the LSA Meeting, Oakland.

Kim, Y. (2008). *Topics in the Phonology and Morphology of San Francisco del Mar Huave*. PhD thesis, UC Berkeley.

Kiparsky, P. (1973). Phonological Representations. In O. Fujimura, editor, *Three Dimensions of Linguistic Theory*. Tokyo: TEC.

Kiparsky, P. (1981). Vowel Harmony. Ms., MIT.

Kiparsky, P. (1993). Blocking in nonderived environments. In Sharon Hargus and Ellen Kaisse, editor, *Studies in Lexical Phonology*, pages 277–313. New York: Academic Press.

Kiparsky, P. and Pajusalu, K. (2003). Towards a Typology of Disharmony. *The Linguistic Review*, 20:217–241.

Koo, H. and Cole, J. (2006). On Learnability and Naturalness as Constraints on Phonological Grammar. In A. Botinis, editor, *Proceedings of ISCA Tutorial and Research Workshop on Experimental Linguistics*, pages 174–177. University of Athens.

Korn, D. (1969). Types of Labial Vowel Harmony in the Turkic Languages. *Anthropological Linguistics*, 11.3:98–106.

Kowalski, T. (1929). *Karaimische texte im dialekt von Troki*. Prace Komisji Orjentalistycznej Polskiej Akademji Umiejętności nr. 11. Cracow: Nakładem Polskiej Akademji Umiejętności.

Krämer, M. (2001). Yucatec Maya vowel alternations – Harmony as Syntagmatic Identity. *Zeitschrift für Sprachwissenschaft*, 20:175–217.

Krämer, M. (2003). *Vowel Harmony and Correspondence Theory*. Mouton de Gruyter.

Kroch, A. (1989). Reflexes of grammar in patterns of language change. *Language Variation and Change*, 1:199–244.

Krueger, J. R. (1962). *Yakut Manual*. Indiana University Publications in Uralic and Altaic Series, Volume 21.

Kuroda, S.-Y. (1967). *Yawelmani Phonology*. Cambridge, MA: MIT Press.

Ladefoged, P. (1968). *A Phonetic Study of West African Languages*. Cambridge University Press.

Langacker, R. W. (1987). *Foundations of Cognitive Grammar*. Stanford University Press.

Leben, W. R. (2006). Rethinking Autosegmental Phonology. In Mugane, J., Hutchison, J., and Worman, D., editors, *Selected Proceedings of the 35th Annual Conference on African Linguistics*, pages 1–9. Somerville, MA: Cascadilla Press.

Leitch, M. (1996). *Vowel harmonies of the Congo Basin: An Optimality Theory analysis of variation in the Bantu zone C*. PhD thesis, University of British Columbia.

Leiwo, M., Kulju, P., and Aoyama, K. (2006). The Acquisition of Finnish Vowel Harmony. *SKY Journal of Linguistics*, 19:149–161.

Levergood, B. (1984). Rule Governed Vowel Harmony and the Strict Cycle. In *The Proceedings of NELS 14*, pages 275–293. Amherst, MA: GLSA.

Levi, S. V. (2004). *The representation of underlying glides: a cross-linguistic study*. PhD thesis, University of Washington.

Li, B. (1996). *Tungusic Vowel Harmony: Description and Analysis*. Holland Institute of Generative Linguistics.

Lieberman, P. (2000). *Human Language and our Reptilian Brain*. Cambridge, MA: Harvard University Press.

Lightfoot, D. and Westergaard, M. (2007). Language Acquisition and Language Change: Inter-relationships. *Language and Linguistics Compass*, 1.5:396–415.

Lindblad, V. (1990). Neutralization in Uyghur. Master's thesis, University of Washington.

Linker, W. (1982). *Articulatory and acoustic correlates of labial activity in vowels: A cross-linguistic study*. PhD thesis, UCLA.

Lodge, K. (1995). Kalenjin phonology and morphology: A further exemplification of underspecification and non-destructive phonology. *Lingua*, 96:29–43.

Longobardi, G. (2005). A minimalist program for parametric linguistics? In Broekhuis, H., Corver, N., Huybregts, R., Kleinheinz, U., and Koster, J., editors, *Organizing Grammar*, pages 407–414. Berlin: Mouton de Gruyter.

Łubowicz, A. (2002). Derived environment effects in Optimality Theory. *Lingua*, 112:243–280.

Magdics, K. (1969). *Studies in the Acoustic Characteristics of Hungarian Speech Sounds*. Indiana University Publications in Uralic and Altaic Series.

