Modular Interactions in Phonology

by

Ezer Rasin

Submitted to the Department of Linguistics and Philosophy
in partial fulfillment of the requirements for the degree of

Doctor of philosophy

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 2018

Signature redacted

Author ....
Department of Linguistics and Philosophy
September 7, 2018

Signature redacted

Certified by ....
Donca Steriade
Professor of Linguistics
Thesis Supervisor

Signature redacted

Accepted by ....
Adam Albright
Professor of Linguistics
Linguistics Section Head

Modular Interactions in Phonology

by

Ezer Rasin

Submitted to the Department of Linguistics and Philosophy
on September 7, 2018, in partial fulfillment of the
requirements for the degree of
Doctor of philosophy

Abstract

This thesis makes two separate claims about the architecture of phonology:

(1) The computation of stress takes place in a distinct cognitive module from segmental phonology. This module is informationally encapsulated from segmental features.

(2) Phonological generalizations over underlying representations can be captured in the lexicon.

The claim in (1) suggests a departure from a consensus view in generative phonology since the 1950's. According to this view, multiple phonological computations, including the computation of word stress and segmental processes, are carried out in a single cognitive module known as *phonology*. In Chapter 1 I challenge this view in two steps. I first argue for a new phonological universal based on the stress patterns of around 400 languages:

(3) STRESS-ENCAPSULATION UNIVERSAL: the distribution of stress is never directly conditioned by segmental features.

After reanalyzing reported counterexamples to the universal, I argue for an account of the universal in terms of a modular decomposition of phonology along the lines of (1).

The claim in (2) suggests a return to the architecture of early generative phonology, in which phonological generalizations could be captured in the lexicon (using constraints on underlying representations) as well as in the mapping from underlying representations to surface forms. Most recent work in phonology has abandoned that architecture, taking the lexicon to be merely a storage place for lexical items. Chapter 2, written jointly with Roni Katzir, presents an argument for constraints on underlying representations from learnability. In Chapter 3 I develop a new theory of blocking in non-derived environments, a phenomenon that has posed a long-standing puzzle for phonological theory since the 1970's. I argue that the new theory, which relies on constraints on underlying representations, offers a better account of the phenomenon than its predecessors.

Thesis Supervisor: Donca Steriade
Title: Professor of Linguistics

Acknowledgments

I am primarily indebted to my thesis committee: Donca Steriade (my committee chair), Adam Albright, Naomi Feldman, Roni Katzir, and Michael Kenstowicz. This incredible committee has provided me with more guidance and support than I could wish for. They have set professional standards that I will always aspire to, with respect to attention to data, clarity of presentation, respect to old or unpopular ideas, detailed knowledge of an unreasonable number of phonological phenomena, consideration of diverse perspectives on a problem, and so much more.

I am particularly grateful to Donca Steriade. I first met Donca in March 2013, when I was a prospective student visiting MIT. What was supposed to be a brief introductory meeting turned into a few-hour-long argument about certain aspects of phonological theory. That meeting made it clear to me that I should be advised by Donca, and it suggested what kind of advisor Donca would be: she would forcefully object to anything I say, but also temporarily take my side to help advance my ideas (even when they seem atrocious to her), and she would spend a significant amount of her free time doing the above. Fortunately for me, our arguments lasted for five more years, and our first meeting was indicative of the years to come. If this thesis had met Donca's standards for research, it would have had more value. Without her objections, insistence that I get everything right, and her continued support, this thesis would not have taken form at all.

Though David Pesetsky and Edward Flemming were not on my committee, they provided valuable feedback on various parts of this thesis. Many of the ideas presented in Chapter 1 and Chapter 3 were developed during meetings with David, especially ideas about architectural issues. Edward provided several crucial and insightful comments on Chapter 3.

This thesis consists of three largely-independent chapters, which benefited from the help and feedback of many people in the linguistics community, both inside and outside MIT. For discussions or feedback on Chapter 1, I thank Rafael Abramovich, Daniel Asherov, Outi Bat-El, Michael Becker, Jonathan Bobaljik, Sylvain Bromberger, Danny Fox, Morris Halle, Bruce Hayes, Paul de Lacy, Tim O'Donnell, Kevin Ryan, Shu-hao

Shih, Benjamin Storme, Bert Vaux, Eva Zimmermann, Sam Zukoff, and the audiences at GLOW 2017, Leipzig University, NAPhC 9, Rutgers University, Stony Brook University, and Tel Aviv University.

Chapter 2 was jointly written with Roni Katzir. For that reason, it is the most thoughtful and well-written chapter in this thesis. Working with Roni has taught me more than I thought was possible about every aspect of scientific work, and knowing that he is always there to provide the most helpful and insightful advise about anything has kept me going. I owe him much more than that, but properly thanking him would have doubled the length of this thesis. Instead, I will make one wish for the future of my career, which is that Roni will still be there.

For help with the Bengali data and judgments presented in Chapter 2, thanks to Neil Banerjee and Ishani Guha. For help with Dutch, thanks to Loes Koring, Marc van Oostendorp, and Coppe van Urk.

For useful feedback on Chapter 3, I thank Arto Anttila, Eric Baković, Michel DeGraff, Morris Halle, Bill Idsardi, Joan Mascaró, Roger Schwarzschild, Benjamin Storme, Jochen Trommer, and Sam Zukoff.

Various linguists have influenced me by providing counsel, by advising me on projects that did not make it into the thesis, or through their teaching. I especially thank Bob Berwick, Danny Fox, Martin Hackl, Irene Heim, David Pesetsky, and Roger Schwarzschild at MIT, and Outi Bat-El and Tal Siloni at Tel Aviv University.

I thank Chrissy Graham, Mary Grenham, Jen Purdy, and Matt Sikorski for turning the administrative part of academia from intimidating to friendly and welcoming.

For keeping me sane, well-fed, and happy during my time at MIT, I thank Daniel Asherov, Itai Bassi, Isa Bayırlı, Ömer Demirok, Sarah Edri, Suzana Fong, Sophie Moracchini, Milena Sisovics, Benjamin Storme, and especially Athulya Aravind, Aron Hirsch, Daniel Margulis, and Michelle Yuan. Daniel Margulis deserves special thanks for creating a place in a foreign country that I was able to call home.

Finally, for supporting me from a distance, I thank Iddo Berger, Dror Levin, Osnat Nixon, and my close family: my brother Shahaf Rasin and my parents Elana and Yaakov Rasin, who allowed this thesis to happen.

Contents

Chapter 1

The Stress-Encapsulation Universal and Phonological Modularity

1.1 Overview

1.1.1 The Stress-Encapsulation Universal

The distribution of segmental features is often conditioned by the position of stress. In American English, for example (and simplifying), /t/ is flapped between a preceding stressed vowel and a following unstressed vowel (políɾical, politícian), voiceless stops are aspirated at the onset of a stressed syllable, (opp^{h}óse, opposítion), stressless vowels undergo reduction (átəm, ətómic), and /h/ is deleted before an unstressed, non-initial vowel (véhicle, vehícular) (see Chomsky and Halle 1968, Kahn 1976, Borowsky 1986, and Davis and Cho 2003, among many others). Such stress-sensitive segmental processes are commonly attested across the world's languages, and they are many and diverse, as shown by the list in (4).

(4) Types of stress-sensitive segmental processes (Gonzalez 2003, Giavazzi 2010, and references therein)

a. Processes affecting consonantal features: affrication, aspiration, deletion, devoicing, flapping, fricativization, glottalization, glottalization-attraction, metathesis, occlusivization, voicing

b. Processes affecting vocalic features: lowering, reduction, vowel harmony (including metaphony, umlaut)

c. Other processes: nasal harmony

As noted by de Lacy (2002) and Blumenfeld (2006), stress-segmental interactions in the opposite direction are almost non-existent. While stress is sensitive to suprasegmental features such as length, syllable structure, and tone, it is arguably never sensitive to segmental features such as aspiration, continuancy, stridency, anteriority, place of articulation, laterality, rhoticity, nasality, rounding, and so on. For example, no language is known to have stress patterns like the following:

(5) a. Stress the leftmost round vowel

b. Stress the penultimate syllable, but if it has an unaspirated onset, stress the antepenultimate syllable

The segmental property that stands apart from the rest is vowel sonority. A literature on so-called 'sonority-driven stress' that goes at least back to Kenstowicz (1997) has documented multiple stress patterns in which the position of stress is determined by the hierarchy in (6). According to this hierarchy, lower vowels are more sonorous than higher vowels and peripheral vowels are more sonorous than central vowels. Kobon (Kenstowicz 1997, Davies 1981) provides an example of a stress pattern that reportedly makes full use of the sonority hierarchy and displays a five-way distinction between vowels in determining stress placement (7).

(6) Vowel sonority hierarchy (Kenstowicz 1997)

a > o, e > u, i > ə > ɨ

(7) Kobon stress in Kenstowicz (1997)

Stress falls on the more sonorous vowel among the final two vowels, according to the sonority hierarchy in (6)

Assuming the existence of sonority-driven stress as an exception, Blumenfeld (2006) treated the universal asymmetry between stress and segmental features as a list of specific universals, one for every segmental feature but sonority:[1]

(8) Blumenfeld's list of universals:

a. The distribution of stress is never conditioned by aspiration

b. The distribution of stress is never conditioned by continuancy

c. The distribution of stress is never conditioned by stridency

d. ...

Continuing a line of work by Hargus (2001), Blumenfeld (2006), Canalis (2007), de Lacy (2013), and Shih (2016), I re-evaluate the evidence for sonority-driven stress. My main claim in this chapter is that reported patterns of sonority-driven stress do not in fact require direct reference to sonority, either because they have been mis-analyzed or because they can be reanalyzed without reference to sonority. If this claim is correct, the result is that Blumenfeld's list of universal asymmetries between stress and segmental features becomes a generalization over all segmental features. This generalization is given in (9) as the *Stress-Encapsulation Universal.*

(9) *The Stress-Encapsulation Universal*

The distribution of stress is never conditioned by segmental features

1.1.2 The Modularity Hypothesis

Apart from establishing the Stress-Encapsulation Universal, my second goal in this chapter is to propose a phonological architecture from which the universal can be derived.

Note, first, that the universal is surprising under existing theories of phonology. Rule-based theories of stress (e.g., Halle and Vergnaud 1987a, Idsardi 1992, Hayes 1995) have assumed a representational separation between stress and segmental features following Liberman and Prince (1977), who argued that the principles that govern the distribution of stress are fundamentally different from those that govern the distribution of segmental features. This view is illustrated in Figure 1-1 in which stress is represented on a separate

[1] A few potential counterexamples to Blumenfeld's universals are discussed in section 1.5.7.

plane and the planes intersect. The planar architecture does not predict an asymmetry between stress and segmental features: regardless of what the content of the planes is and regardless of how one interprets intersection, this architecture is completely symmetric. Intersection is a symmetric relation – if A intersects with B then B intersects with A – so there is no reason to expect any sort of asymmetric encapsulation given this architecture. Indeed, rule-based theories of stress have used rules that make direct reference to segment quality, and even if reference to segment quality can be avoided, the fact that stress rules would consistently ignore the same information in their input would be left as an accident.

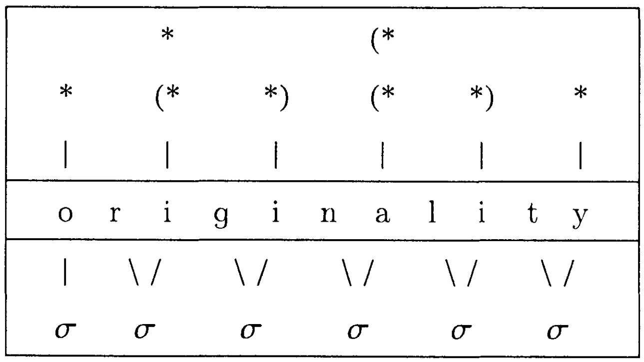

Figure 1-1: Planar architecture of phonology (modeled after a diagram in Halle 1998). The stress plane (top) intersects with the syllable plane (bottom) at the level of segmental representation (middle).

In Optimality Theory (OT; Prince and Smolensky 1993), stress and segmental processes are computed in parallel, and markedness constraints that trigger stress-sensitive segmental processes are symmetric and may be used to trigger quality-sensitive stress. An example is the markedness constraint in (10), which is a simplified version of the constraint that would trigger aspiration in English. This constraint can be satisfied by aspirating a prevocalic voiceless stop, but it can alternatively be satisfied by shifting stress away to a vowel that is not preceded by an unaspirated voiceless stop. Given such constraints, OT has no general way of banning quality-sensitive stress processes, as I discuss later in more detail.

(10) *tV́ = *unaspirated voiceless stop before a stressed vowel

I will show that the Stress-Encapsulation Universal can be derived in an architecture where the computation of stress has no access to segmental features. Information encapsulation of this kind is a hallmark of modular cognitive architectures, and I will use it to argue for a simple decomposition of phonology into modules that can capture the universal (cf. Scheer 2016). The hypothesis, which I refer to as *The Modularity Hypothesis*, is given in (11). The stipulation in (11a) is meant to ensure that computations carried out in the stress module do not refer to segmental features. But (11a) is not enough. Segmental processes that rely on the position of stress require access to stress representations, implying that stress representations must be available wherever segmental processes are computed. The stipulation in (11b) will ensure that access to stress is not exploited outside of the stress module to manipulate stress representations with reference to segmental features.[2] As we will see shortly, the main component of the modular architecture that restricts the interaction between stress and segmental features is the interface. A concrete theory of the interface to the stress module that specifies what information stress can access will determine the range of possible stress-segmental interactions.

(11) *The Modularity Hypothesis*

Stress is computed in an informationally encapsulated module with the following properties:

a. The input to the stress module excludes representations of segmental features

b. Outside of the stress module, stress representations cannot be changed

The move from Blumenfeld's list of universals in (8) to the Modularity Hypothesis in (11) would be a desirable theoretical result, since it eliminates a list of specific stipulations from the theory and replaces them with a simple statement about information encapsulation. It thus achieves greater restrictiveness through a significant simplification of the theory. Given this theoretical advantage, it would make sense to take the Modularity Hypothesis as the null hypothesis regarding the interaction between stress and segmental features and reject it in favor of a list of specific universals only given sufficient evidence to the contrary.

[2]How access to stress can be exploited to change the location of stress depends on the formalism. Suppose that the component responsible for stress-sensitive segmental processes is rule-based, using rules of the form $A \rightarrow B/X\_Y$. Then at least XAY should be able to refer to stress information, and nothing in principle prevents B from doing so as well. If the component in question is implemented using OT and its input contains stress information, nothing in principle prevents GEN from generating candidates with unfaithful stress.

This methodological principle is stated in (12):

(12) Methodological principle

Without evidence to the contrary, prefer the Modularity Hypothesis to a list of specific universals, since the Modularity Hypothesis is both simpler and more restrictive.

Given the methodological principle in (12), my approach to evaluating counterexamples to the Stress-Encapsulation Universal is to require conclusive evidence against it: I will take a tie between a sonority-driven analysis and an alternative that respects encapsulation to be sufficient to reject the evidence for sonority-driven stress in a given language.

1.1.3 Outline of the chapter

The claim that the computation of stress is encapsulated from segmental features can only be evaluated given a concrete phonological architecture. My first step is therefore to develop the basic properties of a modular architecture – the interaction between the stress module and the rest of the grammar, and the theory of the interface to the stress module (section 1.2). After developing the modular architecture, I present some of its predictions regarding possible stress patterns (section 1.3). Then, using the perspective provided by that architecture, I take a closer look at patterns of sonority-driven stress reported in the literature. I first provide a general overview of those patterns (section 1.4) and then re-evaluate individual cases in more detail (section 1.5). Finally, I discuss non-modular accounts of information encapsulation and show that they face non-trivial challenges in accounting for the Stress-Encapsulation Universal (section 1.6).

1.2 A modular architecture

1.2.1 Interaction between stress and the rest of the grammar

The Modularity Hypothesis in (11) only concerns the relationship between stress and segmental features. It has nothing to say about other aspects of phonology or where they are computed. If, for example, the distribution of tone can be conditioned by segmental features (see Tang 2008 and references therein), then at least some aspects of the computation of tone would have to take place outside of the module in which stress is computed. As it currently stands, the Modularity Hypothesis does not preclude non-stress computation from taking place in the stress module as long as it makes no reference to segmental features. It is conceivable, then, that a process like final-vowel lengthening would be computed in the same module as stress. In what follows, I will tentatively name the modules *Stress* and *Segmental phonology*, where *Segmental phonology* minimally includes segmental computation, and I will leave open the question of where other aspects of phonology are computed.

In discussing the interaction between stress and segmental phonology, it would be helpful to make use of the terms *Interactionist* and *Non-Interactionist* sometimes used in the literature to describe models of modular interaction. An interactionist architecture for stress and segmental phonology would be one where stress and segmental processes are interspersed and the grammar goes back and forth between stress and non-stress computation given some ordering, as illustrated in (13). Examples of interactionist architectures for the interaction between morphology and phonology include Lexical Phonology and Morphology (Pesetsky 1979, Kiparsky 1982) and Stratal OT (Kiparsky 2000).

(13) Interactionist architecture (stress and segmental processes are freely interspersed)

- Stress module
  1. Stress rule
- Segmental module
  2. Segmental rule
- Stress module
  3. Stress rule
  4. Stress rule
- Segmental module
  5. Segmental rule

In a non-interactionist architecture, stress computation would precede segmental computation in every cycle, as illustrated in (14).[3] For the interaction between morphology and phonology, a non-interactionist architecture was adopted in SPE (Chomsky and Halle 1968) and later work in Distributed Morphology following Halle (1990).

(14) Non-Interactionist architecture

- Stress module
  1. Stress rule
  2. Stress rule
  3. Stress rule
- Segmental module
  4. Segmental rule
  5. Segmental rule

The non-interactionist architecture for stress and segmental phonology is both simpler and more restrictive than the interactionist architecture. It is simpler since the grammar includes just one instruction to move once from stress to segmental phonology as opposed to multiple instructions to move back and forth between the modules; and it is more restrictive since the requirement that all stress processes precede all segmental processes in every cycle reduces the range of possible orderings. It makes sense, then, to take the non-interactionist architecture as the null hypothesis and abandon it only in the face of sufficient evidence to the contrary.

The final architecture is given in (15). First, underlying phonological representations are inserted using an operation like Vocabulary Insertion (Halle and Marantz 1993). The interface representation is computed based on the phonological representation (which includes segmental information) and is sent off to the stress module. The output of the stress module is sent back and the derivation proceeds to the segmental phonology. Note that

[3]The reverse order is untenable because stress assignment can feed segmental processes within the same cycle. For example, in Noyer's (2013) analysis of Huave, stress assignment must strictly precede a rule of /e/-lowering under stress that crucially applies in the same cycle as stress assignment.

segmental features are not sent off to the stress module but are accessible again in the segmental phonology, so this is not a classical feed-forward architecture. The operations in (15) can be read as a sequence of instructions to a central processor. The stress module serves as a function that receives a representational chunk as an input from the central processor and returns an output.

(15) **Modular architecture (non-interactionist): order of operations**

1. Insert underlying phonological representation
2. Construct interface representation
3. Send interface representation to the stress module (without segmental features)
4. Receive interface representation from the stress module
5. Send phonological representation (interface + segmental representation) to the segmental module
6. In the segmental module, stress representations cannot be changed

1.2.2 The role of the interface

According to the Modularity Hypothesis in (11), the stress module has no access to segmental features. Stress can only see other suprasegmental information, which serves as the interface between the stress module and the rest of phonology. In this architecture, segmental features can only affect stress indirectly through the interface. To illustrate the role of the interface, consider the representation of the made-up word in Figure 1-2. At the top, a stress representation is given in a grid-based theory of stress (Liberman and Prince 1977) where asterisks indicate prominence, as in Prince (1983) and Halle and Vergnaud (1987a). Below stress, a skeletal representation is given which encodes the distinction between consonants and vowels (the CV tier of McCarthy 1979b and Clements and Keyser 1983). The segmental representation at the bottom is connected to the skeletal representation using association lines.

		*								line 2
Stress representation		*				*	*			line 1
		*				*	*		*	line 0
Skeletal representation	C	V	C	C	C	V	V	C	V	
	|	|	|	|	|	\	/	|	|	
Segmental representation	l	i	n	g	k	a		r	ɔ	

Figure 1-2: Representation of the made-up word língkaːrɔ

Suppose now that the stress module has access to the skeletal CV tier (and to association lines) but not to segmental representations (this assumption is only used for illustration and will be replaced below with a concrete proposal). This assumption about the interface separates possible statements that could be made in the stress module from impossible statements. Stated informally in grid-theory terms, examples of possible statements are that 'every vowel projects an asterisk to line 0' and that 'the leftmost vowel projects an asterisk to line 2', as neither statement makes reference to segmental features. Examples of impossible statements are that 'every *low* vowel projects an asterisk to line 0' and that 'every vowel followed by a *flap* projects an asterisk to line 1', as both reference segmental features (*low* and *flap* respectively). In contrast, since a property like *length* is represented at a suprasegmental level – a long vowel is associated with two V slots in Figure 1-2 – stress may be sensitive to length. More generally, if stress is conditioned by some phonological distinction, that distinction must be represented at some suprasegmental level. With this background in hand, I proceed to propose a concrete theory of the interface.

1.2.3 A theory of the interface

My strategy in constructing the theory of the interface is to start with the bare minimum assumptions regarding the information that stress can access and complicate the theory incrementally only when necessary. Simple patterns of quantity-sensitive stress suggest that vowel length and the distinction between consonants and vowels are important for determining stress placement. For example, in Classical Arabic and some of its colloquial dialects, a word-final CVVC sequence (where VV stands for a long vowel) always receives primary stress, but a final CVC sequence does not; similarly, a final CVCC sequence is always stressed but a final CVCV is not (McCarthy 1979a, Watson 2002). Since the CV tier encodes those two properties as suprasegmental, it makes sense to take it as an initial hypothesis regarding interface representation. My first version of the theory of the interface, given in (16), is that interface representations are a subset of the set of strings that can be written using the symbols C and V. The asterisk in (16) stands for the Kleene Star Operator.[4]

(16) *Theory of the interface* (to be updated below in (19))
Interface representations are a subset of Σ^* where $\Sigma = \{C, V\}$

1.2.3.1 Syllable structure

A CV tier is not enough to capture all attested stress patterns. In some languages, segmental features determine syllable structure which in turn affects the position of stress. A simple example comes from Latin (17) (see Allen 1973 and Mester 1994 for general analyses of Latin stress and Lahiri 2001 for a discussion of the significance of syllable structure to Latin stress). In (17a), the penultimate syllable is a heavy CVC syllable which attracts stress, and stress is penultimate. In (17b), the penultimate syllable is a light CV syllable, and stress is antepenultimate. The only relevant difference between the two words is the underlined consonant. In (17b), that consonant is the liquid [r], which allows the preceding consonant to join it into the complex onset of the final syllable, which in turn makes the preceding syllable light. In (17a), that consonant is the non-liquid [t], which cannot function as the second member of a complex onset and thus forces the preceding consonant to be parsed as a coda consonant.

(17) Indirect effect of liquidity on stress in Latin

a. [vo.lúp.tas] (non-liquid)

b. [vó.lu.kris] (liquid)

[4]This version of the theory of the interface omits association lines and thus does not distinguish a long vowel from a sequence of two vowels. That distinction will be made by the updated version of the theory introduced next, using syllable structure and without association lines.

To accommodate such patterns, the input to the stress module should include information about syllable structure. Assuming a CV tier, information about syllable boundaries (without internal syllable structure) will be enough. (18) shows that the difference between the two words can be captured through a distinction in the position of the dot, which indicates a syllable boundary.

(18) a. [voluptas] ↔ [CV.CVC.CVC]

b. [volukris] ↔ [CV.CV.CCVC]

The second version of the theory of the interface, given in (19), includes the new symbol '.' (dot) in the set of interface symbols.

(19) *Theory of the interface* (to be updated below in (23))

Interface representations are a subset of Σ^* where $\Sigma = \{C, V, .\}$

1.2.3.2 Empty vowels

In section 1.4, we will see stress patterns in which stress avoids reduced vowels like schwa ([ə]). A simple example comes from some dialects of French:

(20) *French stress* (violates encapsulation given (19))

Stress is final unless the final vowel is *schwa*, in which case stress is penultimate

This statement makes reference to vowel quality – it mentions schwa – so it is a direct counterexample to the Stress-Encapsulation Universal given my current assumptions about the interface. Since word-final schwas are not epenthetic in French (Anderson 1982), a simple solution that assigns final stress before epenthesis is untenable. The present section introduces a representational mechanism proposed elsewhere in the literature that would allow me to encode the distinction between reduced and full vowels at the interface and avoid reference to vowel quality in the analysis of stress patterns like that of French.

Vowels like schwa exhibit special distributional properties that have motivated various representations of them as structurally deficient segments. In Dutch, for example, Kager (1990) notes that schwa is unstressable and that it is invisible to some syllable-sensitive processes and phonotactic restrictions: some segmental combinations (/h/, /ŋχ/, and /diphthong+r/) occur before full vowels but are banned syllable-finally and before schwa; consonant clusters are broken up by epenthesis syllable-finally and before schwa but not before full vowels; and so on. Kager argues that a structural representation of schwa as a defective vowel that cannot be the nucleus of a syllable provides the best account of its behavior: if stress is a property of syllables, then schwa's inability to be the head of a syllable accounts for its unstressability; and if consonants immediately preceding schwa have no choice but to close the preceding syllable, it follows that schwa is preceded by a syllable boundary.[5] While Kager's original generalizations have been challenged in later literature, his insight that the distributional properties of schwa follow from its structural deficiency has remained (van Oostendorp 1997). In a similar vein, Anderson (1982) argues that the distribution of schwa in French involves an alternation between [œ], [ɛ], and ∅ (it is not pronounced in some environments). He shows that /œ/ and /ɛ/ are not possible underlying representations for schwa and is left to conclude that its underlying representation is ∅. Since the position in which schwas occur is unpredictable, schwa cannot be epenthetic. Consequently, Anderson develops an autosegmental analysis of schwa as a skeletal V slot that lacks any association to segmental features. That V slot is assigned segmental features in some environments in the course of the derivation; otherwise, it is not pronounced. I will refer to V slots that are not associated to any segmental features as *empty vowels*. The representation of empty vowels is given in (21) and a sample spell-out rule for empty vowels is given in (22). Empty vowels or other implementations of structural deficiency have been defended by Levin (1985), Rubach (1986), Szpyra (1992), Zoll (1996), van Oostendorp (1997), and Kiparsky (2003), among others, and have played a central role in the literature on Government Phonology (see especially Lowenstamm 1996 and Scheer 2004).

(21) Representation of empty vowels

[5] Kager's original argument is stated within a moraic framework, where the structural deficiency of schwa is implemented as weightlessness. I restated the argument here in mora-free terms without affecting its force, as far as I can tell.

	Empty vowel	Low central vowel
Skeletal representation	V	V
	|	|
Segmental representation	[ ]	[a] = [+low,+back,. . .]

(22) Sample spell-out rule for empty vowels

$$\begin{array}{ccc} V & & V \\ | & \Rightarrow & | \\ [\ ] & & [\text{ə}] \end{array}$$

If reduced vowels like [ə] are structurally distinct from full vowels, it is a natural move to assume that the stress module can be sensitive to that distinction. I will adopt empty-vowel representations along with the assumption that the stress module can see the binary distinction between an empty vowel and a non-empty vowel at the interface. Formally, empty vowels receive the special skeletal symbol $V_\emptyset$ which I add to the set of interface symbols:

(23) *Theory of the interface* (final)

Interface representations are a subset of Σ^* where $\Sigma = \{C, V, V_\emptyset, .\}$

The updated theory of the interface enables a restatement of French stress that ignores schwa and does not violate encapsulation:

(24) *French stress* (respects encapsulation given (23))

Stress the final V

At present, I do not impose any restrictions on empty-vowel representations other than what is already implied by their definition – namely, that there is a one-to-one mapping between the symbol $V_\emptyset$ and its segmental content (25). I also do not posit any restrictions on empty-vowel spell-out rules.[6]

(25) $V_\emptyset \leftrightarrow []$

[6]This means, for example, that any vowel could be empty, including sonorous vowels like [a]. In my analyses below, empty vowels will always be associated with low-sonority vowels such as [ɔ] and [i]. The theory is compatible with restrictions on the realization of empty vowels and they can be added if needed.

Below I will show that the theory of the interface in (23) can take us quite far in reanalyzing sonority-driven stress patterns, and I will discuss some typological consequences of representing reduced vowels as empty vowels at the interface.

1.3 Predictions regarding possible patterns

With a concrete modular architecture in hand, my next goal is to explore its predictions regarding possible stress patterns. Before doing so, I would like to mention an open issue for this approach that I do not resolve in this chapter.

A modular architecture with encapsulation can sometimes derive patterns that are extensionally equivalent to quality-sensitive stress patterns (derived in architectures with no encapsulation). In such cases, translating encapsulation to predictions regarding possible patterns is not straightforward. To see why, recall that the modular architecture is necessarily serial, because stress and segmental processes are not computed together and, for example, stress can feed segmental processes. In a serial architecture, quality-sensitive stress can be mimicked in indirect ways, such as using a suprasegmental property as a diacritic for the sole purpose of determining stress placement. The grammar in (26) follows a general rule schema that lengthens vowels in some segmental environment only to shorten them back after stress assignment. The result is equivalent to quality-driven stress.

(26) Grammar:

1. $V_{[+F]} \rightarrow$ long / A _ B
2. Assign stress to {every long vowel / the rightmost long vowel / ...}
3. $V_{[+F]} \rightarrow$ short / A _ B

This grammar combines two properties whose existence has been long debated in the literature. First, it involves so-called 'Feeding Duke-of-York' derivations, where a process that changes A into B feeds some process P, before another rule changes B back into A and removes the environment of P (McCarthy 2003b). The second property is a version of 'Absolute Neutralization' where a feature (*long* in the example above) is eliminated from surface representations completely (see Kiparsky 1968, Hyman 1970, McCarthy 2005). To my knowledge, grammars like (26) that combine both properties are unattested independently of stress-segmental interactions, but I am not aware of a satisfying account of their absence within serial architectures. If such grammars are unavailable, though, encapsulation could derive interesting predictions regarding possible patterns which I would like to explore in this section. The predictions I discuss next are therefore conditional on grammars like (26) being unavailable: I will assume that using suprasegmental features as diacritics as in (26) is not an option, but at present I leave as a black box a formal explanation for why this is so.

1.3.1 Prediction regarding vowel invisibility to stress

The theory of the interface in (23) predicts that distributional differences between distinct vowels with respect to stress should be limited to the binary distinction between non-empty vowels and the empty vowel. In some languages, a distinction has been reported between multiple full vowels and multiple reduced vowels, such that the latter are invisible to stress. Since the empty vowel as defined in section 1.2.3.2 is unique (the symbol $V_\emptyset$ corresponds to no segmental features), the theory makes the following prediction regarding invisibility to stress in such languages:

(27) *Prediction regarding invisibility to stress*

All vowels that are invisible to stress must be either epenthetic or (underlyingly) empty[7]

To illustrate this prediction, consider a hypothetical language where stress falls on the final vowel but shifts left when the final vowel is a schwa or an [a], but only when [a] is followed by a glottal stop ([ʔ]). Some examples are given in (28). If epenthesis is not involved, the only way to account for the data systematically is by deriving [a] from schwa precisely where [a] is skipped. In other words, the modular architecture forces the existence of a vowel lowering process that turns schwa into [a] before [ʔ], a process familiar from Semitic languages. In such cases we would expect lowering to leave some distributional signature. For example, the sequence [əʔ] could be unattested in the language and rejected

[7] Leaving aside other options, like underlying glides undergoing vocalization or extrametrical suffixes.

by speakers (29a), or, if lowering only applies before coda glottal stops, adding a suffix with a vowel could reveal a schwa before the glottal stop (29b).

(28) Hypothetical pattern: final stress skips [ə] and [aʔ]

a. kogá

b. kógə

c. kogɨ́ʔ

d. kógaʔ

(29) Possible distributional signatures of lowering

a. *əʔ

b. kógaʔ ~ kógəʔ-i

The theory rules out stress patterns where stress skips two distinct vowels whose distribution is unpredictable. In section 1.5.1 I will discuss the stress pattern of Mari, where stress skips multiple surface-distinct vowels and the prediction in (27) is borne out: all skipped vowels can be traced back to an underlying schwa. I would like to note that even if this prediction turns out to be false, the revision required from the theory would not necessarily be dramatic. The prediction results from a particular implementation of empty-vowel representations that enforces a one-to-one mapping between the interface symbol $V_\emptyset$ and the vocalic features that it is associated to (namely, no features). We could imagine a less restrictive variant of the theory that allows a many-to-one mapping between vowels and the symbol $V_\emptyset$ which would not make the prediction in (27). Instead, stress would be able to skip a set of derivationally unrelated vowels (corresponding to $V_\emptyset$) as long as it treats them in the same way. As a matter of methodology, it makes sense to retreat to the less restrictive variant only given sufficient evidence against (27).

1.3.2 Prediction regarding indirect effects of segmental features on stress

If the interface only allows segmental features to affect stress indirectly through syllable structure, we make the prediction in (30) regarding indirect effects of segmental features on stress placement:

(30) *Prediction regarding indirect effects of segmental features on stress*
Indirect effects of segmental features on stress should have a distributional signature expressed in terms of syllable structure

Consider again the Latin stress pattern, where the presence of a liquid affects stress (31). This effect is mediated by syllable structure: [pt] is broken up by a syllable boundary but [kr] is not. There is an independent restriction on complex onsets in Latin such that a consonant-liquid complex onset like [kr] is allowed but other consonant-stop complex onsets like [pt] are not. What is ruled out is a language that has the same stress pattern as Latin but without the distributional restriction on complex onsets.

