To appear in Steve Parker, (ed.) *Phonological Argumentation: Essays on Evidence and Motivation*. London: Equinox.

Morpheme-Specific Phonology: Constraint Indexation and Inconsistency Resolution*

Joe Pater
University of Massachusetts, Amherst

Abstract. This paper argues that exceptions and other instances of morpheme-specific phonology are best analyzed in Optimality Theory (OT) in terms of lexically indexed markedness and faithfulness constraints. This approach is shown to capture locality restrictions, distinctions between exceptional and truly impossible patterns, distinctions between blocking and triggering, and distinctions between variation and exceptionality. It is contrasted with other OT analyses of exceptions, in particular those that disallow lexically indexed markedness constraints and those that invoke lexically specified rankings (that is, cophonologies). The data discussed are from Assamese, Finnish and Yine (formerly Piro). A learnability account of the genesis of lexically indexed constraints is also provided, in which indexation is used to resolve inconsistency detected by Tesar and Smolensky's (1998, 2000) Recursive Constraint Demotion algorithm.

1. Introduction

Morphemes often behave differently phonologically in ways that cannot be explained purely phonologically: one morpheme undergoes or triggers a process while another morpheme fails to undergo or trigger that process, even though the two are in all relevant respects indistinguishable. Syncope in Yine (formerly known as Piro; Matteson 1965, Kisseberth 1970, Lin 1997) provides an example of such morpheme-specific phonology. Morphemes differ in whether they cause the preceding vowel to delete (/heta+nu/ [hetanu] 'going to see' vs. /heta+lu/ [hetlu] 'see it'), and in whether they undergo deletion themselves (/meyi+wa+lu/ [meyiwlu] 'celebration' vs. /heta+wa+lu/ [hetawalu] 'going to see him yet'). As the behavior of the homophonous pair of /-wa/ morphemes illustrates, morphemes that fail to condition syncope can differ in whether they undergo the process.

The distinction between exceptional triggering and blocking exemplified by Yine is captured straightforwardly in Optimality Theory (OT) if markedness and faithfulness constraints can be lexically indexed (Pater 2000). Morphemes that trigger a process are indexed for the application of a lexically specific markedness constraint, and morphemes that block a process are indexed for the application of a lexically specific faithfulness constraint. However, this distinction is not expressed in either of two alternative approaches to exceptionality in OT: a theory in which morphemes select constraint rankings (the 'cophonology' approach; e.g. Anttila 2002, Inkelas and Zoll 2003), or a theory in which only faithfulness constraints can be lexically

* Discussion in Ling 730 (UMass, Fall 2004), co-taught with John McCarthy, and Ling 606 (UMass, Spring 2005) contributed much to the development of the material herein. For helpful comments in those classes and elsewhere, I thank Eric Bakovic, Michael Becker, Adrian Brasonoveau, Andries Coetzee, Naz Merchant, Alan Prince, Bruce Tesar, and Matt Wolf. I am especially grateful to Arto Anttila for discussion of Finnish, to John McCarthy, Anne-Michelle Tessier and Nicholas Winslow for extensive discussion of the learnability proposal, and to Sara Finley, Kathryn Flack, Maria Gouskova, Peter Jurgec, Shigeto Kawahara, Michael Key, Shakuntala Mahanta, Marc van Oostendorp, Steve Parker and Hideki Zamma for comments on a draft of the paper.

indexed (e.g. Fukazawa 1999, Ito and Mester 1999, 2001). These three theories of exceptionality are introduced in section 1 of the paper. Section 2 shows how indexation of markedness and faithfulness constraints deals with the Yine data, and discusses the difficulties faced by the alternatives.

Anttila (2002) uses data from Finnish /a+i/ allomorphy to argue for a version of the cophonology approach. In this theory, morphemes can only be specified for rankings that are left unspecified in a partially ordered grammar. In the third section of the paper, I show that the types of patterns Anttila captures with this approach can in fact be straightforwardly analyzed with indexed constraints. In addition, I show that constraint indexation captures generalizations that escape the partial ordering theory of cophonologies. One generalization is that the Finnish alternation applies only to a string that includes a portion of the exceptional morpheme. A simple locality convention for the interpretation of indexed constraints accounts for this restriction, and similar ones in other languages, and rules out a range of implausible non-local morpheme-specific processes. This locality restriction is unstateable under the cophonology approach, or under the faithfulness-only indexation theory. A second type of generalization involves distinctions between variation and exceptionality, which are conflated in the partial ordering/cophonology theory, but which can be separated if exceptionality involves indexation.

Section 4 provides a learnability account that uses the inconsistency detection properties of the Recursive Constraint Demotion Algorithm (Tesar and Smolensky 1998, 2000, Tesar 1998, Prince 2002) to trigger the creation of indexed constraints. This account of the genesis of indexed constraints resolves the incongruity of having morpheme-specific constraints in a theory that assumes constraint universality, and also ensures that learners will seek a phonological generalization before resorting to an analysis in terms of exceptionality. In section 5 representational approaches to exceptionality are briefly discussed, with a focus on learnability considerations. Section 6 further discusses the power of the indexation and cophonology approaches, and section 7 concludes.

2. Yine Syncope and Constraint Indexation

2.1 Constraint indexation and cophonologies

Morphologically indexed constraints make their first appearance in the foundational work on Optimality Theory. Prince and Smolensky (1993/2004) propose Edgemost constraints that apply to specific morphemes in order to distinguish prefixes, suffixes, and edge-oriented infixes from one another. McCarthy and Prince (1993) reformulate Edgemost constraints in terms of Generalized Alignment, which they also use for cases of prosodic subcategorization, in which a morpheme is placed next to an instance of a prosodic category.

Fukuzawa (1999), Ito and Mester (1999, 2001), Kraska-Szelenk (1997, 1999) and Pater (2000) extend morphological indexation from Alignment to other constraints. Under this view, a single constraint can be multiply instantiated in a constraint hierarchy, and each instantiation may be indexed to apply to particular set of lexical items. These indexed constraints are universal markedness and faithfulness constraints, whose application is relativized to a set of lexical items (cf. the ‘parochial’ constraints in Hammond 1995 and Green 2005, which are not markedness constraints, since they can demand marked structures, and are not faithfulness constraints, since they apply directly to surface representations; see Russell 1995 and Golston 1996 for related proposals).

To take a simple hypothetical example, a language might have coda deletion (e.g. 1a-b), which is blocked in some lexical items (e.g. 1c):

(1) a. /pak/ → [pa] /pak+a/ → [paka]
b. /lot/ → [lo] /lot+a/ → [lota]
c. /tak/ → [tak] /tak+a/ → [taka]

Coda deletion requires a ranking of NOCODA >> MAX. The exceptional items are targeted by a morphologically indexed MAX constraint. This version of MAX (MAX-L) ranks above NOCODA, and applies only to those lexical items indexed for its application (here with an 'L' for 'lexical').

(2) Grammar: MAX-L >> NOCODA >> MAX
Lexicon: /pak/ /lot/ /tak_L/

The tableaux in (3) show the results of applying this grammar to a form that lacks the index (/pak/), and one that bears it (/tak_L/).

(3)

Input	Output	MAX-L	NOCODA	MAX
pak	pak		* !	
	☞ pa			*
tak_L	☞ tak		*	
	ta	* !		*

An important attribute of this approach to morpheme-specific phonology is that it captures the distinction between an exceptional form and an impossible one. Let us further assume that onset clusters are entirely absent from our hypothetical language. *COMPLEX ("no consonant clusters") would dominate MAX-L, since there is no evidence to contradict the default Markedness >> Faithfulness ranking (Smolensky 1996, Hayes 2004, Prince and Tesar 2004). If under Richness of the Base (Prince and Smolensky 1993/2004) an underlying form with a cluster is given a lexical diacritic, the cluster is reduced, as shown in (4). This is not to say a language could not have both exceptional codas and clusters, but rather that in the absence of evidence of a structure, a learner creates a grammar that rules it out completely.

(4)

Input	Output	*COMPLEX	MAX-L	NOCODA	MAX
CCV	CCV	* !			
	☞ CV				*
CCV_L	CCV	* !			
	☞ CV		*		*

A closely related approach to morpheme-specific phonology is that of cophonology theory (Kirchner 1993, Nouveau 1994, Ito and Mester 1995a,b, Orgun 1996, Inkelas 1999, Anttila 2002, Inkelas and Zoll 2003, Caballero 2005 and Zamma 2005). In this view, a grammar has only a single instantiation of each constraint, but individual morphemes can demand a different ranking of some of them. The term "cophonology" might appear to imply a formalization in which morphemes that are subject to different rankings are submitted to entirely different grammars. I will retain this term, but will adopt a notation similar to that of Anttila (2002), in which there is a single constraint hierarchy, and morphemes impose rankings of only a subset of the constraints.

Our hypothetical example could be analyzed in cophonology theory with a grammatical hierarchy of *COMPLEX, NOCODA >> MAX, and with exceptional lexical items selecting the reverse ranking of MAX and NOCODA. The exceptional ranking is included as part of the lexical entry for the morpheme.

(5) Grammar: *COMPLEX, NOCODA >> MAX
Lexicon: /pak/ /lot/ /tak$_{\text{Max>>NoCoda}}$/

Input	Output	*COMPLEX	NOCODA	MAX
pak	pak		* !	
	☞ pa			*
Input	Output	*COMPLEX	MAX	NOCODA
tak$_{\text{Max>>NoCoda}}$	☞ tak			*
	ta		* !	

Without further elaboration, this theory cannot express the distinction between exceptional and impossible patterns. If lexical items are specified for constraint rankings, then by Richness of the Base any constraint ranking should be available for lexical specification. Returning to our example of a language that has exceptional codas but bans consonant clusters, nothing would rule out the specification of a lexical item for the reverse ranking of *COMPLEX and MAX, as illustrated in the following Richness of the Base tableau:

(6)

Input	Output	*COMPLEX	NOCODA	MAX
CCV	CCV			*
	☞ CV	*!		
Input	Output	MAX	NOCODA	*COMPLEX
CCV$_{\text{Max >>*Complex}}$	☞ CCV			*
	CV	* !		

To overcome this sort of problem, Anttila (2002) extends the partial ordering theory of variation to morpheme-specific phonology. Under the Subregularity Interpretation (Anttila 2002: 22), only pairs of constraints whose ranking is unspecified in the grammar can have lexically specified rankings. Our hypothetical language would have a ranking of *COMPLEX over MAX, and the ranking of NOCODA would be left unspecified. Lexical items are then specified for a ranking of the conflicting constraints, here MAX and NOCODA.