Mahanta, S. (2007). *Directionality and Locality in Vowel Harmony with special reference to Assamese*. PhD thesis, Utrecht University.

Mailhot, F. and Reiss, C. (2007). Computing Long-Distance Dependencies in Vowel Harmony. *Biolinguistics*, 1:28–48.

Marcus, G. (2008). *Kluge*. Houghton Mifflin Company.

McCarthy, J. (1984). Theoretical consequences of Montañes vowel harmony. *Linguistic Inquiry*, 15:291–318.

McCarthy, J. (1985). *Formal Problems in Semitic Phonology and Morphology*. New York: Garland.

McCarthy, J. (1988). Feature Geometry and Dependency: A Review. *Phonetica*, 43:84–108.

McCarthy, J. (2007). *Hidden Generalizations: Phonological Opacity in Optimality Theory*. London: Equinox.

McCarthy, J. and Taub, A. (1992). Review of Paradis & Prunet (eds.), The Special Status of Coronals. *Phonology*, 9:363–370.

Milligan, M. (2000). A New Look at Menominee Vowel Harmony. In John D. Nichols, editor, *Papers of the 31st Annual Algonquian Conference*, pages 237–254. Winnipeg: University of Manitoba.

Miner, K. L. (1979). Theoretical implications of the great Menominee vowel shift. *Kansas Working Papers in Linguistics*, 4.1:7–25.

Mintz, T. and Walker, R. (2006). Infants' Sensitivity to Vowel Harmony and its Role in Word Segmentation. Paper presented the Linguistic Society of America meeting, Albuquerque.

Mithun, M. (1998). The regression of sibilant harmony through the life of Barbareño Chumash. In Hill, J. H., Mistry, P., and Campbell, L., editors, *The life of language: Papers in linguistics in honor of William Bright*, pages 221–242. Berlin: Mouton de Gruyter.

Mohanan, K. P. (1991). On the Bases of Radical Underspecification. *Natural Language and Linguistic Theory*, 9:285–325.

Moro, A. (2008). *The Boundaries of Babel*. Cambridge, MA: MIT Press.

Müller, G. (2008). On Deriving CED Effects from the PIC. Ms., Universität Leipzig.

Müller, M. (1864/1994). *Lectures on the science of language, Series II.* Reprinted as Volumes 3 & 4 of *British linguistics in the nineteenth century*, London: Routledge/Thoemmes Press.

Nash, D. (1980). *Topics in Warlpiri Grammar*. PhD thesis, MIT.

Nespor, M., Peña, M., and Mehler, J. (2003). On the Different Roles of Vowels and Consonants in Speech Processing and Language Acquisition. *Lingue e Linguaggio*, 2:203–231.

Nevins, A. (2004). *Conditions on (Dis)Harmony*. PhD thesis, MIT.

Nevins, A. (2007). The representation of third person and its consequences for person-case effects. *Natural Language and Linguistic Theory*, 25:273–313.

Nevins, A. and Chitoran, I. (2008). Phonological representations and the variable patterning of glides. *Lingua*, 118.12:1979–1997.

Nevins, A. and Vaux, B. (2003). Consonant Harmony in Karaim. In Aniko Csirmaz and Youngjoo Lee and MaryAnn Walter, editor, *MITWPL 46: The Workshop on Altaic in Formal Linguistics*, pages 175–194. Cambridge, MA: MITWPL.

NiChiosain, M. and Padgett, J. (2001). Markedness, segment realization, and locality in spreading. In Linda Lombardi, editor, *Segmental phonology in Optimality Theory*, pages 118–156. Cambridge University Press.

Noske, M. (2001). [ATR] Harmony in Turkana: A Case of Faith Suffix >> Faith Root. *Natural Language and Linguistic Theory*, 18:771–812.

Odden, D. (1980). The irrelevancy of the relevancy condition: evidence for the feature specification constraint. *Linguistic Analysis*, 6:261–304.

Odden, D. (1991). Vowel Geometry. *Phonology*, 8:261–289.

Odden, D. (1994). Adjacency Parameters in Phonology. *Language*, 10.2:289–330.

Odden, D. (2003). *Introducing Phonology*. Cambridge University Press.

Ohala, J. (1981). The Listener as a Source of Sound Change. In *Chicago Linguistic Society 17: Papers from the Parasession on Language and Behavior*, pages 178–203. .

Ohala, J. (1994). Towards a universal, phonetically-based theory of vowel harmony. *3rd International Conference on Spoken Language Processing (ICSLP)*, Yokohama:491–494.