(31) a. [volúp.t̲as] (non-liquid)
b. [vólu.kr̲is] (liquid)

1.3.3 Prediction regarding segmental restrictions on stress alignment

If the computation of stress has no access to segmental features, the assignment of stress to the rightmost or leftmost vowel in some segmental environment is impossible. A general statement of the class of patterns that is ruled out is given in (32).

(32) *Segmental restrictions on stress alignment*
'stress the rightmost/leftmost vowel V such that $f(\mathrm{V})$',
where $f(\mathrm{V})$ is a description of the identity or environment of V that makes reference to segmental features

Examples of unattested stress patterns in this class are the following:

(33) Stress the leftmost round vowel

(34) Stress the penultimate syllable, but if it has an unaspirated onset, stress the antepenultimate syllable

(35) Stress the rightmost vowel not preceded by an unaspirated obstruent

The patterns in (33) and (34) are simple and do not require further elaboration. According to (35), stress seeks the rightmost vowel but shifts left whenever a vowel is preceded by an unaspirated obstruent (like [t]). This pattern is illustrated in (36). In (36a), stress is final since the final vowel is preceded by an aspirated stop. In (36b), the final consonant is unaspirated, so stress shifts once to the left. It remains on the penultimate vowel since the preceding consonant is an aspirated stop. In (36c), the penultimate consonant is an unaspirated stop as well. Stress is antepenultimate since the antepenultimate vowel is preceded by another vowel (and not by an unaspirated stop).

(36) a. [titatut^{h}ó]

b. [titat^{h}úto̱]

c. [tiáṯuto̱]

Patterns with a stress shift along the lines of (34) and (35) can be easily generated in OT using the markedness constraint *tV́ to trigger stress shift.[8]

1.3.4 Prediction regarding destressing

If the stress module has no access to segmental features, feature-specific destressing processes cannot be stated. A general statement of the class of patterns that is ruled out is given in (37), followed by some examples of patterns in this class.

(37) *Feature-specific destressing*

'Delete stress from a vowel V such that f(V)',

where f(V) is a description of the identity or environment of V that makes reference to segmental features

(38) a. Pre-stress destressing of low or front vowels

b. Pre-stress destressing of vowels preceded by an unaspirated obstruent

c. Destressing of high vowels

[8]The precise nature of the shift in (35) will vary depending on the constraints used to generate rightmost and leftmost stress effects. For example, OT with gradient alignment constraints will be able to generate precisely the pattern in (35).

To illustrate (38a), imagine a language that assigns stress to the final vowel of the stem regardless of the identity of the vowel (39). Then, a lexically-stressed suffix is added and creates a sequence of two stressed vowels (40). Finally, only non-low back vowels maintain stress (41).

(39) Stem-final stress

a. [CVCáC]

b. [CVCíC]

c. [CVCúC]

(40) Lexically-stressed suffix creates a clash

a. /CVCáC-ó/

b. /CVCíC-ó/

c. /CVCúC-ó/

(41) Only non-low back vowels maintain stress

a. [CVCaC-ó]

b. [CVCiC-ó]

c. [CVCúC-ó]

Similarly, an example of (38b) is a language that assigns stress to the final vowel of the stem regardless of its segmental environment (42). Then, as before, a lexically-stressed suffix is added and creates a sequence of two stressed vowels (43). Finally, only vowels preceded by a unaspirated obstruent lose stress (44).

(42) Stem-final stress

a. [CVt^{h}áC]

b. [CVCnáC]

c. [CVCtáC]

d. [CVCuáC]

(43) Lexically-stressed suffix creates a clash

a. /CVt^{h}áC-ó/

b. /CVCnáC-ó/

c. /CVCtáC-ó/

d. /CVCuáC-ó/

(44) Only vowels preceded by an unaspirated obstruent lose stress

a. [CVt^{h}áC-ó]

b. [CVCnáC-ó]

c. [CVCtaC-ó]

d. [CVCuáC-ó]

The destressing process in (38c) can create unattested vowel-specific gaps in alternating stress. Suppose that a language assigns alternating stress as in (45a) and deletes stress from every high vowel (45b).

(45) a. Stress every second vowel from the left

b. Destress a high vowel

The result is a pattern where words with only non-high vowels have stress on every second vowel from the left (46) but words with high vowels have gaps in alternating stress such that a stressed vowel may be preceded or followed by three unstressed vowels (47).

(46) Words with only non-high vowels: alternating stress

a. [CaCóCoCáCa]

b. [CaCóCoCáCaCó]

(47) Words with high vowels: gaps in alternating stress

a. [CaC$\underline{i}$CoCáCa]

b. [CaCóCoC$\underline{u}$CaCó]

1.4 Sonority-driven stress in the literature

Previous studies on the phonology of stress include analyses of stress patterns that make direct reference to vowel sonority, thus violating the Stress-Encapsulation Universal. The present section provides a brief history of sonority-driven stress in the literature and of its role in the development of theories of stress.

My starting point is Halle and Vergnaud (1987a), whose grid-based theory of stress explicitly allows vowel quality to influence the distribution of stress through prominence. Halle and Vergnaud (1987a) is by no means the first work to discuss patterns of sonority-driven stress – see references to earlier work in Gordon (1999/2006) – but it will be a convenient point of departure for discussing the role of vowel quality in stress theory. On Halle and Vergnaud's theory (following Liberman and Prince 1977; Prince 1983), asterisks indicate prominence and higher lines on the grid correspond to greater prominence (see Figure 1-2). Metrical constituents are constructed based on the lines of asterisks. Importantly, an element's degree of prominence can be determined by its quality. Stress rules explicitly refer to quality in Halle and Vergnaud's analysis of the default-to-opposite pattern in (48a) that distinguishes full from reduced vowels, reportedly found in 6 languages. Halle and Vergnaud's rule in (48b) is the one that refers to quality, and their system imposes no restrictions on how quality can be used in the description of stress rules.

(48) Sonority-driven stress pattern in Halle and Vergnaud (1987: 51)

a. Stress falls on the last syllable that has a full vowel, but in words where all syllables have only reduced vowels, stress falls on the first syllable

b. Rule: Assign line 1 asterisks to **full** vowels

A more fine-grained sensitivity to vowel quality was considered by Hayes (1995), who developed a theory of stress based on the stress patterns of more than 150 languages. Asheninca, as described by Payne (1990), is the only language in Hayes (1995) whose stress pattern is sensitive to vowel quality. Drawing on Payne's description, Hayes' analysis of Asheninca associates syllables with different degrees of prominence based on vowel length and quality:

(49) Asheninca hierarchy of prominence in Hayes (1995)

*** CVV

** Ca, Co, Ce, CiN (N = nasal consonant)

* Ci

The rhythmic aspect of Asheninca stress is not sensitive to vowel quality: on both Payne's and Hayes' analyses, metrical constituents are built on the basis of quantity alone. The basic rhythmic pattern can be perturbed by processes such as destressing that are sensitive to the prominence hierarchy in (49). Hayes divided stress rules into two subsets, foot construction rules and rules like destressing, end rules (which refer to edges), and extrametricality. He suggested that foot construction is encapsulated from vowel quality but that other rules are not (without developing the architecture responsible for semi-encapsulation in much detail).[9]

In the early OT literature, Kenstowicz (1997) claimed that stress is sensitive to vowel sonority on the basis of the distribution of stress in several languages (Kobon, Chukchi, Aljutor, Mari, and Mordwin). He proposed a hierarchy of markedness constraints that makes more sonorous vowels better stress-bearers. Notably, Kenstowicz offered the fine-grained sonority hierarchy in (50) for Kobon stress. On his analysis, Kobon stress is sensitive to a five-way distinction in terms of vowel quality. Following Kenstowicz's analysis, Kobon has become a showcase pattern of sonority-driven stress. The markedness theory of sonority-driven stress was further developed in a series of works by de Lacy (2002, 2004, 2007) with support from several more languages.

(50) Kobon stress in Kenstowicz (1997)

a. Stress falls on the more sonorous vowel among the final two vowels, according to the sonority hierarchy in (50b)

b. a/au/ai > o/e > u/i > ə > ɨ

Gordon's (1999/2006) survey of 388 languages provided cross-linguistic support for Kenstowicz's small survey, reporting 28 languages with sonority-sensitive stress patterns.

[9]Hayes also allowed segmental features to project directly into the prominence grid in the analysis of Pirahã (Everett and Everett 1984; Everett 1988), where stress assignment has been claimed to be sensitive to the [voice] feature of the onset. See discussion of consonantal features in section 1.5.7.

A rough classification of those patterns according to their type of sonority-sensitivity is given in (51). Type I is of languages that show a distinction between full and reduced vowels and where stress often skips reduced vowels. Out of 28 languages with sonority-driven stress in the survey, 20 are of Type I. Type II is of languages where the low vowel attracts stress as opposed to every other vowel (5/28). Finally, Type III is of languages where stress is sensitive to a fine-grained sonority hierarchy based on vowel height or peripherality (3/28).

(51) Sonority-driven stress in Gordon (1999/2006) (my classification)

- **Type I**: Full vs. reduced vowels (20/28)
  - Aljutor, Au, Chuvash, Javanese, Karo Batak, Lamang, Lillooet, Lushootseed, Malay, Mari, Mordvin, Moro, Nankina, Ngada, Patep, Sarangani Manobo, Sentani, Siraiki, Vach Ostyak, Yil
- **Type II**: Low vowel vs. other vowels (5/28)
  - Gujarati, Kara, Komi, Mayo, Yimas
- **Type III**: Fine-grained sonority hierarchy based on vowel height or peripherality (3/28)
  - Asheninca, Chukchi, Kobon

Following the works of Hayes, Kenstowicz, de Lacy, and Gordon, the existence of sonority-driven stress has been taken for granted in the literature and the theoretical apparatus introduced in those works has influenced later studies on stress. Later works that introduce sonority-driven stress patterns include Crowhurst and Michael (2005), Vaysman (2008), Trommer (2013), and Moore-Cantwell (2016).

Some of the reported cases have already been reanalyzed in the literature. Hargus (2001) suggested that sonority-driven stress can be reduced to quantity-driven stress based on the durational properties of reduced vowels in two languages, Sahaptin and Witsuwit'en. Shih (2016) conducted a phonetic experiment on Gujarati, a Type II language, and showed that low vowels claimed to attract stress do not in fact correlate with stress-related phonetics, suggesting that properties like length may have been misinterpreted as stress (see also

Bowers 2016). Canalis (2007) showed that the correlation between stress and vowel quality in Albanian (Type III, see Trommer 2013) is due to morphological factors. Chukchi, another Type III language, was discussed by de Lacy (2013), who argued that descriptions of Chukchi stress as sonority-sensitive had been based on insufficient evidence from conflicting sources. More generally, de Lacy (2013) rejected the evidence for Type II and Type III patterns of sonority-driven stress in his own work altogether.

In the next section I re-evaluate the evidence for all of the remaining sonority-driven stress patterns in Halle and Vergnaud (1987a), Hayes (1995), Kenstowicz (1997), Gordon (1999/2006), and patterns I have been able to find in later work (Nanti, Crowhurst and Michael 2005; English, Moore-Cantwell 2016). I will offer a general recipe for reanalyzing Type I patterns using empty-vowel representations at the interface, and I will claim that there is no convincing evidence for any Type II or Type III patterns. The tables in (52)-(56) summarize the list of sonority-driven stress languages in Halle and Vergnaud (1987a), Hayes (1995), Kenstowicz (1997), Gordon (1999/2006), and later work and state where each language is re-evaluated. In the columns labeled 'Status', I use the word 'Reanalysis' for cases where an alternative analysis that does not make direct reference to sonority is presented. 'Discussion' is used for cases where a convincing alternative is not presented but a critical discussion of the evidence is provided that I believe weakens the case for sonority-sensitivity.

(52) Sonority-driven stress in Halle and Vergnaud (1987a)

Language	Type	Status
6 languages	I	Recipe for reanalysis in section 1.5.1

(53) Sonority-driven stress in Hayes (1995)

Language	Type	Status
Asheninca	III	Discussion in section 1.5.5

(54) Sonority-driven stress in Kenstowicz (1997)

Language	Type	Status
Kobon	III	Reanalysis in section 1.5.2
Chukchi	III	Data re-evaluated in de Lacy (2013)
Aljutor	I	Recipe for reanalysis in section 1.5.1
Mari	I	Reanalysis in section 1.5.1
Mordwin	I	Recipe for reanalysis in section 1.5.1

(55) Sonority-driven stress in Gordon (1999/2006)[10]

Language	Type	Status
20 languages	I	Recipe for reanalysis in section 1.5.1
Gujarati	II	Data re-evaluated in Shih (2016) and Bowers (2016)
Kara	II	Reanalysis in Blumenfeld (2006)
Komi	II	Re-evaluation in footnote 10
Mayo	II	Reanalysis in section 1.5.4
Yimas	II	Discussion in section 1.5.5
Asheninca	III	Discussion in section 1.5.5
Chukchi	III	Data re-evaluated in de Lacy (2013)
Kobon	III	Reanalysis in section 1.5.2

(56) Sonority-driven stress in later literature

Language	Source	Status
Nanti	Crowhurst and Michael (2005)	Reanalysis in section 1.5.3
Albanian	Trommer (2013)	Reanalysis in Canalis (2007)
English	Moore-Cantwell (2016)	Reanalysis in section 1.5.6

1.5 Re-evaluation of sonority-driven stress patterns

1.5.1 Reanalysis of Mari stress

This section provides a general recipe for reanalyzing Type I patterns of sonority-driven stress, where stress is sensitive to the distinction between full and reduced vowels. The key

[10]The generalization regarding stress in Komi is described in the sources as a diachronic pattern, so I do not discuss it further (see Hausenberg 1998).

ingredient in the analysis is the representation of reduced vowels as empty vowels at the interface. The language that I reanalyze is Eastern Mari (henceforth Mari) as described in Vaysman (2008). Mari was chosen over other Type I languages for two reasons. First, it appears to be a challenging case to the binary distinction between empty and non-empty vowels at the interface. Mari stress often skips schwas, but there is no one-to-one correspondence between schwas and vowels skipped by stress, in both directions (some full vowels are skipped and some schwas are stressed). Mari thus makes a good test case for the prediction in (27). The second reason is that the claims in Vaysman (2008) are supported by rich data controlled for lexical category, morphosyntactic environment, and other factors, so the generalizations regarding stress placement are quite clear. I will begin by discussing stress in mono-morphemic words, all of which are underived nouns, and then proceed to discuss stress in morphologically complex words.

1.5.1.1 Mono-morphemic words

In mono-morphemic words, stress normally falls on the rightmost full vowel – the rightmost vowel that is not a schwa ([ə]):

(57) a. koŋgá 'oven'
b. sérəʃ 'letter'
c. jóŋələs 'mistake'
d. paréŋə 'potato'

Stress also skips vowels that alternate with schwa and are the result of vowel harmony:[11]

(58) a. pőrʃö ~ pőrʃə-m 'frost'~'frost ACC'
b. ʃóʃo ~ ʃóʃə-m 'spring'~'spring.ACC'

Finally, when every vowel in a word is a schwa, stress is initial:

(59) βə́nər 'canvas'

[11] Vaysman presents evidence for vowel harmony over a process that goes in the other direction (vowel reduction). The environment of application of vowel harmony is somewhat complicated; the precise details are not important for the analysis so I will not discuss them here.

The pairs in (60) suggest that schwa is not epenthetic in Mari, as it is not possible to state a general schwa epenthesis rule that would insert the schwa in the second member of each pair without also inserting a schwa in the first member.

(60) a. kučém 'stress'
üréma 'handle'

b. meráŋ 'hare'
paréŋə 'potato'

The basis of a modular analysis is that schwa and vowels derived from it through vowel harmony are underlyingly empty. The analysis is given informally using a serial rule-based formalism with rule ordering and cyclicity as in Halle and Vergnaud (1987a) and assuming the architecture in (15), where stress rules precede segmental rules in every cycle. As far as I can tell, the choice of rules over constraints will not affect the analysis in any meaningful way, but serialism will be needed for a proper treatment of the opacity of Mari stress. A grammar for stress in mono-morphemic words is given in (61). The horizontal line marks the end of the stress rule block.

(61) A fragment of Mari grammar (to be revised below)

1. If no vowel is stressed, stress the rightmost V
2. If no vowel is stressed, stress the leftmost $V_{\emptyset}$

---

3. Vowel harmony
4. Empty-vowel spell-out ([] → [ə])

Here are some sample derivations. (62) shows a derivation of a word with a final schwa and penultimate stress. Rightmost stress applies and targets the penultimate vowel. Then leftmost stress and vowel harmony do not apply and the empty vowel is spelled out as schwa. (63) is an example with vowel harmony. As before, rightmost stress targets the penultimate vowel and leftmost stress does not apply. Then, vowel harmony applies and rewrites the final vowel as the full vowel [ö]. Since the final vowel is no longer empty, empty-vowel spell-out does not apply. Finally, (64) is a word that only contains schwas.

Here all vowels are initially empty, so rightmost stress applies vacuously. Then leftmost stress applies and assigns initial stress.

(62) Derivation of [paréŋə]

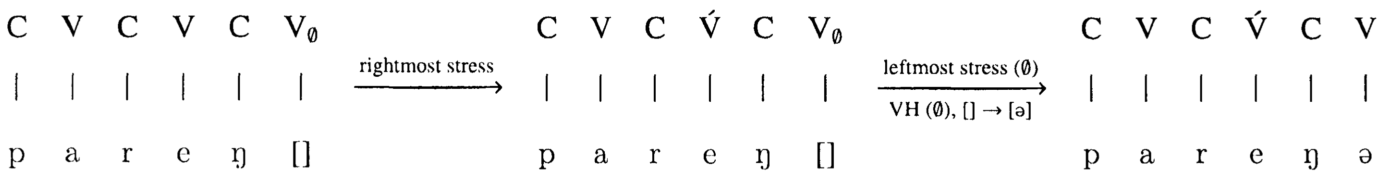

(63) Derivation of [pö́rʃö]

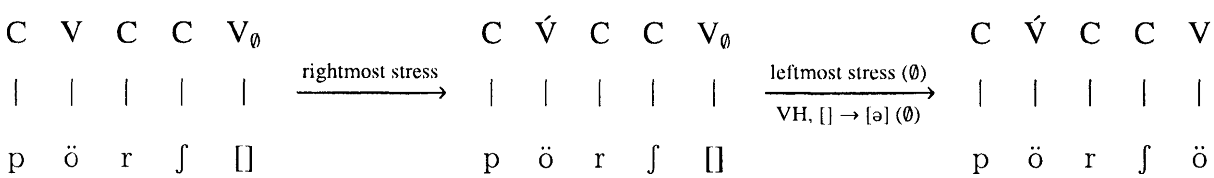

(64) Derivation of [βə́nər]

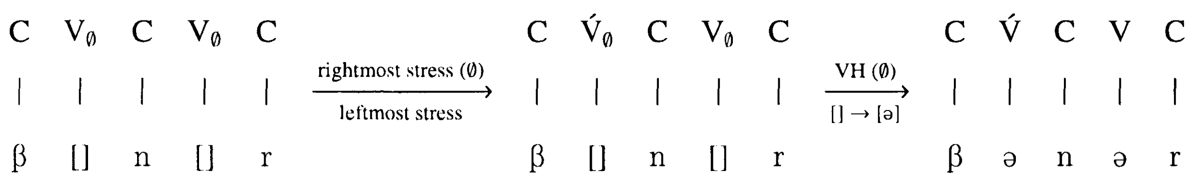

1.5.1.2 Multi-morphemic words

The distribution of stress in suffixed words will be demonstrated using two suffixes, -lan (DATIVE case) and -ge (COMITATIVE case). First, when the root only contains full vowels, stress in the suffixed form is root-final:

(65) a. paʃá ~ paʃá-lan 'work'~'work.DAT'

b. paʃá ~ paʃá-ge 'work'~'work.COM'

When the root only contains schwas, stress falls on the suffix:

(66) a. rə́wəz ~ rəwəz-lán 'fox'~'fox.DAT'

b. rə́wəz ~ rəwəz-gé 'fox'~'fox.COM'

Finally, when the root has non-final stress, the two suffixes behave differently. -lan attracts stress from the root, but -ge does not:

(67) a. sérəʃ ~ serəʃ-lán 'letter' ~ 'letter.DAT'

b. sérəʃ ~ sérəʃ-ge 'letter' ~ 'letter.COM'

Vaysman takes stress attraction to be a general property of suffixes with the vowel [a] (as opposed to suffixes with the vowel [e]). However, the number of suffixes is very small: Vaysman reports 4 suffixes with [a] and 3 suffixes with [e], and it is possible that some idiosyncratic property of the morphemes is what causes their different behavior rather than the quality of the vowel. This property could be lexical stress or, in the cyclic framework of Halle and Vergnaud (1987a) that I have been assuming here, the feature [±cyclic].[12] In the absence of evidence for choosing one option over the other, I will go with lexical stress. I will show that the assumption that suffixes like -lan are lexically stressed (whereas suffixes like -ge are not) is enough to derive the distribution of stress in suffixed words. Respecting encapsulation here comes with a price – a memorization of 4 instances of stress in the lexicon – but I believe that it is a small price to pay.[13] A way to argue against the lexical-stress analysis and in favor of the sonority-driven analysis is to show that speakers of Mari generalize the stress pattern to nonce suffixes with [a] (contrary to the prediction of lexical stress).

The final version of the grammar in (68) includes the assumption regarding stress marking in the lexicon. The rules are divided into a cyclic component, which applies once whenever a morpheme is added in the derivation, and a post-cyclic component which applies once at the end of the derivation. Rightmost stress is a now a cyclic rule and the post-cyclic component includes two new destressing rules.

(68) A fragment of Mari grammar (final)

- Assumptions about the lexicon:

[12][+cyclic] suffixes trigger a pass through the cyclic rule block and can trigger stress rules that would not apply with [-cyclic] suffixes.

[13]Vaysman states that verbal suffixes behave like nominal suffixes in that [a] attracts stress but [e] does not. The data are not provided in Vaysman (2008), but the existence of additional [a]-suffixes would weaken the present analysis. Other sources on Mari morphology (e.g., Kangasmaa-Minn 1998) distinguish two verbal declensions – *-am* and *-em* – but without stress data or a morphosyntactic analysis of those verbs it is difficult to determine whether more than one additional suffix (corresponding to *-am*) would have to be marked as lexically stressed on the present analysis. My attempts to locate Vaysman to ask for more information have been unsuccessful.

- The suffix -lan bears stress
- The suffix -ge does not bear stress

- Cyclic rules:

1. If no vowel is stressed, stress the rightmost V

- Post-cyclic rules:

2. If there are two consecutive stressed Vs, destress the rightmost V
3. If there are two stressed Vs, destress the leftmost V
4. If no vowel is stressed, stress the leftmost $V_\emptyset$

---

5. Vowel harmony
6. Empty-vowel spell-out ([] → [ə])

I will now show how this grammar accounts for the distribution of stress in (65)-(67), starting with the derivation of the two suffixed words in (65), given in (69). I will go through the derivation one rule at a time, considering the effect of each rule on both the lán-derivation and the ge-derivation. In the first cycle, the stem /paʃa/ is evaluated by itself and receives final stress. In the second cycle, the suffixes are added, -lán with lexical stress and -ge without any stress marking. Rightmost stress does not apply again since both representations are already marked for stress, and the representation is sent off to the post-cyclic component. In the post-cyclic component, post-stress destressing resolves the stress clash created by the addition of -lán by removing stress from the suffix. Post-stress destressing does not apply with -ge since only one vowel is marked for stress. None of the remaining rules applies: the environment of pre-stress destressing includes two stressed vowels, vowel harmony is irrelevant here, initial stress does not apply since both representations are marked for stress at the time of its application, and empty-vowel spell-out is irrelevant. The result is stem-final stress in both words.

(69) Derivation of the suffixed words in (65)

Word	[paʃá-lan]	[paʃá-ge]
Cycle I	paʃa	paʃa
Rightmost stress	paʃá	paʃá
Cycle II	paʃá-lán	paʃá-ge
Rightmost stress	-	-
Post-cycle	paʃá-lán	paʃá-ge
Post-stress destressing	paʃá-lan	-
Pre-stress destressing	-	-
Vowel harmony	-	-
Leftmost stress	-	-
$V_\emptyset$ spell-out	-	-
Output	[paʃá-lan]	[paʃá-ge]

Next, (70) shows the derivation of the two suffixed words in (66) along with their unsuffixed variant. Here, square brackets indicate an empty vowel. In the first cycle, rightmost stress does not apply: it only targets full V's, but all vowels are empty ($V_\emptyset$). In the second cycle, the suffixes are added. Rightmot stress again does not apply to the unsuffixed stem. It does not apply in the lán-derivation because stress is already present, but it does apply in the ge-derivation and assigns final stress. This is how the difference between the two suffixes is neutralized when the stem only contains schwas. Next, destressing rules and vowel harmony do not apply, but initial stress targets the first vowel of the unsuffixed word. Then, the empty vowels are spelled out as schwas.

(70) Derivation of the words in (66)

Word	[rə́wəz]	[rəwəz-lán]	[rəwəz-gé]
Cycle I	r[]w[]z	r[]w[]z	r[]w[]z
Rightmost stress	-	-	-
Cycle II	-	r[]w[]z-lán	r[]w[]z-ge
Rightmost stress	-	-	r[]w[]z-gé
Post-cycle	r[]w[]z	r[]w[]z-lán	r[]w[]z-gé

Post-stress destressing	-	-	-
Pre-stress destressing	-	-	-
Vowel harmony	-	-	-
Leftmost stress	r[́]wəz	-	-
$V_\emptyset$ spell-out	rə́wəz	rəwəz-lán	rəwəz-gé
Output	[rə́wəz]	[rəwəz-lán]	[rəwəz-gé]

Finally, (71) shows the derivation of the two suffixed words in (67). In the first cycle, rightmost stress targets the penultimate vowel, which is the rightmost V. In the second cycle, rightmost stress does not apply. In the post-cyclic component, post-stress destressing cannot apply in the lán-derivation since the two stressed vowels are not adjacent. Pre-stress destressing does apply (since it does not require adjacency) and removes stress from the stem. Otherwise, only empty-vowel spell-out applies.

(71) Derivation of the suffixed words in (67)

Word	[serəʃ-lán]	[sérəʃ-ge]
Cycle I	ser[]ʃ	ser[]ʃ
Rightmost stress	sér[]ʃ	sér[]ʃ
Cycle II	sér[]ʃ-lán	sér[]ʃ-ge
Rightmost stress	-	-
Post-cycle	sér[]ʃ-lán	sér[]ʃ-ge
Post-stress destressing	-	-
Pre-stress destressing	ser[]ʃ-lán	-
Leftmost stress	-	-
Vowel harmony	-	-
$V_\emptyset$ spell-out	serəʃ-lán	sérəʃ-ge
Output	[serəʃ-lán]	[sérəʃ-ge]

As far as I can tell, the proposed analysis correctly derives the distribution of stress in all of the data in Vaysman (2008) without reference to vowel quality. Beyond Mari, the analysis demonstrates how the distinction between empty and non-empty vowels at the interface can be used to reanalyze Type I languages without reference to segmental features, even when stress skips multiple surface-distinct vowels.

1.5.2 Reanalysis of Kobon stress

The stress pattern of Kobon has been a showcase of sonority-driven stress and is famous for its fine-grained sonority hierarchy. The source of the claim regarding sonority-driven stress in Kobon is Davies (1981), which states:

> "The rules for positioning stress in two-syllable words have yet to be determined. Relative vowel strength is almost certainly a conditioning factor since *stress is almost always placed on the syllable which is strongest according to the following hierarchy*:
>
> (72) a/au/ai > o/e/u/i > ə/ɨ
>
> Almost all three-syllable words manifest [a] as the vowel of the penultimate syllable and all of these words carry stress on that penultimate syllable. The few words which do not manifest [a] as the vowel of the penultimate syllable also carry stress on the penultimate syllable unless the final syllable manifests a stronger vowel than the penultimate syllable, in which case stress falls on the final syllable. Such cases are very few."

The low vowel [a] and diphthongs containing the low vowel ([au] and [ai]) are at the top of Davies' sonority hierarchy. Lower in the hierarchy are the non-low non-central vowels [o], [e], and [u]. The central vowels [ə] and [ɨ] are the least sonorous.

Based on Davies' (1981) description, Kenstowicz (1997) proposed the following hypothesis for Kobon stress (which assumes a more fine-grained sonority hierarchy than Davies'):

(73) Kobon stress in Kenstowicz (1997)

a. Stress falls on the more sonorous vowel among the final two vowels, according to the sonority hierarchy in (73b)

b. a/au/ai > o/e > u/i > ə > ɨ

The data in (74)-(75) from Kenstowicz (1997) (citing Davies 1981) are given to illustrate the sensitivity of stress to sonority. In (74), the final two vowels in each word differ in their sonority level. (74a) shows that [a] is a better stress-bearer than [o]: when [a] is the penultimate vowel and [o] is the final vowel, the penultimate vowel receives stress, but when the order of the two vowels is reversed it is the final vowel that receives stress. The remaining examples in (74) show that stress tracks sonority when other vowels are involved. When the final two vowels are of equal sonority, stress is penultimate (75).

(74) Vowels that differ in their sonority level

a. [a > o]: alágo vs. kɨdolmáN

b. [a > i]: ki.á vs. háu.i

c. [o > u]: mó.u

d. [o > i]: si.óg

e. [i > ə]: wí.ər

f. ...

(75) Vowels of equal sonority

a. [u ~ u]: dúbu-dúbu

b. [u ~ i]: jínup-jínup

c. ...

Another source on Kobon is Davies (1980), a book on Kobon phonology (written by the same author) where the description of stress does not mention sonority. According to Davies (1980), Kobon stress is normally penultimate:

> "Although the rules for the placement of stress cannot be stated comprehensively at this stage, it appears that stress is not phonemic. In phonological words of more than one syllable *stress normally falls on the penultimate syllable*."

Davies (1980) is a book dedicated to Kobon phonology that includes around 500 examples marked for stress, compared to around 50 examples marked for stress in Davies (1981), which is a general Kobon grammar. Since the description of stress in Davies (1980) as normally penultimate was based on a larger corpus, it raises the question of whether the correlation between sonority and stress placement observed in Davies (1981) generalizes to the entire body of data in both sources. To answer that question, I have reorganized the data from both sources according to lexical category, morphosyntactic environment, and syllable structure, with the goal of comparing the sonority hypothesis in (73) with the penultimate-stress hypothesis. The first observation is that the data include examples that pose a challenge to both hypotheses. Each pair of nouns in (76)-(77) is a near-minimal pair that differs in the location of stress. Aside from stress, the only difference between words in each pair is the place of articulation of some nasal consonants.

(76) a. ɸə́nʌm 'wind' [33]

b. ɸəŋʌ́n 'sweet potato sp.' [33]

(77) a. ambáɲ 'platform' [34]

b. ámbaŋ 'a river name' [34]

There were also examples in the data that posed a challenge to the penultimate-stress hypothesis – words with a final stressed syllable that has a diphthong or a complex coda and words with final stress whose penultimate vowel is schwa. Based on these examples, and based on an examination of the entire data, I have revised Davies' (1980) penultimate-stress hypothesis as follows:

(78) Revised penultimate-stress hypothesis

a. Stress falls on the final syllable if it is heavy (has a diphthong or a complex coda)

b. If the penultimate vowel is $V_\emptyset$ and the final vowel is V, stress falls on the final vowel. ($V_\emptyset \rightarrow$ [ə])

c. Otherwise, stress is penultimate

There are at least two types of examples that could distinguish the revised penultimate-stress hypothesis in (78) from the sonority hypothesis in (73). Consider first words that have a final light syllable with a vowel that is more sonorous than the penultimate (non-schwa) vowel, such as [kɨdolmán] and [gían]. Here the sonority hypothesis predicts final stress but the revised penultimate-stress hypothesis predicts penultimate stress. Consider now words that have a final heavy syllable with a vowel that is not more sonorous than the penultimate vowel, such as [rálemp^{h}]. Here the sonority hypothesis predicts penultimate stress but the revised penultimate-stress hypothesis predicts final stress.

The result is that the two hypotheses are nearly equally successful, with 6 examples that support the sonority hypothesis and 7 examples that support the revised penultimate-stress hypothesis (79).[14] As far as I can tell, each theory would have to mark the counterexamples to it as exceptions.

(79) a. Examples supporting the sonority hypothesis:[15]
kɨdolmán, uréf, k^{h}uám, báwunt, rálemp^{h}, wáimant (6)

b. Examples supporting the revised penultimate-stress hypothesis:
gíaŋ, múmon, kíe, wúse, mímor, gúlo, gúḷo (7)

Since there is a successful alternative to the sonority hypothesis – the hypothesis in (78) with 6 exception marks – I conclude that there is no decisive evidence for sonority-driven stress in Kobon. Given the methodological principle in (12), this tie between the sonority-driven analysis and the analysis that respects encapsulation is a positive result for modularity.

[14] Despite the very different predictions that the two hypotheses make, there were not many distinguishing examples (13/550). One reason is that surprisingly many words in the data have [a] as their penultimate vowel (and such examples are usually unhelpful in distinguishing between the two hypotheses). Another reason is that verbs behave differently (stress is determined based on the identity of the suffixes) and so were not considered by Kenstowicz or in my examination.

[15] I have omitted the following examples that appear in Kenstowicz's paper as support for the sonority hypothesis: [kiá], because the same word is given with penultimate stress in Davies (1980) (of course, the version with penultimate stress was not considered as supporting the penultimate-stress hypothesis); and [sióg], because its surface form is reported to be [sióŋk^{h}], which is a heavy syllable.

1.5.3 Reanalysis of Nanti

1.5.3.1 Introduction

Crowhurst and Michael (2005; C&M) argued that verbal stress in Nanti is sensitive to the fine-grained sonority scale in (80).