(7) Grammar: *COMPLEX >> MAX NO CODA
Lexicon: /pak$_{\text{NoCoda>>Max}}$/ /lot$_{\text{NoCoda>>Max}}$/ /tak$_{\text{Max>>NoCoda}}$/

At first glance, this analysis seems to be a notational variant of the constraint indexation analysis. However, there are two differences that we can see even in this simple example.

First, indexed constraints account for impossible patterns in the way that is standard in OT: by using the grammar, constructed on the basis of positive evidence, to filter out a rich base. On the other hand, the partial ordering/cophonology theory adds a stipulation that lexical rankings are limited to those that are unspecified in the grammar.

The second difference stems from the fact that in the partial ordering theory all lexical items must be specified for a ranking of the grammatically unranked constraints. If a lexical item were left unspecified, then it would show variation, with a ranking chosen randomly each time it is submitted to the grammar. Anttila (2002) presents this connection between variation and exceptionality as a positive attribute of the model, but there are clearly many cases of exceptionality without accompanying variation, and variation without exceptionality. This is evident in that in lexical phonology, variation is seen as a characteristic of post-lexical rules, and exceptionality of lexical rules (see e.g. Kaisse and Shaw 1985). In section 3.2, I will show that this conflation of variation and exceptionality leads to missed generalizations in the Finnish data Anttila (2002) analyzes. However, the problem goes far beyond the Finnish case, since the partial ordering/cophonology model predicts that all variation should be accompanied by exceptionality, and vice versa (see Anttila 2002: 29 for some discussion).

A further difference between lexical rankings and indexed constraints is that only the latter can distinguish between a lexically indexed markedness constraint, which produces exceptional triggering of a process, and a lexically indexed faithfulness constraint, which produces exceptional blocking. This distinction is also unavailable to a theory in which only faithfulness constraints can be indexed (e.g. Fukuzawa 1999, Ito and Mester 1999, 2001, Kraska-Szelenk 1997, 1999). In the next section, I show that the distinction between exceptional blocking and triggering is necessary to capture the facts of Yine syncope.

2.2 Yine and Indexed Constraints

Kisseberth (1970) draws attention to Yine syncope, as described in Matteson (1965), for its implications for a theory of exceptions in rule-based phonology. Unless indicated otherwise, the examples here are ones Kisseberth supplies. The forms in (8) show syncope applying before the nominalizing suffixes /-lu/ and /-nu/ ([ru] in (8c) is an allomorph of /-lu/; /nu/ is used for abstract nouns as in (8d)), before the indirective suffix /-ya/ (8e), and before the third person singular pronominal suffix /-lu/ (8f).

(8)	a. /yimaka+lu/	[yimaklu]	‘teaching’
	b. /kama+lu/	[kamlu]	‘handicraft’
	c. /kakonu+lu/	[kakonru]	‘a shelter in which a hunter hides’
	d. /hata+nu/	[hatnu]	‘light, shining’
	e. /heta+ya/	[hetya]	‘see there’
	f. /heta+lu/	[hetlu]	‘see it’

As the examples in (9) illustrate, syncope fails to apply before the verbal theme formative /-ta/, the anticipatory suffix /-nu/, and the intransitive verb theme suffix /-wa/.

(9)	/meyi+ta/	[meyita]	'to celebrate'
	/hata+ta/	[hatata]	'to illuminate'
	/heta+nu/	[hetanu]	'going to see'
	/meyi+wa+ta/	[meyiwata]	'to celebrate' (Matteson: 303)

All of the suffixes that fail to trigger syncope in (9) do undergo it when placed before one of the syncope-triggering suffixes, as the examples in (10) show.

(10)	/meyi+wa+lu/	[meyiwlu]	'celebration'
	/heta+nu+lu/	[hetanru]	'going to see him'
	/yona+ta+na+wa/	[yonatnawa]	'to paint oneself'

However, there is a further 'exceptional' suffix /-wa/ ('yet, still') which neither conditions syncope nor undergoes it:

(11)	/heta+wa+lu/	[hetawalu]	'going to see him yet'
	/n+hiʃinika+wa+lu/	[nuʃinikawalu]	'I'm still thinking about it' (Matteson: 74)

There is no phonological property that distinguishes the morphemes that trigger syncope from those that do not: the homophones with the shape /-nu/ fall into the two classes. Similarly, morphemes that block syncope and those that do not have no distinguishing property, as clearly illustrated by the two morphemes /-wa/. There is also no apparent morphosyntactic distinction between the different classes of morpheme, and no restriction that triggers must appear closer to the root than non-triggers (or vice versa), as might be expected under an interpretation of the data in terms of lexical or prosodic phonology. In her grammar, Matteson (1965) treats the distinction between triggers and non-triggers as idiosyncratic.

To deal with the Yine data, Kisseberth (1970: 57) proposes a theory of exceptionality in which lexical items are categorized as "either undergoing a rule or not, and as either serving as the context for a rule or not" (see Zonneveld 1978, 1985 for further discussion and a critique). This distinction is expressed naturally in a version of OT with indexed markedness and faithfulness constraints.

As the constraint driving syncope, I will make use of an Alignment constraint requiring a suffix to follow a consonant (cf. Lin 1997).[1] 'Suffix' in (12) is a convenient cover term for the set of morphemes that follow the root.

[1] Lin's (1997) syncope constraint requires the stem to end in a consonant. Lin does not analyze the absence of syncope in unsuffixed stems, or with non-triggering suffixes. Under the locality condition discussed in section 2, the scope of the indexed constraint must include the triggering morpheme; Lin's constraint would not meet this criterion. One might also invoke an Alignment constraint that requires a suffix to follow a heavy syllable; Matteson (1965: 24) describes post-vocalic consonants as variably

(12) ALIGN-SUF-C Align(Suffix, L, C, R)
The left edge of a suffix coincides with the right edge of a consonant

The motivation for this constraint may be to phonologically mark morphological boundaries; the onset of a suffix-initial consonant could well be more prominent when it is preceded by a consonant than when it is preceded by a vowel. Yine syncope does not seem to be driven by a STRESS-TO-WEIGHT constraint (see Gouskova 2003 on prosodically driven syncope); it also occurs before bisyllabic suffixes, in which case the resulting CVC would occupy a stressless syllable, given Yine's pattern of penultimate main stress and clash-avoiding secondary stress.

The ranking ALIGN-SUF-C >> MAX produces vowel deletion, as shown in the tableau in (13):

(13)

Input	Output	ALIGN-SUF-C	MAX
heta+ya	hetaya	* !	
	☞ hetya		*

To distinguish the suffixes that trigger syncope from the non-triggers, the suffixal argument of the ALIGN constraint that dominates MAX is indexed to the set of morphemes that trigger syncope. The general ALIGN constraint rests beneath MAX. The result is shown in (14), which includes in the lexicon all of the suffixes found in examples (8) – (10) above.

(14) Grammar: ALIGN-SUF(L)-C >> MAX >> ALIGN-SUF-C
Lexicon: /-lu_L/ /-nu_L/ /-lu_L/ /-ya_L/ /-ta/ /-nu/ /-wa/ /-wa/

Input	Output	ALIGN-SUF(L)-C	MAX	ALIGN-SUF-C
heta+ya_L	hetaya	* !		*
	☞ hetya		*	
heta+wa	☞ hetawa			*
	hetwa		* !	

Similarly, to distinguish the suffixes that undergo syncope from those that don't, MAX appears in both a lexically indexed and a general version. The lexically indexed version ranks above the indexed markedness constraint; the general one ranks beneath it. The index for MAX is given as 'L2', and for ALIGN as 'L1'.

closing the preceding syllable (cf. Lin 1997: 425). However, in deference to Matteson's claim that pre-consonantal consonants are invariably syllabic, I retain the formulation in (12).

(15) Grammar: MAX-L2 >> ALIGN-SUF(L1)-C >> MAX >> ALIGN-SUF-C
Lexicon: /-lu$_{L1}$/ /-nu$_{L1}$/ /-ya$_{L1}$/ /-ta/ /-nu/ /-wa/ /-wa$_{L2}$/

Input	Output	MAX-L2	ALIGN-SUF(L1)-C	MAX
heta+nu+lu$_{L1}$	hetanulu		* !	
	☞ hetanru			*
heta+wa$_{L2}$+lu$_{L1}$	☞ hetawalu		*	
	hetawlu	* !		*

Syncope is also blocked when it would create a triconsonantal cluster (Matteson 1965: 36, Lin 1997), which indicates that a constraint against such clusters dominates the indexed alignment constraint. I will use the simple constraint *CCC for this purpose; see Lin (1997: 420) for an alternative formulation. The result of this ranking is shown in (16), with an example Matteson (1965: 36) glosses as 'she washes it'.

(16)

Input	Output	*CCC	ALIGN-SUF(L1)-C
terka+lu$_{L1}$	terklu	* !	
	☞ terkalu		*

A hierarchy for Yine syncope that incorporates indexed markedness and faithfulness constraints is thus as in (17).

(17) *CCC, MAX-L2 >> ALIGN-SUF(L1)-C >> MAX >> ALIGN-SUF-C

This analysis distinguishes morpheme-specific triggering from morpheme-specific blocking, and also accounts for phonological blocking.

2.3 Yine and Cophonologies

In a cophonological analysis of Yine, morphemes would select rankings between the markedness constraint causing syncope and the conflicting faithfulness constraint. Under this approach, a morpheme causing syncope could be distinguished from one that does not as follows (compare 14):

(18)

Input	Output	ALIGN-SUF-C	MAX
heta+ya	hetaya	* !	
	☞ hetya		*
Input	Output	MAX	ALIGN-SUF-C
heta+wa	☞ hetawa		*
	hetwa	* !	

One issue is the determination of the outcome with morphemes that demand opposite rankings of the constraints. In /heta+nu+lu/ [hetanru], for example, /-nu/ requires MAX >> ALIGN-SUF-C to block deletion in the stem-final syllable, while /-lu/ requires ALIGN-SUF-C >> MAX to force deletion in /-nu/. However, given some kind of cyclic evaluation, which Lin (1997) argues is independently necessary for Yine, the outcome for /heta+nu/ could be calculated before, and independently of, the outcome for the entire string (see Orgun 1996 and Inkelas and Zoll 2005 on cyclicity and cophonologies). This is not a complete solution, since the syncope-producing constraint must also be limited to apply only to the environment of the outermost suffix of the complete /heta+nu+lu/ (perhaps using bracket erasure), but the problem does seem resolvable.