Öhman, S. (1966). Coarticulation in VCV utterances: Spectrographic measurements. *Journal of the Acoustic Society of America*, 39:151–168.

O'Keefe, M. (2003). Akan Vowel Harmony. Bachelors Thesis, Swarthmore College.

Ola Orie, O. (2001). An alignment-based account of vowel harmony in Ifẹ Yoruba. *Journal of African Languages and Linguistics*, 22:117–143.

Osborn, H. (1966). Warao I: Phonology and Morphophonemics. *International Journal of American Linguistics*, 32:108–123.

Owren, M. and Cardillo, G. (2006). The relative roles of vowels and consonants in discriminating talker identity versus word meaning. *Journal of the Acoustic Society of America*, 119.3:1727–1739.

Padgett, J. (2002). Feature Classes in Phonology. *Language*, 78.2.

Palmer, F. (1970). *Prosodic Analysis*. London: Oxford University Press.

Parker, S. (2002). *Quantifying the Sonority Hierarchy*. PhD thesis, UMass, Amherst.

Pensalfini, R. (1997). *Jingulu Grammar, Dictionary, and Texts*. PhD thesis, MIT.

Pensalfini, R. (2002). Vowel Harmony in Jingulu. *Lingua*, 112:561–586.

Peperkamp, S., Skoruppa, K., and Dupoux, E. (2007). Implicit phonological learning in artificial language learning paradigms. Paper presented the Linguistic Society of America meeting, Anaheim.

Pesetsky, D. (1989). The Earliness Principle. Paper presented GLOW; Ms., MIT.

Piggott, G. (1988). A Parametric Approach to Nasal Harmony. In van der Hulst, H. and Smith, N., editors, *Features, segmental structure, and harmony processes*, pages 131–167. Dordrecht:Foris.

Piggott, G. (1992). Variability in Feature Dependency: The Case of Nasality. *Natural Language and Linguistic Theory*, 10:33–77.

Piggott, G. (1996). Implications of Consonant Nasalization for a Theory of Harmony. *Canadian Journal of Linguistics*, 41.2:141–174.

Piggott, G. (2003). Theoretical implications of segment neutrality in nasal harmony. *Phonology*, 20:375–424.

Pöchtrager, M. (2006). *The Structure of Length*. PhD thesis, Universität Wien.

Polgárdi, K. (1998). *Vowel Harmony: An account in terms of Government and Optimality*. PhD thesis, Leiden University.

Poliquin, G. (2006). *Canadian French Vowel Harmony*. PhD thesis, Harvard University.

Poppe, N. (1960). *Vergleichende Grammatik der Altaischen Sprachen*. Wiesbaden: Harrassowitz.

Poppe, N. (1964). *Bashkir Manual*. Indiana University Publications in Uralic and Altaic Series.

Poser, W. (1982). Phonological Representation and Action-at-a-Distance. In van der Hulst, H. and Smith, N., editors, *The Structure of Phonological Representations, Part II*, pages 121–158. Dordrecht: Foris.

Poser, W. (1993). Are Strict Cycle Effects Derivable? In Sharon Hargus and Ellen Kaisse, editor, *Studies in Lexical Phonology*, pages 315–321. New York: Academic Press.

Prince, A. and Smolensky, P. (1993). Constraint Interaction in Generative Grammar. ROA-537.

Przezdziecki, M. (2005). *Vowel Harmony and Coarticulation in Three Dialects of Yoruba: Phonetics Determining Phonology*. PhD thesis, Cornell.

Pycha, A., Nowak, P., Shin, E., and Shosted, R. (2003). Phonological Rule-Learning and Its Implications for a Theory of Vowel Harmony. In Garding, G. and Tsujimura, M., editors, *WCCFL 22*, pages 423–435. Cascadilla Press.

Raimy, E. (2000). *The Phonology and Morphology of Reduplication*. Mouton de Gruyter.

Recasens, D. (1984). An electropalatographic investigation of coarticulation in VCV sequences. *Journal of the Acoustic Society of America*, 76:1624–1635.

Reinhart, T. and Reuland, E. (1993). Reflexivity. *Linguistic Inquiry*, 24.4:657–720.

Reiss, C. (2003). Deriving the Feature-Filling/ Feature-Changing Contrast: An Application to Hungarian Vowel Harmony. *Linguistic Inquiry*, 34.2:199–224.

Rice, K. (1993). A Reexamination of the Feature [Sonorant]: The Status of 'Sonorant Obstruents'. *Language*, 69.2:308–344.