(80) Vowel height scale for Nanti

a > e, o, u > i

Nanti's vowel inventory consists of the five vowels in (80), and each of them contrasts for length (e.g., [aa] vs. [a]). The basic verbal stress pattern of Nanti is rhythmic iambic: in verbs with only simple CV syllables, stress falls on every second syllable starting from the second syllable of the word. Final CV syllables are never stressed. The examples in (81) illustrate C&M's foot-based analysis, where ']' marks the right edge of the prosodic word.

(81) Basic rhythmic pattern

a.	no.né.he.ro	(no.né).he].ro	'I will see it'
b.	o.kò.wo.gó.te.ro	(o.kò)(wo.gó).te].ro	'she harvests it'

The basic iambic pattern is reportedly overridden by several factors, including syllable weight, stress clash, and vowel height. The effect of vowel height according to the scale in (80) is partly illustrated in (82), where the underlined feet show surprising trochaic stress. Since all Nanti vowels contrast for length, attraction to the short [a] is not due to length. According to C&M, sonority is the cause: the vowel [a] can attract stress from its less sonorous sisters.

(82) Surprising trochaic stress

a. (<u>nà.bo</u>)(bu.tái)].ro 'I re-sew it'

b. (pi.pò)(<u>ká.kse</u>)].na 'you came to me'

I revisit the status of Nanti's verbal stress as a sonority-driven system by examining the effects of Nanti's morphosyntax and phonological structure on stress. I will develop a cyclic analysis of Nanti in which apparent sonority-sensitivity is due to a sonority-blind underlying system that is rendered opaque on the surface. For example, the surface forms in

(82) are underlyingly (83a)-(83b). In (83a), vowel deletion applies after stress assignment and makes an underlying iambic foot look trochaic on the surface. In (83b), the suffix -ak is lexically stressed and keeps its stress on the surface.

(83) a. /no-abobu-áh-i=ro/ → (no.à).bo... → nà.bo
     b. /pi-pok-ák-e=na/

My analysis of Nanti will replace the scale in (80) with the binary scale ({a, e, o, u} > i). On the assumption that [i] is the empty vowel in Nanti, the new scale would make Nanti consistent with the Stress-Encapsulation Universal. In addition to simplifying C&M's scale, I will argue that my analysis has two advantages over C&M's analysis, stated in (84) and described below.

(84) Advantages of the present analysis over C&M
     a. Supporting evidence
     b. No scale reversal

The first advantage of the analysis is that it accounts for the distribution of stress in several words which C&M's analysis makes incorrect predictions for. As we will see, the distribution of stress in those words is surprising if stress is sensitive to the sonority scale in (80). The second advantage is that the analysis eliminates an unusual property of Nanti's stress system under C&M's account. According to C&M, main stress and secondary stress are sensitive to different, partly opposing weight scales. For example, while [eN] is heavier than [a] for secondary stress, [a] is heavier than [eN] for main stress. Though scales can reportedly differ between processes in a given language, scale reversal is otherwise unattested (see, for example, Gordon 2006, who reports no scale reversals in a survey of the stress systems of around 400 languages). We will see that once morphology is considered, scale reversal for Nanti can be avoided.

Before turning to the analysis, a note on data. For the purposes of this analysis, I created a corpus of Nanti that combines the stress data from C&M with the several stress-marked verbs in Michael 2008. The corpus consists of 122 words, all collected by Lev Michael. The stress data in C&M are morphologically unanalyzed, and no analyzed stress corpus of Nanti is publicly available (Lev Michael, p.c.). To segment the data morphologically, I relied on the description of Nanti's verbal morphosyntax and morphophonology in Lev Michael's grammatical sketch of Nanti (Michael 2008). This sketch included sufficient information for segmenting most of the verbs, but unsegmentable verbs remained, either because a root or an affix did not appear in Michael 2008. Since my theory of Nanti stress crucially relies on morphology, testings its predictions using unsegmented words is difficult. For this reason, I will first focus on words which I was able to segment exhaustively. I will present the main ingredients of the analysis, which is able to correctly account for the distribution of stress in all such words. After presenting the main ingredients of the analysis, I will discuss its advantages over C&M's sonority-driven analysis. Finally, I will discuss some refinements to the analysis that are required to account for stress in some additional examples.

1.5.3.2 Analysis: main ingredients

The starting point of my analysis is Michael's (2008) observation that the Nanti verb is divided into two morphological domains: the stem domain, which consists of prefixes and the root, and the suffix domain, which consists of all suffixes, as illustrated in (85). The two domains are distinguished by their phonologies. For example, vowel hiatus is resolved by glide formation or deletion in the stem domain but by [t]-epenthesis in the suffix domain, and a consonant cluster is resolved by consonant-deletion in the stem domain but by [a]-epenthesis in the suffix domain. I propose a similar division of labor for stress: the stem domain is constructed first with its own phonology, then suffixes are added one-by-one, each triggering a pass through the cyclic phonology.

(85) $\underbrace{\text{[SUBJECT-IRREALIS-CAUSATIVE-ROOT]}}_{\text{Stem domain}}\underbrace{\text{-DERIVATION-INFLECTION=OBJECT}}_{\text{Suffix domain (cyclic)}}$

In the stem domain, stress is simple. First, bisyllabic feet are constructed from left to right. Then, iambic main stress is assigned to the rightmost foot and iambic secondary stress is assigned to all other feet. In the examples I will consider first, all syllables are of equal weight and the suffix domain (developed later) does not alter stress assigned in the stem domain. Following C&M, I assume that object clitics (which follow '=') are extra-metrical, and I will ignore them when presenting the derivations. To see how stress is assigned in the stem cycle, consider first the following word:

(86) no.né.he.ro
no-neh-e=ro
1.SG-SEE-IRREAL=OBJ
'I will see it'

Here, the stem consists of the subject marker /no-/ and the root /neh/, and there is a single cyclic suffix, the irrealis marker /-e/. As shown in (87), the input to the stem cycle is /no-neh/. In the stem cycle, foot construction applies and is followed by stress assignment. The output of the stem cycle is [(no.néh)]. In the next cycle, the irrealis suffix is added. Stress is transferred to the next cycle but, for reasons that will become clear below, I assume that foot structure gets destroyed and must be reconstructed in each cycle (for the current word, foot reassignment creates the same foot structure that had been erased). Assuming that the extra-metrical object clitic /=ro/ is added in the end with no effect, the final output is [(no.né).he=ro].

(87)

Word	[no.né.he.ro]
Structure	/no-neh-e=ro/
Cycle I (stem)	noneh
Foot construction	(no.neh)
Stress	(no.néh)
Cycle II	no.néh-e
Foot construction	(no.né).he
Output	[(no.né).he=ro]

The next step is to introduce syllable-structure rules (still considering words in which suffixes do not alter stress). Nanti's surface syllable structure is restricted to syllables of the forms CV(V)(N), where VV is a long vowel or a diphthong. Nasal consonants are the only permissible codas and they must be homorganic with a following stop. Onsets are required except in word-initial position. When illicit syllables would arise through affixation, various processes apply. As mentioned above, different processes apply in the stem domain and in the suffix domain:

(88) Syllable-structure rules (stem)

a. Glide formation (i $\rightarrow$ j / # _ V)

b. Vowel deletion (V $\rightarrow \emptyset$ / _ V)

c. Consonant deletion (C $\rightarrow \emptyset$ / _ C)

(89) Syllable-structure rules (suffix)

a. Vowel epenthesis ($\emptyset \rightarrow$ a / C _ C)

b. Consonant epenthesis ($\emptyset \rightarrow$ t / V _ V)

c. h-deletion (h $\rightarrow \emptyset$ / V _ V)

I group the processes in each domain under 'Syllable rules', and I assume that they apply in the order specified in (88) and (89). My preliminary grammar for Nanti is the following:

(90) A fragment of Nanti grammar (to be revised below)

- Stem rules:

  1. Foot construction
  2. Stress
  3. Syllable rules

- Suffix (cyclic) rules:

  4. Syllable rules
  5. Foot construction

In the suffix domain, at least one syllable rule – vowel epenthesis – will need to apply before foot construction, since epenthesis changes the number of syllables in a way that is relevant to stress assignment, as we will see below. In my description of the modular architecture in (15), I remained agnostic about which processes count as interface-construction processes and can precede stress processes in a given cycle. Since interface representations consist of syllable structure, it seems harmless to assume that a process like vowel epenthesis counts as an interface-construction process, since vowel epenthesis directly modifies syllabic representations. On this assumption, the ordering of processes in the suffix domain where epenthesis precedes stress assignment is consistent with the non-interactionist architecture in (15).

We can now see the first main difference between my analysis and C&M's analysis. Consider the word in (91), in which stress is surprisingly word-initial. According to C&M, the foot (ja.nu) is trochaic because [a] attracts stress from a less sonorous vowel within the foot (which would have been assigned iambic stress otherwise).

(91) já.nu.ti
i-anu-i
3.MASC.SG-WALK-REAL
'he walked'

The derivation of this word under the grammar in (90) is given in (92). In the stem cycle, stress is assigned to /a/, which heads an iambic foot at this stage of the derivation. Glide formation (as one of the syllable rules) applies after stress assignment and reduces the number of syllables, making main stress appear on the initial syllable. Stress (but not foot structure) is sent to the following cycle, in which the realis suffix /-i/ is added. After the application of t-epenthesis and foot construction (which still has no observable effect), the final output is [(já.nu).ti].

(92) Derivation of (91)

Word	[já.nu.ti]
Structure	/i-anu-i/
Cycle I (stem)	i-anu
Foot construction	(i.a).nu
Stress	(i.á).nu
Syllable rules	(já).nu
Cycle II	já.nu-i

Syllable rules	já.nu.ti
Foot construction	(já.nu).ti
Output	[(já.nu).ti]

This analysis of [(já.nu).ti] is very different from C&M's analysis. On this analysis, trochaic stress is the result of an opaque interaction between stress and glide formation. Importantly, the first syllable gets stress for reasons that have nothing to do with the sonority of the second vowel of the word. It is easy to construct examples of hypothetical roots in the same morphosyntactic environment that would distinguish the two analyses. For example, the present analysis predicts the hypothetical root /una/ (as opposed to /anu/) to surface as [(jú.na).ti] in the same environment, whereas C&M predict it to surface as [(ju.ná).ti]. None of the four other examples in the corpus that show glide formation are informative.[16]

While glide formation applies in a small number of words, other processes with a similar effect are more widespread. Consider vowel deletion, which applies in at least 20/122 words. Vowel deletion removes a pre-vocalic vowel, as illustrated in (93). The first-person singular marker is /no-/, and it loses its vowel before the vowel-initial stem /aʃint/. The derivation of this word is provided in (94).[17]

(93) ná.ʃin.tem.pa.ra
no-aʃint-empa=ra
1.SG-OWN-IRREAL=SUBORD
'I will own'

[16][já.mu.ta.kòi.ga.kse.na] (with the underlying stem /i-amu/) is similar to [(já.nu).ti]; for [jóo.ga.kse.ro], both analyses predict initial stress: my analysis predicts initial stress because the root is monosyllabic (/i-oog/) and C&M predict initial stress because the initial vowel is longer (thus heavier, for sonority-independent reasons) from the second vowel. The two other examples, [jo.bíi.kài.ga.kse] and [ja.máa.ta.kòi.ga.nà.kse], are also unhelpful, because the initial vowel is followed by a long vowel. C&M predict a long vowel to win over a short vowel, and my analysis will predict initial stress to get deleted before a long vowel in those cases for independent reasons discussed below.

[17]Feet in word-final position are never assigned secondary stress, as discussed later. For now, I will ignore the possibility of assigning a foot to the last two syllables of the prosodic word.

(94)

Word	[ná.ʃin.tem.pa.ra]
Structure	/no-aʃint-empa=ra/
Cycle I (stem)	no-aʃint
Foot construction	(no.a).ʃint
Stress	(no.á).ʃint
Syllable rules	(ná).ʃint
Cycle II	ná.ʃint-empa
Syllable rules	-
Foot construction	(ná.ʃin).tem.pa
Output	[(ná.ʃin).tem.pa=ra]

For C&M's analysis, which ignores morphophonological structure and assigns stress to surface representations, trochaic stress in (93) requires a complication. VN syllables seem heavier than V syllables in various other examples, such as (95), where secondary stress falls on the initial syllable rather than the second.

(95) [pìŋ.ka.mo.sói.ga.kse]

According to C&M, what distinguishes (95) from (93) is secondary stress as opposed to main stress: VN syllables are heavier than V syllables for secondary stress, which is why stress shifts to the first syllable in (95). For main stress, in contrast, V_i wins over V_jN if V_i is more sonorous than the V_j. Since [a] is more sonorous than [i], it attracts stress from the syllable [in] in (93). In my analysis, V can never win over VN. What explains the surprising trochaic stress in (93) is the opaque interaction between stress and vowel deletion (my own analysis of examples like (95) will come later). Importantly, (93) illustrates one way in which opposing scales for main and secondary stress can be avoided. As in the case of glide formation, the two analyses make very different predictions for examples with vowel deletion, and all other examples in the corpus with vowel deletion are consistent with my analysis. Since some of those examples are more complicated, it would be better to postpone that discussion until after we have seen the effects of stressed suffixes on stress, to which I now turn.

As mentioned above, the verbal stem is followed by a sequence of suffixes, which Michael (2008) classifies into derivational and inflectional suffixes (this classification will not play any role in my analysis). Derivational suffixes (which are optional) are followed by inflectional suffixes. The inflectional reality-status marker is obligatory and always comes last:

(96) Order of inflectional suffixes

VERB QUANTIFIER- ARGUMENT NUMBER - DIRECTIONAL - ASPECT - REALITY STATUS

I assume that Nanti suffixes are divided into two stress classes: lexically stressed suffixes, which come with initial main stress, and lexically unstressed suffixes. As a working intensional characterization, I assume that all suffixes are lexically stressed, except for /-an/, /-ut/, and the four reality-status markers (/-a/, /-e/, /-i/, and /-eNpa/).[18] A list of the suffixes discussed below is provided in (97). While most suffixes contain the low vowel [a], some unstressed suffixes contain [a] and not all stressed suffixes contain [a].

(97) a. Stressed suffixes: /-áNt/, /-híg/, /-áko/, /-ápah/, /-áh/, /-ák/, /-há/

b. Unstressed suffixes: /-an/, /-ut/, {/-a/, /-e/, /-i/, /-eNpa/}

Since main stress is always assigned to the stem, when the first stressed suffix is added in the derivation, the result is two main-stressed syllables. In a situation like this, the rule of STRESS REDUCTION in (98) reduces one of the two main stresses to secondary stress. After STRESS REDUCTION creates a secondary stress, another rule removes that secondary stress if it ends up in a doubly-stressed foot or in a word-final foot.

(98) STRESS REDUCTION (to be revised below)

Given two main-stressed syllables:

[18] An alternative formulation of the distinction is that one class of suffixes attracts stress and requires main stress right after the preceding morpheme boundary (rather than being lexically stressed). This alternative formulation would make different predictions for consonant-initial suffixes (e.g., the frustrative suffix /-be/) in situations where they follow a consonant. In such situations, [a] would be epenthesized, possibly into a position right after the morpheme boundary (e.g., /C+be/ → [C+abe]). If suffixes are stress-attracting and require stress after the morpheme boundary, epenthetic [a] would get stress. However, if they are lexically stressed, as in the formulation in the main text, the suffix itself and not the epenthetic vowel would get stress. My choice of the lexical-stress formulation was arbitrary. I have not been able to find examples with sufficient information for distinguishing between the two formulations.

- If either has a shorter nucleus, reduce its stress;[19]
- Otherwise (if the two are of equal length), reduce the stress on the stem.

(99) 2-STRESS DELETION

Delete secondary stress from doubly-stressed or final feet.

To illustrate the effect of these two rules, consider the suffix /-ák/ in the following word:

(100) pi.pò.ká.kse.na

pi-pok-ak-e=na

2.PL-COME-PERF-REAL=1.OBJ

'You came to me'

As shown in the derivation below, main stress ends up on the second syllable in the stem cycle. When the stressed suffix is added, the two equally-short stressed syllables compete for main stress. According to STRESS REDUCTION, the suffix wins in this case and stress in the stem is reduced to secondary stress. The result is secondary stress on the stem and main stress on the suffix. On C&M's analysis, main stress falls on the rightmost strongest syllable in the word (simplifying). Since [a] is the most sonorous vowel in (100), it gets main stress. On the other hand, on my analysis main stress falls on [a] because this vowel is a suffixal vowel that comes with lexical stress. We will see examples that distinguish between the two analyses later on.

(101) Derivation of (100)

Word	[pi.pò.ká.kse.na]
Structure	/pi-pok-ak-e=na/
Cycle I (stem)	pi-pok
Foot construction	(pi.pok)
Stress	(pi.pók)
Syllable rules	-

[19]A nucleus is short if it consists of a short vowel. It is long if it consists of a long vowel or a diphthong.

Cycle II	pi.pók-ák
Syllable rules	-
Foot construction	(pi.pó).(kák)
Stress reduction	(pi.pò).(kák)
2-stress deletion	-
Cycle III	pi.pò.kák-e
Syllable rules	-
Foot construction	(pi.pò)(ká.ke)
Stress reduction	-
2-stress deletion	-
Output	[(pi.pò)(ká.kse)=na]

When the stem contains a long vowel or a diphthong that gets main stress, as in (102), the suffix /-ák/ loses, as shown in the derivation in (103). In this case, secondary stress on the suffix ends up in a word-final foot and thus gets deleted through 2-STRESS DELETION.

(102) i.pái.ta.kʃi.ri
i-pait-ak-i=ri
3.SG-NAME-PERF-REAL=3.MASC.OBJ
'He named him'

(103) Derivation of (102)

Word	[i.pái.ta.kʃi.ri]
Structure	/i-pait-ak-i=ri/
Cycle I (stem)	i-pait
Foot construction	(i.pait)
Stress	(i.páit)
Syllable rules	-
Cycle II	i.páit-ák
Syllable rules	-

Foot construction	(i.pái).(ták)
Stress reduction	(i.pái).(tàk)
2-stress deletion	-
Cycle III	i.pái.tàk-i
Syllable rules	-
Foot construction	(i.pái).(tà.ki)
Stress reduction	-
2-stress deletion	(i.pái).(ta.ki)
Output	[(i.pái).(ta.kʃi)=ri]

We are now ready to see more complicated examples where vowel deletion makes an underlying iambic foot trochaic on the surface. These example will also demonstrate the necessity of re-footing. Consider the word in (104), given by C&M as an indication that the greater sonority of [a] relative to [o] can affect stress. (The suffix /-áh/ is a regressive perfective suffix.) On C&M's analysis, the word is footed as in (105). The first foot gets trochaic stress due to sonority and the second gets iambic stress due to vowel length. This word poses a new challenge for the vowel deletion analysis: if [a] gets stress because it heads an underlying iambic foot, we would expect the two syllables that follow [a] to form another foot and get at least one stress. The fact that [a] is followed by two stressless syllables is surprising. This is where re-footing comes into play.

(104) nà.bo.bu.tái=ro

no-abobu-ah-i=ro

1.SG-SEW-REG.PERF-REAL=3.NONMASC.OBJ

'I re-sew it'

(105) (nà.bo)(bu.tái)=ro

In the derivation of (104) below, iambic secondary stress is assigned to /a/ in the stem cycle and iambic main stress is assigned to /bu/. Vowel deletion eliminates a syllable, making the first foot monosyllabic. Foot structure is not transferred to the next cycle, in which the stressed suffixed is added. When bisyllabic feet are re-constructed, the syllable whose nucleus is /a/ is now the left member of a foot. The syllable /bu/, which received iambic stress as the right member of a foot in the previous cycle, now joins another stressed syllable as the left member of a foot. Both stresses are equally short and the first of them belongs to the stem, so it gets reduced to secondary stress. This secondary stress cannot be tolerated in the same foot as another main stress and it gets deleted by secondary-stress deletion.

(106)

Word	[nà.bo.bu.tái.ro]
Structure	/no-abobu-ah-i=ro/
Cycle I (stem)	no-abobu
Foot construction	(no.a)(bo.bu)
Stress	(no.à)(bo.bú)
Syllable rules	(nà)(bo.bú)
Cycle II	nà.bo.bú-áh
Syllable rules	nà.bo.bú.táh
Foot construction	(nà.bo)(bú.táh)
Stress reduction	(nà.bo)(bù.táh)
2-stress deletion	(nà.bo)(bu.táh)
Cycle III	nà.bo.bu.táh-i
Syllable rules	nà.bo.bu.tái
Foot construction	(nà.bo).(bu.tái)
Stress reduction	-
2-stress deletion	-
Output	[(nà.bo).(bu.tái)=ro]

The interaction between re-footing and secondary-stress deletion makes the following prediction. A main stress that gets reduced to secondary stress before an adjacent main-stressed syllable will survive as secondary stress depending on the parity of the number of syllables that precede it. If that secondary stress is preceded by an even number of syllables, as in (106), it gets deleted. In a situation where it is preceded by an odd number of syllables, it would not be parsed into the same foot as the following main stress and would therefore survive. This prediction is illustrated by the minimal pair of verbs in (107), which differ morphologically only with respect to their reality-status markers. In the second verb, the irrealis prefix /r-/ blocks glide formation. The stressed suffix /-áko/ is preceded by an even number of syllables in (107a) but by an odd number of syllables in (107b). In both cases it loses to a following main-stressed long vowel, but it only survives as secondary stress in (107b). The full derivations of both words are provided side-by-side in (108).

(107) a. jà.mu.ta.kói.ga.kse.na
i-amu-áko-hiǵ-ák-e=na
3-HELP-APPL-PL-PERF-REAL=1.OBJ
'they helped us with something else'

b. i.rà.mu.tà.kói.ga.ksem.pa
i-r-amu-áko-hiǵ-ák-empa
3-IRREAL-HELP-APPL-PL-PERF-IRREAL
'they will help someone with something'

(108) Interaction of re-footing and secondary-stress deletion

Word	[jà.mu.ta.kói.ga.kse.na]	[i.rà.mu.tà.kói.ga.ksem.pa]
Structure	/i-amu-áko-hiǵ-ák-e=na/	/i-r-amu-áko-hiǵ-ák-empa/
Cycle I (stem)	i-amu	i-r-amu
Foot construction	(i.a).mu	(i.ra).mu
Stress	(i.á).mu	(i.rá).mu
Syllable rules	(já).mu	-
Cycle II	já.mu-áko	i.rá.mu-áko
Syllable rules	já.mu.tá.ko	i.rá.mu.tá.ko
Foot construction	(já.mu)(tá.ko)	(i.rá)(mu.tá).ko
Stress reduction	(jà.mu)(tá.ko)	(i.rà)(mu.tá).ko
2-stress deletion	-	-
Cycle III	jà.mu.tá.ko-hiǵ	i.rà.mu.tá.ko-hiǵ
Syllable rules	jà.mu.tá.koiǵ	i.rà.mu.tá.koiǵ
Foot construction	(jà.mu)(tá.koiǵ)	(i.rà)(mu.tá)(koiǵ)

Stress reduction	(jà.mu)(tà.koíg)	(i.rà)(mu.tà)(koíg)
2-stress deletion	(jà.mu)(ta.koíg)	-
Cycle IV	jà.mu.ta.koíg-ák	i.rà.mu.tà.koíg-ák
Syllable rules	-	-
Foot construction	(jà.mu)(ta.koí)(gák)	(i.rà)(mu.tà)(koí.gák)
Stress reduction	(jà.mu)(ta.koí)(gàk)	(i.rà)(mu.tà)(koí.gàk)
2-stress deletion	-	(i.rà)(mu.tà)(koí.gak)
Cycle V	jà.mu.ta.koí.gàk-e	i.rà.mu.tà.koí.gak-empa
Syllable rules	-	-
Foot construction	(jà.mu)(ta.koí)(gà.ke)	(i.rà)(mu.tà)(koí.ga)(kem.pa)
Stress reduction	-	-
2-stress deletion	(jà.mu)(ta.koí)(ga.ke)	-
Output	[(jà.mu)(ta.koí)(ga.ke)=na]	[(i.rà)(mu.tà)(koí.ga)(kem.pa)]

The final ingredient of the analysis is TROCHAIC SHIFT, a group of processes that shift iambic stress the to first syllable of the foot under certain conditions. Here I inherit from C&M a strength scale that determines whether TROCHAIC SHIFT applies. My scale is a simplification over C&M in two ways. First, it reduces the sonority hierarchy to the binary distinction (V > i). Second, as mentioned above, C&M propose partly opposing scales for main and secondary stress: V_i is stronger than V_jN for main stress if V_i is more sonorous than V_j, but VN is always stronger than V for secondary stress. My proposal avoids this reversal. For now, I will use the unified scale in (109) for both main and secondary stress, though a refinement (which still avoids reversal) will be introduced below. According to (110), stress shifts left to a stronger syllable, or to an equally strong syllable in case of a clash with a following stressed short vowel.

(109) Strength scale

VV > VN > iN > V > i

(110) TROCHAIC SHIFT (to be updated below)

In the environment $(\sigma_1\sigma_2)\sigma_3$ where σ_2 is stressed, shift stress from σ_2 to σ_1 if one of the following conditions holds:

- $\sigma_1 > \sigma_2$
- $\sigma_1 = \sigma_2$ and σ_3 has a stressed short vowel

Trochaic shift is illustrated by the examples in (111). For these examples, there is sufficient morphological information to tell that no processes apply that can create trochaic feet from underlying iambic feet (like glide formation or vowel deletion). In (111a), the syllables of the first foot are equally strong and there is no clash, so trochaic shift does not apply. In (111b), stress shifts from the second syllable to the stronger first syllable even though there is no clash. In (111c), the first two syllables are equally strong and trochaic shift applies due to clash with a following stressed short vowel. In (111d), the first syllable is weaker than the second syllable, so trochaic shift does not apply despite a clash with a short stressed vowel. In (111e), the first two syllables are equally strong but the following stressed nucleus is long and does not trigger trochaic shift.

(111) a. no.nè.bi.tá.kse
no-nebi-ak-e
1-REQUEST-PERF-REAL
'I requested'

b. pìŋ.ka.mo.sói.ga.kse
pi-ŋ-kamoso-hig-ak-e
2-IRREAL-VISIT-PL-PERF-IRREAL
'you.PL will have visited'

c. nòŋ.ksen.tá.kse.ro
no-ŋ-kent-ak-e=ro
1-IRREAL-PIERCE-PERF-IRREAL=3.NONMASC.OBJ
'I will have pierced it (with an arrow)'

d. iŋ.ksèn.tá.kse.ro
i-ŋ-kent-ak-e=ro
3-IRREAL-PIERCE-PERF-IRREAL=3.NONMASC.OBJ
'he will have pierced it (with an arrow)'

e. noŋ.kàn.tái.ga.kse

no-ŋ-kant-hig-ak-e

1-IRREAL-SAY-PL-PERF-IRREAL

'we will have said'

As before, this analysis is substantially different from C&M's analysis and it is easy to imagine distinguishing examples. For example, C&M take the pair (111c) and (111e) to be an indication that the sonority difference between [a] and [o] is a factor that affects stress: there is no shift in (111e) because stress does not shift to weaker vowels and [o] is weaker than [a]. However, the two examples do not form a minimal pair. They differ with respect to various factors, such as the identity of the root, the plurality of the subject, and the length of the clashing syllable's nucleus. The motivation behind the restriction to short vowels in my formulation of trochaic shift is that length could play a role in other examples, such as [no.tsà.róo.ga.kse], [mu.tà.kói.ga.kʃim.pi.ra], and perhaps also [o.sà.ráan.tai.ga.kse] and [no.kà.rái.ga.kse] (though I have not been able to establish the morphology of the latter two examples).[20] The two analyses diverge in their predictions for examples like the hypothetical verb 'we will have pierced it (with an arrow)', the plural variant of (111c):

(112) Divergent predictions for hypothetical 'we will have pierced it (with an arrow)'

no-ŋ-kent-hig-ak-e=ro

1-IRREAL-PIERCE-PL-PERF-IRREAL=3.NONMASC.OBJ

a. Prediction of my analysis (long vowel does not trigger shift):

[(noŋ.kèn)(tái.ga).ke=ro]

b. Prediction of C&M (long vowel triggers shift):

[(nòŋ.ken)(tái.ga).ke=ro]

Unfortunately, this example was not available in the corpus. The only similar example I was able to find is the following:

(113) [nò.ne.háem.pi], 'I will see you again'

[20] An alternative to length that is also consistent with these examples is that certain suffixes like /-hig/ and long stem-vowels do not trigger shift.

At first sight, this example seems to corroborate C&M's analysis: a trochaic foot with two equally-strong syllables is followed by a long nucleus ([ae]), and it is possible to verify that no processes like vowel deletion apply that can create trochees. For trochaic shift to apply to the foot (no.nè) on my analysis, the following syllable must have a short vowel. It turns out that the vowel [a] in [ae] is the short vowel of the regressive suffix /-áh/. It is only in a later cycle that it forms a diphthong with the irrealis suffix [e]. If trochaic shift applies cyclically, the right prediction is made.[21] The derivation is as follows:

(114)

Word	[nò.ne.háem.pi]
Structure	/no-neh-ah-e=mpi/
Cycle I (stem)	no-neh
Foot construction	(no.neh)
Stress	(no.néh)
Syllable rules	-
Cycle II	no.néh-áh
Syllable rules	-
Foot construction	(no.né)(háh)
Stress reduction	(no.nè)(háh)
2-stress deletion	-
Trochaic shift	(nò.ne)(háh)
Cycle III	nò.ne.háh-e
Syllable rules	nò.ne.háe
Foot construction	(nò.ne)(háe)
Stress reduction	-
2-stress deletion	-
Trochaic shift	-
Output	[(nò.ne)(háem)=pi]

This concludes the main ingredients of the analyses. The grammar we have so far is

[21] In previous examples, long nuclei where introduced in a single cycle, so the application of trochaic shift in a cyclic fashion would not make any difference there.

summarized in (136), and it accounts for most of the examples that have motivated the sonority hierarchy for Nanti, even before the refinements presented later. A detailed discussion of additional individual examples will come later. Before that, I would like to discuss the advantages of the present analysis over C&M's analysis.

(115) A fragment of Nanti grammar (to be revised below)

- Stem rules:
  1. Foot construction
  2. Stress
  3. Syllable rules
- Suffix (cyclic) rules:
  4. Syllable rules
  5. Foot construction
  6. Stress reduction
  7. 2-stress deletion
  8. Trochaic shift

1.5.3.3 Advantages over the sonority-driven stress analysis

Supporting evidence

On C&M's analysis, main stress is computed globally. Its default position is the rightmost strongest syllable in the word: in words with a unique strongest syllable, that syllable gets main stress. Otherwise – if there are multiple strongest syllables – the rightmost of them wins.[22]

The minimal pair in (116) supports a role for morphology in stress assignment and poses a problem for C&M's morphology-blind analysis. In both examples in (116), the final syllable is a unique strongest syllable. C&M predict final stress in both cases. In particular, they incorrectly predict final stress in (116b): *[(i.pò)(ka.pái)]. On my analysis,

[22]This default can be overridden by several factors, which are not relevant in the discussion to follow.

the difference between the two words comes from their morphology. In (116b), main stress falls on the stressed bisyllabic suffix /-ápah/, which wins over an equally-long nucleus in the stem. In (116a), the suffix /-an/ is unstressed (in fact, it is never stressed in the corpus) and main stress falls on the stressed suffix /-áh/ which wins over an equally-long nucleus in the stem. Without reference to morphology, C&M have no way to distinguish between these two words in terms of the location of main stress. The divergent predictions are shown in (117).

(116) a. i.ʃì.ga.nái

i-ʃig-an-ah-i

3-RUN-ABL-REG-REAL

'he ran away'

b. i.pò.ká.pai

i-pok-apah-i

3-COME-ADL-REAL

'he is coming towards'

(117) Divergent predictions for (116b)

a. Prediction of my analysis (correct):

[(i.pò)(ká.pai)]

b. Prediction of C&M (incorrect):

*[(i.pò)(ka.pái)]

The minimal pair in (118) poses another problem for C&M's analysis. The only phonological difference between the two words on the surface (aside from stress) is the irrealis prefix /N-/ in (118b), which creates a VN syllable.[23] On C&M's analysis, a VN syllable is strictly stronger than a V syllable with an identical vowel. Since main stress falls on the rightmost strongest syllable, C&M incorrectly predict initial main stress in (118b). On my analysis, main stress is computed locally and normally only shifts rightwards from the stem. In the stem cycle, main stress is assigned to the final syllable of the quadrisyllabic stem (o(ŋ).kò)(wo.gó) in both words, regardless of the presence of the irrealis prefix,

[23]The realis suffix /-i/ lowers to [e] in some environments. This is presumably what happens in (118a).

which does not change the syllable count. Since there are no stressed suffixes, main stress remains on that syllable. Trochaic shift creates the initial trochaic foot in (118b). The divergent predictions are shown in (119). This minimal pair is a reason to think that C&M's 'rightmost strongest' condition for main stress is incorrect. Many of the examples that motivated their condition were examples with suffixes with the vowel [a], which on my analysis are lexically stressed and can attract stress rightwards from the stem.