The more serious problem is how to distinguish the two forms of /-wa/ noted by Kisseberth (1970) and discussed in section 2.2. Both suffixes fail to trigger syncope, which should indicate that MAX dominates the markedness constraint. The problem is that only one of the suffixes fails to undergo syncope: in constrast to the /-wa/ meaning 'yet, still' in [hetawalu], the intransitive verb theme suffix in /meyi+wa+lu/ does syncopate ([meyiwlu]). As it stands, the account predicts that both should block (or that both should undergo, if the following suffix determines the ranking, as in the cyclic account mooted in the last paragraph). In other words, cophonology theory fails to distinguish morpheme-specific triggering from morpheme-specific blocking. This is an instance of a more general locality problem for cophonologies: when morphemes impose a ranking on the grammar, they do not specify where in the string that ranking should apply. When /-wa/ demands a ranking MAX >> ALIGN-SUF-C, it does not specify whether that ranking stops /-wa/, the preceding syllable, or some other syllable from undergoing syncope.

3. Finnish /a+i/ Allomorphy

*3.1 Locality and *[ai]*

Anttila (2002) presents a pair of morphologically conditioned alternations in Finnish as providing evidence for a choice of lexical rankings over indexed constraints. In this section, I discuss the constraint that triggers these alternations, and show that constraint indexation allows for an analysis of the local nature of this and other cases of exceptional triggering. I also show that extant versions of cophonology theory and faithfulness-only indexation fail to capture such locality effects. Except where indicated, all of the data come from Anttila (2002), who based his study on an electronic version of a dictionary of Modern Finnish (Sadeniemi 1973).

The alternations affect a stem-final low vowel /a/ that precedes one of two homophonous suffixes /-i-/, which indicate either past or plural. The /a/ either deletes, or mutates to [o]. The examples in (19) show that the choice between the alternations is sometimes lexically determined; the stems are identical in all relevant phonological respects, yet one undergoes final vowel mutation (19a), one undergoes final vowel deletion (19b), and one varies between mutation and deletion (19c).

(19)	a. /tavara+i+ssa/	[tavaroissa]	'thing (plural-inessive)'
	b. /jumala+i+ssa/	[jumalissa]	'God (plural-inessive)'
	c. /itara+i+ssa/	[itaroissa] ~ [itarissa]	'stingy (plural-inessive)'

The alternations do not apply stem internally (20a), nor do they apply in all derived environments; (20b) is an example of non-application with the conditional suffix /-isi/ (Arto Anttila p.c.).

(20) a. /taitta-i/ [taittoi] ~ [taitti] *[toittoi] *[titi] 'break (past)'
b. /anta-isi/ [antaisi] 'give (conditional)'

Anttila's (2002) analysis focuses on the interplay between morphological and phonological conditioning in the choice between mutation and deletion, and does not include the constraint that drives the alternation. To penalize the [ai] sequence in the appropriate morphological context, we can index the relevant constraint to the plural and past tense morphemes. For expository ease, I will adopt the straightforward but stipulative *[ai] as the active constraint. The indexed version of the constraint ranks above MAX and IDENT, while the general version of it ranks beneath the faithfulness constraints. The past and the plural /-i-/ are both indexed for the application of the specific constraint:

(21) Grammar: $*[ai]_L$ >> MAX, IDENT >> *[ai]
Lexicon: $/\text{-i-}/_L$ $/\text{-i-}/_L$ /-isi-/ /taitta/

In the analysis of Yine, the indexed markedness constraint was an ALIGN constraint that provides an argument for the suffix. As such, the alignment constraint automatically specifies the context in which it applies. The following schema specifies how other indexed constraints assess violation marks:

(22) $*X_L$
Assign a violation mark to any instance of X that contains a phonological exponent of a morpheme specified as L

This formulation serves as a locality convention for indexed constraints: they apply if and only if the locus of violation contains some portion of the indexed morpheme. Note especially that "phonological exponent" in (22) is fulfilled by any subpart of the morpheme; there is no requirement that the entire morpheme be contained within the scope of the constraint. The locality convention in (22) also provides an explicit formulation of how a constraint indexed to a general morphological category applies: $*X_{SUFFIX}$, for example, would apply to all instances of X including an element of a suffix (see Pater 2006 for an application to derived environment effects). Along these lines, the Finnish *[ai] constraint could also have been specified to apply to Plural and Past tense morphemes.

The following tableau shows how the ranking in (21) applies to /taitta-i/:

(23)

Input	Output	$*[ai]_L$	MAX	IDENT	*[ai]
$/\text{taitta-i}_L/$	taittai	* !			**
	☞ taitti		*		*
	☞ taittoi			*	*
	titti		** !		
	toittoi			** !	

The indexed constraint *[ai]$_L$ assigns a violation only to an [ai] sequence containing the indexed morpheme /-i/ and not to the root internal one. Because the indexed constraint does not apply to the root-internal [ai] sequence, the faithfulness constraints protect it. Because the ranking between MAX and IDENT is unspecified, either deletion or mutation of the stem-final vowel can emerge as the optimal outcome (see Casali 1997 on constraints that block changes to the suffix-initial vowel). Section 3.4 shows how the choice between deletion and mutation is made for stems that allow only one outcome. But first, I will provide another example of a case of exceptional triggering that is similarly local, and discuss the locality problems faced by alternative theories.

3.2 Exceptional triggering in Assamese

Mahanta (to appear) provides an example of local exceptional triggering in Assamese [ATR] harmony. In this language, [+ATR] spreads regressively from the high vowels [i] and [u], yielding an allophonic alternation in the mid vowels, which appear as [+ATR] [e] and [o] when they precede [+ATR] vowels, and as [-ATR] [ɛ] and [ɔ] elsewhere. Underlying [-ATR] [ʊ] also neutralizes with [+ATR] [u] when it precedes a [+ATR] vowel. Harmony is iterative, in that the targets of harmony spread [+ATR] to preceding vowels. Iterativity is illustrated in examples (24a-c), where harmony is triggered by the [+ATR] vowels in the initial position of the suffixes /-iyɑ/ and /-uwɑ/.

(24)	a. [bɔyɔx]	‘age’	[boyoxiyɑ]	‘aged’
	b. [tɔlɔt]	‘below’	[tolotiyɑ]	‘subordinate’
	c. [gʊbɔr]	‘dung’	[guboruwɑ]	‘fly (with dung-like smell)'

The low [-ATR] vowel [ɑ] is normally opaque to vowel harmony, as shown in examples (25a) and (25b); the /-i/ suffix regularly triggers harmony in other [-ATR] vowels.

(25)	a. [kɔpɑh]	‘cotton’	[kɔpɑhi]	‘made of cotton’
	b. [zɔkɑr]	‘shake’	[zɔkɑri]	‘shake’ (infinitive)

The exceptional triggering pattern arises when the morphemes /iyɑ/ and /uwɑ/ in (24) are added to stems whose final vowel is [ɑ]. With these morphemes, /ɑ/ raises to a mid [+ATR] vowel, usually [o] (26a-g), except when [e] precedes, in which case /ɑ/ also surfaces as [e] (26h-i). In examples (26e-i), we see again that harmony is iterative; the raised /ɑ/ yields the regular harmony pattern to its left.

(26)	a. [sɑl]	‘roof’	[soliyɑ]	'roof’
	b. [dɑl]	‘branch’	[doliyɑ]	‘branch-ed’
	c. [d^hɑr]	‘debt’	[d^horuwɑ]	‘debtor’
	d. [mɑr]	‘beat’(verb)	[moruwɑ]	‘beat’(causative)

e. [bɔzɑr]	'marketplace'	[bozoruwɑ]	'cheap'
f. [kɔpɑl]	'destiny'	[kopoliyɑ]	'destined'
g. [gʊlɑp]	'rose'	[gulopiyɑ]	'pink'
h. [ɛŋɑr]	'charcoal'	[eŋeruwɑ]	'black as charcoal'
i. [dʰɛmɑli]	'play/amusing'	[dʰemeliyɑ]	'playful/amusing'

The local nature of the exceptional triggering pattern is shown in the following examples, in which non-stem-final /ɑ/ fails to raise.

(27)

a. [pɑtɔl]	'light'	[pɑtoliyɑ]	'lightly'
b. [ɑpɔd]	'danger'	[ɑpodiyɑ]	'in danger'
c. [ɑbɔtɔr]	'bad time'	[ɑbotoriyɑ]	'bad timed'
d. [ɑlɑx]	'luxury'	[ɑloxuwɑ]	'pampered'
e. [ɑdʰɑ]	'half'	[ɑdʰoruwɑ]	'halved'

Mahanta's analysis involves a lexically indexed version of the following constraint, whose definition I have slightly elaborated:

(28) *[-ATR][+ATR]
Assign a violation mark to the minimal string containing a [-ATR] vowel followed by a [+ATR] vowel

'Minimal' in this definition refers to the smallest string of segments that contains the specified sequence, and which would fail to contain the sequence if a segment were removed. This is included to make explicit the locus of violation (thanks to Colin Wilson for related discussion). The indexed constraint would then be formulated as in (29).

(29) *[-ATR][+ATR]$_L$
Assign a violation mark to the minimal string containing a [-ATR] vowel followed by a [+ATR] vowel, if that string contains a phonological exponent of a morpheme indexed as L

As Mahanta shows, a ranking in which these constraints straddle a faithfulness constraint generates /ɑ/ raising only in the immediate context of the exceptional morpheme:

(30)

Input	Output	*[-ATR][+ATR]$_L$	IDENT-LOW	*[-ATR][+ATR]
/ɑlɑx-uwɑ$_L$/	ɑlɑx-uwɑ	*!		**
	☞ ɑlox-uwɑ		*	*
	olox-uwɑ		**!	

For further details of the analysis, see Mahanta (to appear).

In the Assamese example, the indexed constraint is well-motivated both in the phonology of the language and cross-linguistically. It is only the ranking that is exceptional. The locality of the exceptional pattern provides strong support for the claim that indexed constraints are interpreted as in (22) (see Finley 2006 for further discussion of locality in exceptional vowel harmony).

There are some instances of exceptional triggering in which the exceptional pattern is perpetuated beyond the purely local environment (thanks to Matt Wolf for discussion). Terena, for example, has iterative nasal spreading in the context of the first person morpheme (Bendor-Samuel 1960), but no nasal harmony elsewhere. All cases of this type, however, seem to involve a "floating feature", that is, a morpheme that is realized only as a featural change on the stem. McCarthy (2002: 204) uses comparative markedness to analyze a tonal pattern of this type in which a floating feature docks and triggers iterative spreading. Interestingly, comparative markedness can only generate iterative spreading (it would fail to capture a case like Assamese exceptional triggering), and indexed markedness can only generate non-iterative spreading (that is, where the spreading is exceptional). Why the scenarios traditionally described in terms of floating features should be associated with the constraint generating iterative spreading, and the ones that involve triggering by context should involve indexed markedness or the like, is a question for further research.