Ringen, C. (1975). *Vowel Harmony: Theoretical Implications*. PhD thesis, Indiana University.

Ringen, C. and Heinämäki, O. (1999). Variation in Finnish Vowel Harmony: An OT Account. *Natural Language and Linguistic Theory*, 17:303–337.

Ringen, C. and Kontra, M. (1989). Hungarian Neutral Vowels. *Lingua*, 78:181–191.

Ringen, C. and Vago, R. (1998). Hungarian vowel harmony in Optimality Theory. *Phonology*, 15:393–416.

Rizzi, L. (1982). Violations of the Wh- Island Condition and the Subjacency Condition. In *Issues in Italian Syntax*, pages 49–76. Dordrecht: Foris.

Rizzi, L. (1990). *Relativized Minimality*. Cambridge, MA: MIT Press.

Rizzi, L. (2001). Relativized Minimality Effects. In *The Handbook of Syntactic Theory*, pages 89–110. Blackwell.

Roeper, T. (1999). Universal bilingualism. *Bilingualism: Language and Cognition*, 2:169–186.

Rose, S. and Walker, R. (2004). A typology of consonant agreement as correspondence. *Language*, 80:475–531.

Ross, H. (1984). Inner Islands. In Claudia Brugman and Monica Macaulay, editor, *Proceedings of the Tenth Berkeley Linguistics Society*, pages 258–265. Berkeley.

Ross, J. R. (1967). *Constraints on Variables in Syntax*. PhD thesis, MIT.

Sagey, E. (1986). *The Representation of Features and Relations in Nonlinear Phonology*. PhD thesis, MIT.

Sagey, E. (1987). Non-constituent spreading in Barra Gaelic. Ms., UC Irvine.

Sagey, E. (1988). On the Ill-formedness of Crossing Association Lines. *Linguistic Inquiry*, 19:1:109–118.

Schein, B. and Steriade, D. (1986). On Geminates. *Linguistic Inquiry*, 17.4:691–744.

Schlindwein, D. (1987). P-bearing units: a study of Kinande vowel harmony. In *NELS 17*, pages 551–567. .

Sedivy, J. C., Tanenhaus, M. K., Chambers, C. G., and Carlson, G. N. (1999). Achieving incremental semantic interpretation through contextual representation. *Cognition*, 71:109–147.

Senghas, A. (1995). *Children's contribution to the birth of Nicaraguan Sign Language*. PhD thesis, MIT.

Shahin, K. (2002). *Postvelar Harmony*. Amsterdam: John Benjamins.

Siptár, P. and Törkenczy, M. (2000). *The Phonology of Hungarian*. Oxford University Press.

Smith, N. and Tsimpli, I.-M. (1995). *The Mind of a Savant: Language, Learning and Modularity*. Blackwell.

Sohn, H.-M. (1971). a-raising in Woleaian. *University of Hawaii Working Papers in Linguistics*, 3.8:15–36.

Steriade, D. (1981). Parameters of Metrical Harmony Rules. Ms., MIT.

Steriade, D. (1987). Redundant values. In Bosch, A., Need, B., and Schiller, E., editors, *Papers from the Parasession on Autosegmental and Metrical Phonology*, pages 339–362. Chicago: Chicago Linguistic Society.

Steriade, D. (1995). Underspecification and markedness. In Goldsmith, J., editor, *The Handbook of Phonological Theory*, pages 114–174. Blackwell.

Stevens, K. N. (1998). *Acoustic Phonetics*. Cambridge, MA: MIT Press.

Suomi, K. (1983). Palatal Vowel Harmony: A Perceptually Motivated Phenomenon? *Nordic Journal of Linguistics*, 6:1–35.

Suomi, K., McQueen, J., and Cutler, A. (1997). Vowel Harmony and Speech Segmentation in Finnish. *Journal of Memory and Language*, 36:422–444.

Surendran, D. and Niyogi, P. (2003). Measuring the Usefulness (Functional Load) of Phonological Contrasts. Technical Report TR-2003-12., Department of Computer Science, University of Chicago.

Suzuki, K. (1998). *A Typological Investigation of Dissimilation*. PhD thesis, University of Arizona.

Svantesson, J.-O. (1985). Vowel Harmony Shift in Mongolian. *Lingua*, 67:283–327.

Svantesson, J.-O., Tsendina, A., Karlsson, A., and Franzen, V. (2005). *The Phonology of Mongolian*. Oxford University Press.

Terbeek, D. (1977). *A cross-language multidimensional scaling study of vowel perception*. PhD thesis, UCLA.