(118) a. o.kò.wo.gó.te.ro
o-kowogo-i=ro
3.F-HARVEST-REAL=3.NONMASC.OBJ
'she harvests it'

b. òŋ.ko.wo.gó.te=ro
o-ŋ-kowogo-e=ro
3.F-IRREAL-HARVEST-IRREAL=3.NONMASC.OBJ
'she will harvest it'

(119) Divergent predictions for (118b)

a. Prediction of my analysis (correct):
[(òŋ.ko)(wo.gó).te=ro)]

b. Prediction of C&M (incorrect):
*[(óŋ.ko)(wo.gò).te=ro)]

The pair in (120) illustrates another problem with C&M's analysis for main stress. We already saw (120a) and its derivation above in (104)-(106): according to my analysis, main stress on the final syllable comes from the lexically stressed suffix. (120b) can be analyzed in a similar fashion, assuming that the suffix /-há/ is lexically stressed. My assumption that (120b) includes that suffix is based on the following reasoning. The root *kamoso* means 'visit' while the meaning of the verb in (120b) is 'visit by water'. The inflectional morphemes /no/ and /i/ could also be identified, and Nanti has a suffix /-há/ which means 'water'. I do not have independent evidence that /h/ could be pronounced as [w] intervocalically, but this seems like a reasonable assumption given that /h/ is sometimes pronounced as the front glide [j] intervocalically after a front vowel (e.g. /i-neh-ako-e=ro/

→ [i.nè.já.ko.te.ro]). Assuming that the suffix /-há/ is stressed, it attracts main stress from the stem and nothing special needs to be said beyond what I have already said for (120a).[24] For C&M, main stress on the final foot in (120b) is problematic. On their analysis, the default location of main stress is on the rightmost strongest syllable. But this default can be overridden by a dispreference for stressing word-final feet: in words with multiple strongest syllables, if the rightmost strongest is in a word-final foot, a syllable of equal strength further to the left gets the main stress. C&M invoke that principle to explain why main stress is not penultimate in words like (121) (on my analysis, this principle is not needed for (121): main stress is assigned to the stem and there are no stressed suffixes to attract it (other examples used by C&M to support the principle will be discussed below). This principle predicts that penultimate stress should be avoided in (120b) in favor of the second syllable, which is of equal strength. The divergent prediction is shown in (122).

(120) a. nà.bo.bu.tái=ro
no-abobu-ah-i=ro
1.SG-SEW-REG.PERF-REAL=3.NONMASC.OBJ
'I re-sew it'

b. no.kà.mo.sò.wá.ti
no-kamoso-?ha?-i
1-VISIT-?WATER?-REAL
'I visited (by water)'

(121) no.né.he=ri=ro
no-neh-e=ri=ro
1-SEE-IRREAL=3.MASC.OBJ=3.NONMASC.OBJ
'I will definitely see him'

(122) Divergent predictions for (120b)

a. Prediction of my analysis (correct):
[(no.kà)(mo.sò)(wá.ti)]

[24]To account for the location of secondary stress on [sò] in (120b), I assume a complication to the conditions of application of trochaic shift that prevents it from applying in case of it would create another clash.

b. Prediction of C&M (incorrect):

*[(no.ká)(mo.sò).wa.ti]

For the same reasons, C&M incorrectly predict pen-initial stress in (123), as illustrated in (124). Here, however, I have no sufficient morphological information to test the prediction of my anaysis.

(123) no.sà.me.rè.há.ka

no-?????-ak-a

1-?????-PERF-REAL

'I got a blister'

(124) Prediction of C&M (incorrect):

*[(no.sá)(me.rè).ha.ka]

No scale reversal

As mentioned above, C&M's analysis involves another unusual property in addition to its fine-grained sonority hierarchy for stress. C&M propose partly opposing scales for main and secondary stress: V_i is stronger than V_jN for main stress if V_i is more sonorous than V_j, but VN is always stronger than V for secondary stress. This makes Nanti an exception to the following universal (as also mentioned above, Gordon's 1999/2006 survey of stress in around 400 languages reports no scale reversals):

(125) In a given language, weight scales for stress cannot be reversed

As evidence for scale reversal, C&M provide the examples in (126) (which I present by order of morphological certainty). In all of these examples, main stress falls on [a] and there is another VN syllable in the word. According to my analysis, in the first five examples [a] wins because it is part of the lexically-stressed perfective suffix /-ak/. In the sixth example, [a] gets main stress in the stem cycle before vowel deletion, and there are no stressed suffixes that attract main stress. In sum, an advantage of the sonority-blind analysis is that it eliminates an exception to what otherwise seems like a robust universal.

(126) a. a > oN nòŋ.ksen.tá.kse.ro

no-ŋ-kent-ak-e=ro

1-IRREAL-PIERCE-PERF-IRREAL=3.NONMASC.OBJ

'I will pierce it (with an arrow)'

b. a > oN om̀.pi.tá.kse

o-ŋ-?pi?-ak-e

3-IRREAL-STAY-PERF-IRREAL

'she will stay'

c. a > oN nòm.pi.tá.kse

no-ŋ-?pi?-ak-e

1-IRREAL-STAY-PERF-IRREAL

'I will stay'

d. a > iN pìŋ.ki.sá.kse.ra

pi-ŋ-?kis?-ak-e=ra

2-IRREAL-?HIT?-PERF-IRREAL-SUBORD

'you will hit'

e. a > eN tso.tèŋ.ka.ná.ka.ra

tso-?eŋkani?-ak-a=ra

FINISH-?PASSIVE?-PERF-REAL=SUBORD

'(it) was finished up'

f. a > eN á.pje.ʃi.jèm.pa.ra

o-?ape?-ʃi-empa=ra

3-?DECAY?-LEAF-IRREAL=SUBORD

'it (thatch) will decay'

1.5.3.4 Refinements

The analysis as described above can be improved by minor modifications required by a few words in the data. After I mention a few of them, I will quantitatively evaluate the ability of my analysis to cover the examples in C&M given to support sonority-driven stress.

In the analysis above I proposed the following unified strength scale for main and secondary stress:

(127) Strength scale

VV > VN > iN > V > i

In particular, I proposed that [iN] syllables are stronger than [V] syllables to account for the shift of secondary stress in feet like [(pìŋ.ka)]. There is a single example in the data that suggests that main stress does not obey the same scale:

(128) iŋ.kó.ge

i-ŋ-kog-e

3-IRREAL-WANT-IRREAL

'he will want'

Here, main stress does not shift from [V] to [iN]. This example might suggest that [V] and [iN] count as equally strong for main stress:

(129) Strength scale for main stress

VV > VN > iN, V > i

Another example which the analysis currently does not account for is the following:

(130) o.gó.te.ro

o-ogo-e=ro

3-KNOW-IRREAL=3.NONMASC.OBJ

'she will know it'

This example involves vowel deletion (two underlying /o/'s surface as a short vowel). If iambic feet are assigned to the stem before vowel deletion as I suggested above, my analysis incorrectly predicts stress to be word-initial: (o.ó).go → (ó).go. There are two possible responses to this issue. The first is that when two consecutive vowels are identical, it is the second vowel that gets deleted and not the first one. If this is the case, stress might fall on [gó] because of a constraint that requires the root to be stressed in the stem cycle. This constraint would ensure that stress falls on [gó], which is the only remaining syllable of the root after post-vocalic deletion. An alternative to the root stress-constraint, which is more in line with the serial architecture I have been assuming, is that while vowel deletion normally applies after stress assignment, the deletion of a vowel following an identical vowel applies first, before stress assignment. Post-stress pre-vocalic deletion then takes care of the remaining cases of vowel hiatus and deletes the first vowel among two non-identical vowels. If post-vocalic deletion applies before stress assignment in the example above, stress would correctly fall on the second syllable of the word.

Consider now the following example:

(131) ja.máa.ta.koì.ga.nà.ksc
i-?am?-?ha?-ako-hig-an-ak-i
3-BRING-WATER-APPL-PL-ABL-PERF-REAL
'they floated it away'

Here, main stress falls on [máa] even though it is followed by a stressed suffix with a long nucleus (/-híg/, which becomes part of a diphthong). The rule of stress reduction I presented above said nothing about a situation in which two suffixes compete for main stress. In the example above, it is possible that main stress falls on the stressed suffix /-há/ 'water'. If this is true, stress reduction can be modified so that when two stressed suffixes compete for main stress, the leftmost suffix wins (see Halle & Vergnaud 1987 for a similar analysis of stress in languages like Lithuanian):

(132) STRESS REDUCTION (final)
Given two main-stressed syllables:

- If either has a shorter nucleus, reduce its stress;
- Otherwise (if the two are of equal length), reduce the stress which is not on the leftmost suffix.

With the revised stress reduction, the suffix /-há/ which ends up with a long stressed vowel after [a]-epenthesis and h-deletion, will win over the stressed suffix /-híg/ which surfaces as a diphthong.

Consider now the following word, which shows a surprising misapplication of trochaic shift under clash:

(133) no.gà.pá.ku.ti

no-?ogi?-?ap?-ak-ut-i

1-CAUS-GO-PERF-REV-REAL

'I let it go'

Stress remains on the second syllable even though it clashes with a following short stressed vowel, contrary to the prediction of my rule of trochaic shift. While I have not been able to identify all of the morphemes in this word, I would like to suggest a possible response to trochaic-shift misapplication. The meaning of the word is 'I let it go', which has a causative component. One of the causative prefixes of Nanti is /ogi-/. When this prefix follows another vowel-final prefix and precedes a vowel-initial stem, it triggers two applications of vowel deletion. This creates a situation that we have not seen yet, where two secondary stresses end up in the same foot. The relevant part of the derivation is as follows:

(134)

Word	[no.gà.pá.ku.ti]
Structure (uncertain)	/no-ogi-ap-ak-ut-i/
Cycle I (stem)	no-ogi-ap
Foot construction	(no.o)(gi.ap)
Stress	(no.ò)(gi.áp)
Syllable rules	(nò)(gáp)
Cycle II	nò.gáp-ák
Syllable rules	-
Foot construction	(nò.gá)(pák)
Stress reduction	(nò.gà)(pák)
...	...

If secondary-stress deletion applies next, it would remove stress from the initial syllable, incorrectly feeding trochaic shift:

(135)

Word	[no.gà.pá.ku.ti]
...	...
Stress reduction	(nò.gà)(pák)
2-stress deletion	(no.gà)(pák)
Trochaic shift	(nò.ga)(pák)
Output	*[(nò.ga)(pák)]

While we have not seen reasons to order trochaic shift after secondary-stress deletion, reversing the order in the cycle to block trochaic shift will not help, since it would simply get another chance to apply in the following cycle. Suppose, however, that secondary-stress deletion applies post-cyclically, after all cyclic rules have had their chance to apply. If this is the case, the conditions of application of trochaic shift would never be met by this word, and the initial secondary stress would be correctly removed by secondary-stress deletion. Concretely, the proposed grammar for Nanti can be updated as follows:

(136) A fragment of Nanti grammar (final)

- Stem rules:
  1. Foot construction
  2. Stress
  3. Syllable rules
- Suffix (cyclic) rules:
  4. Syllable rules
  5. Foot construction
  6. Stress reduction
  7. Trochaic shift
- Post-cyclic rules:
  8. 2-stress deletion

With the final proposal for the grammar of Nanti in hand, here is a summary of its coverage of examples of sonority-driven stress in C&M. My goal has been to reduce the sonority scale (a > e, o, u > i) to the binary scale (V > i). The crucial examples are those which C&M account for by appealing to the greater sonority of [a] relative to the mid vowels (e, o, u). There are 37 such examples in C&M, which can be classified according to two factors: the availability of morphological information, and the ingredients of my analysis that replace sensitivity to sonority:

(137) Summary: coverage of crucial examples in C&M by my analysis[25]

# of ex.	Morphological analysis	Coverage	Alternative to sonority
24/37	Full	24/24	Stressed suffixes (15), syllabic processes (5), no clash before a long vowel (4)
1/37	Full	1/1	Specific stipulation
3/37	Almost full	3/3	Syllabic processes (2), stressed suffixes (1)
1/37	Partial	Depends on parse	Stressed suffixes
4/37	Partial	Depends on parse	Leftmost-suffix stipulation
2/37	Partial	Depends on parse	Specific stipulations
2/37	Minimal	Unclear	N/A

The column 'Morphological analysis' indicates the availability of information about the morphological decomposition of the words: 'Full' means that there was sufficient information for an exhaustive morphological analysis, 'Almost full' means that only one or two morphemes remained unidentified, and so on. Whenever the morphology is understood, the analysis correctly derives the distribution of stress, without exception. There are 7 examples where the morphology is only partially understood, and where the success of the analysis depends on the correct parse (examples (19a-i), (23b-i), (24-ii), (25-iii), (30-v), (30-vii), and (31a-iv) in C&M). Finally, there are two examples where information about the morphology is very minimal (examples (9a-i) and (30-vi) in C&M). For these two examples, it is difficult to estimate what the analysis would predict. Given the coverage of the analysis and the supporting evidence discussed above, I tentatively conclude that there is a

[25]The specific stipulation referred to in the second line is that the stressed suffix -ánt never loses its stress through stress reduction. The specific stipulations referred to un the sixth line is that the stem kamant 'tell' consists of the suffix -ánt and that the suffix eNkani is part of the stem.

sufficiently successful alternative to sonority-driven stress in Nanti.

1.5.4 Reanalysis of Mayo stress

Mayo (Foreman and Marten 1973) is cited in Gordon (1999/2006) as a Type II language, where the low vowel attracts stress as opposed to other vowels. The stress generalization provided in the source is the following:

(138) Sonority generalization in Foreman and Marten (1973)

1. The first syllable (of a word) which contains /a/ is stressed
2. When there is no syllable containing /a/ in a word, the first syllable of the word is stressed

The source includes around 400 examples marked for stress, most of which have initial stress. There are two types of examples that could distinguish a naive initial-stress generalization from Foreman and Marten's sonority generalization. Examples that would support initial stress are words with initial stress on a vowel other than [a] that have an [a] in a non-initial syllable; examples with non-initial stress on the first [a] in the word would support the sonority generalization. My count of distinguishing examples resulted in a near-tie between the two generalizations: 13 examples in the source support the sonority generalization but 11 examples support initial stress:

(139) a. Examples supporting the sonority generalization (13)
t^{h}owknátɨ, k^{h}ʌnákʌm, ngɨláng^{w}ow, t^{h}ɔpátɨ, t^{h}eyá, tɔrámsɨ, t^{h}ɨtánʌ, rɨmá, rɨmbá, kɔrándʌ, wiyákʌ, sɨpá, t^{h}ʌk^{h}námbʌ

b. Examples supporting initial stress (11)
ə́rank^{h}, ə́rowk^{w}atɨ, ə́rangiy, sɨ́ngampk^{h}, ə́rastɨ, lówanɨm, ə́rangʌrmbʌ, ə́ras, ə́raw, wúswar, lέt^{h}lanʌ

Moreover, some of the counterexamples to initial stress may be due to morphosyntactic factors. For example, two of the examples that support the sonority generalization are infinitival forms with penultimate stress (t^{h}ɔpátɨ 'to buy', t^{h}ɨtánʌ 'to be'). All but one of the 9 infinitival forms in Foreman and Marten (1973) have penultimate stress. If infinitival forms are exceptions to initial stress and receive penultimate stress, the two hypotheses would be tied with 11 counterexamples each which would have to be marked as exceptions.[26]

Since initial stress is at least as successful as the sonority generalization, I conclude that the data do not support sonority-driven stress in Mayo. Given the methodological principle in (12), a tie between the sonority-driven analysis and an analysis that respects encapsulation is a positive result for modularity.

1.5.5 Discussion of Asheninca and Yimas

The present section discusses the reported patterns of sonority-driven stress in Asheninca (Type III; Payne 1990) and Yimas (Type II; Foley 1991). Both are cited in Gordon (1999/2006) as stress patterns that are sensitive to vowel height. As mentioned in section 1.4, Asheninca played a special role in Hayes (1995) as the only stress pattern analyzed using reference to vowel quality. Both cases involve an optional process that either shifts or deletes stress. While I do not provide sufficient support for alternative analyses of the data, I will discuss possible ways out for a modular analysis that make use of the fact that in both languages the processes are optional, and I will explain what hypothetical evidence would provide evidence against the modular analysis.

Consider first Asheninca, as presented in Payne (1990). The basic vowel inventory of Asheninca is given in (140).

(140) Vowel inventory of Asheninca

	front	central	back
high	i		
mid	e		o
low		a	

According to Payne (1990), the basic stress pattern of Asheninca is Left-to-Right Iambic where CVV(C) syllables are heavy and CV(C) are light and where the final syllable is ex-

[26]The single infinitival form with final stress is rimá 'to strengthen', which is already included in the count of counterexamples to initial stress in (139a).

trametrical. The examples in (141) illustrate Payne's analysis (for presentational purposes, Payne ignores the distinction between primary and secondary stress). In (141a), all syllables are light and binary feet are constructed from left to right. The final syllable is extrametrical and does not receive stress. The penultimate syllable is assigned a degenerate foot (not marked in the example) that loses its stress due to clash with the preceding stressed syllable. Example (141b) demonstrates that heavy syllables always carry stress and can form their own foot.

(141) Basic quantity-sensitive rhythmic pattern

a. (pa.mé).(na.kó).(weŋ.tá).ke.ro — 'take care of her'

b. (no.má).(ko.ryáa).(wái).(ta.páa).ke — 'I rested a while'

Payne presents four processes that perturb that basic rhythmic pattern based on vowel quality. Three of them are sensitive to the sonority scale in (142a) and the fourth to the more fine-grained scale in (142b).

(142) Payne (1990)'s sonority scales for Asheninca stress[27]

a. a, e, o > i

b. a > e, o > i

An example of a process that relies on (142a) is prestress destressing, which removes stress from a CV syllable before a heavy syllable. Destressing applies obligatorily to Ci syllables but only optionally to Ce, Co, and Ca syllables. In (143), expected secondary stress on the second syllable is absent from a Ci syllable before a heavy CVV syllable. In contrast, (144) shows two variants of a word with a Ca syllable before a heavy CVV syllable, one with stress and one without it.

(143) Ci syllable *obligatorily* loses stress before a heavy syllable

kan.<u>ti</u>.mái.ta.cya — 'however' — (no expected secondary stress)

[27]Payne's hierarchy includes a further distinction between Ci with a strident onset (realized allophonically as Cɨ) and Ci with a non-strident onset. Feet with a second Cɨ syllable are unexpectedly trochaic. Hayes (1995:289-290) shows that most examples of this sort can be analyzed using an /i/-deletion rule that triggers stress shift, though as noted by Hayes, several problematic examples remain where deletion is of no help. These examples would have to be treated as memorized exceptions to the general stress pattern.

(144) Ca syllable *optionally* loses stress before a heavy syllable

a. a.tì.ri.<u>pà</u>.yée.ni 'people'

b. a.tì.ri.<u>pa</u>.yée.ni 'people' (no expected secondary stress)

A binary distinction as in (142a) can be captured using empty-vowel representations, assuming that the vowel [i] is the empty vowel of Asheninca (for a similar assumption about Asheninca, see Gordon 1999/2006). Prestress destressing would apply obligatorily to $CV_{\emptyset}$ but optionally to CV. We can analyze in a similar manner the two other aspects of stress that show the binary distinction in (142a), main-stress assignment and destressing in rapid-speech, so I do not discuss them here.

More problematic is another process of prestress destressing which is sensitive to the scale in (142b). Here, Ci obligatorily loses stress before Ca (145) but optionally before Ce, Co, and Ci, illustrated in (146) using Ce. In (146a), where destressing does not apply, the penultimate syllable (which forms a degenerate foot) loses its stress due to clash. This process is problematic because the low vowel [a] behaves differently from the mid vowels [e] and [o], and the empty-vowel has been reserved for the representation of [i]. Payne provides 3 examples like (145) and 4 examples like (146).

(145) Ci syllable *obligatorily* loses stress before a Ca syllable

o.<u>pi</u>.ná.ta 'it costs' (no expected secondary stress)

(146) Ci syllable *optionally* loses stress before a Ce syllable

a. i.<u>kí</u>.te.ti 'people'

b. i.<u>ki</u>.té.ti 'people'

The examples provided by Payne are consistent with an alternative formulation of the optional process which omits reference to vowel quality. A grammar that optionally omits stress from a $CV_{\emptyset}$ syllable before *any* CV syllable using the rule in (147) will never be contradicted by the data – it could be an accident that we have not yet encountered an example where stress is found on a Ci syllable that precedes a Ca syllable.

(147) Optional rule: destress a $CV_{\emptyset}$ syllable before a CV syllable

Of course, more data along the lines of (145)-(146) would make an accident less plausible, but the number of examples is quite small – there are 3 examples like (145) and 4 examples like (146).[28] One way to argue in favor of quality-driven stress and against (147) is to show that speakers of Asheninca reject forms like [o.pí.na.ta] (where destressing does not apply before a Ca syllable), which is unexpected given (147). Until such evidence is provided, I will assume the stress rule in (147) as an account of the data in (145)-(146), which makes the process compatible with the Stress-Encapsulation Universal: on this account, Payne's two sonority scales for Asheninca in (142a) are reduced to the following, quality-insensitive scale:

(148) Scale for Asheninca given the rule in (147)

$V > V_0$

A similar challenge for the Stress-Encapsulation Universal comes from Yimas (Foley 1991). In Yimas, more examples seem to support a quality-sensitive analysis, but the author chose an analysis that uses a general quality-insensitive rule. Foley (1991: 78) reports that stress in Yimas can optionally shift from the first to the second syllable in the word, and that this shift "*is found with many disyllabic or trisyllabic words with underlying vowels in the first two syllables, especially when these vowels are /a/*". Two examples out of 11 provided by Foley are given below. In all 11 cases the second vowel is [a].

(149) a. yúan ~ yuán 'good'

b. yánara ~ yanára 'bark of clove tree'

Foley's actual analysis of Yimas stress does not make reference to segmental features. Default stress assignment is optional (and quality-insensitive); when it does not apply, a second, obligatory stress rule assigns stress to the second syllable regardless of its quality.

Since the proposed quality-insensitive analyses for Asheninca and Yimas can be easily refuted, I will treat both languages as potential counterexamples to the universal, noting that the discussion in this section at least provide a loophole for analyses that respect encapsulation.

[28]There is no indication in the paper that Asheninca provides additional examples, and my attempts to locate Payne to ask for more information have been unsuccessful.

1.5.6 Reanalysis of English i-extrametricality

As noted by Chomsky and Halle (1968), English stress normally ignores the final syllable of the word if its nucleus is the vowel [i]. Thus, for example, in words like résidency and éfficacy, main stress falls on the pre-antepenultimate syllable. These facts are surprising given the rules that govern the distribution of stress elsewhere in English (the expected forms are *residény, and *effícacy), but they are immediately explained if the final vowel in those words is invisible to stress assignment (e.g., the stress contour of résidency is identical to that of résidence). Moore-Cantwell (2016) tested English speakers' preference regarding stress placement in trisyllabic nonce words that end in [i] or [ə]. Speakers showed a strong preference for antepenultimate over penultimate stress with [i] but just a slight preference with [ə], reflecting a similar asymmetry in the English lexicon. Moore-Cantwell proposed a constraint that makes word-final [i] extrametrical and thus violates encapsulation. Here is a version of the problematic constraint, stated as a rule of extrametricality that refers to vowel quality:

(150) Mark word-final [i] as extrametrical

Halle (1998) proposed a different treatment of the distributional facts. On his analysis, it is the suffix *-y* rather than the final vowel that is extrametrical. He notes that other English suffixes, such as the suffix *-ure*, show a similar behavior (e.g., main stress is surprisingly initial in words like músculature, cándidature, and líterature). On Halle's interpretation, we can restate (150) as a rule that refers to the morphological identity of the suffix rather than the quality of its vowel:

(151) Mark word-final [$-Y$] as extrametrical

If Halle is right, a plausible interpretation of Moore-Cantwell's results is that participants had a strong preference for parsing nonce words with a final [i] as morphologically complex (there is no comparable parse for words with a final [ə]), and that the grammatical statement in (151) was responsible for antepenultimate stress in those words. A way to argue against (151) and in favor of (150) is to show that speakers show a preference for earlier stress in [i]-final nonce words that cannot be parsed using [$-Y$].

1.5.7 A remaining challenge: consonantal features

The Stress-Encapsulation Universal states that stress is never conditioned by any segmental features, including consonantal features. While the empirical focus of the present chapter is on the relationship between vowel sonority and stress, effects of consonantal features on stress are potential counterexamples to the universal. The literature reports four rare types of such effects (see Davis 2011 for a summary): 1) *Variable coda weight.* $CVC_{[+son]}$ syllables are reportedly heavier than $CVC_{[-son]}$ syllables for stress in three languages: Kwak'wala, the closely related Nuuchahnulth, and Inga Quechua (see Zec 1995 and references in Gordon, 1999/2006). 2) *Vowel - glottal stop is heavy.* Gordon (1999/2006) lists three languages in which a vowel followed by a coda glottal stop ([Vʔ]) is reportedly heavier than other vowel-coda sequences (Kamchadal, Mundari, Mam). 3) *Onset voice.* Syllables with a voiceless onset have been claimed to be heavier than ones with a voiced onset in Pirahã (Everett and Everett 1984; Everett 1988), Karo (Blumenfeld 2006, citing Gabas 1999), and Arabela (Topintzi 2005, citing Payne and Rich 1988). 4) *Coda place.* In Ngalakgan, CVC is heavy unless the postvocalic consonant is a glottal stop, the first part of a geminate consonant, or the first part of a homorganic nasal-stop sequence (Baker 2008).

Those cases have already been analyzed in the literature as indirect effects of consonantal features on stress through syllable structure, as in Latin, making them consistent with the universal. I will provide references to the relevant analyses, though I leave a closer examination of the assumptions needed for those analyses and their consequences for the universal for a separate occasion. Analyses of variable coda weight in terms of syllable structure can be found in Levin (1985) and Hulst and Ritter (1999) (see also Zec 1995). Gordon (1999/2006) proposes an analysis of heavy vowel - glottal stop sequences in which stress makes reference to vowel length rather than to the quality of the coda. See Everett (1988) for an analysis of onset voice cases in terms of syllable structure and see Baker (2008) and Davis (2011) for two different interpretations of the Ngalakgan data as an indirect effect of [place] on stress.

1.6 Alternatives to modularity

My next goal is to discuss alternative explanations to the Stress-Encapsulation Universal that do not involve modularity. The reason is that information encapsulation by itself is not a sufficient argument for modularity: encapsulation can be emulated in non-modular architectures, whether serial or parallel. My claim, however, is that the Stress-Encapsulation Universal poses a special problem for non-modular accounts of encapsulation. Consider the diagram in (152), which shows the picture regarding attested phonological interactions that I have argued for. The bottom arrow indicates that stress is visible to segmental features and the dotted top arrow indicates that segmental features are not visible to stress. There are bidirectional interactions between stress and syllable structure and between syllable structure and segmental processes.

(152) *Attested phonological interactions* (a full arrow from A to B indicates that A is visible to B)

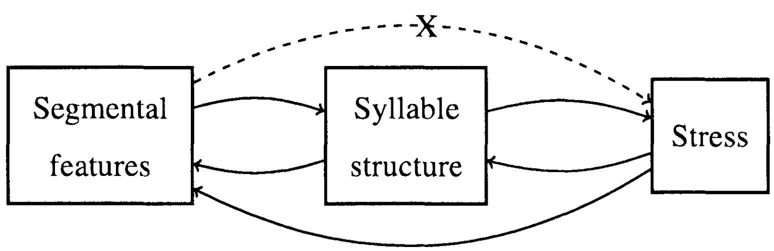

The modular architecture captures the asymmetry in (152) by removing segmental features from the input to the stress module and allowing segmental features to only affect stress through the interface. We will see that a main prediction made by non-modular accounts of encapsulation is that visibility is transitive: if A is visible to B and visible to C, then A should be visible to C. The challenge to that prediction comes from indirect effects of segmental features on stress (as in Latin, discussed above). Since segmental features are visible to syllable structure and syllable structure is visible to stress, non-modular accounts of encapsulation incorrectly predict that segmental features should be visible to stress as well and thus over-generate quality-driven stress. Blocking quality-driven stress comes at the cost of under-generating attested indirect effects of segmental features on stress.

1.6.1 An ordering account

The first non-modular account of encapsulation to consider is an ordering account within a serial phonological architecture. The main idea behind an ordering account is that stress is universally assigned *before* the insertion of segmental features: stress can never see segmental features because they are universally inserted later. The first issue with implementing such an account is that all working theories of phonology assume that stored phonological information (including segmental features) is present in URs. For example, the place of articulation of the first consonant in the English word [k^hæt] 'cat' has to be memorized and present when stress applies. But perhaps the Stress-Encapsulation Universal suggests that phonology should be reconceptualized such that segmental features, including memorized ones, are inserted late. Here is how this reconceptualization would work. We can impose a universal ordering on phonological processes as in (153). According to (153), stress processes apply before the insertion of segmental information in the derivation. A word like [k^hæt] would be derived by inserting the information as two separate tiers, shown in (154).

(153) Universal ordering

a. Insert CV tier and syllable structure

b. Apply stress processes

c. Apply non-stress processes and insert segmental tier (in any order)

(154) Representation of [k^hæt]:

a. CV tier and syllable structure: /.CVC./

b. Segmental tier: /kæt/

The problem with (153) is that the input to stress computation should be determined based on segmental features. This is particularly easy to see with the Latin example in (17), repeated below: segmental features must be available for the computation of syllable structure before the application of stress.

(155) Segmental feature → syllable structure → stress

Latin: [vo.lúp.t̠as] (non-liquid) vs. [vó.lu.kr̠is] (liquid)

This is a general problem posed by indirect effects of segmental features on stress through syllable structure. For segmental features to determine syllable structure, they must be present in the derivation whenever syllable structure is computed. And for syllable structure to affect stress assignment, it must be present before stress is computed. By transitivity, segmental features are present in the derivation before stress applies.

I conclude that an ordering account does not provide a viable alternative to modularity.

1.6.2 Universal constraint rankings within a parallel architecture

Another alternative to modularity is to fix constraint rankings within a parallel architecture such as OT. Here the strategy would be to impose a universal ranking relation between disjoint sets of OT constraints (e.g., Prince and Smolensky 1993, de Lacy 2002). As an illustration of this strategy, one can define two sets of constraints $\mathbb{C}_1$ and $\mathbb{C}_2$ as in (156) and impose the universal ranking relation in (157), which means that every constraint in $\mathbb{C}_1$ outranks every constraint in $\mathbb{C}_1$ in every language. This strategy emulates encapsulation because, intuitively, $\mathbb{C}_2$ constraints will never be strong enough to affect $\mathbb{C}_1$-computation.

(156) a. $\mathbb{C}_1 = \{m : m \text{ is a prosodic markedness constraint}\}$

b. $\mathbb{C}_2 = \{f : f \text{ is a faithfulness constraint of the form } \textsc{ident}[\text{F}]\}$

(157) $\mathbb{C}_1 \gg \mathbb{C}_2$

To see how this strategy can be implemented as an account of the Stress-Encapsulation Universal, consider again English aspiration and the following constraint:

(158) *tV́ = *unaspirated voiceless stop before a stressed vowel

The constraint can be satisfied by shifting stress away from a syllable with an unaspirated voiceless onset – an unattested quality-driven stress pattern – as shown by the tableau in (159). Candidate (b) violates a faithfulness constraint that penalizes deviations in aspiration between URs and surface forms (ASP is used here as an abbreviation of [spread glottis]); candidate (a) violates *tV́, so candidate (c) with shifted stress wins.

(159) Satisfying *tV́ by shifting stress

	/datá/	IDENT[ASP]	*tV́	FINAL STRESS
a.	datá		*!	
b.	datʰá	*!		
c. ☞	dáta			*

There are two ways to block such patterns of quality-driven stress using the universal-ranking strategy. The first is to impose a universal ranking of stress constraints over segmental faithfulness constraints (which would translate into the ranking FINAL STRESS ≫ IDENT[ASP] in the example above). The second is to impose a ranking of stress over segmental markedness (which would translate into the ranking FINAL STRESS ≫ *tV́). In both cases, if FINAL STRESS is ranked higher, the problematic candidate (c) will be blocked. I will only discuss the 'stress over segmental faithfulness' approach since the logic of the heart of the argument against the 'stress over segmental markedness' approach is similar.

As shown in (160), forcing the ranking FINAL STRESS ≫ IDENT[ASP] blocks quality-driven stress shift:

(160) FINAL STRESS ≫ IDENT[ASP]

	/datá/	*tV́	FINAL STRESS	IDENT[ASP]
a.	datá	*!		
b. ☞	datʰá			*
c.	dáta		*!	

Implementing this restriction as a universal can be done by enforcing the ranking $\mathbb{C}_1 \gg \mathbb{C}_2$ where the constraint sets $\mathbb{C}_1$ and $\mathbb{C}_2$ are defined as follows:

(161) a. $\mathbb{C}_1 = \{m : m$ is a markedness constraint that mentions stress$\}$[29]

b. $\mathbb{C}_2 = \{f : f$ is a faithfulness constraint that mentions segmental features$\}$

I discuss two problems for this account, an under-generation problem and an over-generation problem.

[29]This definition is overly simplified. To explain why languages that have contrastive aspiration do not show aspiration in response to *tV́, the definition has to be changed so as to allow IDENT[ASP] to outrank *tV́. The definition should be complicated as follows: $\mathbb{C}_1 = \{m :$ m is a markedness constraint that mentions stress but not segmental features$\}$.

1.6.2.1 Problem #1: under-generation

Indirect effects of segmental features on stress as in Latin pose an under-generation problem for the universal-ranking approach. To see why, we will need to look at such patterns in more detail. I will discuss an oversimplified version of the Latin stress pattern. As far as I can tell, the simplification does not affect the argument.

In Latin, the penultimate syllable is stressed if it is heavy; otherwise, the antepenultimate syllable is stressed. For the analysis of Latin, I will use the default-stress constraint in (162a), a cover constraint that penalizes words with non-antepenultimate stress, and the weight-to-stress constraint in (162b).