3.3 The challenge of local exceptional triggering

In this section, I discuss the challenges posed by local exceptional triggering for cophonology theory and for indexed faithfulness theory, using the Finnish case to illustrate. Recall that /-i-/, but not /-isi-/, triggers mutation. A cophonology analysis would allow /-i-/ and /-isi-/ to select different rankings of *[ai] and the faithfulness constraints:

(31) a. /-i-/ *[ai] >> MAX, IDENT
b. /-isi-/ MAX, IDENT >> *[ai]

An indexed faithfulness analysis would also leave the *[ai] constraint in its general form, but have faithfulness constraints indexed to morphemes that do not trigger the alternation:

(32) Grammar: MAX_L, IDENT_L >> *[ai] >> MAX, IDENT
Lexicon: /-i-/ /-i-/ /-isi-/_L

It is not clear, on either account, how to deal with a form like /taitta-i/, in which one /ai/ sequence undergoes the alternation, and the other one does not. In the cophonology analysis, we might put /taitta/ in the list of morphemes with the FAITH >> MARK ranking:

(33) a. /-i-/ *[ai] >> MAX, IDENT
b. /-isi-/ /taitta/ MAX, IDENT >> *[ai]

But which ranking should /taitta-i/ select? Either ranking will yield the wrong result: either both /ai/ sequences will surface, or neither will.

In the morpheme-specific faithfulness account, we could similarly index /taitta/ to the faithfulness constraints:

(34) Grammar: MAX_L, IDENT_L >> *[ai] >> MAX, IDENT
Lexicon: /-i-/ /-i-/ /-isi-/$_L$ /taitta/$_L$

The problem here is that this would protect both instances of /a/, rather than just the second one.

A possible solution for the Finnish case would be to relativize *[ai] to the derived context. With this constraint, which I will label *[a+i], analyses are available in either framework; (35) sketches the lexical ranking analysis, and (36) the indexed faithfulness one.

(35) a. /-i-/ /taitta/ *[a+i] >> MAX, IDENT
b. /-isi-/ MAX, IDENT >> *[a+i]

(36) Grammar: MAX_L, IDENT_L >> *[a+i] >> MAX, IDENT
Lexicon: /-i-/ /-i-/ /-isi-/$_L$ /taitta/

One issue with this approach is that it requires an additional theory of derived environment effects: with morpheme-specific constraints, these are captured by simply indexing a constraint to a general morphological category (see Pater 2006). A bigger issue is that this solution does not address the general locality problem for morpheme-specific phonology (see Horwood 1999 and Wolf 2006 for discussion of locality problems in the Antifaithfulness theory of Alderete 2001 and the Realize Morpheme theory of Kurisu 2001).

For the analyses of Finnish in (35) and (36) to be successful, a certain amount of non-locality must be countenanced. The morpheme that fails to trigger (/-isi-/) must demand higher faithfulness not for itself, but for the immediately adjacent segment, so that the stem-final /a/ neither mutates nor deletes. A similar scenario will obtain for any situation in which triggering morphemes must be distinguished from non-triggers, and the alternation takes place outside of the morpheme; the Assamese case discussed in the previous section provides another clear example. The problem is in defining how much non-locality is allowed.

If the ranking introduced by a morpheme holds over the entire string, clearly undesirable results follow (see also Horwood 1999). For example, a language could have a general ranking ONSET >> DEP, which produces epenthesis in vowel-initial stems. If a suffix could introduce a DEP >> ONSET ranking that holds over the entire string, then epenthesis would be blocked word-initially only in the presence of that suffix, as in (37), where /ba/ is the exceptional morpheme.

(37) /amana/ [ʔamana] /amana+da/ [ʔamanada]
/amana+ba/ [amanaba]

A legion of similarly implausible cases could be constructed; I leave this to the reader's imagination.

Kiparsky (1993), Inkelas (2000), Mascaró (2003) and Kurisu (2006) discuss several cases in which affixation has effects that are somewhat similar to (37); it leads to a change that is not phonologically conditioned by the affix, and occurs at a distance. In Catalan, for example, exceptions to unstressed vowel deletion are regularized in derivation (e.g. the exceptional unreduced [e] of [tótem] 'totem' is lost in [tutəmízmə]). As Kiparsky (1993) notes, this can be characterized as the loss of exception features in the derived form. Note that this is different in at least two ways from the hypothetical example in (37), in which derivation *induces* exceptionality, and only one morpheme introduces the exceptional alternation. Possible analyses for cases of the Catalan type include an exceptionality analogue of bracket erasure, or indexation of faithfulness to a category that identifies the bare stem, but not the derived one. It does not seem that the full power of cophonology theory is needed (cf. Inkelas 2000), since this would also generate unattested cases like those in (37).

A position between the extremes of the ranking holding only of the morpheme itself, and of it holding for the whole string, is that it holds of a string that contains some portion of that morpheme. This restriction is straightforwardly captured by the interpretive schema for indexed constraints in (24). In this theory, an indexed constraint cannot have the unattested long-distance triggering effect shown in (37). If indexation were limited to faithfulness constraints, it is difficult how to see how the intended restriction could be stated. It would also be difficult to formalize in cophonology theory, especially since it is taken as a fundamental and distinguishing assumption of the cophonology program that phonological constraints are not specified for morphological context:

> All constraints are fully general, but morphological class or lexical class are potentially associated with distinct rankings of those constraints. (Inkelas and Zoll 2003: 1)
>
> The alternative [to indexed constraints - JP] is to keep phonological constraints purely phonological, but posit a range of distinct COPHONOLOGIES, that is, different constraint rankings for different morphological categories. (Anttila 2002: 2)

3.4 Morphological and Phonological Conditions on Repair Choice

Anttila (2002) shows that both morphological idiosyncrasy, and the phonological environment, can affect the choice of mutation or deletion as the repair for *[ai]. He uses the partial ordering/lexical ranking theory to analyze the interplay between morphological and phonological conditioning. Here I replicate a portion of Anttila's analysis with indexed constraints to show that they are capable of expressing these sorts of generalizations. I also provide an analysis of generalizations that can be captured with indexed constraints, but not under partial ordering/cophonology theory.

As mentioned in section 3.1, stems of the same phonological shape can show three patterns. They can either select mutation, deletion, or vary between the two. The examples are repeated in (38).

(38) a. /tavara+i+ssa/ [tavaroissa] 'thing (plural-inessive)'
b. /jumala+i+ssa/ [jumalissa] 'God (plural-inessive)'
c. /itara+i+ssa/ [itaroissa] ~ [itarissa] 'stingy (plural-inessive)'

I will follow Anttila (1997, 2002) in analyzing variation as the result of conflicting constraints being unranked with one another, with a ranking being randomly selected each time the grammar derives an output. However, it is worth noting that the present account of morpheme-specific phonology is compatible with other approaches to variation, including that of Boersma and Hayes (2001).

The grammar in (39) deals with the three stem types. The general MAX and IDENT are left unranked, so that unindexed stems show variation. Morpheme-specific versions of the constraints are ranked above the general ones; stems indexed to one of them will show consistent deletion or mutation.

(39) Grammar: *[ai]-L1 >> MAX-L2, IDENT-L3 >> MAX, IDENT >> *[ai]
Lexicon: /-i-/$_{L1}$ /-i-/$_{L1}$ /-isi-/ /tavara/$_{L2}$ /jumala/$_{L3}$ /itara/

The result of applying this grammar to each of the three stem types is shown in (40). The unindexed stem shows variation (two optimal candidates in this tableau), while the indexed stems show either mutation or deletion, depending on whether they are indexed to MAX or IDENT.

(40)

Input	Output	*[ai]-L1	MAX-L2	IDENT-L3	MAX	IDENT
/itara-i$_{L1}$-ssa/	itaraissa	* !				
	☞ itarissa				*	
	☞ itaroissa					*
/tavara$_{L2}$-i$_{L1}$-ssa/	tavaraissa	* !				
	tavarissa		* !		*	
	☞ tavaroissa					*
/jumala$_{L3}$-i$_{L1}$-ssa/	jumalaissa	* !				
	☞ jumalissa				*	
	jumaloissa			* !		*

Antilla (2002) presents data showing that in some phonological contexts, not all these options are observed. With stems that consist of an even number of syllables (the stems in (40) are trisyllables), the generalization he uncovers is relatively straightforward: mutation occurs unless the preceding vowel is round, in which case deletion occurs instead. Anttila analyzes deletion after round vowels as an effect of an OCP constraint against a sequence of round vowels, which is violated when mutation creates an [o]. Since Finnish constructs trochees from left-to-right, Anttila (2002: 17) derives the syllable count generalization by restricting the constraint to the foot-internal context. To rule out mutation for this type of stem, Anttila ranks OCP/V[rd]$_\phi$ above MAX (which he labels *DEL). In his account, the presence of this ranking in the grammar bans lexical items from choosing the reverse order of the constraints. We can achieve the same effect with morpheme-specific constraints by ranking OCP/V[rd]$_\phi$ above the indexed version of MAX:

(41) OCP/V[rd]$_\phi$, IDENT-L3 >> MAX-L2 >> MAX, IDENT

With this ranking, even if a lexical item with a final round vowel is indexed to lexically specific MAX, it will undergo deletion rather than mutation, since mutation conflicts with OCP/V[rd]$_\phi$. The difference between disyllabic and trisyllabic stems is demonstrated in the following Richness of the Base tableaux, in which foot boundaries are indicated by parentheses. These tableaux show the result of indexation to MAX-L2 for both stem shapes. In disyllabic forms, where OCP/V[rd]$_\phi$ applies, mutation is ruled out, even with this indexation. In trisyllabic forms, mutation does occur with indexation. Hypothetical forms are used in both cases; Anttila (2002: 10) states that 1/3 of the forms like /itota/ undergo categorical mutation, though he does not provide any examples.

(42)

Input	Output	OCP/V[rd]$_\phi$	IDENT-L3	MAX-L2	MAX	IDENT
/tota$_{L2}$-i$_{L1}$/	☞ (toti)			*	*	
	(totoi)	* !				*
/itota$_{L2}$-i$_{L1}$/	(ito)ti			* !	*	
	☞ (ito)(toi)					*

The other side of this even-numbered stem generalization, that mutation applies in the absence of a preceding round vowel, is not dealt with in Anttila's (2002) analysis. As it stands, the analysis predicts that these stems should behave just like trisyllables of the same phonological shape. To some extent, this is borne out: as the example /taitta-i/ [taitt-oi] ~ [taitt-i] discussed in the last section shows, this side of the generalization is not iron-clad. Anttila (2002: 5) cites Karlsson (1982) as noting approximately 35 verb stems with variation in this context. However, unlike the trisyllabic stems, none of this type undergo categorical deletion.