Thorpe, W. (1958). The learning of song patterns by birds with special reference to the song of the chaffinch. *Fringilla coelebs. Ibis*, 100:535–570.

Tolskaya, I. (2008). Oroch Vowel Harmony. Master's thesis, University of Tromsɔ.

Traunmüller, H. (1994). Conventional, biological and environmental factors in speech communication: a modulation theory. *Phonetica*, 51:170–183.

Travis, L. (1984). *Parameters and Effects of Word Order Variation*. PhD thesis, MIT.

Trigo, L. (1991). On pharynx-larynx interactions. *Phonology*, 8:113–136.

Trubetzkoy, N. S. (1931). Die phonologischen Systeme. *TCLP*, 4:96–116.

Ultan, R. (1973). Some reflections on vowel harmony. *Working Papers on Language Universals*, 12:37–67.

Urbanczyk, S. (1995). *Patterns of Reduplication in Lushootseed*. Garland; Outstaning Dissertations in Linguistics.

Vago, R. (1975). *Hungarian Generative Phonology*. PhD thesis, Harvard University.

Vago, R. (2006). Hungarian: Phonology. In Keith Brown, editor, *The Encyclopedia of Language and Linguistics, 2nd Ed.* Oxford: Elsevier.

Välimaa-Blum, R. (1999). A featore-geometric description of Finnish vowel harmony covering both loans and native words. *Lingua*, 108:247–268.

van der Hulst, H. and Smith, N. (1989). Tungusic and Mongolian vowel harmony: a minimal pair. *Linguistics in the Netherlands*, 12:79–88.

van der Hulst, H. and van de Weijer, J. (1995). Vowel Harmony. In Goldsmith, J., editor, *The Handbook of Phonological Theory*, pages 495–534. Blackwell.

Vaux, B. (1993). The Origins of Altaic Labial Attraction. In Thráinsson, H., Calabrese, A., Carrier, J., Hale, M., and Watkins, C., editors, *Harvard Working Papers in Linguistics Volume 2*, pages 228–237. Cambridge, MA: HWPL.

Vaux, B. (1999). A Note on Pharyngeal Features. In Kuno, S. and Vaux, B., editors, *Harvard Working Papers in Linguistics 7*, pages 39–63. Cambridge, MA: HWPL.

von Soden, W. (1969). *Grundriss der Akkadischen Grammatik*. Rome: Pontificum Institutum Biblicum.

Vroomen, J., Tuomainen, J., and deGelder, B. (1998). The Roles of Word Stress and Vowel Harmony in Speech Segmentation. *Journal of Memory and Language*, 38:133–149.

Walker, R. (1999). Guaraní voiceless stops in oral vs. nasal contexts: an acoustical study. *Journal of the International Phonetic Association*, 29.1:63–94.

Walker, R. (2001). Round Licensing, Harmony, and Bisyllabic Triggers in Altaic. *Natural Language and Linguistic Theory*, 19:827–878.

Wechsler, S. and Zlatíc, L. (2003). *The Many Faces of Agreement*. Stanford and Chicago: CSLI Publications and Chicago University Press.

Wilson, C. (2003). Analyzing unbounded spreading with constraints: marks, targets, and derivations. Ms., UCLA.

Wolf, M. (2007). For an autosegmental theory of mutation. In Leah Bateman and Michael O'Keefe and Ehren Reilly and Adam Werle, editor, *UMOP 32: Papers in Optimality Theory III*, pages 315–404. Amherst, MA: GLSA.

Yamada, N. (1983). A Constraint on Phonological Variables. *Linguistic Analysis*, 12.1:29–84.

Yang, C. (2003). *Knowledge and Learning in Natural Language*. Oxford University Press.

Yip, M. (2006). The search for phonology in other species. *Trends in Cognitive Sciences*, 10.10:442–446.

Zhang, X. (1996). *Vowel systems of the Manchu-Tungus languages of China*. PhD thesis, University of Toronto.

Zhang, X. and Dresher, B. E. (1996). Labial Harmony in Written Manchu. *Saksaha: A Review of Manchu Studies*, 1:13–24.

Zoll, C. (1996). *Parsing Below the Segment in a Constraint-Based Framework*. Stanford: CSLI Publications.

Zoll, C. (2003). Optimal Tone Mapping. *Linguistic Inquiry*, 34.2:225–268.

Zsiga, E. (1997). Features, Gestures, and Igbo Vowels: An Approach to the Phonology-Phonetics Interface. *Language*, 73.2:227–274.