(162) Constraints for an OT analysis of Latin stress

a. DEFAULT STRESS: assign * if the antepenultimate syllable is not stressed

b. WEIGHT-TO-STRESS PRINCIPLE (WSP): assign * for every unstressed heavy syllable

Assuming the ranking WSP ≫ DEFAULT STRESS, the tableau in (163) shows that a heavy syllable attracts stress. Candidate (a) with antepenultimate stress violates WSP; candidate (b) wins even though it violates the lower ranked DEFAULT STRESS constraint.

(163)

/voluptas/	WSP	DEFAULT STRESS
a. vó.lup.tas	*!	
b. ☞ vo.lúp.tas		*

The challenge for this approach is to block candidate (c) in (164), where the underlying consonant /t/ is changed into a liquid on the surface to avoid a violation of DEFAULT STRESS. Candidate (c) violates neither constraint and is thus more optimal than the desired winner candidate (b).

(164)

/voluptas/	WSP	DEFAULT STRESS
a. vó.lup.tas	*!	
b. ➽ vo.lúp.tas		*!
c. ☞ vó.lu.pras		

Candidate (b) should be selected as a winner because in Latin, violating default stress is better than changing the liquidity of a consonant. The following tableau includes the new segmental faithfulness constraint IDENT[LIQUID], which rules out candidate (c).

(165)

	/voluptas/	WSP	IDENT[LIQUID]	DEFAULT STRESS
a.	vó.lup.tas	*!		
b. ☞	vo.lúp.tas			*
c.	vó.lu.pras		*!	

For candidate (b) to win, IDENT[LIQUID] must outrank DEFAULT STRESS. But since IDENT[LIQUID] is in $\mathbb{C}_2$ and DEFAULT STRESS is in $\mathbb{C}_1$, this ranking violates the universal ranking $\mathbb{C}_1 \gg \mathbb{C}_2$. Note that replacing IDENT[LIQUID] with a markedness constraint like *COMPLEX to penalize candidate (c) would incorrectly allow changing a liquid to an obstruent in /volukris/, favoring *[vo.lúk.tis] over the correct output [vó.lu.kris]. To block *[vo.lúk.tis], IDENT[LIQUID] would have to outrank *COMPLEX and, by transitivity, DEFAULT STRESS.

The argument does not depend on the choice of markedness constraints for the analysis of Latin stress. This is easy to see using the pair of words [vó.lu.kris] and (hypothetical) [vo.lúk.tis] which, stress aside, differ in the quality of a single consonant. No choice of stress markedness constraints could prefer [vó.lu.kris] to [vo.lúk.tis] as the output of /volukris/ while simultaneously preferring [vo.lúk.tis] to [vó.lu.kris] as the output of /voluktis/. To block the undesirable candidates that surface with an unfaithful consonant, a faithfulness constraint must outrank at least one markedness constraint. Since stress or syllable faithfulness would be of no help (stress and syllable structure are predictable), that faithfulness constraint must be a segmental faithfulness constraint. The problem, then, is quite general. As long as segmental features and stress are computed in parallel and segmental features indirectly affect stress, there will be a candidate that changes the feature instead of moving stress. To block that candidate, segmental faithfulness will have to outrank stress markedness. If this ranking is made impossible, as in the universal-ranking approach, stress patterns as in Latin cannot be generated.

1.6.2.2 Problem #2: over-generation

While a universal ranking of stress constraints over segmental faithfulness constraints blocks stress shift as in (160), it does not block all effects of segmental features on stress. Consider the following ranking:

(166) *CLASH ≫ *í, *ú ≫ *é, *ó

This ranking is not blocked by the universal ranking, but it can create the following quality-driven clash-resolution: given two adjacent stressed vowels where one is a high vowel and the other is a mid vowel, the high vowel will lose stress regardless of the order of the vowels (e.g., /íó/ → [ió], /óí/ → [ói]). The modular architecture blocks such patterns, which to my knowledge are unattested.

1.6.2.3 Summary: universal ranking

We have seen that a universal ranking of stress markedness constraints over segmental faithfulness constraints under-generates attested patterns (indirect effects of segmental features on stress) and over-generates unattested ones (quality-driven clash-resolution). The argument can be replicated for a universal ranking of stress markedness over segmental markedness constraints: such a ranking would not address the over-generation problem; regarding the under-generation problem, the argument from Latin can be restated using the candidate *[vó.lu.ptas] as a potential output of /voluptas/. Blocking such a candidate while generating [vó.lu.kris] would require ranking a segmental markedness constraint that blocks $CC_{[+liquid]}$ complex onsets over a stress constraint. I conclude that the universal-ranking approach is less successful than modularity.[30]

[30]Here is a direction for a response to the under-generation problem faced by the 'stress over segmental markedness' approach. In addition to the set of constraints that mention stress but not segmental features, one could split segmental markedness constraints into two subsets – segmental markedness constraints that mention stress and segmental markedness constraints that do not – and force a universal ranking of constraints that mention stress but not segmental features over the former subset. This response would still not address the over-generation problem (which would require further commitments in order to block quality-driven clash-resolution), so I do not develop it here further.

1.7 Conclusion

I started this chapter with de Lacy's and Blumenfeld's observation that the interaction between stress and segmental features is asymmetrically restricted: while the distribution of segmental features is often conditioned by the position of stress, the distribution of stress is never conditioned by any segmental feature but sonority. I reviewed the literature on sonority-driven stress and showed that reference to vowel sonority can be avoided if stress is allowed to see syllable structure and the binary distinction between empty vowels and non-empty vowels: the distribution of stress seems to never be conditioned by segmental features. I referred to this generalization as the Stress-Encapsulation Universal and argued that it supports a modular architecture of grammar, repeated in (167), where stress is severed from the rest of phonology. This is a welcome result: modularity provides a simple account of information encapsulation and makes various typological predictions regarding stress patterns and their interaction with other aspects of phonology. Testing these predictions is a task I leave for future work.

(167) **Hypothesis about the architecture of grammar**

1. Insert underlying phonological representation
2. Construct interface representation
3. Send interface representation to the stress module (without segmental features)
4. Receive interface representation from the stress module
5. Send phonological representation (interface + segmental representation) to the segmental module
6. In the segmental module, stress representations cannot be changed

Chapter 2

A conditional learnability argument for constraints on underlying representations (joint with Roni Katzir)

2.1 Where are phonological generalizations captured?

As noted by Halle (1962) and Chomsky and Halle (1965), speakers judge some nonce forms as nonexistent but possible – that is, as *accidental gaps* – and other nonce forms as nonexistent and impossible – that is, as *systematic gaps*. In Dutch, for example, the distribution of the voiceless alveolar strident [s] and its palatalized variant [ʃ] is restricted such that the palatalized variant occurs precisely before the palatal glide [j] (Booij 1995). Thus, forms such as [ɔstər] and [oʃjər] are accidental gaps, while *[ɔʃtər] and *[osjər] are systematic gaps.[1] Capturing this distinction in speakers' judgments is a central task of phonological theory, and it involves answering two questions. First, how is the distinction between the two kinds of gaps represented? And second, since the judgments of speakers regarding nonce forms differ between languages, how is the relevant knowledge acquired? In what follows, we point out a dependence between the two questions: on certain assumptions

[1]Our focus in this chapter is on the distribution of the stridents [s] and [ʃ]. We will set aside other systematic properties of Dutch surface forms, such as the distribution of tense vowels (like [o]) and lax vowels (like [ɔ]). As far as we can tell, this does not affect our argument.

about learning, the phonological component must follow one of several specific representation schemes discussed below in order to ensure that the acquisition process leads to the judgments that actual speakers make.

To set the stage for our argument, let us briefly review the two main views in the literature on the representations behind phonological well-formedness judgments. Early generative approaches relied on a combination of two factors: constraints on underlying representations (CURs) in the lexicon;[2] and phonological rules. In the example above, an early generative account might use a CUR such as (168) and a phonological rule such as (169) as the basis for capturing the distribution of stridents in Dutch:[3]

(168) CUR IN DUTCH: No ʃ in the lexicon

(169) [+*strident*] → ʃ /__ j

(168) ensures that stridents will be alveolar underlyingly, while (169) ensures that they will become palatalized in exactly the right environment. The combination of (168) and (169) handles the distinction between the accidentally missing [ɔstər] and [oʃjər] on the one hand and the systematically missing *[ɔʃtər] and *[osjər] on the other, on the assumption that accidental gaps are those forms that can be derived by a new UR and without changing the rest of the grammar and that systematic gaps are those forms that would require a change to the rest of the grammar. The accidentally missing [ɔstər] and [oʃjər] could be added to Dutch with the URs /ɔstər/ and /oʃjər/; the palatalizing rule in (169) would then turn the former into its surface form. For *[ɔʃtər] and *[osjər] the situation is different. Since (168) prohibits the storing of /ʃ/ in the lexicon of Dutch, [ʃ] must follow from rule application; but the palatalization rule in (169) does not apply before /t/, which leaves no way to derive

[2]Halle (1959, 1962) proposed to capture the relevant generalizations through rules that apply to URs. Stanley (1967) argued that these should be constraints rather than rules. In the generative tradition these became known as morpheme-structure conditions or morpheme-structure constraints. We use CURs as a cover term for rules or constraints of this kind.

[3]To simplify the presentation, here and below we use *strident* to refer to the voiceless coronal stridents [s] and [ʃ] only, excluding other Dutch stridents such as [z] and [x]. As far as we can tell, this simplification is orthogonal to the argument we will present.

We will assume that [s] and [ʃ] are distinguished from each other by the feature *anterior*: [s] is specified as [+*anterior*] and [ʃ] as [−*anterior*]. A different yet equivalent formulation of (169) using *anterior* would be [+*strident*] → [−*anterior*] /__ j. To simplify the presentation, we will use symbols (like s and [ʃ]) instead of features whenever possible.

*[ɔʃtər]. For *[osjər], on the other hand, obligatory palatalization ensures that this surface form cannot appear. Both gaps are thus correctly treated as systematic.

CURs, then, offer one way in which patterns such as the distribution of stridents can be captured. A different way to capture the same pattern forgoes CURs and relies on phonological rules alone. For example, instead of stating that stridents are alveolar by default using a CUR, we could accomplish the same by a rule such as (170) below, which makes stridents alveolar regardless of their underlying specification or of their environment:

(170) [+*strident*] → s

If (170) is ordered before (169), any UR would first have its stridents made alveolar ([s]), after which its pre-palatal stridents will be made palatalized ([ʃ]). This would make the URs /osjər/ and /oʃjər/ surface as [oʃjər], while the URs /ɔʃtər/ and /ɔstər/ will surface as [ostər]. The systematically missing *[ɔʃtər] and *[osjər] will correctly be predicted to be impossible to derive.

We thus have two different ways to represent the distinction between accidental and systematic gaps. The first involves a combination of CURs and phonological processes, and the second relies on phonological processes alone. The former approach was the one favored in early generative phonology: while the architecture assumed at the time allowed for both kinds of analysis, CURs were taken to be preferred by the simplicity metric (for a simplicity-based argument for CURs, see Halle 1962, pp. 59–60). The latter approach has been adopted within Optimality Theory (OT; Prince and Smolensky, 1993), where a representational principle, *Richness of the Base*, prevents CURs from being stated:[4]

(171) Richness of the Base (ROTB; Prince and Smolensky 1993, p. 191,Smolensky 1996, p. 3):

a. All systematic language variation is in the ranking of the constraints.

b. In particular, there are no language-specific CURs.

[4]The discussion above uses phonological rules, but both approaches can just as easily be stated using OT constraints (which will be the main representation used in section 2.3) or even more neutrally, as a reviewer notes, using mapping statements as in Tesar 2014. Stated in terms of constraints, the first approach would combine the CUR in (168) with a constraint ranking such as *sj ≫ IDENT[ANT], while the second approach would avoid (168) and instead add a mid-ranking constraint banning [ʃ], as in *sj ≫ *ʃ ≫ IDENT[ANT]. The question of whether to use CURs is thus separate from the choice between rules and constraints, and we will focus exclusively on the former question in what follows.

Clearly, the two representational choices for how to handle the distributional pattern of stridents are meaningfully different. For example, the use of CURs distributes the knowledge of such patterns between two distinct components of the grammar – CURs versus phonological rules or constraints – while ROTB leads to a unitary treatment of such patterns. This difference can lead to different ways in which various phenomena can be accounted for – for example, in loanword adaptation – but to date it has been hard to find empirical arguments for one view or the other. Below, we will show how considerations of learnability can be brought to bear on the choice.

Turning to the question of how the relevant knowledge is acquired, we will rely on a general approach to learning, following the principle of Minimum Description Length, that has much in common with the evaluation metric of early generative phonology but is quite different from much of the literature on learning within OT. The following section briefly introduces the learning approach that we will be using and motivates our choice of using it instead of familiar suggestions in the literature on learning in OT.

2.2 Learning

Our discussion below relies on a general approach to learning according to which the child attempts to make inductive inferences by balancing the simplicity of the grammar (or its prior probability) against its fit to the data (or the likelihood of the data). A preference for simplicity favors general grammars that do not overfit the data. However, simplicity on its own, as in the evaluation metric of early generative grammar, leads to grammars that are overly general. By balancing simplicity against tightness of fit, the learner can hope to find an intermediate level of generalization that is appropriate given the data.

Our argument, which we present in section 2.3, can be followed at the informal level of balancing the simplicity of grammar against tightness of fit. As far as we can tell, nothing in our discussion will depend on the specifics of how this balancing is formalized. For concreteness, however, we will frame our discussion of balanced generalization in terms of one particular formalization: namely, the principle of Minimum Description Length (MDL; Solomonoff 1964, Rissanen 1978), which we now briefly sketch, along with references for details and further discussion. For this sketch, it will be convenient to think of both grammars and their encoding of the data as sitting in computer memory according to a given encoding scheme. Using $|\cdot|$ to notate length, we can write $|G|$ for the length of the grammar G as measured in bits. The encoding of the data D using G will be written $D : G$, and the tightness of fit will be the length of this encoding, $|D : G|$. Using this notation, MDL can be stated as follows:[5]

(172) MDL EVALUATION METRIC: If G and G' can both generate the data D, and if $|G| + |D : G| < |G'| + |D : G'|$, prefer G to G'

The balancing of economy and tightness of fit has made MDL – and the closely related Bayesian approach to learning – helpful across a range of grammar induction tasks, in works such as Horning 1969, Berwick 1982, Ellison 1994, Rissanen and Ristad 1994, Stolcke 1994, Grünwald 1996, de Marcken 1996, Brent 1999, Clark 2001, and Goldsmith 2001, among others. Recently, this approach to learning has been used to provide learners for both constraint-based phonology (Rasin and Katzir 2016) and rule-based phonology (Rasin et al. 2017, To appear) that acquire a lexicon, the phonological processes involved, and their interactions, all from distributional evidence alone.

The MDL view predicts that the child will invest in grammatical statements only when the cost of the investment (in terms of increase in $|G|$) will be offset by the increase in tightness of fit to the data (in terms of decrease in $|D : G|$). Applied to the case of the distribution of stridents in Dutch, the fact that *[ɔʃtər] and *[osjər] are judged as ill-formed teaches us that the investment in ruling out these forms, through the relevant statements in G, has been justified by the shortening of $D : G$. The acceptability of [ɔstər] and [oʃjər], on the other hand, teaches us that the benefits for $|D : G|$ in ruling out these forms are too small to justify an investment in the requisite grammatical statements.[6]

Before proceeding, we note that the view of the child as inductive learner is not the only view on phonological learning in the literature. There is a prominent alternative,

[5]Here and below the grammar G will be taken to be not just the phonological rules and their ordering (or the constraints and their ranking) but also the lexicon. Thus, by saying that a grammar G generates the data D, we mean that every string in D can be derived as a licit surface form from some UR in the lexicon and the relevant rules (or constraints).

[6]The above assumes that the grammatical statements under consideration need to be acquired (rather than being given to the learner in advance) and that they are allowed by UG in the first place.

which we will refer to as *restrictive consistency seeking* (RCS), according to which the child attempts to find the most restrictive grammar consistent with the data. On this view, common within OT, the child starts with a maximally restrictive hypothesis about the world (typically assuming a finite number of innate markedness constraints penalizing various surface patterns); this hypothesis is gradually relaxed, with individual prohibitions being eliminated or demoted, in the face of conflicting evidence. Representative proposals of RCS include an initial ranking of markedness over faithfulness (M≫F; Smolensky 1996, Tesar and Smolensky 1998) from which a search for a consistent ranking begins, as well as a more sustained bias for M≫F (Hayes 2004, Prince and Tesar 2004) throughout the search for a consistent ranking. On this view, the child hypothesizes in advance that they should ban *[osjər] and *[ɔʃtər], and since Dutch provides no counter evidence, they never need to change their mind. What the acceptability of [oʃjər] and [ɔstər] teaches us, on this view, is that UG simply neglects to provide the means to ban these forms without banning attested forms.[7] Had it done so, the absence of [oʃjər] and [ɔstər] from the child's input data would have allowed the child to maintain the more restrictive hypothesis that these forms are impossible.

In the literature, the representational principle of ROTB has often been bundled together with learning models that follow RCS. This bundling does not follow logically – ROTB does not imply RCS, nor is it implied by it – but given the connection made in the actual literature, we wish to explain why we set aside RCS and instead propose to use the inductive approach to learning, as in MDL, for our probing of the choice between CURs and ROTB.

Our first reason for setting aside RCS and focusing exclusively on MDL is that, to date, only the latter has actually provided working distributional learners for phonology (that is, implemented algorithms that take raw surface data and induce a phonology and a lexicon). The MDL learner of Rasin and Katzir 2016, for example, takes unanalyzed surface forms and induces a lexicon, often with abstract URs that differ from the surface forms, along

[7] The qualification regarding not banning attested forms is needed. UG may well provide the means to rule out [oʃjər] and [ɔstər], but if these means are blunt instruments such as *ʃ or *ɔ they will also rule out attested forms and will therefore be demoted below the relevant faithfulness constraints and will not play a role for [oʃjər] and [ɔstər]. Such accidental gaps, then, teach us about the absence of more specific constraints such as *[ʃV] (within an RCS model) but are uninformative regarding more sweeping constraints that would be demoted independently.

with a set of markedness and faithfulness constraints and their ranking.[8] The MDL learner of Rasin et al. (2017) accomplishes a similar task but within rule-based phonology: it takes unanalyzed surface forms and induces a lexicon and a set of ordered context-sensitive rewrite rules. Despite active research into RCS in phonology over the past few decades, no currently available learner has been presented that can accomplish similar tasks, making the RCS idea promissory.

While it is conceivable that future work will provide an implemented RCS learner, the path toward such a result is quite unclear at present. For one thing, the relevant notion of restrictiveness has been hard to formulate, with concrete choices such as M≫F giving rise to problems that have been recognized in the literature (see, e.g., Hayes 2004, Prince and Tesar 2004, Tauberer 2009). The induction of the lexicon has also posed a challenge for RCS models. Such models have typically relied on notions such as Lexicon Optimization (Prince and Smolensky 1993, p. 209, Inkelas 1995, Smolensky 1996), which encourages the learner to posit URs that are identical to the corresponding surface forms unless forced to do otherwise by paradigmatic information from alternations. (Variants of this idea have been explored as well; see McCarthy 2005 and Krämer 2012, among others. In McCarthy 2005's variant, the Free Ride Principle, nonidentical URs are posited when an alternation suggests – rather than strictly force – a deviation from identity.) As argued in detail by Alderete and Tesar (2002), McCarthy (2005), Idsardi (2006), Nevins and Vaux (2007), and Krämer (2012), Lexicon Optimization (and its variants) cannot handle the abstract URs that speakers actually posit, often without any supporting alternations to force such URs. For example, Nevins and Vaux consider the case of rhotics in Spanish, which can be realized as [r] or [ɾ] word-medially but only as [r] word-initially. When induced to move a word-initial [r] to a word-medial position as part of a language game, speakers sometimes realized it as [r], in line with a faithful UR, but sometimes as [ɾ], suggesting an unfaithful UR. Crucially, Lexicon Optimization cannot attribute this deviation to any available alternations. While the induction of abstract URs posits a (currently unresolved) challenge for RCS models, MDL learners succeed in acquiring such abstract URs using nothing more than the general

[8]The learner also works with a set of constraints that are given in advance, as is assumed in much work within OT.

principle in (172), as shown in Rasin and Katzir (to appear).

Our second reason to set aside RCS models and focus on MDL is that the latter has been supported empirically by a range of lab experiments on generalization, while similar support is lacking for the former. In a variety of learning tasks in the lab, ranging from word learning (Xu and Tenenbaum 2007) through causal reasoning (Sobel et al. 2004) to sensorimotor control (Körding and Wolpert 2006) and visual scene chunking (Orbán et al. 2008), among many other tasks, human subjects have been argued to balance between the specificity of a hypothesis, corresponding to $|D : G|$, and its independent plausibility, corresponding to $|G|$. If humans indeed use this way of learning across different domains, it seems sensible to consider their use of the same in phonology. We are not aware of similar considerations for RCS models.

If the technical challenges for RCS models are eventually addressed and an implemented distributional RCS learner provided, or if new lab experiments change the currently available evidence for inductive models like MDL over RCS, a reexamination of the choice between CURs and ROTB will of course be warranted. Until then, it strikes us as reasonable to examine the implications of this choice within an MDL model, to which we now turn.

2.3 The MDL-learnability implications of ROTB

The general shape of our MDL-based reasoning about ROTB will be as follows.[9] As we will note, ROTB will usually result in some part of the distribution of stridents being stored faithfully. For that part of the pattern, there will be no MDL motivation to invest in grammatical statements (whether rules and their ordering or constraints and their ranking), and so such statements will not be systematically acquired. Depending on the properties of the initial state, this can result in adults who do not have the knowledge of the relevant part of

[9]For concreteness, we present most of our discussion within the framework of constraint-based phonology and refer to similar considerations within rule-based phonology only occasionally. (As mentioned above, the question of whether ROTB holds is separate from the question of rules versus constraints, though in section 2.3.3 we note one place in which the two choices interact.) We will also stay with the example of stridents in Dutch (though the argument for CURs can be made using any of a wide variety of patterns from different languages).

the pattern as part of the input-output mapping, and – again, due to ROTB – such adults will be able to accept nonce forms with the incorrect kind of strident in that part of the distribution, which does not match the judgments that actual speakers make on such forms. We will point out two possible responses to this predicament. The first response is to abandon ROTB and admit CURs, which, as we will show, lead to the correct pattern of speaker judgments and also have a clear MDL motivation (by supporting a shorter encoding of the lexicon) and will therefore be acquired by the learner. Since the learning problem that we note is caused directly by ROTB, and since ROTB has not been particularly well supported by other evidence in the literature, the re-introduction of CURs strikes us as the most natural response. The second response is to maintain ROTB but adopt special measures to ensure the knowledge of the pattern. For example, the problem we outline obviously does not arise if the full knowledge of the pattern is given to the child in advance (by building the relevant constraints and their ranking into the initial state, as is often assumed within OT, or by doing the same with the rules and their ordering). For rule-based phonology (but not for OT), a more imaginative possibility within the second response is to allow for underspecification in the storage of URs. This choice allows a rule-based learner to store non-faithful stridents throughout, and, on certain assumptions discussed below, it ensures that the full pattern of strident distribution is acquired.

The structure of the argument and the range of responses are intricate, and in what follows we discuss both in some detail. The basic observation, however, is straightforward: ROTB leads to a learnability challenge given the data available to the child and the judgments that speakers have, and one of a small range of representational responses is called for. In the literature to date, ROTB has mostly been left as a matter of theoretical taste, but our observation shows that this need not remain the case: the range of possible responses to the learning challenge amounts to an empirical prediction of ROTB that can be tested, though we will not be able to do so within the present chapter. Beyond the issue of ROTB, our argument illustrates a methodological point that was central in earlier generative phonology but has not received much attention in recent years: that a general evaluation metric for learning can yield architectural predictions about linguistic representations and help choose between competing theories of UG. We return to both the specific implications of our argument for ROTB and the general methodological point in section 2.4.

In order to develop our argument, we will need to examine the MDL implications of the possible choice points under various reasonable representational assumptions. These assumptions include both cases in which the constraints are given to the child in advance and cases in which they are acquired. While the former possibility has been widely assumed within the literature on OT, it will be useful to consider the latter possibility in some detail for several reasons. First, we would like to get a picture of the connection between learnability and the choice between CURs and ROTB, not just in specific configurations that have received attention in the literature but generally. As we will show, this broader examination will allow us to identify an empirical connection between assumptions that have often been bundled together in the literature without argument. Second, language-specific constraints that need to be acquired have occasionally been suggested even within the OT literature (see, e.g., Kager and Pater 2012, Pater 2014 and references therein, as well as the earlier literature on arbitrary phonological rules, e.g., Bach and Harms 1972 and Anderson 1981; of course, language-specific rules that need to be acquired were broadly assumed within earlier generative phonology). Finally, the case of constraints that need to be acquired is somewhat simpler to analyze than the case of constraints that are given in advance. It will thus be convenient presentationally to start from the simpler case, which we discuss in section 2.3.1, and introduce the complications of constraints that are given in advance only later, which is what we do in section 2.3.2. The configurations we discuss in sections 2.3.1 and 2.3.2 are summarized in Figure 2-1. In section 2.3.3 we discuss the special case of a rule-based system with underspecification (and certain additional assumptions), which, as mentioned above, allows a ROTB learner to acquire the full pattern of strident distribution correctly.

2.3.1 Constraints acquired

For the first few configurations that we will consider, suppose that the child, using MDL, needs to acquire the constraints, with each additional constraint costing a positive number of bits, and suppose further that /s/ and /ʃ/ each costs some fixed number of bits to store

	Acquired constraints	Innate constraints
No $M \gg F$	CURs	CURs
$M \gg F$	CURs	

Figure 2-1: Summary of the configurations we discuss in sections 2.3.1 and 2.3.2 as part of the conditional argument for CURs. We consider two conditions: whether constraints are acquired or innate and whether markedness constraints preferably outrank faithfulness constraints ($M \gg F$). Cells labeled as 'CURs' correspond to configurations for which we show that ROTB fails and CURs are required for learning. The empty cell corresponds to the only configuration in which ROTB succeeds, which combines innate constraints with $M \gg F$.

in the lexicon. The exact form of the argument depends on which of the two segments is costlier, if either. We will consider the three different possibilities in turn, followed by an examination of costs that vary between contexts.

2.3.1.1 Globally fixed costs for /s/ and /ʃ/, *Cost*(/ʃ/) > *Cost*(/s/)

Suppose that the cost of storing an instance of /ʃ/ in the lexicon is greater than the cost of storing an instance of /s/, *Cost*(/ʃ/) > *Cost*(/s/). Consider first the situation of the child on the assumption that ROTB holds. Since instances of /ʃ/ are costly to store in the lexicon, it will be preferable in terms of MDL to invest in a markedness constraint that triggers palatalization of /s/ before /j/ (e.g., *sj) and then store all stridents as /s/. Adding the markedness constraint will cost a few bits, but this cost will be outweighed by the savings from not having to store any instances of the costlier /ʃ/ in the lexicon. A faithfulness constraint ensuring that /s/ surfaces faithfully – possibly a general Faith and possibly something more specific such as Ident[ant] – will also be acquired, but it will be outranked by *sj so that the latter will have its desired effect.[10] By following this reasoning,

[10]See Rasin and Katzir 2016 for why faithfulness constraints will be acquired by an MDL learner, quite independently of the palatalization pattern under discussion here. Faithfulness constraints ensuring that /s/ surfaces faithfully will be acquired independently of the palatalization pattern on the assumption that UG allows writing faithfulness constraints that are at least as general as Ident[F] – that is, constraints such as Ident[ant] or Ident. We note that if there is a reason to restrict faithfulness constraints to be very specific (e.g., if Ident-constraints can only take the form Ident[±F]), our claim that a faithfulness constraint protecting [+ant] will be acquired independently of the palatalization pattern will no longer be true. If that turns out to be the case, at least some learners may acquire a markedness constraint such as *ʃ instead of the relevant faithfulness constraint, and if the relative costs of the two kinds of constraints ensure that all learners acquire *ʃ, the challenge for ROTB will be avoided. In the absence of reasons to assume that the only faithfulness

the child has successfully learned that pre-palatal stridents are systematically palatalized in the language. For example, the child will now correctly rule out forms such as *[osjər], with an alveolar [s] before [j].

Unfortunately for ROTB, this is also the extent of the child's acquisition of the pattern: the child will not learn to block forms such as *[ɔʃtər], with [ʃ] in 'elsewhere' environments. The reason is that, for an ROTB child, such forms must be blocked through the input-output mapping, for example through *ʃ or a similar markedness constraint that penalizes [ʃ]. And there is simply no MDL justification for acquiring this kind of constraint. Recall that all stridents are already stored in the lexicon as alveolar ([s]) and that by default they are mapped faithfully to the surface due to the low-ranking faithfulness constraints. Consequently, a markedness constraint such as *ʃ will be of no use in deriving the observed forms in the input data, and the cost of adding such a constraint to the grammar will not be justified. An ROTB child, then, will become an adult who knows only half of the distributional pattern of stridents in Dutch: such an adult will correctly rule out *[osjər] but fail to rule out *[ɔʃtər] (which, due to ROTB, can be stored as is and then mapped faithfully to the surface given the acquired constraints).[11]

If ROTB does not hold, the learning process can succeed in full. The first step is similar to the one with ROTB: the child will invest in a constraint like *sj and then store all stridents as alveolar, which will allow the child to correctly rule out forms like *[osjər]. But with the possibility of stating CURs, a crucial second step becomes available. The first step involved the extensional removal of all instances of /ʃ/ from the lexicon. The child can now conclude that this was no accident, and that /ʃ/ should be eliminated *intensionally*

constraints can be acquired are highly specific, we set aside this possibility and do not consider it further in what follows.

[11]The reasoning above makes no reference to alternations, but it is possible in principle that paradigms will allow the learner to acquire phonological knowledge that is otherwise hard or impossible to obtain. In the present case, one could imagine that alternations would provide MDL justification for learning that stridents are alveolar in 'elsewhere' environments, which, in turn, would allow the learner to acquire the full distributional pattern without abandoning ROTB. As far as we can tell, however, the alternations that are actually available in Dutch do not provide an MDL learner with the relevant information and thus do not help address the challenge to ROTB, either with the present representations or with those discussed below. In section 2.3.1.3 we will mention one case where alternations can help learn part but not all of the Dutch pattern. There, too, the challenge to ROTB remains even if the learner can make use of alternations. Other than that case, since alternations do not help address the challenge to ROTB we will continue to ignore their role in what follows.

from the very alphabet in which the lexicon is written. That is, the child can reach the following conclusion (restating (168) above):

(173) CUR IN DUTCH: No ʃ in the alphabet of the lexicon

Let us first see why (173) is justified in terms of MDL. All things being equal, removing a possible segment from the underlying inventory makes it slightly easier to specify the remaining segments, some of which may now cost fewer bits than before. Consequently, the lexicon will now be encoded with fewer bits, thus providing MDL justification for adopting (173). We can now see why, in a world that allows CURs, the child can go beyond what was possible under ROTB and acquire the second part of the pattern of distribution of stridents. The reason is that with a CUR like (173), the child will now correctly rule out surface forms like *[ɔʃtər]: /ɔʃtər/ is now no longer a possible UR; and given the grammar that has been induced, this UR is the only potential source for this putative surface form. In other words, the impossibility of even stating /ʃ/ in the lexicon, with its MDL justification, means that the learner has correctly learned to block illicit palatalization.

We conclude that, under the representational choice of constraints that need to be acquired and of *Cost*(ʃ) > *Cost*(s), the ability to state CURs allows for the full distributional pattern of stridents to be acquired, while the adoption of ROTB leads to a failure in learning half of the pattern.

2.3.1.2 Globally fixed costs for /s/ and /ʃ/, *Cost*(/ʃ/) < *Cost*(/s/)

Suppose now that *Cost*(/ʃ/) < *Cost*(/s/). In this case, the learner will store /ʃ/ throughout the lexicon and acquire constraints that will ban [ʃ] in 'elsewhere' (that is, non-pre-palatal) environments. The precise constraints and their implications, however, can vary, depending on the complexity of describing the 'elsewhere' environments.

Suppose first that these environments are easy to characterize. In that case, a learner following ROTB will acquire a single markedness constraint banning [ʃ] in those environments – say, *ʃT (no [ʃ] before a stop) – and be done.[12] In particular, *sj will not be learned, and the target grammar will fail to ban impossible nonce words such as *[osjər] (which,

[12]As with the earlier cost scheme (as well as with the configurations below), a relevant faithfulness constraint will also be acquired and ranked below the markedness constraint.

again, can be stored as is due to ROTB and then mapped faithfully to the surface). This is the mirror image of the problem for ROTB in the previous setting, and again the ability to violate ROTB and remove elements from the alphabet of the lexicon will allow the learner to acquire the full pattern.

On the other hand, if the 'elsewhere' environment is difficult to describe, it will be costly to state a constraint that bans [ʃ] directly in these contexts, and it will make more sense for the child to learn to ban [ʃ] in general and allow it only before [j]. That is, it will acquire a low-ranking *ʃ to prevent underlying /ʃ/ from surfacing faithfully and a high-ranking *sj to ensure that stridents surface correctly before [j]. On this scenario, the learner will have correctly acquired the full pattern without requiring a CUR, thus allowing ROTB to be maintained.