To generate the default pattern of mutation, we need a constraint that applies to even-numbered stems, but not odd-numbered ones; ranking MAX over IDENT will not suffice. Drawing on Anttila's proposal that foot structure is responsible for syllable count generalizations, we can note that in the bare form of an odd-numbered stem the final syllable will be unparsed. Given a faithfulness relation between the stem and the suffixed form (see esp. Bakovic's 2000 elaboration of Benua's 2000 proposal), a MAX constraint that protects only footed segments will target the final vowels of only even-numbered stems. This constraint, which I will label OO-MAX$_\phi$, ranks above the general form of IDENT, but beneath OCP/V[rd]$_\phi$, so that deletion still occurs with a preceding round vowel, as in /tota-i/ [toti] (see (42)). To allow for variation between mutation and deletion in exceptional cases like /taitta-i/, the lexically specific version of IDENT ranks evenly with OO-MAX$_\phi$. The tableaux in (43) demonstrate the results of these rankings for an unindexed disyllabic stem (/pala/ 'burn'; Anttila 2002: 3) and an indexed one

(/taita/ from (20)), along with an indexed trisyllabic stem (/jumala+i+ssa/ 'God' (plural-inessive))

(43)

Input	Output	OO-MAX$_{\phi}$	IDENT-L3	MAX-L2	MAX	IDENT
/taita$_{L3}$-i$_{L1}$/	☞ taiti	*			*	
	☞ taitoi		*			*
/pala-i$_{L1}$/	pali	* !			*	
	☞ paloi					*
/jumala$_{L3}$-i$_{L1}$-ssa/	☞ jumalissa				*	
	jumaloissa		* !			*

Because the final vowel of the trisyllabic stem /jumala/ lies outside of a foot in the base form and is hence not subject to OO-MAX$_{\phi}$, indexing it to IDENT-L3 will categorically choose deletion. And as the tableaux in (40) show, indexation of a trisyllable to MAX-L2 chooses consistent mutation, while lack of indexation produces variation. For disyllables like /taita/ and /pala/, however, the options in (43) are the only available ones; consistent deletion requires a preceding round vowel.

This analysis cannot be translated into Anttila's (2002) theory. As noted above, because the partial ordering/lexical ranking theory allows lexical rankings to fix only grammatically unordered constraints, exceptionality and variation are conflated. To get variation for /taita-i/, in partial ordering/lexical ranking theory the constraints picking mutation and deletion would have to be unranked. But if they are unranked, then stems should be able to select a ranking of the constraint picking deletion over the one preferring mutation, resulting in the unattested pattern of consistent mutation in this environment. See Pater (2006) for discussion of further generalizations in the Finnish data that cannot be expressed with cophonologies.

4. Constraint Indexation as Inconsistency Resolution

In this section I propose an account of creation of lexically indexed constraints in terms of inconsistency resolution, and show that it can handle the Yine case, in which both faithfulness and markedness constraints must be appropriately indexed (see also Winslow 2003, Pater 2004, 2006, Becker 2006 and Tessier 2006 on indexation as inconsistency resolution, as well as Ota 2004 for discussion of Japanese postnasal voicing in similar terms).

As Tesar, Alderete, Horwood, Merchant, Nishitani and Prince (2003) point out in the context of lexical stress, exceptions to patterns of phonological alternation give rise to an inconsistent set of data, which can be detected by Tesar and Smolensky's (1998) Constraint Demotion Algorithm (CDA; see Tesar 1998, Prince 2002, McCarthy 2005, and Tesar and Prince 2004 for other applications of inconsistency detection). This is illustrated by the set of winner-loser pairs (Prince 2002) for the hypothetical language discussed in section 2.1, which appear in (44).

(44)

Input	W ~ L	NoCoda	Max
pak	pa ~ pak	W	L
lok	lo ~ lok	W	L
tak	tak ~ ta	L	W

Each row of the table presents the optimal form, or 'winner', paired with a suboptimal competitor or 'loser'. For each pair, the W's and L's in the constraint columns indicate whether a constraint prefers the winner or loser. A constraint ranking will correctly choose the winner if and only if every constraint preferring the loser is dominated by some constraint preferring the winner. The winner-loser pairs in (44) impose incompatible requirements on the ranking; /pak/ → [pa] requires NoCoda >> Max, and /tak/ → [ta] requires Max >> NoCoda. The CDA will fail to rank these constraints, and "declare inconsistency" (Prince 2002).

The recursive version of the CDA (RCD; Tesar and Smolensky 1998, 2000) is a particularly useful tool in diagnosing the locus of inconsistency. It starts by identifying all of the constraints that prefer only winners (that is, that prefer the winner or prefer neither the winner or the loser), and installs them in a stratum. It then eliminates all of the winner-loser pairs from the dataset in which the just-installed constraints prefer the winner (since the optimality of the winner is now guaranteed), and goes on to construct the next stratum in the same fashion. The procedure stops when all of the data are eliminated, or when no constraints prefer only winners. In the latter case, the residue of the dataset contains inconsistency (unless none of the constraints assign Ws to the remaining pairs, in which case the constraint set is inadequate for some other reason).

What happens after RCD detects inconsistency? For lexical stress, Tesar *et al.* (2003) propose that the lexical representation is altered, the winner-loser pairs are updated, and constraint demotion restarts. It is unlikely, however, that all instances of morpheme-specific alternation can be dealt with in terms of differences in lexical representation. And even in those cases in which a structural account is available, the search space of possible lexical changes is extremely large. Tesar *et al.* (2003) abstract from this problem by only considering changes in underlying stress, but if lexical "surgery" is disconnected from constraint ranking, it is not at all clear how the pattern of constraint violations can guide the change in underlying representation (see further section 5).

Here I suggest instead that when the constraint demotion algorithm can no longer find constraints that favor only winners, it seeks a constraint that favors only winners *for all instances of some morpheme*. It then ranks that constraint, indexed to all of the morphemes for which it favors only winners. In the simple case in (45), there are two such constraints: Max and NoCoda. It is perhaps inconsequential which is chosen, since an indexed version of either one will allow inconsistency to be resolved. If, however, it is taken as a goal to lexically index the smaller set of forms (i.e. the 'exceptional' ones), then a bias to a smaller set of indexed morphemes could be built in (see Winslow 2003, Pater 2004, and the Appendix), thus choosing Max to be the indexed constraint, as in (2) above.

The Yine case is more interesting in that the correct choice between constraints must be made for each morpheme in order to get the right results. For example, the following data must lead to morpheme-specific ALIGN for /-lu/ and morpheme-specific MAX for /-wa/, and no marking for the other morphemes

(45)	/heta+lu/	[hetlu]	'see it'
	/heta+nu/	[hetanu]	'going to see'
	/heta+nu+lu/	[hetanru]	'going to see him'
	/heta+wa+lu/	[hetawalu]	'going to see him yet'

When learning commences, there are only unindexed versions of the constraints:

(46) ALIGN-SUF-C, MAX

In (47) we see that the winner-loser pairs are inconsistent with one another.

(47)

Input	W ~ L	MAX	ALIGN-SUF-C
heta+lu	hetlu ~ hetalu	L (heta)	W (lu)
heta+nu	hetanu ~ hetnu	W (heta)	L (nu)
heta+wa+lu	hetawalu ~ hetawlu	W (wa)	L (lu)
heta+nu+lu	hetanlu ~ hetanulu	L (nu)	W (lu)

In this table, the relevant morpheme is indicated alongside the W and L marks. This information is needed to correctly identify the locus of inconsistency. A more explicit statement of the inconsistency resolution routine that makes use of this information appears in (48).

(48) i. Clone a constraint that prefers only Ws in all instances of some morpheme
ii. Index it to every morpheme for which it prefers only Ws

In the first step, we can only index MAX for /-wa/ - MAX prefers Ls and Ws for /heta/, and ALIGN-SUF-C prefers Ls and Ws for /lu/.

(49)

Input	W ~ L	MAX	ALIGN-SUF-C	MAX-L1
heta+lu	hetlu ~ hetalu	L	W	
heta+nu	hetanu ~ hetnu	W	L	
heta+wa_{L1}+lu	hetawalu ~ hetawlu	W	L	W
heta+nu+lu	hetanlu ~ hetanulu	L	W	

Since MAX-L1 prefers only winners, RCD places it in the first stratum. Since RCD eliminates winner-loser pairs after installing a constraint preferring the winner, the third pair of table (49) is removed, yielding the following partial hierarchy and table:

(50) MAX-L1 >>

Input	W ~ L	MAX	ALIGN-SUF-C
heta+lu	hetlu ~ hetalu	L	W
heta+nu	hetanu ~ hetnu	W	L
heta+nu+lu	hetanlu ~ hetanulu	L	W

In the above set of winner-loser pairs ALIGN-SUF-C now assigns only Ws for /lu/; MAX continues to assign a W and an L to /heta/. Table (51) shows the result of cloning and indexation:

(51) MAX-L1 >>

Input	W ~ L	MAX	ALIGN-SUF-C	ALIGN-SUF-C-L2
heta+lu_{L2}	hetlu ~ hetalu	L	W	W
heta+nu	hetanu ~ hetnu	W	L	
heta+nu+lu_{L2}	hetanlu ~ hetanulu	L	W	W

The RCD's stratum formation and winner-loser pair elimination now yields (52):

(52) MAX-L1 >> ALIGN-SUF(L2)-C

Input	W ~ L	MAX	ALIGN-SUF-C
heta+nu	hetanu ~ hetnu	W	L

RCD's next two iterations place all of the constraints in the hierarchy:

(53) MAX-L1 >> ALIGN-SUF-C-L2 >> MAX >> ALIGN-SUF-C

The ranking of the constraint *CCC will be established straightforwardly, given forms like /terkalu/ [terkalu] *[terklu], which show that *CCC dominates ALIGN-SUF-C:

(54)

Input	W ~ L	*CCC	ALIGN-SUF-C-L2
terka+lu_{L2}	terkalu ~ terklu	W	L

In fact, except for the special circumstances discussed in Lin (1997), from which I abstract here, *CCC is unviolated in Yine.[2] Thus, *CCC would be installed in the first step of the constraint demotion algorithm, since it prefers only winners, and the winner-loser pair in (54) would be immediately eliminated from consideration. The grammar constructed for Yine by RCD with inconsistency resolution would thus be as in (55).

(55) *CCC >> MAX-L1 >> ALIGN-SUF-C-L2 >> MAX >> ALIGN-SUF-C

The ranking of *CCC raises an important point. Given morphologically indexed faithfulness constraints, one might worry that the learner wouldn't bother with the phonological

[2] These special circumstances are in the context of an opaque interaction in which vowel deletion feeds consonant deletion and compensatory lengthening, as well in as clusters incorporating a mono-consonantal affix.

generalization at all. In the present case, why doesn't the learner just index /terka/ to MAX, and treat this as another morphological exception? The answer lies in the way that this modified RCDA works. RCD seeks to deal with the data in terms of a constraint ranking. Only when inconsistency prevents this from happening are indexed constraints created. In the case of forms like /terkalu/, the winner-loser pair in (54) will already have been eliminated from consideration by the ranking of *CCC before inconsistency resolution applies; the phonological explanation takes precedence over the lexical one (see also Tesar *et al.* 2003).

The prioritization of ranking over indexation is not in itself sufficient for the learner to capture all generalizations in patterns of alternation (note that this proposal is not meant to handle exceptions to phonotactic generalizations; see Pater 2005, Albright 2006, Coetzee and Pater 2007 and Hayes and Wilson to appear for recent discussion of the learning of gradient phonotactics). At the end of learning, there will be a set of indexed morphemes over which there may be both phonological and morphological generalizations that are not yet expressed. The phonological generalizations would be those that cannot be captured by the posited universal constraint set (see Becker 2006 for related discussion). The morphological generalizations would be ones in which morphemes of a particular lexical or morpho-syntactic class pattern together (see e.g. Ito and Mester 1995a,b, and 1999 on lexical strata, and Smith 1997 on noun faithfulness). For the latter, it may be that learners are biased to index constraints to information already present in the lexicon, since this would avoid the postulation of an arbitrary diacritic, and would presumably render the lexical encoding more robust (see Anttila 2002 on the emergence of morphological conditioning of exceptionality).

There are several other ways in which this learning proposal may be profitably developed. Here the data were presented to the learner all at once; an on-line learner would be more realistic (see Tesar 1998 and subsequent work on how inconsistency can be detected by an on-line learner). The implementation of this proposal in such terms would be straightforward; the only difference would be that the learner might clone more constraints than it would if it had access to all of the data. Another extension would be to variation. As Tessier (2006) discusses in detail, an alternative OT learning model that handles variation, the GLA of Boersma (1998) and Boersma and Hayes (2001), does not deal with the full range of patterns of exception, even with the supplementary proposals of Zuraw (2000). As Andries Coetzee (p.c.) has pointed out to me, one way of extending the current theoretical model to variation would be to treat a morpheme that displays variation as one that bears multiple indices. This may lead to a fairly simple elaboration of the learning model (see Tessier 2006 for an alternative). And finally, the procedure for selecting a constraint to clone likely needs to be further specified; see the Appendix for discussion.

5. Representational Approaches to Exceptionality

The lexical idiosyncrasies of Yine syncope could be reanalyzed in representational terms, without recourse to the lexical diacritics that indexed constraints require. Matteson's (1965: 36) convention for notating non-triggering morphemes in fact suggests such an analysis: that they possess an unspecified vocalic position (e.g. /-Vta/), whose presence blocks the application of syncope (see Wolf 2006 on how to derive this sort of blocking in OT). For the morpheme that fails to undergo deletion, one might lexically specify some aspect of syllable structure (e.g. a mora), which is protected by a faithfulness constraint (e.g. MAX-μ; see Inkelas, Orgun and Zoll 1997 and Inkelas 2000 for related proposals).

The main challenge for such an analysis is to specify how a learner would arrive at the representations needed to make it work. The most difficult cases are ones where the feature that must be specified in the non-alternating forms is not the only feature that changes in the alternating form. Taking the Yine example of exceptional blocking of syncope, it is not just the vocalic mora that is variably present in surface allomorphs of alternating vowels, but the entire vowel. For the analysis mentioned in the above paragraph to work, the learner must somehow be guided to choose the mora as the element to specify in non-alternating /-wa/, and leave absent from the alternating form. Presumably, the guide would be the presence of a faithfulness constraint that specifically targets the mora, and can protect it from deletion. While it is plausible that a learner might restructure an underlying representation so that it allows a constraint to choose the correct output, no extant learnability proposal allows for this. The process of surgery, as Tesar *et al.* (2003) term it, is grammar-blind. Inkelas (1994) proposes a version of lexicon optimization (Prince and Smolensky 1993/2004) that leads to underspecified representations. However, as Bruce Tesar (p.c.) notes, this proposal assumes the existence of a constraint ranking; it has yet to be incorporated into a learning model that also acquires a ranking. One might also limit the representational search space; see Tesar (2006) on contrast analysis. Ultimately, the success of learnability proposals, rather than some unstated aesthetic principle, should choose between representational and diacritic approaches to exceptionality, insofar as they can handle the same range of data. Since the pattern of constraint violations directly guides the required adjustment to the system in the inconsistency resolution account of constraint indexation, this seems a particularly promising approach.

The empirical scope of various approaches to exceptionality remains to be determined. It is unclear whether all cases of exceptionality can, or should, be given a diacritic treatment (see also Ito and Mester 2001, Albright 2002, and Becker 2004 for recent arguments for diacritic analyses; cf. Inkelas *et al.* 1997 on Turkish, though see Becker 2006). Instances where the markedness motivation for the alternation has disappeared, as in many cases of mutation (see Wolf 2006), seem particularly amenable to a representational analysis. And of course, non-surface-true underlying forms are required to deal with opacity (see McCarthy 2005). However, it is equally unclear whether a representational account of morpheme-specific phonology can be made fully general. It should be noted that even proponents of representational approaches in OT see a role for morpheme-specific constraints or rankings (see e.g. Inkelas 1999, Kager to appear).

As well as clarifying the relative scope of purely diacritic and structural approaches, further research should cast light on the viability of various intermediate positions (e.g. indexation limited to a subset of markedness constraints, or structural hypotheses limited to surface observable forms), and of alternative analyses of exceptionality like Alderete's (2001) Antifaithfulness theory, the Realize Morpheme theory of Kurisu (2001) and others, Kager's (to appear) extension of allomorphy theory, and Zuraw's (2000) extension of Boersma (1998) and Boersma and Hayes' (2001) stochastic OT (see also Hayes and Londe 2006).

6. On the Power of Diacritic Theories

It is most likely that representational and diacritic theories of exceptionality do not overlap completely in the set of attested phenomena that they can account for, and so the challenge will be to parcel out the explanatory burden appropriately. However, the three diacritic theories discussed in this paper do account for essentially the same phenomena, and so we can ask which of them seems most likely to account for all and only the attested types of exceptionality and of other morpheme-specific phonology. In earlier parts of the paper I have advanced specific arguments in favor of constraint indexation; here I attempt to address the more general question of whether the theories differ in the extent to which they allow different phonological patterns to co-exist in a single language.

It might seem obvious that a theory that allows both markedness and faithfulness constraints to be indexed is more powerful than one that allows only faithfulness indexation. However, since most phonological processes involve crucial rankings between markedness and faithfulness constraints, it is somewhat difficult to find phenomena that cannot be analyzed in terms of faithfulness indexation (for attested cases that argue for markedness indexation, see Pater 2000, 2006, Ota 2004, and Flack to appear; see also Gelbart 2005). Inkelas and Zoll (2003) provide a cogent demonstration that the restriction of indexation to faithfulness does not impose substantial limits on the set of patterns that can co-exist in a language. I know of just two problems. Indexed markedness constraints can produce the unattested templatic backcopying pattern that McCarthy and Prince (1995, 1999) avoid by eliminating templates from prosodic morphology (see Flack to appear; cf. Inkelas and Zoll 2005, Gouskova 2006). The other major worry with indexed markedness constraints is that if they can be relativized to a category like ROOT, then they can subvert a fixed ROOT-FAITH >> AFFIX-FAITH meta-ranking (McCarthy and Prince 1995), and produce a language that that neutralizes contrasts root-internally, but not in affixes (see Albright 2004, Tessier 2004, Beechey 2005, Urbanczyk 2006 for related discussion). How these issues are best addressed is a matter for further research; completely ruling out markedness indexation does not seem to be a viable approach.

Most versions of cophonology theory can express the same sorts of morphologically specific ranking reversals as a theory with indexation of both markedness and faithfulness constraints. Anttila's (2002) partial ordering/cophonology theory also seems nearly, if not completely equivalent. Insofar as there is no restriction on how many constraints can be left unordered, this theory does not impose any restrictions on what rankings can co-exist in a language. The only exception is that it cannot express the equivalent of the following type of indexed constraint grammar, where the lexically specific version of the constraint is separated by more than one stratum from its general counterpart:

(56) Con1-L >> Con2 >> Con3 >> Con1

As far as I know, the empirical force of this restriction has yet to be demonstrated.

Since all of the diacritic theories seem to allow an extremely wide range of rankings to co-exist in a single language, an obvious question is why we don't find languages that mix very different properties together. The main issue with this line of objection is that it's difficult to measure heterogeneity: how does one quantify the attested phonological differences between morphemes that we do find in a single language, versus those that do not co-exist? However, it does seem that languages maintain a certain degree of homogeneity, for example by regularizing exceptional patterns over time. This would receive a relatively natural explanation in indexation theory, if one assumes that constraint clones and indices add complexity to the system that is sometimes eliminated over the course of time.

7. Conclusions

In this paper, I have shown that several considerations favor a theory of morpheme-specific phonology with indexed markedness and faithfulness constraints over either a theory with only indexed faithfulness constraints, or cophonology theory. Unlike these alternatives, it allows for a distinction between exceptional triggering and blocking, which was shown to be necessary to analyze the Yine data. Given the locality convention introduced here, indexed markedness constraints allow for a straightforward analysis of the locus of Finnish /ai/ allomorphy and Assamese exceptional ATR harmony, and also resolve general locality problems for morpheme-specific phonology. And finally, indexed constraints allow for distinctions between exceptional and impossible patterns, and between exceptionality and variation.