While this scenario provides a way to learn the distribution of stridents without CURs, it is ultimately unsuccessful because the cost assignment makes the mirror image of the Dutch pattern unlearnable – a pattern with stridents that are alveolar in some specific environments but are palatalized elsewhere. Consider Bengali, where the default sibilant is [ʃ], and [s] occurs only in word-initial consonant clusters and word-medially before dental stops (Evers et al. 1998). For example, the nonce forms [t̪uʃka], with an [ʃ] before a velar consonant, and [t̪ust̪a], with an [s] before a dental consonant, were both accepted by two native speakers of Bengali we consulted and are thus accidental gaps, while the nonce form [t̪uska], with an [s] before a velar consonant, was rejected and is therefore a systematic gap.[13] An appropriate constraint ranking for the Bengali pattern (ignoring both optionality and word-initial clusters) would be the following:

(174) Constraint ranking for Bengali (without optionality): *ʃt̪≫*s≫IDENT[ANT]

And paralleling the discussion of the Dutch pattern with the earlier cost assignment of *Cost*(/ʃ/) > *Cost*(/s/), the present cost assignment of *Cost*(/ʃ/) < *Cost*(/s/) will prevent the full Bengali pattern from being acquired. Given the present cost assignment, an ROTB

[13]The speakers' responses were variable with respect to the nonce form [t̪uʃt̪a], with an [ʃ] before a dental consonant: one speaker rejected it as ill-formed; the other speaker accepted it as well-formed, suggesting that the process that turns [ʃ] into [s] applies optionally in her grammar. That variability with respect to [t̪uʃt̪a] is consistent with our argument, which is independent of whether the process applies obligatorily or optionally.

learner will store all stridents as /ʃ/ and then acquire a markedness constraint forcing stridents to surface as [s] in the relevant environments. The same reasoning used earlier will prevent the learner from acquiring a constraint that enforces [ʃ] elsewhere (in this case, since all stridents are already stored as /ʃ/ in the lexicon), which will result in an inability to rule out nonce forms with [s] in 'elsewhere' environments (e.g., before velar consonants, as in *[t̪uska]), contrary to fact. On our current assumptions that constraints are acquired and that /ʃ/ is less costly than /s/, succeeding in learning Bengali requires abandoning ROTB and using CURs.

2.3.1.3 Globally fixed costs for /s/ and /ʃ/, *Cost*(/ʃ/) = *Cost*(/s/)

Consider now the possibility that *Cost*(/ʃ/) = *Cost*(/s/). In this case, compression cannot learn either part of the Dutch pattern in the absence of CURs (and similarly for Bengali): with fixed, equal costs for /s/ and /ʃ/, compression will favor the storing of URs that are identical to their corresponding surface forms in terms of palatalization, along with the acquisition of the relevant faithfulness constraints that will guarantee that the stored values surface faithfully. Any markedness constraints governing palatalization will be superfluous and will therefore not be acquired. An ROTB learner will consequently fail to reject both *[ɔʃtər] and *[osjər].[14]

2.3.1.4 Contextualized costs for /s/ and /ʃ/

The learnability argument against ROTB extends to some other representational possibilities that UG might make available. For example, suppose that UG makes the cost of /ʃ/ lower than that of /s/ before /j/ and higher than it in other environments. In the absence of the ability to state CURs, this cost assignment will make both kinds of markedness con-

[14] In the case of equal costs for /s/ and /ʃ/, alternations may help learn half of the Dutch pattern. The pressure for economy will push the learner to store the stem in surface pairs like *ijs* [ɛis] 'ice cream' and *ijsje* [ɛiʃjə] 'small ice cream' as a single UR – either /ɛis/ or /ɛiʃ/ – and derive the [s]~[ʃ] alternation from the input-output mapping by adding appropriate constraints to the grammar. However, on either choice of UR only half of the pattern will be learned. If the UR of the stem is /ɛis/ (and [ɛiʃjə] is derived through palatalization, using a constraint like *sj), the constraint *ʃ will serve no compressional purpose and thus will not be learned; in this case, an ROTB learner will fail to reject *[ɔʃtər]. If, on the other hand, the UR of the stem is /ɛiʃ/ (and [ɛis] is derived through de-palatalization, using a constraint like *ʃ# which penalizes word-final [ʃ]), the constraint *sj will serve no compressional purpose and thus will not be learned; in this case, an ROTB learner will fail to reject *[osjər].

straints necessary for capturing the Dutch pattern unlearnable by an MDL learner for the same reasons as discussed above: as in the case of identical costs, the learner will store URs that are identical to the corresponding surface forms in terms of palatalization (with /ʃ/ before /j/ and /s/ elsewhere) and, given the faithfulness constraints, will not invest in any markedness constraints for palatalization.

For the opposite weighting scheme, with the cost of /ʃ/ higher than that of /s/ before /j/ and lower than it in other environments, things are different. This scheme will allow both kinds of markedness constraints relevant for the Dutch pattern to be learned by an MDL learner, regardless of CURs. As with *Cost*(/ʃ/) < *Cost*(/s/), however, this scheme makes patterns like the one in Bengali unlearnable by an MDL learner that follows ROTB. Since neither /s/ nor /ʃ/ precedes /j/ in Bengali,[15] the contextualized cost assignment will have the same effect as *Cost*(/ʃ/) < *Cost*(/s/): an ROTB learner will store all stridents as /ʃ/ in the lexicon and, given the relevant faithfulness constraints, will fail to invest in a markedness constraint such as *s; this, in turn, will lead to an inability to rule out forms with [s] in elsewhere environments (as in *[t̪uska]), contrary to fact.

This concludes our discussion of the case of constraints that need to be acquired. We have seen that across various representational choices, the ability to state CURs in the lexicon is necessary for successful learning, assuming that the constraints are not given in advance.

2.3.2 Constraints given in advance

If the constraints are not acquired but rather given to the learner in advance, as is commonly assumed in the OT literature, a slightly more complex situation arises. We now turn to this case, building on the argument in Rasin and Katzir 2015 that, unless a preference for markedness over faithfulness is incorporated, an MDL learner would still need to abandon ROTB and adopt CURs. Suppose that the learner is given the two relevant markedness constraints for the Dutch pattern: *sj, which penalizes alveolar pre-palatal stops; and *ʃ,

[15]Ferguson and Chowdhury (1960) suggest that the glide [j], if it exists in Bengali, is only available as the second member of a diphthong (such as /ai/). The two native speakers we have consulted confirmed that [j] never follows stridents in their dialects. One speaker reported that [j] does not exist in her dialect at all. The second speaker reported that [j] only occurs after vowels.

which penalizes [ʃ] in general. Suppose further, as in 2.3.1.1, that *Cost*(/ʃ/) > *Cost*(/s/).

As in the setting with acquired constraints, the constraint *sj poses no special problem for an MDL learner following ROTB. Ranking this markedness constraint above the relevant faithfulness constraints will serve the compressional purpose of enabling the elimination of /ʃ/ from all URs. As for *ʃ, the learner is now assumed to be given this constraint in advance; differently from the case of a learner that needs to acquire the constraints, the presence of *ʃ will no longer incur costs in the present setting. However, the constraint still offers no compressional advantage. Consequently, the learner will not benefit from ranking this constraint above any faithfulness constraints, such as IDENT[ANT], that penalizes modifications of the feature *anterior*. We would thus expect speakers to vary in the relative ranking of *ʃ and IDENT[ANT]. But this means, on ROTB, that speakers of Dutch should differ in whether they accept forms such as *[ɔʃtər] as possible, contrary to fact.[16] In other words, for an MDL learner following ROTB that is given the constraints in advance, the problem lies not with the possibility of attaining the appropriate constraint ranking but rather with ensuring that this ranking is attained systematically, for all speakers, and not just occasionally.

It is at this point that a preference for M≫F becomes relevant.[17] In the settings discussed in section 2.3.1 above, with binary features and acquired constraints, M≫F does not solve the problem for ROTB, and adopting CURs would be needed to ensure the learning of the distribution of stridents. With constraints that are given in advance, on the other hand, M≫F enables successful acquisition: as we just saw, the challenge in this case is not justifying the constraints (which, in the current setting, are already provided) but rather ensuring that the markedness constraints outrank the faithfulness constraints; by stipulation, a preference for M≫F addresses this challenge.[18] The combination of M≫F with constraints that are given in advance, then, is one way to preserve ROTB in the face of the learnability

[16]We have consulted three native speakers of Dutch, who all rejected *[ɔʃtər].

[17]As discussed earlier, variants of such a preference have been used within RCS approaches, which are not considered in this chapter, to increase the restrictiveness of the grammars arrived at. Within inductive learning approaches, such as the MDL one discussed here, a preference for M≫F can similarly be implemented, most straightforwardly through the cost scheme for the statement of rankings.

[18]The same reasoning applies if *Cost*(/ʃ/) = *Cost*(/s/) or if *Cost*(/ʃ/) < *Cost*(/s/). In both cases, the problematic ranking IDENT[ANT]≫*ʃ can be avoided with a preference for M≫F (but otherwise remains a problem for an ROTB learner).

challenge (in effect, by giving the child knowledge of the pattern as part of the initial state). We now turn to a less stipulative response available within rule-based phonology.

2.3.3 Special case: rule-based phonology with underspecification

So far we have assumed that features are binary. This assumption contributed to the fact that an ROTB child would always store part of the distribution of stridents faithfully, which in turn made it superfluous to acquire that part of the distribution within the input-output mapping, thus leading to the challenge to ROTB. We saw two responses to this challenge: allowing CURs (and thereby rejecting ROTB); and endowing the child with prior knowledge of the pattern (in the shape of constraints that are given in advance combined with $M \gg F$). If underspecification is allowed, a third response suggests itself: if storing an underspecified value such as [0*ant*] is less costly than either of the specified values, the learner might prefer storing all stridents unfaithfully as [0*ant*] and invest in the markedness constraints *sj and *ʃ, along with a high-ranking markedness constraint such as *[0*F*] that blocks underspecified values from surfacing. This response still requires something like $M \gg F$ to ensure that *sj and *ʃ outrank faithfulness (since otherwise inappropriate stridents in nonce forms could be accommodated, as discussed earlier), but otherwise this seems like a way to allow ROTB to be maintained without giving full prior knowledge to the learner (since the markedness constraints are now acquired). However, as we now show, the help that underspecification offers ROTB is considerably more limited than might appear to be the case: within OT, capturing certain simple cases of systematic gaps will still require both innate constraints and $M \gg F$, which means that underspecification leaves the challenge to ROTB without change; and within rule-based phonology, underspecification will enable general learning while maintaining ROTB, but only under specific assumptions.

The problem with using underspecification to succeed in learning while maintaining ROTB is that, while underspecification indeed makes a correct grammar (with underspecified URs and an investment in the requisite markedness constraints) more economical than the kinds of incorrect grammars considered earlier, it sometimes makes a new kind of incorrect grammar more economical still. As a concrete illustration, consider a case of four consonants such as the velar obstruents [k], [g], [x], and [ɣ], which are identical with respect to all features but two (*voice*, which distinguishes the voiced [g] and [ɣ] from the voiceless [k] and [x], and *continuant*, which distinguishes the continuants [x] and [ɣ] from the stops [k] and [g]). Consider now a language that has exactly three of those four consonants – for example, German, which has [k], [g] and [x], but not [ɣ]. To correctly rule out surface forms with [ɣ], the German-learning ROTB child will need to learn a high-ranking markedness constraint such as *ɣ. Earlier, with binary features, a ROTB learner would have had no motivation to posit such a constraint: in analogy with our discussion of Dutch and Bengali, an incorrect grammar storing voicing faithfully (and with no need for *ɣ) would have been optimal. With underspecification, faithful storage of that kind is no longer optimal. In particular, storing the attested [x] as underspecified for voicing in the lexicon will provide an incentive to derive the voicelessness of [x] through the input-output mapping using *ɣ, which, in turn, would correctly prevent URs with underlying /ɣ/ from surfacing faithfully, as in the following grammar G_1:

(175) G_1 (complex; correct)

a. Lexicon: [x] is stored as [0*voice*]; [k],[g] are stored faithfully

b. Constraint ranking: {*[0*voice*],*ɣ} ≫ IDENT[VOICE]

The challenge to ROTB with underspecification and acquired constraints is that the correct G_1 has a simpler but incorrect alternative grammar G_2 which stores not only [x] but also [k] as underspecified for voicing, and which maps only underspecified velars to voiceless:

(176) G_2 (simple; incorrect)

a. Lexicon: [k] and [x] are stored as [0*voice*]; [g] is stored faithfully

b. Constraint ranking: *[0*voice*] ≫ IDENT[VOICE] ≫ *[+*voice*]

G_2 is simpler than G_1 for two reasons. First, its constraint ranking is slightly simpler since it replaces the specific constraint *ɣ (= *[*velar*, +*cont*., +*voice*]) with the more general constraint *[+*voice*]. Second, and much more significantly, its lexicon is more economical since it stores more features as underspecified than G_1 does. As opposed to G_1, the simpler

G_2 does not rule out [ɣ], because the faithfulness constraint IDENT[VOICE] preserves underlying instances of /ɣ/. And crucially, the existence of [g] prevents a preference for $M \gg F$ from being helpful: a ranking with the markedness constraint *[+*voice*] above IDENT[VOICE] would rule out both [ɣ] and [g] and will thus fail to generate the data. In other words, an ROTB learner with underspecification and acquired constraints will converge on G_2 and fail to capture the absence of [ɣ]. Avoiding this problem requires a combination of *ɣ and $M \gg F$, which in turn means that within OT, underspecification offers ROTB no learnability advantage over full specification.

The problem with using underspecification to keep ROTB can be replicated within rule-based phonology, where G_1' and G_2' serve as counterparts of G_1 and G_2:

(177) G_1' (complex; correct)

a. Lexicon: [x] is stored as [0*voice*]; [k],[g] are stored faithfully

b. Rule: [*velar*, +*cont.*] $\rightarrow$ [−*voice*]

(178) G_2' (simple; incorrect)

a. Lexicon: [k] and [x] are stored as [0*voice*]; [g] is stored faithfully

b. Rule: [*velar*, 0*voice*] $\rightarrow$ [−*voice*]

As with OT, the incorrect rule-based solution of G_2' is more economical than the correct one in G_1'. Differently from OT, however, rule-based phonology offers a straightforward way to avoid the problem: if it is impossible to refer to underspecified values within a rule (but still cheaper to state underspecified values in the lexicon, so that the rule in (177) will be learned at all), then the correct G_1' will be acquired. The following summarizes the two assumptions that we relied on here to preserve ROTB in the present setting (by benefiting from underspecification on the one hand and avoiding the trap of G_2' on the other hand):

(179) Assumptions that allow ROTB to be maintained within rule-based phonology:

a. The cost of storing an underspecified feature is strictly lower than that of storing a fully specified one

b. Rules may not refer to underspecified feature values

Within rule-based phonology, then, and assuming the possibility of underspecification along with the specific statements in (179), CURs are not necessary to learn systematic gaps (as in the Dutch, Bengali, and now German pattern above) even if knowledge of the pattern is not given to the learner in advance.

2.3.4 Phonotactic learning does not help ROTB

The discussion in the previous sections illustrated the challenge for ROTB using primarily the pattern of distribution of stridents in Dutch. Similar patterns abound, and the same argument could be made with other instances. In the present section we discuss a potential concern with the challenge. We assumed that acquiring the Dutch pattern is part of full phonological learning, where what is acquired is not just the constraint ranking (possibly along with the constraints themselves) but also the lexicon. It is conceivable, however, that the relevant constraint ranking is established at an early stage of purely phonotactic learning without reference to the lexicon (perhaps along the lines of Hayes and Wilson 2008's Maximum Entropy learner); this ranking could then be passed along to a second, more complete phase of phonological learning, which could proceed without requiring CURs.[19]

We are not aware of actual implementations of the idea just sketched. Regardless of whether such an idea might handle the Dutch pattern and allow ROTB to be maintained in that specific case, there are other patterns in which an earlier phonotactic stage will be of little help and in which our reasoning regarding ROTB can be repeated, which in turn means that the challenge to ROTB will remain unchanged. In particular, a phonotactic first step will run into difficulties with any pattern where the 'elsewhere' part is obscured and cannot be directly observed from surface forms alone. Opaque interactions are a case in point. For example, McCarthy (2007b) discusses a case of opacity in Bedouin Arabic that allows us to replicate the reasoning regarding CURs and ROTB in a setting that is considerably more challenging for phonotactic learning than the Dutch pattern. Bedouin Arabic has a process that palatalizes velar consonants before front vowels. It also has a process that deletes high

[19]Such an approach would need to ensure that in the second, phonological stage, the constraint(s) banning [ʃ] in 'elsewhere' environments will not end up being ranked below the relevant faithfulness constraints. For the purposes of the present section, we will assume that there is a principled way to prevent the phonological-learning stage from ranking the markedness constraints for 'elsewhere' environments too low.

vowels in certain environments. Palatalization precedes deletion, which results in certain velars surfacing as palatalized because of an underlying /i/ that then deletes, an instance of counterbleeding opacity (e.g., /haːkim-iːn/ [haːkʲm-iːn] 'ruling(m.pl.)').

Our reasoning from Dutch can be restated for the Bedouin Arabic palatalization process, which in turn will allow us to restate our challenge to ROTB without the possible escape hatch of earlier phonotactic learning. Restating the discussion in section 2.3.1.1 as an example: if velars are costlier to store as palatalized, an ROTB learner who needs to induce the constraints in the phonology will fail to acquire the part of the pattern that says that velars should not be palatalized in 'elsewhere' environments (here, not preceding a front vowel). The reasoning is familiar by now: velars will be stored as alveolar in 'elsewhere' environments, which will give the child no reason to invest in a markedness constraint that penalizes palatalization; as an adult, they will then fail to rule out palatalized velars in 'elsewhere' environments.

Because of the opacity involved in this case, however, it is hard to see how an earlier phase of phonotactic learning might help: surface velars in 'elsewhere' environments appear sometimes as palatalized (when a following underlying /i/ was deleted) and sometimes as alveolar; consequently, a phonotactic constraint banning palatalized velars in 'elsewhere' environments will clash with the surface pattern and will most likely not be induced. Penalizing palatalization must be left for the lexicon-aware stage of full phonological learning, where, as we just saw, an ROTB learner will fail unless knowledge of the pattern is given to the learner in advance. We conclude that, even if an early phonotactic stage is available and offers a way to acquire some markedness constraints, the challenge to ROTB remains without change.

2.4 Discussion

We started this chapter by asking how the distinction between accidental and systematic gaps is represented and how the relevant knowledge is acquired. We pointed out a connection between the two questions: assuming that phonological knowledge is acquired using MDL (or a similar inductive approach), and across several different representational choices, we must allow CURs to be stated as part of the grammar if we wish to account for speakers' ability to distinguish between the two kinds of gap. The one major exception concerns the possibility that the 'elsewhere' knowledge is guaranteed to be available by some independent principle, such as the combination of constraints given in advance and a preference for markedness outranking faithfulness.

The general shape of the argument was this. A ROTB learner will usually store part of the distribution of stridents faithfully. Since these stridents then surface faithfully (whether through an independently-acquired faithfulness constraint or through the default faithfulness on a rule-based system), stating the knowledge of the relevant part of the pattern of distribution through the input-output mapping will be superfluous and will not be acquired by an MDL learner. But in the absence of such a statement, speakers would be predicted to accept nonce forms in which this material appears in an inappropriate environment, contrary to fact. The solution, then, is to do one of the following: either (a) abandon ROTB and allow the learner to eliminate predictable material not just from the lexicon but, using a CUR, from the very alphabet in which the lexicon is written; or (b) bypass the challenge by either minimizing the learning task (for example, by providing in advance both the constraints and their preferred ranking) or by ensuring that stridents are not stored faithfully (for example, by using underspecification and rule-based phonology, along with certain additional assumptions, as discussed above).

This disjunctive conclusion might seem reassuring for ROTB: after all, the first choice within the (b) option is quite close to the view, common within OT, that all constraints are given in advance and that the markedness constraints are ranked above the faithfulness constraints unless forced otherwise. However, this conclusion also highlights the stakes for the combination of given constraints and markedness over faithfulness. In the OT literature, these assumptions are often bundled together with ROTB, but this bundling is not logically necessary: it is easy to imagine either component being true while the other is false (or that both are true or both are false). What we have shown is that there is an *empirical* dependence between them: the patterns of well-formedness that speakers show are such that, if the combination of constraints given in advance and markedness over faithfulness is false within OT, then ROTB must be abandoned (since otherwise part of the pattern becomes unlearnable).[20] Consequently, any attempt within OT to defend ROTB that does not reject MDL must involve a defense of markedness over faithfulness, along with constraints given in advance. Since language-specific constraints have occasionally been proposed in the literature (see, e.g., Kager and Pater 2012, Pater 2014 and references therein), and since the empirical support for markedness over faithfulness has been thin (though see Davidson et al. 2004), this challenge strikes us as nontrivial, though a proper assessment would clearly take us beyond the scope of the present chapter.

Looking past the specific question of CURs versus ROTB, this note illustrates a way in which a general learning criterion can help evaluate competing representational possibilities. The idea is not new. Works such as Halle 1962, 1978, Baker 1979, and Dell 1981 used the simplicity metric of early generative grammar to argue for specific conclusions about representations. As noted above, however, the simplicity metric lacks a pressure for a tight fit of the data (in terms of the MDL quantity $|G| + |D : G|$, the simplicity metric minimizes $|G|$ but does not include a counterpart for $|D : G|$). Consequently, the simplicity metric leads to overly general hypotheses and is not a suitable metric for learning. From the perspective of architecture comparison, the simplicity metric often leads to incorrect conclusions about what representations are learnable.[21] More recently, Katzir 2014 and Piantadosi et al. 2016 have raised the possibility of using MDL and Bayesian reasoning to evaluate competing architectures, thus returning to the kind of architecture comparison envisioned in early generative grammar but with a better supported approach to learning than the one assumed at the time. The present note offers a concrete application of this idea to an actual architectural question.

[20]Note that things could have been different. For example, if $cost(ʃ) > cost(s)$ and if speakers rejected [osjər] but accepted [ɔʃtər], then ROTB could have been maintained without markedness over faithfulness or even constraints that are given in advance. The observation is thus a contingent connection that happens to be true of humans.

[21]In particular, it incorrectly suggests that restricted optionality of the kind studied in Baker 1979 and Dell 1981 is unlearnable without severe representational limitations. See Rasin et al. 2017 for discussion of this case and a comparison of learnability using the simplicity metric and using MDL. Not all uses of the simplicity metric suffer from this problem. As far as we can tell, Halle 1962, 1978's simplicity argument for feature-based representations stands.

Chapter 3

Morpheme structure constraints and blocking in nonderived environments

3.1 Introduction

In Nonderived Environment Blocking (NDEB), a phonological process applies across morpheme boundaries or morpheme-internally when fed by another phonological process but is otherwise blocked. To illustrate the phenomenon of NDEB, I will use a simple artificial example modeled after the description of Finnish assibilation in Kiparsky (1973, 1993) which I will refer to as *Finnish' assibilation*.[1] In this artificial example, assibilation turns the stop /t/ into the strident [s] before the high vowel /i/. The process applies before suffixes that begin with /i/ (180a). Morpheme-internally, it applies only when the high vowel is the result of final-vowel raising (which raises /e/ to [i] word-finally), as in (180b). Otherwise, assibilation does not apply within morphemes (180c). The underlying sequence /ti/ is often referred to as a *derived environment* in (180a) and (180b) and as a *nonderived environment* in (180c).

[1]The literature on Finnish assibilation following Kiparsky (1973, 1993) has challenged the original description and offered alternative analyses of Finnish assibilation according to which the process is not blocked in nonderived environments. See Hammond (1992), Wolf (2008), and especially Anttila (2006). The present chapter will not contribute to the debate about Finnish assibilation. Since the process as originally described is a familiar and simple case that demonstrates two types of NDEB effects simultaneously, it will be convenient to use a constructed variant of it for illustration before turning to slightly more complex cases from natural languages later in the chapter.

(180) Assibilation in the artificial language Finnish'

a. Assibilation applies across a morpheme boundary:

lut-a ~ lus-i (/lut-i/)

b. Assibilation applies morpheme-internally when fed by final-vowel raising (e $\rightarrow$ i / _ #):

vete-pa ~ vesi (/vete/)

c. Otherwise, assibilation is blocked morpheme-internally:

- tila
- niti

NDEB is an instance of under-application opacity that poses a challenge to both rule-based phonology and OT: in rule-based phonology, a rule of assibilation that turns the stop /t/ into the strident [s] before the high vowel /i/ would incorrectly apply to nonderived /ti/ sequences if no conditions on its application are posited. Similarly, in OT, the markedness constraint *ti would equally penalize derived and nonderived surface sequences of [ti]. And if *ti is allowed to be repaired by assibilation in derived environments (by appropriately ranking it over faithfulness constraints like IDENT[cont]), assibilation would incorrectly apply in nonderived environments as well. More generally, if *P* is a process that is blocked in nonderived environments, the challenge in both frameworks is to partition the set of environments of application of *P* into two subsets – corresponding to derived and nonderived environments – and block the application of *P* precisely in nonderived environments. Previous works that have tried to address the challenge include Mascaró (1976), Kiparsky (1993), Burzio (2000) Inkelas (2000), Łubowicz (2002), McCarthy (2003a) van Oostendorp (2007), Kula (2008), Wolf (2008), and Anttila (2009), among others.

My goal in this chapter is to develop and defend a theory of NDEB which is an extension of Kiparsky's (1993) underspecification theory of NDEB. According to Kiparsky, a process which shows NDEB is structure-building, which means that it can apply to underspecified but crucially not to fully specified structure. On this view, Finnish' assibilation would be a feature-filling rule that applies to underspecified /T/ but cannot apply to fully specified /t/ (181a). /T/ here stands for a variant of the voiceless alveolar stop /t/ in which the feature [continuant] is not specified. A default rule which applies after assibilation would convert any underspecified /T/ that did not undergo assibilation into [t] (181b). (182) shows the derivation of hypothetical [timas-i], which includes two potential environments for assibilation, one which is fully contained in the stem and one which spans a morpheme boundary. According to Kiparsky, the second environment contains an underspeficied /T/ which ends up undergoing assibilation; the first environment contains a fully-specified /t/ which is protected from assibilation.

(181) Grammar of Finnish' according to Kiparsky's (1993) theory

a. T → s / _ i

b. T → t

(182) Derivation of [timasi]

UR	/timaT-i/
T → s / _ i	timas-i
T → t	-
SR	[timasi]

As noted by Burzio (2000) and discussed in detail by Inkelas (2000), Kiparsky's analysis of NDEB is missing a crucial component, as it leaves the underlying distribution of underspecified /T/ and fully-specified /t/ as an accident of the lexicon. Nothing in the analysis prevents fully-specified /t/ from occurring stem-finally and blocking assibilation before a suffix-initial /i/. The grammar thus freely generates ungrammatical forms such as *[rat-i] in which assibilation has not applied across a morpheme boundary, as illustrated in (183). In other words, Kiparsky's theory cannot represent *obligatory* NDEB processes, but such processes are attested (one concrete example from Romanian will be discussed below).

(183)

UR	/rat-i/
T → s / _ i	-
T → t	-
SR	*[rati]

In this chapter, I propose to address Burzio's and Inkelas' challenge for Kiparsky's theory directly, using Morpheme Structure Constraints (MSCs) which regulate the distribution of underspecified /T/ and fully-specified /t/ in the lexicon.[2] In particular, I will propose that the grammar of a language like Finnish' includes the MSC in (184). This MSC bans a fully-specified /t/ from underlying morpheme-final position and thus blocks monomorphemic URs like /rat/ which generate ungrammatical forms under Kiparsky's theory.

(184) MSC IN FINNISH' (informal): /t/ occurs before /i/; /T/ occurs elsewhere

The MSC in (184) seems like a suspicious addition to the grammar, since it makes a distributional statement in the lexicon that applies to exactly the same environments that trigger assibilation in the phonology. I will argue, however, that this MSC plays a key role in the theory of NDEB. I will show, using several case studies, that more than providing a simple theory of NDEB, this theory is also more successful than previous theories of NDEB proposed in the literature in accounting for known cases of NDEB. If this view of NDEB is correct, it would provide further support for a dual-component architecture of phonology as in SPE and against ROTB.

The chapter is structured as follows. First, in section 3.2, I implement an architecture that uses MSCs to regulate the distribution of underspecified and fully-specified structure in the lexicon (3.2.1) and develop an analysis of Finnish' assibilation within this architecture (3.2.2). Then, in section 3.3, I use several case studies to discuss the predictions of the MSC-based approach for NDEB. In section 3.4 I discuss the predictions of previous approaches to NDEB, including Inkelas (2000), the Strict Cycle Condition (Mascaró, 1976), Coloured Containment (van Oostendorp, 2007), Optimal Interleaving with Candidate Chains (Wolf, 2008), and Sequential Faithfulness (Burzio, 2000). I show that none of those approaches can correctly account for all the case studies from section 3.3.

3.2 Proposal

3.2.1 Architecture

This subsection describes the phonological architecture that will be used in 3.2.2 for an account of NDEB. My claim in this chapter is that NDEB supports a component that re-

[2] In this chapter I will use the term Morpheme Structure Constraints to refer to constraints on underlying representations that apply to individual morphemes in the lexicon.

stricts possible URs in the lexicon. I will have nothing to say about the phonological formalism (e.g., rule-based or constraint-based) or the nature of lexical representations (e.g., underspecified or fully specified). To make the proposal explicit, I will present it using a ruled-based formalism and underspecification. [3] The architecture, which I now describe, is schematized in (185).

(185) Architecture

Lexicon	$\{M_1\}$ $\{M_2\}$ $\{M_3\} \leftarrow \Sigma_L$ $\downarrow$ $\downarrow$ $\downarrow$ $/M_1/$ $/M_2/$ $/M_3/$
Phonology	$/M_1 - M_2 - M_3/$ $\downarrow$ $[M_1 - M_2 - M_3]$

A central component of the architecture is the mapping from URs to surface forms, which is implemented here using ordered phonological rules as in SPE. I assume that a phonological grammar includes an alphabet – an inventory of feature bundles Σ – the elements of which can be concatenated. For example, if $k, a, t \in \Sigma$, then {kat} and {takta} are possible concatenations, among many others. I assume that individual languages can

[3]The choices of rules and underspecification are arbitrary: the mappings presented in this chapter using a rule-based formalism can be reformulated using constraints, and one can think of variants of the proposal that do not make use of underspecification. For example, the distinction between unspecified and specified features can be replaced with a distinction between plain specified features and specified features alongside an exception diacritic that prevents a feature from being changed by a particular rule.

restrict Σ to a proper subset, call it Σ_L. For a segment $\sigma \in \Sigma$, we can write $\sigma \notin \Sigma_L$, meaning that σ cannot be used for concatenation in that language. For example, if English rules out /x/ from its alphabet and we write $x \notin \Sigma_L$, then {bax} is not a possible concatenation in English. Negative statement such as $x \notin \Sigma_L$ are used for convenience and should not be taken to be grammatical constraints per se. What I mean by writing $x \notin \Sigma_L$ is that Σ_L, which could be positively stated in the grammar as a set of segments, does not include x. I will refer to representations created by concatenating elements from Σ_L as *initial representations*, and I will mark them using curly brackets, as in {anta}. Morpheme structure rules map initial representations to URs. For example, if {anta} is an initial representation and post-nasal voicing (t → d / n _) is the only morpheme structure rule in the grammar, the result of applying post-nasal voicing to {anta} would be the UR /anda/. Morpheme structure rules have the same format as ordinary rules, but they apply to isolated morphemes in the lexicon before the morphemes are combined. In this framework, then, URs are created in two steps: first, elements from Σ_L are concatenated to form an initial representation. Then, morpheme structure rules apply and map this representation to a UR. Later on, phonological rules map URs to surface forms.

In addition, I assume that lexical representations may be underspecified: segments in Σ (and in Σ_L) may be underspecified for some of their features. See Kiparsky (1982), Archangeli (1988), and Steriade (1995) for relevant discussion. For example, a variant of the voiceless alveolar stop [t] in which the feature [continuant] is not specified may be in Σ. We can refer to this segment as [T] and write $T \in \Sigma$. Underspecified features are filled in either by morpheme structure rules or by phonological rules. Finally, both morpheme structure rules and phonological rules may be *feature filling*. This means that they can target segments underspecified for some feature F and fill in the relevant value but, crucially, without affecting segments that are already specified for F. Example (186) demonstrates the property of feature filling using a version of Finnish assibilation that applies to underspecified [T].

(186) Assibilation: T → s / _ i (feature-filling)

a. /Ti/ $\xrightarrow{\text{assibilation}}$ [si]

b. /ti/ $\xrightarrow{\text{assibilation}}$ [ti]

3.2.2 Analysis

In this subsection I provide an analysis of Finnish' assibilation using the architecture described in 3.2.1. The basic pattern of Finnish' assibilation was presented above in (180a)-(180c) and is repeated here as (187a)-(187c). Following the convention in the literature, I use the term *morphologically-derived environment* to refer to an environment created through affixation, as in (187a), and *phonologically-derived environment* to refer to an environment created through the application of a phonological process, as in (187b).

(187) a. Assibilation applies across a morpheme boundary:

lut-a ~ lus-i (/lut-i/)

b. Assibilation applies morpheme-internally when fed by final-vowel raising (e → i / _ #):

vete-pa ~ vesi (/vete/)

c. Otherwise, assibilation is blocked morpheme-internally:

- tila
- niti

The first ingredient in the analysis is the rule of assibilation (188), which, following Kiparsky (1993), I take to be a feature-filling rule that specifies the voiceless alveolar [T] as [+continuant].[4] The second ingredient is a rule that I refer to as *anti-assibilation* (189). Anti-assibilation is similar to the rule of assibilation: it is a feature-filling rule that applies in the same environment (/Ti/) and fills in a value for the feature [continuant]. The only difference between the two is that anti-assibilation specifies that value as [-continuant] rather than [+continuant]. That is, anti-assibilation specifies /T/ as [t].

(188) Assibilation

T → s / _ i (feature-filling)

[4]For presentational ease, I ignore the feature [strident], which could be filled in by the assibilation rule itself or by a separate rule.

(189) Anti-assibilation

$T \rightarrow t$ / _ i (feature-filling)

To see how assibilation and anti-assibilation interact, consider the UR /Ti/ and a hypothetical grammar in which anti-assibilation is ordered before assibilation. The derivation is provided in (190). First, anti-assibilation applies and specifies /T/ as [t]. Then, assibilation does not apply since its structural description is not met: the rule is feature filling, but /t/ is not underspecified for continuancy. The result is the surface form [ti]. In short, anti-assibilation bleeds assibilation by destroying its environment of application.

(190) Interaction between assibilation and anti-assibilation (hypothetical grammar)

UR	/Ti/
$T \rightarrow t$ / _ i	ti
$T \rightarrow s$ / _ i	-
SR	[ti]

My proposal is that in the grammar of Finnish', anti-assibilation is a morpheme structure rule that applies to isolated morphemes, whereas assibilation is a phonological rule that is part of the mapping from URs to surface forms. Fully-specified /t/ is not part of the Finnish' alphabet.

(191) Morpheme structure component:

a. $t \notin \Sigma_L$

b. $T \rightarrow t$ / _ i

The consequence for the form of URs in Finnish' is that /t/ and /T/ are in complementary distribution in the lexicon: /t/ occurs only before /i/ (following the application of anti-assibilation) and /T/ occurs elsewhere. Here are some examples. 192a shows the derivation of the UR /tila/. Since $t \notin \Sigma_L$, any instance of /t/ in URs must be derived from /T/. The initial representation is therefore {Tila}, which anti-assibilation maps to /tila/. 192b indicates that /lata/ is not a possible UR in Finnish': since $t \notin \Sigma_L$ and the environment for anti-assibilation is not met before /a/, /t/ cannot occur in a a pre-/a/ position.

(192) a. {Tila} $\rightarrow$ /tila/

b. */lata/

c. /laTa/, /luT/

In (192c), anti-assibilation does not apply, and /T/ remains underspecified. The value for [continuant] will be filled in by the mapping from URs to surface forms: the rule of assibilation turns /T/ into [s] before /i/; otherwise – that is, whenever assibilation does not apply – /T/ is specified as [t] through the default rule T → t.

(193) Phonological rules:

a. T → s / _ i

b. T → t

Example (195) demonstrates the application of phonological rules in the derivation of the alternants in (194), assuming the UR /luT/ for the stem.

(194) lut-a ~ lus-i

(195)

UR	/luT-i/	/luT-a/
T → s / _ i	lusi	-
T → t	-	luta
SR	[lusi]	[luta]

This is the grammar of Finnish' we have so far:

(196) a. Morpheme structure component:

- $t \notin \Sigma_L$
- T → t / _ i

b. Phonological rules:

- T → s / _ i
- T → t

I will now show why this grammar applies assibilation in morphologically-derived environments but not in nonderived environments. Consider the derivation of hypothetical [timas-i], which alternates with [timat-a] and includes two potential environments for the application of assibilation: the first is morpheme-internal, and the second spans the morpheme boundary. Assibilation only applies in the latter.

(197) timat-a ~ timas-i

First, morpheme structure rules apply to each morpheme individually (198a). Since t $\notin \Sigma_L$, the initial representation of the stem must be {TimaT}. Anti-assibilation applies to the first instance of T, but not to the second: at this stage of the derivation, the second T is stem-final and the environment for anti-assibilation is not met. The result is the UR /timaT/, where only the second T remains underspecified for continuancy. In the mapping from URs to surface forms (198b), assibilation successfully applies to the sequence /T-i/ which was created through affixation. It does not apply to the stem-initial /ti/, which at this point is already fully specified. The final surface form is [timasi].

(198) Derivation of [timas-i] (alternant: [timat-a])

a. Morpheme structure rules apply to each morpheme individually:

- {TimaT} → /timaT/
- {i} → /i/

b. Phonological rules apply:

UR	/timaT-i/
T → s / _ i	timas-i
T → t	-
SR	[timasi]

The next step is to show why assibilation applies in phonologically-derived environments. Recall that final-vowel raising (199) raises a word-final /e/ to [i] (199a). Assibilation may apply morpheme-internally when fed by final-vowel raising (199b).

(199) e → i / _ #

a. juke-pa ~ juki

b. vete-pa ~ vesi

Here, nothing further has to be said. Final-vowel raising is ordered before assibilation (200). In words like [vesi], alternating /T/ precedes /e/ in the UR, so anti-assibilation does not get to apply. /T/ remains underspecified, which means that assibilation will get to apply after affixation. The full derivation is provided in (201).

(200) a. Morpheme structure component:

- $t \notin \Sigma_L$
- $T \rightarrow t$ / _ i

b. Phonological rules:

- $e \rightarrow i$ / _ #
- $T \rightarrow s$ / _ i
- $T \rightarrow t$

(201) Derivation of [vesi]

a. Morpheme structure rules apply (vacuously):

{veTe} → /veTe/

b. Phonological rules apply:

UR	/veTe#/
e → i / _ #	veTi#
T → s / _ i	vesi#
T → t	-
SR	[vesi]

In sum, a process *P* that is blocked in nonderived environments applies unless its focus is made immune in an earlier stage of the derivation. Foci can be made immune by a feature-filling rule *anti-P* that shares its structural description with *P* and can apply to isolated morphemes in the lexicon. *Anti-P* thus induces the following partition on the set of environment of *P*:

(202) Partition into nonderived and derived environments

a. Environments present when *anti-P* applies (correspond to nonderived environments)

b. All other environments (correspond to derived environments)

Rules of the form *anti-P* are unusual rules. Their formulation seems arbitrary and their environment duplicates the environment of *P*, and at present I have nothing to say to make them seem less arbitrary. In what follows, I will assume that such rules are available without trying to derive their existence from deeper principles. Instead, I will focus on the picture of NDEB that arises from (202) and evaluate the success of the MSC-based theory in accounting for known cases of NDEB compared to previous proposals.

3.3 Predictions and case studies

3.3.1 Blocking is determined before suffixation

The MSC-based theory determines the alternation status of a feature at the individual morpheme level. Consider again the blocking of Finnish' assibilation in morphologically-nonderived environments, using the example [timas-i] (which alternates with [timat-a]). For the MSC-based theory, blocking is exclusively determined according to the environment of each potential input for assibilation in the stem /timaT/: the first consonant – but not the stem-final consonant – precedes /i/ and therefore becomes immune to assibilation. Once a potential input to assibilation is made immune before suffixation, it is predicted to stay immune even if the environment for assibilation is re-created after suffixation. To illustrate this prediction, suppose that Finnish' had a process of vowel deletion that deletes stem-final vowels pre-vocalically and may feed assibilation:

(203) /munte-i/ → [munsi]

The prediction is that assibilation would not apply to suffixed forms if the stem-final vowel is /i/, even though the environment for assibilation is re-created through suffixation and spans the morpheme boundary:

(204) Prediction of MSC-based theory for /munti-i/

/munti-i/ → [munti]

On the MSC-based approach, the morpheme structure component captures the distinction between the two verbs at the UR level: the URs are /munTe/ (with underspecified /T/) and

/munti/ (with fully specified /t/). Assibilation can only apply to the first. /t/ in /munti/ remains immune to assibilation even after the deletion of the stem-final vowel and the addition of an /i/-initial suffix which creates the environment for assibilation. Other theories of NDEB that make this prediction include Kiparsky (1993), Burzio (2000), and Wolf (2008).

In contrast to these theories, much of the previous literature on NDEB has followed the idea that NDEB should be understood through a characterization of the set of derived environments. The guiding intuition is that in both types of environments in which *P* applies – across a morpheme boundary and when part of its environment is the result of another phonological process – part of the environment is 'new', or, stated differently, is introduced in the course of the derivation. In the Finnish' assibilation case, the environment in /lut-i/ is 'new' because it is formed through affixation, and the environment in /veti/ (derived from /vete/ through vowel raising) is 'new' because the high vowel is the result of vowel raising. Theories guided by this idea, like the Strict Cycle Condition (Mascaró, 1976) and Coloured Containment (van Oostendorp, 2007), incorporate a notion of 'new' or 'derived' environments into the grammar and often introduce a licensing condition to allow the application of *P* only in such environments. I will refer to such theories as *derived-environment* theories.

In derived-environment theories, application is determined based on the morphologically-complex form: for [timas-i], the relevant representation would be /timat-i/, the suffixed form before the application of assibilation. Assibilation applies in the second environment (/timat-i/) but not in the first (/timat-i/) since only the second environment is 'derived' and spans a morpheme boundary. Below, I will discuss in more detail some of these approaches and how they enforce application in derived environments. For now, what matters is that they all license application across a morpheme boundary:

(205) *Prediction of derived-environment theories*

Spanning a morpheme boundary is a sufficient condition for licensing.

For the Finnish' example /munti-i/, the prediction is that assibilation would apply after vowel deletion, since an environment that triggers assibilation would span a morpheme boundary:

(206) Prediction of derived-environment theories for /munti-i/

/munti-i/ → [munsi]

In this section I show that the prediction of the MSC-based theory is correct, using Romanian palatalization as a case study, and using data from unpublished notes by Donca Steriade (Steriade 2008b).

In Romanian, a palatalization rule turns a velar stop into a palatal before a front vowel or glide:

(207) a. k → tʃ / _ {e, i, j}

b. g → dʒ / _ {e, i, j}

Palatalization applies across morpheme boundaries (208) and is blocked morpheme-internally (209).[5]

(208) mak ~ matʃ-j 'poppy-SG.'~'poppy-PL.'

(209) a. unkj 'uncle-SG.'

b. rokie 'dress-SG.'

c. paket 'package-SG.'

Vowels are deleted before the plural suffix /-i/, which is sometimes realized as a glide (210). The vowel-glide alternation is irrelevant for our current purposes, so I will leave it as a black box in what follows, assuming that deletion applies pre-vocalically and that a cover rule i → j, which is responsible for the glide-vowel alternation, applies after deletion.

(210) a. metru ~ metr-i 'meter-SG.'~'meter-PL.'

b. bere ~ ber-j 'beer-SG.'~'beer-PL.'

c. popʌ ~ pop-j 'priest-SG.'~'priest-PL.'

[5]For presentational ease, I have omitted secondary palatalization from the examples below. The distribution of secondary palatalization is irrelevant for our purposes.

Crucially, palatalization is blocked exactly when the deleted vowel had been a palatalization trigger: In (211a), the final vowel in the singular is a back vowel and palatalization applies in the plural. In (211b), the final vowel is a front vowel and palatalization in the plural is blocked.

(211) a. mɨnekʌ ~ minetʃ-j 'sleeve-SG.'~'sleeve-PL.'

b. pʌduke ~ pʌdu**k**-j 'louse-SG.'~'louse-PL.'

This contrast is quite general. The following table demonstrates the behavior of palatalization in the plural form of every nominal declension class that takes the plural suffix /-i/: for each class, the two rightmost columns indicate the identity of the stem-final vowel and whether palatalization applies in the plural form.[6]

(212) Palatalization in Romanian nouns that take the plural suffix /-i/

	Noun-SG.	Noun-PL.		Final vowel	Palatalzation applies
MASC					
a.	mak	matʃ-j	'poppy'	/u/	✓
b.	pʌduke	pʌdu**k**-j	'louse'	e	*
c.	dukʌ	dutʃ-j	'duke'	ʌ	✓
d.	flamiŋgo	flamiŋdʒ-j	'flamingo'	o	✓
FEM					
e.	fabrikʌ	fabritʃ-j	'factory'	ʌ	✓
f.	pereke	pere**k**-j	'pair'	e	*

For the sake of concreteness, let us see why the MSC-based theory accounts for this pattern without modification. The grammar, with anti-palatalization (213) as a morpheme structure rule and palatalization as a phonological rule, is given in (214).

(213) Anti-palatalization

$K \rightarrow k \;\; / \_ \; \{e, i, j\}$

[6]The range of possible noun-final vowels in Romanian is restricted, perhaps suggesting that the final vowel should be regarded as an idiosyncratic theme vowel specified on a root by root basis. If this is true, a necessary assumption on the present account is that the theme vowel is present before the application of anti-palatalization.

(214) a. Morpheme structure component:

- $k \notin \Sigma_L$
- K → k / _ {e, i, j}

b. Phonological rules:

- V → ∅ / _ V
- K → tʃ / _ {e, i, j}
- K → k
- i → j

Anti-palatalization applies to individual morphemes in the lexicon and specifies K as /k/ in (211b) but not in (211a):

(215) Derivation of [pʌduk-j] (singular: [pʌduke])

a. Morpheme structure rules apply:

1. {pʌduKe} → /pʌduke/
2. {i} → /i/

b. Phonological processes apply:

UR	/pʌduke-i/
V → ∅ / _ V	pʌduki
K → tʃ / _ {e, i, j}	-
K → k	-
i → j	pʌdukj
SR	[pʌdukj]

(216) Derivation of [minetʃ-j] (singular: [minekʌ])

a. Morpheme structure rules apply (vacuously):

1. {mineKʌ} → /mineKʌ/
2. {i} → /i/

b. Phonological processes apply:

UR	/mɨneKʌ-i/
V → ∅ / _ V	mɨneKi
K → tʃ / _ {e, i, j}	mɨnetʃi
K → k	-
i → j	mɨnetʃj
SR	[mɨnetʃj]

Romanian palatalization, then, supports the prediction of the MSC-theory. In section 3.4.2, I discuss specific derived-environment theories, which make the prediction in (205) and thus incorrectly predict application of palatalization in /pʌduke-i/.

3.3.2 Blocking within suffixes

The MSC-based analysis of Finnish' assibilation relies on MSCs to restrict the distribution of /t/ and /T/ in the lexicon: non-alternating /t/ occurs before /i/, and alternating /T/ occurs elsewhere. MSCs apply to each morpheme individually and ensure that stems like /timaT/ have the desired specification before suffixation. The same result could be alternatively achieved in a cyclic architecture where MSCs are replaced with first-cycle evaluation and the distributional restriction applies once, before suffixation. My goal in this section is to discuss the cyclic variant of the MSC-based proposal, which can successfully capture the Finnish' assibilation pattern, and show that it faces a challenge in accounting for cases of NDEB where application is blocked not only within stems, but also within suffixes. As we will see, accounting for blocking within suffixes requires a level of representation in which phonological restrictions apply to suffixes in isolation from the rest of the string – a level available in MSC-based architectures but crucially not in cyclic architectures that reject MSCs.

Cyclic architectures allow phonological processes to be interleaved with affixation. Examples of cyclic architectures are Lexical Phonology and Morphology (Kiparsky, 1982 et seq.), its implementation within OT known as Stratal OT (Kiparsky, 2000), and Halle and Vergnaud's (1987b) theory of the cycle. I will first show that a cyclic variant of the MSC-based analysis can account for Finnish' assibilation without MSCs. In this variant, there are no restrictions on the alphabet, which means that both /t/ and /T/ can be used in writing URs. Moreover, since anti-assibilation is not an MSC, /t/ and /T/ may occur anywhere within URs: URs like /rat/ (with fully-specified /t/ in final position) and /Tila/ (with underspecified /T/ before /i/) can be generated. A cyclic grammar is provided in (217). It contains two rule blocks separated by suffixation. To keep the discussion general and compatible with various cyclic architectures, I will not name the rule blocks and will refer to them simply as Rule block A and Rule block B.[7] Rule block A contains two rules, which mirror the effects of MSCs in the MSC-based analysis. The first rule turns every /t/ to [T], which has a similar effect to the constraint $t \notin \Sigma_L$ in banning /t/ from initial representations. The second rule is the anti-assibilation rule. The remaining rules, including assibilation, are part of Rule block B.

(217) a. Rule block A:

$t \rightarrow T$

$T \rightarrow t \ / \ \_\ i$

b. Add the suffix [-i]

c. Rule block B:

$T \rightarrow s \ / \ \_\ i$

$T \rightarrow t$

The derivation of [timas-i] (which alternates with [timat-a]) using this grammar is given in (218). As there are no MSCs, multiple URs for the stem may lead to the same output. To see the rules in working, I have chosen the UR /Tilat/. Notice that the correct output is derived. The analysis straightforwardly extends to phonologically-derived environments if final-vowel raising is placed in Rule Block B.

(218) Cyclic derivation of [timas-i]

[7]The distinction between Rule block A and Rule block B may correspond to the following distinctions made by cyclic approaches: cyclic vs. post-cyclic, stem-level vs. word-level, lexical vs. post-lexical, etc.

Rule block A	/Tilat/
t → T	TimaT
T → t / _ i	timaT
Suffixation	/timaT-i/
Rule block B	/timaT-i/
T → s / _ i	timasi
T → t	-
	[timasi]

A cyclic architecture, then, can capture the Finnish' assibilation pattern without using MSCs since it can impose the same distributional restriction on the stem before suffixation. More generally, the cyclic architecture succeeds because every nonderived environment is introduced into the derivation before every derived environment. This allows anti-*P* to be ordered at a stage in the derivation after every nonderived environment has been created and before any derived environment has been created, which, in turn, allows anti-*P* to apply exclusively to nonderived environments and *P* to apply later to derived environments.

The two architectures diverge in their predictions when the derivational precedence between nonderived and derived environments required by the cyclic approach breaks down. This may happen when a phonological process that is blocked in nonderived environments is also blocked within suffixes. Cases of such blocking mentioned in the literature are consonant gradation in Finnish (Kiparsky 1993, 2003), spirantization in Luganda (Wolf 2008, citing Odden 1990), and palatalization in Meskwaki (Wier, 2004). The challenge for the cyclic approach from Luganda spirantization is discussed in Wolf (2008, pp. 443-447), and I will present another version of the argument from Finnish consonant gradation.

Finnish consonant gradation (CG) de-geminates a double stop at the onset of a closed syllable and yields alternations like the following:

(219) tten → ten, ttain → tain

Example (220), taken from Kiparsky (1993), is a single example that contains three environments for CG. CG is blocked in the first, nonderived, environment (underlined) and applies in the other two, derived, environments (bold). The second geminate (/...totti.../)

and the third geminate (/...**tt**oma.../) undergo CG since they are onsets of closed syllables at some level of representation.[8]

(220) /hottentotti-ttoma-ta/ → [hottentoti-ton-ta] 'Hottentotless-PART.SG.'

CG is blocked when its environment is fully contained within the suffix -tten, an allomorph of the genitive plural (221).

(221) maa-i-tten *maa-i-ten 'country-PL.GEN'

Non-application in (221) is not yet a problem for the cyclic analysis, since the suffix -tten might be added only after CG gets its last chance to apply. The crucial example is the contrast between (222a), where CG applies optionally, and (222b), where it applies obligatorily:

(222) a. /ullakko-i-hin/ → [ul.la.koi.hin] ~ [ul.lak.koi.hin] 'attic-PL.ESS.'

b. /ullakko-i-tten/ → [ul.la.koit.ten] ~ *[ul.lak.koit.ten] 'attic-PL.GEN.'

This contrast suggests that the suffix -tten itself creates an environment for the application of CG to a preceding geminate (/kk/). In (222a), with a suffix that begins with a non-geminate, CG applies optionally. On one account, optional application is due to two available structures for [oi], one of which triggers CG and one which does not (Keyser and Kiparsky 1984, Kiparsky 2003). Kiparsky (2003: 121) notes that when a diphthong like [oi] is followed by a geminate-initial suffix, CG applies obligatorily rather than optionally. That is, the geminate /tt/ of the suffix -tten plays a role in triggering CG and eliminating the illicit representation *[ul.lak.koit.ten] in (222b), where CG does not apply. To eliminate the illicit variant with non-application in (222b) while not eliminating it in (222a), CG needs to apply again after -tten is added in the derivation. This leads to an ordering paradox for the cyclic approach. On the one hand, CG must be able to apply *after* the addition of -tten to make sure that /kk/ undergoes gradation. On the other hand, anti-CG could not

[8]For the second geminate, the syllable is closed by the third, suffix-initial geminate. For the third geminate, the syllable is closed after deletion of the suffix-final vowel /a/ triggered by the following suffix. In a derivational approach, an explicit analysis could either order vowel deletion before a directional CG rule that applies left-to-right, or (in a cyclic architecture) apply CG after deletion in every cycle. See Kiparsky (1993) for further discussion on CG in an underspecification-based account.

have applied to -tten at any prior level of representation, so, paradoxically, CG must not be able apply once -tten is added (otherwise, it would incorrectly apply to -tten).

The problem for the cyclic approach is that there is no level of representation in which phonological restrictions apply to suffixes in isolation from the rest of the string. Whenever the nonderived environment in -tten is present in the derivation, a derived environment (the hetero-morphemic closed syllable kko-i-t) is present as well. This is why anti-CG cannot apply to -tten without causing trouble elsewhere. MSCs address this problem directly: if anti-CG applies to individual morphemes in the lexicon before they are combined with other morphemes, it can apply to -tten before any derived environment is created.

One way out for a cyclic approach without MSCs is to mark the suffix -tten as an exception to CG. As Wolf (2008) notes, however, this solution will not work for a similar problem in Luganda, where a single suffix that behaves like -tten contains two potential targets for the application of a process that shows NDEB. Only one of the two targets is included in an environment that is fully contained within the suffix, so marking the suffix as an exception will incorrectly block both targets (rather than just one) from undergoing the process.

Citing an example from Odden (1990), Wolf (2008) notes that morpheme-final consonants in Luganda undergo spirantization before /i/. Otherwise, consonants do not undergo spirantization before a tauto-morphemic /i/. This is illustrated using the following example, where the morpheme-final /k/ and /r/ become [s] and [z] respectively, but the first /r/ in /-irir/ does not change.

(223) /lamuk-irir-i/ → [lamus-iriz-i] 'greet without ceasing'

Since the the suffix /-irir/ triggers the application of spirantization to /k/ and fully contains an environment for spirantization itself, it poses the same problem for the cyclic approach as the Finnish suffix /-tten/: spirantization is predicted to apply within the suffix, producing the incorrect *[lamus-iziz-i]. In this case, marking /-irir/ as an exception to spirantization would incorrectly block spirantization from applying to the morpheme-final /r/, incorrectly producing *[lamus-irir-i]. To account for blocking within suffixes, then, the cyclic approach will have to mark as exceptions precisely those suffix-internal targets that are a part of underlying morpheme-internal environments, while the MSC-based approach avoids arbitrary exception-marking in these cases altogether.

3.3.3 Blocking in partially-nonderived environments

3.3.3.1 Blocking in Romanian and Armenian

The MSC-based approach identifies nonderived environments as environments present at a particular level of representation: the level at which *anti-P* applies. Other approaches in the literature that follow a similar path include Wolf's (2008) Optimal Interleaving with Candidate Chains and Burzio's (2000) Sequential Faithfuless. In these approaches, the presence of an environment at some privileged level leads to blocking, but the relevant level is identified without using MSCs. In Wolf (2008), application to environments present before suffixation may be blocked in suffixed forms due to violation of a precedence constraint. In Burzio (2000), environments present at the UR of individual morphemes are subject to a faithfulness constraint. In this section I discuss a pattern of NDEB in which part of the environment of *P* is predictable. In particular, the application of *P* depends on the position of stress, but the distribution of stress is determined by the grammar. Given ROTB, underlying stress can be generated anywhere: output constraints enforce its correct output position. This leads to an over-generation problem for Wolf's and Burzio's approaches: if stress is not in its correct position at the relevant level of representation (the level subject to the blocking constraint), the environment for *P* is not met at that level, so the blocking constraint is avoided and *P* incorrectly applies. Examples of such blocking patterns are vowel raising in Romanian (Steriade 2008a) and vowel reduction in Armenian (Khanjian 2009). I will describe the Romanian case.

In Romanian, unstressed /a/ raises to [ə] in suffixed forms (224), but only if [a] is stressed in the unsuffixed form (225).

(224) Raising

a.	bárbə	'beard'	bərb-ós	'bearded-MASC'
b.	fáur	'artisan'	fəur-í	'to fashion'
c.	ispràvə	'brave deed'	isprǎv-nik	(nobleman's title)

(225) No raising

a.	mazíl	'deposed official'	mazil-í	'depose'
b.	kartóf	'potato'	kartof-jór	'potato-DIM'

Stress is predictable: it is penultimate by default, but falls on the final syllable on the surface in words that undergo final-/u/ deletion or have stress-attracting suffixes. There is independent distributional evidence for an underlying /u/ in words like [mazíl]: this /u/ surfaces before suffixes, i.e., in some environments where it is not word-final; singular nouns may only end in a surface [u] when this [u] follows an otherwise impermissible complex coda (as in the word [metru]), suggesting that deletion does not apply in these cases; and except for those singular nouns that end in a consonant and show a surface [u] elsewhere, singular nouns must end in a vowel. Final stress in consonant-final singular nouns makes sense if penultimate stress is assigned to the pre-deletion representation. Assuming this description to correctly reflect speakers' grammars, I will proceed to present an MSC-based analysis of Romanian. In section 3.4 I show why ROTB leads to a problem for Wolf's and Burzio's proposals.

3.3.3.2 An MSC-based account

The challenge posed by the blocking pattern of Romanian raising is that part of the conditioning environment is predictable: the vowel /a/ raises if it is unstressed, and stress is assigned by the grammar. For the MSC-based approach, an account of blocking would require the following ingredients. First, a variant of /a/ that is underspecified for the feature [low] would be referred to as /A/. Raising would be stated as in (226) and anti-raising as in (227).

(226) Raising: $A_{[\text{-stress}]} \rightarrow ə$

(227) Anti-raising: $A_{[\text{-stress}]} \rightarrow a$

If we follow the same recipe as in previous sections, the basic grammar would be (228), with a cover stress rule preceding raising in the phonology.

(228) Grammar for Romanian raising (to be revised below)

a. Morpheme structure component:

- $a \notin \Sigma_L$
- $A_{[\text{-stress}]} \rightarrow a$

b. Phonological rules:

- Stress
- $A_{[\text{-stress}]} \rightarrow ə$
- $A \rightarrow a$

The problem with (228) is that anti-raising must be able to protect underlying unstressed /a/'s in the unsuffixed form, but stress is only assigned later: anti-raising cannot make the necessary distinction between stressed and unstressed vowels and thus fails to capture the distinction in (229).

(229)	a.	bárbə	'beard'	bərb-ós	'bearded-MASC'
	b.	mazíl	'deposed official'	mazil-í	'depose'

The MSC-based approach allows for a remedy: since anti-raising is part of the grammar, it can interact in non-trivial ways with stress. In particular, stress can be assigned to the unsuffixed form before anti-raising. There are two ways to implement this solution. The first would be to relegate anti-raising to a cyclic phonology and apply stress and anti-raising, in this order, in the first cycle. In the second cycle, stress would apply again, followed by raising. This account, which is consistent with the proposed architecture, would assign a more important role than before to the mapping from URs to surface forms in accounting for blocking, and would leave the morpheme-structure component with the minor role of banning /a/ from URs. Another way to achieve the same result is to keep anti-raising as a morpheme-structure rule and minimally modify (228) so as to assign stress in the morpheme structure component before the application of anti-raising. At present, I am not aware of any good reason to choose between the two variants. For concreteness, I will use the second. Here is the final grammar, followed by derivations of the forms in (229) (for simplicity, I omit vowel-deletion rules from the grammar and drop stem-final vowels when convenient):

(230) Grammar for Romanian raising (final)

a. Morpheme structure component:

- a, $\acute{\sigma} \notin \Sigma_L$
- Stress
- $A_{[\text{-stress}]} \rightarrow$ a

b. Phonological rules:

- Stress
- $A_{[\text{-stress}]} \rightarrow$ ə
- A → a

(231) Derivation of [bárbə]

a. Morpheme structure rules apply:

1. {bArbə} → /bÁrbə/

b. Phonological processes apply:

UR	/bÁrbə/
Stress	-
$A_{[\text{-stress}]} \rightarrow$ ə	-
A → a	bárbə
SR	[bárbə]

(232) Derivation of [bərb-ós] (unsuffixed form: [bárbə])

a. Morpheme structure rules apply:

1. {bArbə} → /bÁrbə/
2. {os} → /ós/

b. Phonological processes apply:

UR	/bÁrbə-ós/
Stress	bArbós
$A_{[\text{-stress}]} \rightarrow$ ə	bərbós
A → a	-
SR	[bərbós]

(233) Derivation of [mazíl]

a. Morpheme structure rules apply:

1. {mAzilu} → /mazílu/

b. Phonological processes apply:

UR	/mazílu/
STRESS	-
$A_{[\text{-stress}]} \rightarrow ə$	-
A → a	-
SR	[mazíl]

(234) Derivation of [mazil-í] (unsuffixed form: [mazíl])

a. Morpheme structure rules apply:

1. {mAzilu} → /mazílu/

2. {i} → /í/

b. Phonological processes apply:

UR	/mazílu-í/
STRESS	mazilí
$A_{[\text{-stress}]} \rightarrow ə$	-
A → a	-
SR	[mazilí]

3.4 Previous approaches

In this section I use the case studies from the previous section to take a critical look at alternative approaches to NDEB. Competing approaches proposed in the literature include Mascaró (1976), Kiparsky (1993), Burzio (2000) Inkelas (2000), Łubowicz (2002), McCarthy (2003a) van Oostendorp (2007), Kula (2008), Wolf (2008), and Anttila (2009), among others. As the literature on NDEB is quite vast, I will not be able to do justice to all of the relevant approaches. Instead, I will discuss what I take to be a representative sample of the literature and refer the reader to critical reviews of approaches not discussed here directly. See, in particular, Kiparsky (1993) for a review of the literature prior to 1993, Inkelas (2000) for a review of the early OT literature (1993-2000), and Wolf (2008) for a more comprehensive critical review of the literature, including the literature following Inkelas (2000). As we will see, none of the previous approaches is able to account for all case studies at once.

3.4.1 Inkelas (2000): lexical typing and analogy

Above we have seen Kiparsky's 1993 proposal, in which the mapping from URs to SRs is the same as in the MSC-based account, but the morpheme structure component is absent:

(235) Grammar of Finnish' according to Kiparsky's (1993) theory

a. T → s / _ i

b. T → t

(236) Derivation of [timasi]

UR	/timaT-i/
T → s / _ i	timas-i
T → t	-
SR	[timasi]

Without MSCs, this grammar leaves the underlying distribution of underspecified /T/ and fully-specified /t/ as an accident of the Finnish' lexicon. This leads to over-generation, which can be avoided by adopting MSCs. In principle, however, it may be possible to combine Kiparsky's grammar with a mechanism other than MSCs to rule out undesirable URs like /rat/, where fully-specified /t/ (which blocks assibilation) occurs in a position to which assibilation should be able to apply. This is what Inkelas (2000) proposes. In particular, she proposes to extend underspecification with a position-based mechanism of lexical typing and analogy designed to rule out fully-specified /t/ from UR-final positions. The mechanism of lexical typing and analogy is left mostly unspecified, but I will show that any position-based mechanism would lead to incorrect predictions.

To see how URs like /rat/ would be ruled out under this proposal, consider the tree in (237), which is supposed to represent the internal organization of the lexicon. The assumption is that the lexicon keeps track of the identity of the final segment, including its probability of occurrence in the lexicon: in the Finnish' lexicon, every final voiceless coronal plosive is underspecified for continuancy. To determine the UR of a stem such as [rat], a mechanism of analogy scans the lexicon and finds that final voiceless coronal plosives are always underspecified. As a consequence, /rat/ is never selected as a UR (even though it can be represented in principle), and *[rati] is blocked.

(237)

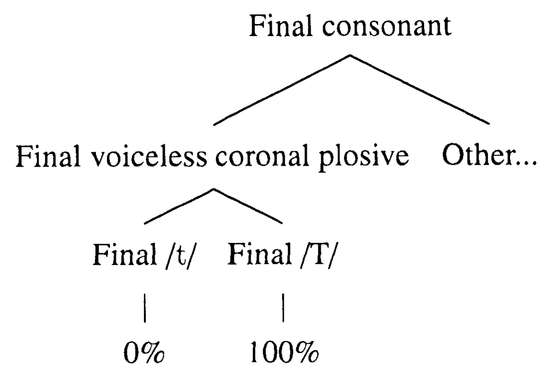

The position-based mechanism fails once we move from segments in final position to segments in penultimate position. Recall the hypothetical Finnish' process of vowel deletion that deletes stem-final vowels before another vowel. Vowel deletion may feed assibilation, which then targets stem-penultimate segments:

(238) /munte-i/ → [munsi]

Assibilation does not apply morpheme-internally to a /t/ that precedes an underlying /i/, and such /t/'s may occur in penultimate position:

(239) niti

The conclusion is that /t/'s in stem-penultimate position may be either fully specified (as in /tila/) or underspecified (as in /munTe/, the UR of the stem in [munsi]). Crucially, their specification depends on whether they precede an underlying /i/ (as enforced by anti-assibilation) and not on their position within the UR. While Finnish' is of course an artificial language, Romanian palatalization discussed above is an attested process with the same profile that poses a problem to Inkelas (2000).

3.4.2 Derived-environment theories

3.4.2.1 Strict cycle condition

Mascaró (1976) argues that NDEB provides evidence for a phonological analog of Chomsky's Strict Cycle Condition (SCC; Chomsky, 1973). The phonological version is given in (240).[9]

(240) **Strict Cycle Condition.** For a cyclic rule *R* to apply properly in any given cycle *j*, it must make specific use of information proper to (i.e., introduced by virtue of) cycle *j*.
This situation obtains if either of the following conditions is met:

1. The rule makes crucial reference to the information in the representation that spans the boundary between the current cycle and the preceding one.

2. The rule applies solely within the domain of the previous cycle but crucially refers to information supplied by a rule operating on the current cycle.

Application of a cyclic rule is licensed in morphologically-derived environments by the first condition and in phonologically-derived environments by the second condition. The following table illustrates the analysis of Finnish' assibilation using the SCC. Final-vowel raising and assibilation are both assumed to be cyclic rules. Cyclic rules cannot apply in the first cycle by stipulation. A word boundary is inserted in the final cycle. The leftmost column demonstrates application in a morphologically-derived environment, the middle column application in a phonologically-derived environment, and the rightmost column blocking in a nonderived environment. The number of the SCC condition that licenses each rule application is given in brackets next to the outcome of the rule.