I have also provided a learnability account in which indexed constraints are created to deal with inconsistency in the learning data. This is a relatively straightforward extension of Tesar and Smolensky's (1998, 2000) RCDA, and avoids some of the issues faced by learnability analyses posited for representational theories of exceptionality. In addition, this learnability proposal provides an explicit resolution of an apparent theoretical incongruity. The existence of morpheme-specific constraints is sometimes seen as incompatible with a theory in which constraints are universal (see e.g. Green 2005): is the Tagalog constraint ALIGN([UM]AF, L, STEM, L) proposed by McCarthy and Prince (1993) present in all grammars? Under this proposal, morpheme-specific constraints are constructed from universal constraints in the course of learning.

Appendix: Analyzing Inconsistency

The learnability proposal in sec. 4 remains to be fully elaborated and implemented. The main issue that arises is how to choose between constraints to clone when more than one is selected by the criterion in (48i.), repeated in (A1i.).

(A1) i. Clone a constraint that prefers only Ws in all instances of some morpheme
ii. Index the cloned constraint to every morpheme for which it prefers only Ws

Recursive Constraint Demotion (RCD) will usually winnow down the constraint set considerably by the time inconsistency is encountered. However, as members of the Rutgers Optimality Research Group have pointed out to me, RCD may stall at a point when constraints not involved in the inconsistency pattern remain to be installed. Such a situation is illustrated schematically in (A2):

(A2)

Winner~Loser	CON1	CON2	CON3	CON4
a~b	W	L		
c~d	L	W	L	
e~f			W	L

The problem for the constraint indexation proposal is the following. Assuming that the inputs underlying 'a~b', 'c~d', and 'e~f' are all separate morphemes, CON1, CON2 and CON3 could all be cloned under the definition in (A1), but cloning CON3 would not resolve inconsistency. One possibility would be to simply allow the learner to choose randomly, since it would eventually find a useful constraint to clone. However, a guided approach is likely preferable. Here I address only the problem of finding useful constraints; see Becker (2006) and Pater (2006; sec. 3.3) on choosing between useful constraints to ensure restrictiveness. Before proceeding, I should emphasize that the analysis of inconsistency is important not only for this learnability proposal, but more generally for diagnosing the source of failure when a constraint set fails to capture a data pattern.

The task of identifying the locus of inconsistency is complicated by the fact that more than two constraints, and more than two winner-loser pairs, can be involved in creating an inconsistent pattern of W and L marks. The following table illustrates.

(A3)

Winner~Loser	CON1	CON2	CON3
a~b	W	L	
c~d		W	L
e~f	L		W

The pattern of marks in (A3) can be extended to create an "irreducibly inconsistent system" (IIS; van Loon 1981) with a set of winner-loser pairs of any size. IIS is a term used in the analysis of inconsistent systems in Linear Programming; see Pater, Potts and Bhatt (2006) on the connection between Linear Programming and the constraint demotion algorithm. A definition adapted to OT from Chinneck and Dravnieks (1991) appears in (A4) (readers that consult the original definition should note that the term "constraint" in Linear Programming refers to a linear inequality, not a constraint in the linguistic sense).

(A4) *Irreducibly Inconsistent System (in OT)*
An irreducibly inconsistent system is a minimal inconsistent set of Winner-Loser pairs, that is, an inconsistent set that becomes consistent when one Winner-Loser pair is removed.

To find an IIS, one can apply Chinneck and Dravnieks' (1991) technique of Deletion Filtering, in which the following procedure is applied to each Winner-Loser pair until they are all tested (see Chinneck and Dravnieks for a proof that Deletion Filtering yields an IIS):

(A5) *Deletion Filtering (in OT)*
i. Temporarily remove a Winner-Loser pair from the inconsistent set
ii. If the remaining set is consistent, return the Winner-Loser pair to the set; if the remaining set in inconsistent, remove the Winner-Loser pair permanently

This procedure would eliminate the final row of the table in (A2), leaving W marks only in the columns of the constraints involved in the inconsistency pattern, thus providing the desired input for (A1).

However, it may be that the goals of finding useful constraints to clone, and of analyzing inconsistent OT systems for other purposes, do not require this brute force method. The procedure I term Analytic Filtering (A6) takes as its first step the elimination of constraints that cannot be contributing to inconsistency (many other such constraints - those that prefer only winners - would have already been installed by RCD). The second step pares down the resulting winner-loser set. Analytic Filtering reduces the table in (A2) to an IIS containing just the first two winner-loser pairs, and will leave the table in (A3) untouched.

(A6) *Analytic Filtering*
i. Remove constraint columns containing only Ls
ii. Remove Winner-Loser pair rows containing only Ws

Though Deletion Filtering and Analytic Filtering yield similar results for cases (A2) and (A3), there are at least two ways in which the results of these two procedures will diverge. Analytic Filtering will not eliminate Winner-Loser pairs that duplicate the same portion of a pattern of inconsistency, as in (A7). Since either of the last rows could be removed without yielding a consistent set, this is not yet an IIS.

(A7)

Winner~Loser	CON1	CON2
a~b	W	L
c~d	L	W
e~f	L	W

Another situation in which Analytic Filtering will fail to yield an IIS is when there are multiple patterns of inconsistency in the residue of RCD, as in the following example:

(A8)

Winner~Loser	CON1	CON2	CON3
a~b	W	L	
c~d	L	W	L
e~f		L	W

Here either the first or third row would have to be removed to yield an IIS.

It is not clear that the further reduction performed by Deletion Filtering in these cases is either necessary, or useful, for an OT learner or analyst. For example, if it is taken as a goal to index the smallest set of morphemes, then the information that Con1 assigns Ws to fewer morphemes than Con2 in (A7) is crucial, and would be lost through Deletion Filtering. And if it is taken as a goal to clone the smallest number of constraints, then (A8) contains the information that Con2 participates in both patterns of inconsistency, which would also be lost after Deletion Filtering. However, other cases of inconsistency may well be better analyzed using Deletion Filtering or some other method; how to best perform inconsistency analysis in OT remains an open research question.

References

Albright, A. (2002) A restricted model of UR discovery: Evidence from Lakhota. Ms, University of California at Santa Cruz.

Albright, A. (2004) The emergence of the marked: Root-domain markedness in Lakhota. Handout from the meeting of the Linguistic Society of America.

Albright, A. (2006) Gradient phonotactic effects: lexical? grammatical? both? neither? Handout from the meeting of the Linguistic Society of America.

Alderete, J. (2001) Dominance effects as Transderivational Anti-Faithfulness. *Phonology* 18: 201--253.

Anttila, A. (1997) Deriving variation from grammar. In F. Hinskens, R. van Hout and W. L. Wetzels (eds.) *Variation, Change and Phonological Theory* 35--68. Amsterdam: John Benjamins.

Anttila, A. (2002) Morphologically conditioned phonological alternations. *Natural Language and Linguistic Theory* 20: 1: 1--42.

Bakovic, E. (2000) *Harmony, Dominance and Control*. Doctoral dissertation, Rutgers University, New Brunswick, New Jersey. ROA-360

Becker, M. (2004) Lexical exceptions: putting the grammar to work. Handout of a talk presented at HUMDRUM, Rutgers University.

Becker, M. (2006) Learning the lexicon through the grammar - Evidence from Turkish. Ms., University of Massachusetts, Amherst.

Beechey, T. (2005) A non-representational theory of geminate inalterability and integrity. Ms., University of Massachusetts, Amherst.

Bendor-Samuel, J. T. (1960) Some problems of segmentation in the phonological analysis of Terena. *Word* 16: 348--355.

Benua, L. (2000) *Phonological Relations Between Words*. New York: Garland Press. ROA-259

Boersma, P. (1998*) Functional Phonology: Formalizing The Interactions Between Articulatory And Perceptual Drives*. Doctoral dissertation, University of Amsterdam, the Netherlands.

Boersma, P. and B. Hayes (2001) Empirical tests of the Gradual Learning Algorithm. *Linguistic Inquiry* 32: 45--86. ROA-348

Caballero Hernández, G. (2005) The stress system of Central Raramuri: Root privilege, prosodic Faithfulness and Markedness reversals. Ms, University of California, Berkeley. ROA-706

Casali, R. (1997) Vowel elision in hiatus contexts: Which vowel goes? *Language* 73: 493--533.

Chinneck, J.W. and E.W. Dravnieks (1991) Locating minimal infeasible constraint sets in Linear Programs. *ORSA Journal on Computing*. 3/2: 157--168.

Coetzee, A. and J. Pater (2007) Weighted constraints and gradient restrictions on place co-occurrence in Muna and Arabic. Ms, University of Michigan and University of Massachusetts, Amherst.

Finley, S. (2006) Exceptions in vowel harmony are local. Ms., Johns Hopkins University, Baltimore, Maryland.

Flack, K. (To appear) Templatic morphology and indexed markedness constraints. In *Linguistic Inquiry*.

Fukazawa, H. (1999) *Theoretical Implications of OCP Effects on Features in Optimality Theory*. Doctoral dissertation, University of Maryland, College Park. ROA-307

Gelbart, B. (2005) *The Role of Foreignness in Phonology and Speech Perception.* Doctoral dissertation, University of Massachusetts, Amherst.

Golston, C. (1996) Direct OT: Representation as pure markedness. *Language* 72, 713--48. ROA-71.

Gouskova, M. (2003) *Deriving Economy: Syncope in Optimality Theory*. Doctoral dissertation, University of Massachusetts, Amherst. ROA-610.

Gouskova, M. (2006) A templatic constraint in Tonkawa. Ms, New York University.

Green, A. Dubach. (2005) *Phonology Limited*. Ms, Centre for General Linguistics, Typology and Universals Research (ZAS), Berlin. ROA-745

Hammond, M. (1995) There is no lexicon! Ms, University of Arizona. ROA-45.

Hayes, B. (2004) Phonological acquisition in Optimality Theory: the Early Stages. In R. Kager, J. Pater and W. Zonneveld (eds.) *Constraints in Phonological Acquisition*. Cambridge University Press. ROA-327

Hayes, B., and Z. Londe (2006) Stochastic phonological knowledge: the case of Hungarian vowel harmony. *Phonology* 23: 59--104.

Hayes, B. and C. Wilson. (To appear) A Maximum Entropy Model of Phonotactics and Phonotactic Learning. In *Linguistic Inquiry*. ROA-858

Horwood, G. (1999) Anti-faithfulness and subtractive morphology. Ms, Rutgers University, New Brunswick, NJ. ROA-466.