(241) Finnish' assibilation using the SCC

[9] The presentation of the SCC in this section is based on Kenstowicz (1994).

First cycle	[lut]	[vete]	[tila]
e → i / _ #	-	-	-
t → s / _ i	-	-	blocked
Second cycle	[#[lut]i#]	[#[vete]#]	[#[tila]#]
e → i / _ #	-	veti (1)	-
t → s / _ i	lusi (1)	vesi (2)	blocked
	[lusi]	[vesi]	[tila]

Condition (1) of the SCC dictates that spanning a morpheme boundary is a sufficient property for licensing. This wrongly predicts obligatory application of Romanian palatalization after vowel deletion:

(242) Prediction for Romanian [pʌduk-j] (incorrect):
/pʌduke-i/ → *[pʌdutʃ-j]

3.4.2.2 Coloured Containment

van Oostendorp (2007) proposes an account of NDEB that makes use of a mechanism of morpheme indexing called "colouring". The assumption is that every morpheme is annotated with its own "color" – a morpheme-specific index which is distributed over all segments and other material (features, moras, etc) which make up the morpheme. For example, in the representation of Finnish' /timat-i/, the first morpheme would be associated with the color α and the second morpheme with the color β, as shown in (243) using a simplified linear representation.

(243) /$t_\alpha i_\alpha l_\alpha a_\alpha t_\alpha$-$i_\beta$/

Blocking in nonderived environments arises from a proposed constraint against monochromatic feature spreading, which I have simplified using the following statement (see the original chapter for more details about the mechanics of colouring and spreading):

(244) Do not associate a feature and a segment of the same colour.

Finnish' assibilation would presumably involve spreading of the feature [continuant] from /i/ to /t/, but only if /i/ and /t/ are not of the same color:

(245) a. $t_\alpha i_\alpha \rightarrow t_\alpha i_\alpha$

b. $t_\alpha i_\beta \rightarrow s_\alpha i_\beta$

This account makes the right prediction that /timat-i/ should become [timas-i], but it fails for the vowel deletion case in Romanian, as demonstrated in (246): spreading is incorrectly licensed across a morpheme boundary.

(246) Prediction for Romanian [pʌduk-j] (incorrect):

/$p_\alpha ʌ_\alpha d_\alpha u_\alpha k_\alpha e_\alpha$-$i_\beta$/ → $p_\alpha ʌ_\alpha d_\alpha u_\alpha k_\alpha$-$i_\beta$ → *[pʌdutʃ-j]

3.4.3 Wolf (2008): Optimal Interleaving with Candidate Chains

Wolf's (2008) architecture is a cyclic implementation of Optimality Theory with Candidate Chains, a serial variant of OT (OT-CC; McCarthy, 2007a). Wolf's account of NDEB is guided by the following intuition: P is blocked in some environment if it can apply in this environment before the application of some other process P_0. For morphologically-derived environments, P_0 is set as affixation; for phonologically-derived environments, P_0 is set as the relevant phonological process that precedes P. Blocking is enforced by a precedence constraint that requires P_0 to crucially precede P.

Let us see how this account correctly derives Finnish' [timasi] from /timat-i/, where there are two potential environments for application. The first step is a precedence constraint that requires affixation to crucially precede assibilation. Informally, the first sequence [ti] (/timat-i/) is present before affixation: assibilation can apply to this sequence before or after affixation, so it is not crucially preceded by affixation, in violation of the precedence constraint. Application to the second sequence /tilat-i/ is not blocked since the process can only apply after affixation. More formally, the derivation starts with an abstract morphosyntactic structure (/ROOT-AF/) and morpheme exponents are inserted in the phonology in violation of the faithfulness constraints INSERT-ROOT and INSERT-AFFIX. Here are the constraints relevant for NDEB:

(247) a. *ti

b. IDENT[cont]

c. PREC(INSERT-AFFIX,IDENT[cont]): assign a violation mark for each time that:

- A process that violates IDENT[cont] applies without having been preceded by a process that violates INSERT-AFFIX
- A process that violates IDENT[cont] applies and is followed by a process that violates INSERT-AFFIX

The markedness constraints *ti triggers assibilation, but only when the higher ranked precedence constraint is satisfied:

(248) PREC(INSERT-AFFIX,IDENT[cont]) $\gg$ *ti $\gg$ IDENT[cont]

The tableau in (252) demonstrates the derivation of [timasi]. A candidate consists of a chain in which each member differs from the preceding member by one atomic change, like a feature change, epenthesis, deletion, affix insertion, and so on. The final member of the chain is the output. Candidate (a) includes the chain that outputs [timasi]: first the root is inserted, then the affix, then assibilation applies. Whenever multiple distinct chains lead to the same output, they are merged into a single candidate, as in (b), which represents the output candidate [simasi]. Precedence is evaluated based on this merged candidate. In the first chain in (b), assibilation is applied to the first /ti/ sequence after suffixation, but in the second chain it applies before suffixation. This means that this application of assibilation is not crucially preceded by suffixation, incurring a violation of PREC. To ensure that multiple applications of assibilation are distinguished from one another, precedence is not directly evaluated on the candidates themselves, but rather on tuples of faithfulness violations that the candidates induce, called LUMSeqs (250), and each violation is indexed with respect to the position in the word which is the source of the violation. In both LUMSeqs for (b), the violation ID[cont]$_{@5}$ (which corresponds to the application of assibilation to the fifth segment of the word) follows INSERT-AF (which corresponds to affixation), which means that this instance of assibilation is crucially preceded by affixation and so does not incur a PREC violation. Since candidate (b) violates the highest ranked PREC constraint and candidate (a) does not, candidate (a) is the winner.

(249) Tableau for [timasi]

		/ROOT-AF/	PREC(INSERT-AF,ID[cont])	*ti	ID[cont]
a.	☞	<timat-AF, timati, timasi>		*	*
b.		<timat-AF, timati, simati, simasi> <timat-AF, simat-AF, simati, simasi>	*!		**

(250) a. LUMSeq: <INSERT-AF, ID[cont]$_{@5}$>

b. LUMSeq: <INSERT-AF, ID[cont]$_{@1}$, ID[cont]$_{@5}$>

LUMSeq: < ID[cont]$_{@1}$, INSERT-AF, ID[cont]$_{@5}$>

Here is how an analysis of Romanian raising would work in this architecture. The ranking, given in (251), is of the following constraints: a cover constraint STRESS stands for whatever constraints enforce correct surface stress in Romanian; the constraint PREC(INSERT-AFFIX,IDENT[low]) requires that raising is crucially preceded by affixation; the markedness constraint *a$_{[\text{-stress}]}$ is responsible of triggering raising, in violation of the faithfulness constraint IDENT[low].

(251) STRESS , PREC(INSERT-AFFIX,IDENT[low]) ≫ *a$_{[\text{-stress}]}$ ≫ IDENT[low]

The tableau in (252) demonstrates the correct derivation of [bərb-ós] assuming the UR /bárbə/ for the root (for simplicity, the tableau ignores the deletion of stem-final /ə/). Notice that, crucially, stress is underlyingly penultimate: /a/ is stressed from the outset, so raising cannot apply before suffixation and there is no PREC violation. Hence, raising is (correctly) not blocked.

(252) Correct derivation of [bərb-ós] (assuming the UR /bárbə/)

		/ROOT-AF/	STRESS	PREC(AF,IDENT[low])	*a$_{[\text{-stress}]}$	ID[low]
a.	☞	<bárbə-AF, bárbə-ós, barbə-ós, bərbə-ós>				*
b.		<bárbə-AF, bárbə-ós, barbə-ós>			*!	

Given ROTB, the predictability of stress allows for URs in which stress is placed in arbitrary positions and output constraints enforce surface penultimate stress (253). The problem is that for such URs, the environment for raising is met before or after suffixation, so raising is not crucially preceded by suffixation and is incorrectly blocked (254a).

(253) a. /barbə́/ → [bárbə]

b. /barbə/ → [bárbə]

(254) a. /barbə́-ós/ → *[barb-ós]

b. /barbə-ós/ → *[barb-ós]

The tableau in (255) is a concrete tableau for (254a), where stress is underlyingly final. The conclusion is that given ROTB, the grammar over-generates pairs of nonderived-derived forms where raising is incorrectly blocked in the derived form.

(255) Incorrect derivation of *[barb-ós] (assuming the UR /barbə́/):

	/ROOT-AF/	STRESS	PREC(AF,IDENT[low])	$*a_{[-stress]}$	ID[low]
a.	<barbə́-AF, bárbə-ós, barbə-ós, bərbə-ós> <barbə́-AF, bərbə́-AF, bərbə́-ós, bərbə-ós>		*!		*
b. ✕	<barbə́-AF, barbə́-ós, barbə-ós>			*	

3.4.4 Burzio (2000): Sequential faithfulness

The intuition behind Burzio's 2000 approach to NDEB is that sequences that are present in the UR of a single morpheme have a privileged status in terms of faithfulness compared to other sequences. He proposes a new type of faithfulness constraints to account for NDEB: as opposed to traditional faithfulness constraints which typically protect individual features, Burzio's constraints penalize modifications of sequences or combinations of features. I will refer to these constraints as Sequential Faithfulness constraints. An example of a Sequential Faithfulness constraint is FAITH[ti], which penalizes any output deviation from the input sequence /ti/. Burzio assumes that such constraints do not protect sequences that are separated by morpheme boundaries, presumably because morphemes are not concatenated in the input. This assumption creates a distinction between the two /ti/ sequences in the Finnish' /timat-i/: modifying the first sequence (for instance, by applying assibilation) would incur a violation of FAITH[ti], but modifying the second sequence will not. The following tableau shows how Sequential Faithfulness successfully accounts for the derivation /timat-i/ → [timasi]:

(256) Tableau for [timasi]

	/timat-i/	FAITH[ti]	*ti	ID[cont]
a.	timat-i		**!	
b.	silat-i	*!	*	*
c.	silas-i	*!		**
d. ☞	timas-i		*	*

Burzio's theory shares much of its predictions with the MSC-based theory – if an environment for *P* is present in the UR of some morpheme, application in this environment will be avoided – but the connection between presence in a UR and blocking is made without MSCs. Blocking in nonderived environments that are partially predictable poses an over-generation problem for Sequential Faithfulness since ROTB allows (given a single morpheme) both for URs in which the environment is present (and is therefore protected by a Sequential Faithfulness constraint) and URs in which it is not (and is therefore not protected). *P* will incorrectly apply to the latter. The following simplified tableau demonstrates the (correct) derivation of Romanian [bərb-ós], assuming the UR /bárbə/ for the stem: FAITH[$a_{[\text{-stress}]}$] is not violated since [a] is stressed in the input.

(257) Tableau for [bərb-ós], UR: /bárbə-ós/

	/bárbə-ós/	FAITH[$a_{[\text{-stress}]}$]	*$a_{[\text{-stress}]}$	ID[low]
a.	barb-ós		*!	
b. ☞	bərb-ós			*

Given ROTB, other possible URs for the stem are /barbə/ and /barbə́/. Here vowel raising would incur a violation of FAITH[$a_{[\text{-stress}]}$] since [a] is unstressed in the input. The result is that raising is incorrectly blocked in the derived form:

(258) Tableau for [bərb-ós], UR: /barbə/

	/barbə-ós/	FAITH[$a_{[\text{-stress}]}$]	*$a_{[\text{-stress}]}$	ID[low]
a. ✕	barb-ós		*	
b.	bərb-ós	*!		*

3.5 Conclusion

OT dispensed with MSCs for reasons of theoretical simplicity: a single-component architecture seemed more appealing than a dual-component one, and output constraints unified MSCs and the input-output mapping. In this chapter, I examined the implications of MSCs for the phenomenon of NDEB and claimed that NDEB can be characterized as an opaque interaction between MSCs and the input-output mapping. I discussed a few cases of NDEB that receive a simple account using MSCs but cannot all be accounted for using any previous theory of NDEB. This provides support for a dual-component architecture of phonology with MSCs over architectures that adopt the principle of ROTB.

Bibliography

Alderete, John, and Bruce Tesar. 2002. Learning covert phonological interaction: an analysis of the problem posed by the interaction of stress and epenthesis. Technical Report RuCCS-TR-72, Rutgers Center for Cognitive Science, Piscataway, NJ. ROA 543.

Allen, William Sidney. 1973. *Accent and rhythm*. Cambridge University Press.

Anderson, Stephen R. 1981. Why phonology isn't "natural". *Linguistic Inquiry* 12:493–539.

Anderson, Stephen R. 1982. The analysis of French schwa: or, how to get something from nothing. *Language* 58.

Anttila, Arto. 2006. Variation and opacity. *Natural Language and Linguistic Theory* 24:893–944.

Anttila, Arto. 2009. Derived environment effects in colloquial Helsinki Finnish. In *The nature of the word: Essays in honor of paul kiparsky*, ed. Kristin Hanson and Sharon Inkelas, 433–460. Cambridge: MIT Press.

Archangeli, Diana. 1988. Aspects of underspecification theory. *Phonology* 5:183–207.

Bach, Emmon, and Robert T Harms. 1972. How do languages get crazy rules. *Linguistic change and generative theory* 1:21.

Baker, Brett. 2008. *Word structure in Ngalakgan*. Stanford: CSLI Publications.

Baker, C. L. 1979. Syntactic theory and the projection problem. *Linguistic Inquiry* 10:533–581.

Berwick, Robert C. 1982. Locality principles and the acquisition of syntactic knowledge. Doctoral Dissertation, MIT, Cambridge, MA.

Blumenfeld, Lev A. 2006. Constraints on phonological interactions. Doctoral Dissertation, Stanford University.

Booij, Geert. 1995. *The phonology of Dutch*. Oxford: Clarendon Press.

Borowsky, T. 1986. Topics in the lexical phonology of English. Doctoral Dissertation, University of Massachusetts, Amherst.

Bowers, Dustin. 2016. Acoustic correlates of Gujarati stress: A failure to replicate. University of Alberta ms.

Brent, Michael. 1999. An efficient, probabilistically sound algorithm for segmentation and word discovery. *Computational Linguistics* 34:71–105.

Burzio, Luigi. 2000. Cycles, non-derived-environment-blocking, and correspondence. In *Optimality theory: Phonology, syntax, and acquisition*, ed. Joost Dekkers, Frank van der Leeuw, and Jeroen van de Weijer. Cambridge: Cambridge University Press.

Canalis, Stefano. 2007. To what extent is Albanian word stress predictable? In *Proceedings of the XXXIII Incontro di grammatica generativa*, ed. Antonietta Bisetto and Francesco E. Barbieri, 1–14. Bologna: Università di Bologna.

Chomsky, Noam. 1973. Conditions on transformations. In *A festschrift for Morris Halle*, ed. Stephen Anderson and Paul Kiparsky. Holt, Rinehart and Winston, Inc.

Chomsky, Noam, and Morris Halle. 1965. Some controversial questions in phonological theory. *Journal of Linguistics* 1:97–138.

Chomsky, Noam, and Morris Halle. 1968. *The sound pattern of English*. New York: Harper and Row Publishers.

Clark, Alexander. 2001. Unsupervised language acquisition: Theory and practice. Doctoral Dissertation, University of Sussex.

Clements, G. N., and Samuel Jay Keyser. 1983. *CV phonology: A generative theory of the syllable.*. Cambridge MA: MIT Press.

Crowhurst, Megan J., and Lev D. Michael. 2005. Iterative footing and prominence-driven stress in Nanti (Kampa). *Language* 81:47–95.

Davidson, Lisa, Peter Jusczyk, and Paul Smolensky. 2004. The initial and final states: theoretical implications and experimental explorations of richness of the base. In *Constraints in phonological acquisition*, ed. René Kager, Joe Pater, and Wim Zonneveld, chapter 10, 321–368. Cambridge University Press.

Davies, H.J. 1980. *Kobon phonology*. Australian National University, Canberra.

Davies, H.J. 1981. *Kobon*, volume 3 of *Lingua Descriptive Series*. Amsterdam, North Holland.

Davis, Stuart. 2011. Quantity. In *The handbook of phonological theory*, ed. John Goldsmith, Jason Riggle, and Alan Yu, 103–140. Oxford: Blackwell, second edition edition.

Davis, Stuart, and Mi-Hui Cho. 2003. The distribution of aspirated stops and /h/ in American English and Korean: an alignment approach with typological implications. *Linguistics* 41:607–652.

Dell, François. 1981. On the learnability of optional phonological rules. *Linguistic Inquiry* 12:31–37.

Ellison, Timothy Mark. 1994. The machine learning of phonological structure. Doctoral Dissertation, University of Western Australia.

Everett, Daniel, and Keren Everett. 1984. On the relevance of syllable onsets to stress placement. *Linguistic Inquiry* 15:705–711.

Everett, Daniel L. 1988. On metrical constituent structure in Pirahã phonology. *Natural Language and Linguistic Theory* 6:297–246.

Evers, Vincent, Henning Reetz, and Aditi Lahiri. 1998. Crosslinguistic acoustic categorization of sibilants independent of phonological status. *Journal of Phonetics* 345–370.

Ferguson, Charles A, and Munier Chowdhury. 1960. The phonemes of Bengali. *Language* 36:22–59.

Foley, William A. 1991. *The Yimas language of New Guinea.* Stanford University Press.

Foreman, Velma, and Helen Marten. 1973. Yessan-Mayo phonemes. *Workpapers in Papua New Guinea Languages* 2:79–108.

Gabas, Nilson. 1999. A grammar of Karo (Tupi). Doctoral Dissertation, University of California, Santa Barbara.

Giavazzi, Maria. 2010. The phonetics of metrical prominence and its consequences for segmental phonology. Doctoral Dissertation, Massachusetts Institute of Technology.

Goldsmith, John. 2001. Unsupervised learning of the morphology of a natural language. *Computational Linguistics* 27:153–198.

Gonzalez, Carolina. 2003. The effect of stress and foot structure on consonantal processes. Doctoral Dissertation, USC.

Gordon, Matthew. 1999. Syllable weight: phonetics, phonology, and typology. Doctoral Dissertation, UCLA.

Gordon, Matthew. 2006. *Syllable weight: phonetics, phonology, typology.* New York: Routledge.

Grünwald, Peter. 1996. A minimum description length approach to grammar inference. In *Connectionist, statistical and symbolic approaches to learning for natural language processing*, ed. G. S. S. Wermter and E. Riloff, Springer Lecture Notes in Artificial Intelligence, 203–216. Springer.

Halle, Morris. 1959. *The sound pattern of Russian.* Walter de Gruyter.

Halle, Morris. 1962. Phonology in generative grammar. *Word* 18:54–72.

Halle, Morris. 1978. Knowledge unlearned and untaught: What speakers know about the sounds of their language. In *Linguistic theory and psychological reality*, ed. Morris Halle, Joan Bresnan, and George A. Miller, 294–303. MIT Press.

Halle, Morris. 1990. An approach to morphology. In *Proceedings of NELS*, volume 20, 150–184.

Halle, Morris. 1998. The stress of English words 1968-1998. *Linguistic Inquiry* 29:539–568.

Halle, Morris, and Alec Marantz. 1993. Distributed morphology and the pieces of inflection. In *The view from Building 20*, ed. Kenneth Hale and Jay Keyser, 111–176. Cambridge, MA: MIT Press.

Halle, Morris, and Jean-Roger Vergnaud. 1987a. *An essay on stress*. MIT Press, Cambridge MA.

Halle, Morris, and Jean-Roger Vergnaud. 1987b. Stress and the cycle. *Linguistic Inquiry* 18:45–84.

Hammond, Michael. 1992. Deriving the strict cycle condition. *CLS* 28:126–140.

Hargus, Sharon. 2001. Quality-sensitive stress reconsidered. In *University of Washington Working Papers in Linguistics*, volume 20, 25–56.

Hausenberg, Anu-Reet. 1998. Komi. *Abondolo (1998b)* 305–326.

Hayes, Bruce. 1995. *Metrical stress theory: principles and case studies*. Chicago, The University of Chicago Press.

Hayes, Bruce. 2004. Phonological acquisition in Optimality Theory: The early stages. In *Constraints in phonological acquisition*, ed. René Kager, Joe Pater, and Wim Zonneveld, 158–203. Cambridge, UK: Cambridge University Press.

Hayes, Bruce, and Colin Wilson. 2008. A maximum entropy model of phonotactics and phonotactic learning. *Linguistic Inquiry* 39:379–440.

Horning, James. 1969. A study of grammatical inference. Doctoral Dissertation, Stanford.

Hulst, H.G. van der, and Nancy Ritter. 1999. Theories of the syllable. In *The syllable: views and facts*, ed. H.G. van der Hulst and N. Ritter, 13–52. Berlin: Mouton de Gruyter.

Hyman, Larry M. 1970. How concrete is phonology? *Language* 46:58–76.

Idsardi, William. 2006. A Bayesian approach to loanword adaptations. Poster presented at the annual meeting of the Linguistic Society of America, Albuquerque, NM, January 2006.

Idsardi, William James. 1992. The computation of prosody. Doctoral Dissertation, Massachusetts Institute of Technology.

Inkelas, Sharon. 1995. The consequences of optimization for underspecification. In *Proceedings of the North East Linguistic Society 25*, ed. Jill Beckman, 287–302. University of Pennsylvania: Graduate Linguistic Student Association.

Inkelas, Sharon. 2000. Phonotactic blocking through structural immunity. In *Lexicon in focus*, ed. Barbara Stiebels and Dieter Wunderlich, Studia Grammatica, 7–40. Akademie Verlag.

Kager, René. 1990. Dutch schwa in moraic phonology. In *Papers from the 26th Regional Meeting of the Chicago Linguistic Society*, ed. M. Noske M. Ziolkowski and K. Deaton, volume II: The parasesssion on the syllable in phonetics and phonology, 189–214.

Kager, René, and Joe Pater. 2012. Phonotactics as phonology: Knowledge of a complex restriction in dutch. *Phonology* 29:81–111.

Kahn, Daniel. 1976. Syllable-based generalizations in English phonology. Doctoral Dissertation, MIT.

Kangasmaa-Minn, Eeva. 1998. Mari. In *The uralic languages*, ed. Daniel Abondolo, 219–248. Routledge.

Katzir, Roni. 2014. A cognitively plausible model for grammar induction. *Journal of Language Modelling* 2:213–248.

Kenstowicz, Michael. 1994. *Phonology in Generative Grammar*. Blackwell.

Kenstowicz, Michael. 1997. Quality-sensitive stress. *Rivista di Linguistica* 9:157–188.

Keyser, Samuel Jay, and Paul Kiparsky. 1984. Syllable structure in Finnish phonology. In *Language sound structure*, ed. Mark Aronoff and Richard T. Oehrle, 7–31. Cambridge, Mass: MIT Press.

Khanjian, Hrayr. 2009. Stress dependent vowel reduction. In *Annual Meeting of the Berkeley Linguistics Society*, volume 35, 178–189.

Kiparsky, Paul. 1968. How abstract is phonology? Indiana University Linguistics Club.

Kiparsky, Paul. 1982. Lexical phonology and morphology. In *Linguistics in the Morning Calm*. Linguistics Society of Korea, Seoul, Korea: Hanshin Publishing Company.

Kiparsky, Paul. 1993. Blocking in nonderived environments. In *Phonetics and phonology 4: Studies in lexical phonology*, ed. Ellen M. Kaisse and Sharon Hargus, 277–313. San Diego: Academic Press.

Kiparsky, Paul. 2000. Opacity and cyclicity. *The Linguistic Review* 17:351–367.

Kiparsky, Paul. 2003. Finnish noun inflection. In *Generative approaches to finnic linguistics*. CSLI.

Körding, Konrad P., and Daniel M. Wolpert. 2006. Bayesian decision theory in sensorimotor control. *Trends in Cognitive Sciences* 10:319–326.

Krämer, Martin. 2012. *Underlying representations*. Cambridge University Press.

Kula, Nancy C. 2008. Derived environment effects: a representational approach. *Lingua* 118:1328–1343.

de Lacy, Paul. 2004. Markedness conflation in optimality theory. *Phonology* 21:145–199.

de Lacy, Paul. 2007. The interaction of tone, sonority, and prosodic structure. *The Cambridge handbook of phonology* 281–307.

de Lacy, Paul. 2013. What counts as evidence for universals? Questioning sonority-driven stress. Talk given at Language Universals Workshop, Harvard University, November 7, 2013.

de Lacy, Paul V. 2002. The formal expression of markedness. Doctoral Dissertation, University of Massachusetts Amherst.

Lahiri, Aditi. 2001. Metrical patterns. *Language typology and language universals* 1347–1367.

Levin, Juliette. 1985. A metrical theory of syllabicity. Doctoral Dissertation, Massachusetts Institute of Technology.

Liberman, Mark, and Alan Prince. 1977. On stress and linguistic rhythm. *Linguistic inquiry* 8:249–336.

Lowenstamm, Jean. 1996. CV as the only syllable type. *Current trends in Phonology. Models and Methods, edited by Jacques Durand & Bernard Laks* 419–441.

Łubowicz, Anna. 2002. Derived environment effects in optimality theory. *Lingua* 112:243–280.

de Marcken, Carl. 1996. Unsupervised language acquisition. Doctoral Dissertation, MIT, Cambridge, MA.

Mascaró, Joan. 1976. Catalan phonology and the phonological cycle. Doctoral Dissertation, MIT.

McCarthy, John. 1979a. Formal problems in semitic morphology and phonology. Doctoral Dissertation, MIT.

McCarthy, John. 2003a. Comparative markedness. *Theoretical Linguistics* 29:1–51.

McCarthy, John. 2007a. *Hidden generalizations: Phonological opacity in optimality theory*. London: Equinox.

McCarthy, John J. 1979b. On stress and syllabification. *Linguistic Inquiry* 10:443–465.

McCarthy, John J. 2003b. Sympathy, cumulativity, and the Duke-of-York gambit. In *The syllable in optimality theory*, ed. Caroline Fery and Ruben van de Vijver, 23–76. Cambridge: Cambridge University Press.

McCarthy, John J. 2005. Taking a free ride in morphophonemic learning. *Catalan Journal of Linguistics* 4:19–56.

McCarthy, John J. 2007b. Derivations and levels of representation. In *The cambridge handbook of phonology*, ed. Paul de Lacy, 99–117. Cambridge University Press.

Mester, R. Armin. 1994. The quantitative trochee in Latin. *Natural Language and Linguistic Theory* 12:1–62.

Moore-Cantwell, Claire. 2016. The representation of probabilistic phonological patterns: Neurological, behavioral, and computational evidence from the English stress system. Doctoral Dissertation, University of Massachusetts Amherst.

Nevins, Andrew, and Bert Vaux. 2007. Underlying representations that do not minimize grammatical violations. In *Freedom of analysis?*, ed. Sylvia Blaho, Patrik Bye, and Martin Krämer, 35–61. Mouton de Gruyter.

Noyer, Rolf. 2013. A generative phonology of San Mateo Huave. *International Journal of American Linguistics* 79:1–60.

Odden, David. 1990. Phonology and its interaction with syntax and morphology. *Studies in the Linguistic Sciences* 20.2:69–108.

van Oostendorp, M. 2007. Derived environment effects and consistency of exponence. In *Freedom of analysis*, ed. Sylvia Blaho, Patrik Bye, and Martin Krämer. Berlin: Mouton de Gruyter.

van Oostendorp, Marc. 1997. Vowel quality and phonological projection. Doctoral Dissertation, Katolieke Universiteit Brabant.

Orbán, Gergő, József Fiser, Richard N Aslin, and Máté Lengyel. 2008. Bayesian learning of visual chunks by human observers. *Proceedings of the National Academy of Sciences* 105:2745–2750.

Pater, Joe. 2014. Canadian raising with language-specific weighted constraints. *Language* 90:230–240.

Payne, David, and Furne Rich. 1988. Sensitivity to onset in Arabela stress. *Ms, Instituto Lingüıstico de Verano, Pucallpa, Peru* .

Payne, Judith. 1990. Asheninca stress patterns. In *Amazonian linguistics*, ed. Doris Payne, 185–212. Austin: University of Texas Press.

Pesetsky, David. 1979. Russian morphology and lexical theory. Ms., MIT.

Piantadosi, Steven T., Joshua B. Tenenbaum, and Noah D. Goodman. 2016. The logical primitives of thought: Empirical foundations for compositional cognitive models. *Psychological review* 123:392–424.

Prince, Alan. 1983. Relating to the grid. *Linguistic Inquiry* 14:19–100.

Prince, Alan, and Paul Smolensky. 1993. Optimality theory: Constraint interaction in generative grammar. Technical report, Rutgers University, Center for Cognitive Science.

Prince, Alan, and Bruce Tesar. 2004. Learning phonotactic distributions. In *Constraints in phonological acquisition*, ed. René Kager, Joe Pater, and Wim Zonneveld, 245–291. Cambridge University Press.

Rasin, Ezer, Iddo Berger, Nur Lan, and Roni Katzir. 2017. Learning rule-based morphophonology. Ms., MIT and TAU.

Rasin, Ezer, Iddo Berger, Nur Lan, and Roni Katzir. To appear. Rule-based learning of phonological optionality and opacity. In *Proceedings of NELS 48*.

Rasin, Ezer, and Roni Katzir. 2015. Compression-based learning for OT is incompatible with Richness of the Base. In *Proceedings of NELS 45*, volume 2, 267–274.

Rasin, Ezer, and Roni Katzir. 2016. On evaluation metrics in optimality theory. *Linguistic inquiry* 47.

Rasin, Ezer, and Roni Katzir. To appear. Learning abstract URs from distributional evidence. In *Proceedings of NELS 48*.

Rissanen, Jorma. 1978. Modeling by shortest data description. *Automatica* 14:465–471.

Rissanen, Jorma, and Eric Sven Ristad. 1994. Language acquisition in the MDL framework. In *Language computations: DIMACS Workshop on Human Language, March 20-22, 1992*, 149. Amer Mathematical Society.

Rubach, Jerzy. 1986. Abstract vowels in three dimensional phonology: the yers. *The Linguistic Review* 5:247–280.

Scheer, Tobias. 2004. *A lateral theory of phonology: What is cvcv, and why should it be?*, volume 1. Walter de Gruyter.

Scheer, Tobias. 2016. Melody-free syntax and phonologically conditioned allomorphy. *Morphology* 26:341–378.

Shih, Shu-hao. 2016. Sonority-driven stress does not exist. In *Proceedings of the Annual Meetings on Phonology*, volume 3.

Smolensky, Paul. 1996. The initial state and 'richness of the base' in Optimality Theory. Technical Report JHU-CogSci-96-4, Johns Hopkins University.

Sobel, D.M., J.B. Tenenbaum, and A. Gopnik. 2004. Children's causal inferences from indirect evidence: Backwards blocking and bayesian reasoning in preschoolers. *Cognitive science* 28:303–333.

Solomonoff, Ray J. 1964. A formal theory of inductive inference, parts I and II. *Information and Control* 7:1–22, 224–254.

Stanley, Richard. 1967. Redundancy rules in phonology. *Language* 43:393–436.

Steriade, Donca. 1995. Underspecification and markedness. In *The handbook of phonological theory*, ed. John Goldsmith, 114–174. Oxford: Blackwell.

Steriade, Donca. 2008a. A pseudo-cyclic effect in Romanian morphophonology. *Inflectional identity* 18:313–360.

Steriade, Donca. 2008b. Unpublished lecture notes, MIT.

Stolcke, Andreas. 1994. Bayesian learning of probabilistic language models. Doctoral Dissertation, University of California at Berkeley, Berkeley, California.

Szpyra, Jolanta. 1992. Ghost segments in nonlinear phonology: Polish yers. *Language* 277–312.

Tang, Katrina Elizabeth. 2008. The phonology and phonetics of consonant-tone interaction. Doctoral Dissertation, UCLA.

Tauberer, Joshua. 2009. Goldilocks meets the subset problem: Evaluating error driven constraint demotion (RIP/CD) for OT language acquisition. In *Proceedings of the 32nd Annual Penn Linguistics Colloquium*, volume 15 of *University of Pennsylvania Working Papers in Linguistics*, 25.

Tesar, Bruce. 2014. *Output-driven phonology: Theory and learning*. Cambridge University Press.

Tesar, Bruce, and Paul Smolensky. 1998. Learnability in Optimality Theory. *Linguistic Inquiry* 29:229–268.

Topintzi, Nina. 2005. On sonority, onsets and stress in Arabela. In *Proceedings of the University of Cambridge 3rd Postgraduate Conference in Language Research 2005*, ed. Gerasimos Chatzidamianos et al., 279–286.

Trommer, Jochen. 2013. Stress uniformity in Albanian: Morphological arguments for cyclicity. *Linguistic inquiry* 44:109–143.

Vaysman, Olga. 2008. Segmental alternations and metrical theory. Doctoral Dissertation, Massachusetts Institute of Technology.

Watson, Janet CE. 2002. *The phonology and morphology of arabic*. Oxford: Oxford University Press.

Wier, Thomas R. 2004. Comparative markedness and opacity in Meskwaki palatalization. In *Proceedings of the Thirtieth Annual Meeting of the Berkeley Linguistics Society*, ed. Marc Ettlinger, Nicholas Fleisher, and Mischa Park-Doob, 426–434. Berkeley: BLS.

Wolf, Matthew. 2008. Optimal interleaving: serial phonology-morphology interaction in a constraint-based model. Doctoral Dissertation, University of Massachusetts, Amherst.

Xu, F., and J.B. Tenenbaum. 2007. Word learning as bayesian inference. *Psychological review* 114:245.

Zec, Draga. 1995. Sonority constraints on syllable structure. *Phonology* 85–129.

Zoll, Cheryl. 1996. Parsing below the segment in a constraint based framework. Doctoral Dissertation, University of California at Berkeley.