Inkelas, S. (1994) The consequences of optimization for underspecification. In *Proceedings of NELS 25*, ed. Jill Beckman, 287--302. Amherst, Mass: GLSA Publications. ROA-40

Inkelas, S. (1999) Exceptional Stress-Attracting Suffixes in Turkish: Representations vs. the Grammar. In *The Prosody-Morphology Interface*. ed. Harry van der Hulst, René Kager, and Wim Zonneveld, 134--187, Cambridge University Press. ROA-39

Inkelas, S. (2000) Phonotactic blocking through structural immunity. In B. Stiebels and D. Wunderlich (eds.), *Lexicon in Focus. Studia grammatica 45*, 7-40. Berlin: Akademie Verlag. ROA-366.

Inkelas, S., O. Orgun and C. Zoll. (1997) The implications of lexical exceptions for the nature of grammar. In I. Roca ed. *Derivations and Constraints in Phonology*, 393--418. New York: Oxford University Press.

Inkelas, S. and C. Zoll (2003) Is Grammar Dependence Real? Ms, UC Berkeley and MIT. ROA-587.

Inkelas, S. and C. Zoll (2005) *Reduplication: Doubling in Morphology*. Cambridge University Press.

Ito, J. and A. Mester (1995a) The core-periphery structure in the lexicon and constraints on re-ranking. In J. Beckman, L. Walsh Dickey and S. Urbanczyk (eds.) *Papers in Optimality Theory*, 181--210. GLSA, University of Massachusetts, Amherst.

Ito, J. and A. Mester (1995b) Japanese phonology. In J. Goldsmith (ed.) *The Handbook of Phonological Theory*, 817--38. Oxford: Blackwell.

Ito, J. and A. Mester (1999) The Phonological Lexicon. In N. Tsujimura (ed.) *The Handbook of Japanese Linguistics*, 62--100. Oxford: Blackwell.

Ito, J. and A. Mester (2001) Covert generalizations in Optimality Theory: the role of stratal faithfulness constraints. *Studies in Phonetics, Phonology, and Morphology* 7: 273--299.

Kager, R. (To appear) Lexical irregularity and the typology of contrast. In *The Structure of Words*, ed. K. Hanson and S. Inkelas. MIT Press.

Karlsson, F. (1982) *Suomen kielen äänne- ja muotoraknne [The Phonological and Morphological Structure of Finnish]*. Helsinki: Werner Söderström Osakeyhtiö.

Kaisse, E. and P. Shaw (1985) On the theory of Lexical Phonology. *Phonology Yearbook* 2*:10-–30.*

Kiparsky, P. (1993) Blocking in non-derived environments. In S. Hargus and E. Kaisse (eds.) *Studies in Lexical Phonology*. San Diego: Academic Press.

Kirchner, R. (1993) Turkish vowel harmony and disharmony: An Optimality Theoretic account. Ms., UCLA. ROA-4

Kisseberth, C. (1970) The treatment of exceptions. *Papers in Linguistics* 2: 44--58.

Kraska-Szlenk, I. (1997) Exceptions in phonological theory. *Proceedings of the 16th International Congress of Linguists.* Paper No. 0173. Oxford: Pergamon.

Kraska-Szlenk, I. (1999) Syllable structure constraints in exceptions. In J. Rennison and K. Kühnhammer (eds.) *Phonologica* 1996: *Syllables !?*, 113--131. The Hague: Thesus.

Kurisu, K. (2001) *The Phonology of Morpheme Realization*. Doctoral Dissertation, University of California Santa Cruz.

Kurisu, K. (2006) Weak Derived Environment effect. Handout from NELS 37.

Lin, Y.-H. (1997) Syllabic and moraic structures in Piro. *Phonology* 14: 403-436.

Lubowicz, A. 2002. Derived environment effects in Optimality Theory. *Lingua* 112: 243-280.

Mahanta, S. (To appear). *Directionality and Locality in Vowel Harmony*. Doctoral dissertation, Utrecht University, the Netherlands.

Mascaró, J. (2003) Comparative markedness and derived environments*. Theoretical Linguistics* 29. 113–-122.

Matteson, E. (1965) *The Piro (Arawakan) language*. Berkeley: University of California Press.

McCarthy, J. (2002) Comparative markedness. In A. Carpenter, A. Coetzee, and P. de Lacy (eds.) *Papers in Optimality Theory II*, 147--246. GLSA, University of Massacusetts, Amherst.

McCarthy, J. (2005) Taking a Free Ride in morphophonemic learning. *Catalan Journal of Linguistics* 4: 19--56. ROA-683.

McCarthy, J. and A. Prince (1993) Generalized Alignment. *Yearbook of Morphology*: 79–-153.

McCarthy, J. and A. Prince. (1995) Faithfulness and reduplicative identity. In J. Beckman, S. Urbanczyk and L. Walsh Dickey. *Papers in Optimality Theory*, 249–-384. GLSA, University of Massacusetts, Amherst. ROA-60

McCarthy, J. and A. Prince (1999) Faithfulness and identity in prosodic morphology. In H. van der Hulst, R. Kager, and W. Zonneveld (eds.) *The Prosody-Morphology Interface*, 218--309. Cambridge University Press. ROA-216

Nouveau, D. (1994) *Language Acquisition, Metrical Theory, and Optimality: A Study of Dutch Word Stress*. Doctoral dissertation, Utrecht University.

Orgun, C. O. (1996) *Sign-based morphology and phonology: with special attention to Optimality Theory.* Doctoral dissertation, University of California, Berkeley. ROA-171

Ota, M. (2004) The learnability of the stratified phonological lexicon. *Journal of Japanese Linguistics* 20, 4: 19-40 ROA-668

Pater, J. (2000) Nonuniformity in English stress: the role of ranked and lexically specific constraints. *Phonology* 17: 237--274.

Pater, J. (2004) Exceptions in Optimality Theory: Typology and Learnability. Handout from the Conference on Redefining Elicitation: Novel Data in Phonological Theory, New York University. (Available at http://people.umass.edu/pater/exceptions.pdf)

Pater, J. (2005) Learning a stratified grammar. In A. Brugos, M. R. Clark-Cotton, and S. Ha. (eds.) *Proceedings of the 29th Boston University Conference on Language Development*, 482-492. Cascadilla Press, Somerville, MA. ROA-739.

Pater, J. (2006) The Locus of Exceptionality: Morpheme-Specific Phonology as Constraint Indexation. In L. Bateman, M. O'Keefe, E. Reilly, and A. Werle (eds.) *University of Massachusetts Occasional Papers in Linguistics 32: Papers in Optimality Theory III*, Amherst: GLSA. (Available at http://people.umass.edu/pater/locus.pdf)

Pater, J. C. Potts, and R. Bhatt (2006) Harmonic Grammar with Linear Programming. Ms., University of Massachusetts, Amherst. ROA-872

Prince, A. (2002) Arguing optimality. In A. Coetzee, A. Carpenter, and P. de Lacy (eds.) *Papers in Optimality Theory II*, 269–-304. Amherst, MA: GLSA. ROA-562.

Prince, A., and P. Smolensky (1993/2004) *Optimality Theory: Constraint interaction in generative grammar*. Technical Report, Rutgers University and University of Colorado at Boulder, 1993. Revised version published by Blackwell, 2004.

Prince, A. and B. Tesar (2004) Learning phonotactic distributions. In R. Kager, J. Pater and W. Zonneveld (eds.) *Constraints in Phonological Acquisition,* 245--291. Cambridge University Press.

Russell, K. (1995) Morphemes and candidates in Optimality Theory. Ms, University of Manitoba. ROA-44.

Sadeniemi, M., ed. (1973) Nyksuomen sanakirja [Dictionary of Modern Finnish]. Porvoo, Finland, Werner Söderström Osakeyhtiö.

Smith, J. (1997) Noun faithfulness: On the privileged behavior of nouns in phonology. Ms., University of Massachusetts, Amherst. ROA-242.

Smolensky, P. (1996) The initial state and 'Richness of the Base' in Optimality Theory. Technical Report, Department of Cognitive Science, Johns Hopkins University. ROA-154.

Tesar, B. (1998) Using the mutual inconsistency of structural descriptions to overcome ambiguity in language learning. In P. Tamanji and K. Kusumoto (eds.) *Proceedings of NELS 28*, 469--483. Amherst, MA: GLSA, University of Massachusetts.

Tesar, B. (2006) Learning from Paradigmatic Information. *Wondering at the Natural Fecundity of Things: Essays in Honor of Alan Prince*. Paper 14. Linguistics Research Center. available from http://repositories.cdlib.org/lrc/prince/14

Tesar, B., J. Alderete, G. Horwood, N. Merchant, K. Nishitani, and A. Prince. (2003) Surgery in language learning. In G. Garding and M. Tsujimura (eds.) *Proceedings of the Twenty-Second West Coast Conference on Formal Linguistics*, 477--490. Somerville, MA: Cascadilla Press. ROA-619

Tesar, B., and A. Prince. (2004) Using Phonotactics to Learn Phonological Alternations. In *Proceedings of CLS 39, Vol. II: The Panels*.

Tesar, B. and P. Smolensky. (1998) Learnability in Optimality Theory. *Linguistic Inquiry* 29. 229--268. ROA-155

Tesar, B. and P. Smolensky. (2000) *Learnability in Optimality Theory.* Cambridge, MA: MIT Press. ROA-156

Tessier, A.-M. (2004) Root-Restricted Markedness and Morpho-Phonological Domains. Handout from the Montreal-Ottawa-Toronto Phonology Workshop.

Tessier, A.-M. (2006) *Biases and stages in phonological acquisition*. Doctoral dissertation, University of Massachusetts, Amherst. ROA-883

Urbanczyk, S. (2006) Reduplicative form and the root-affix asymmetry. *Natural Language and Linguistic Theory*. 179--240.

Winslow, N. (2003) *Incorporating Exceptions in Optimality Theoretic Learnability*. Honor's Thesis, University of Massachusetts, Amherst.

Wolf, M. (2006) For an autosegmental theory of mutation. In *University of Massachusetts Occasional Papers in Linguistics 32: Papers in Optimality Theory III*, ed. L. Bateman, M. O'Keefe, E. Reilly, and A. Werle. Amherst: GLSA. ROA-754

Zamma, H. (2005) Predicting varieties: Partial orderings in English stress assignment. Ms, Kobe City University of Foreign Studies/University College London.

van Loon, J. (1981) Irreducibly inconsistent systems of linear inequalities. *European journal of operations research* 8, 263-288.

Zonneveld, W. (1978) *A Formal Theory of Exceptions in Generative Phonology*. Lisse: The Peter de Ridder Press.

Zonneveld, W. (1985) Environmental concerns in the study of phonology. *International Journal of American Linguistics*, 5/4: 626-628.

Zuraw, K. (2000) *Patterned Exceptions in Phonology*, Doctoral dissertation, University of California, Los Angeles. ROA-788