UNIVERSITY OF CALIFORNIA SAN DIEGO

Conspiratorial Exceptionality: A Case Study of Mushunguli

A dissertation submitted in partial satisfaction of the
requirements for the degree
Doctor of Philosophy

in

Linguistics

by

Katherine Hout

Committee in charge:

Professor Eric Baković, Chair
Professor Gabriela Caballero
Professor Sarah Creel
Professor Marc Garellek
Professor Sharon Rose

2020

The dissertation of Katherine Hout is approved, and it is acceptable in quality and form for publication on microfilm and electronically:

Chair

University of California San Diego

2020

DEDICATION

For Aunt Linda

EPIGRAPH

Yet exceptions do exist, so how do they arise? It seems to me that they arise to the extent that we, the grammarians, have got it wrong. We introduce them from outside with rules that are not quite right. If a rule is 100% correct it will have no (unexplained) exceptions whatsoever, if it is almost right it will have smaller number of exceptions, and if it is badly wrong it will have lots of exceptions. (Beedham 2005:153)

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

LIST OF ABBREVIATIONS

ACC: accusative

ANTIC: anticipatory

Appl: applicative

AUG: augment

Cl 1,2, etc.: Noun class 1, 2, etc.

CMPV: comparative

Dem: demonstrative

ELV: elative

FUT: future

Obj: object

OT: Optimality Theory

PB: Proto-Bantu

PFV: perfective

PL: plural

prox: proximal

PST: past

QUOT: quotative

SG: singular

SGF: singular feminine

SGM: singular masculine

SUBJ: subject

TAM: Tense/Aspect/Mood

TNP: total non-participant(s)

VCL: verb theme closure

ACKNOWLEDGEMENTS

I am surprised to find that writing these acknowledgements is nearly as awkward and stressful as writing the dissertation itself. This is not due to a lack of gratitude, but rather an excess of it. I had a non-traditional path to graduate school—I nearly flunked out of both high school and community college, and in my twenties I was besieged by a seemingly never-ending onslaught of personal tragedies and health issues. To be sitting here, misty-eyed, at the tail end of a successfully-completed doctoral program is genuinely surreal, and I'm wondering how the hell I could possibly ever thank everyone who has helped me get here, because I sure as hell couldn't have done it by myself.

First, and perhaps primarily, I have to thank Eric Baković. Over the years, Eric has played the role of advisor, coauthor, and occasional therapist. Eric pushed me to become a better and *bolder* writer. He encouraged me to make strong claims—even if they're wrong—up front, and to stop hedging and apologizing in my writing. I specifically recall getting my dissertation proposal back from him with what seemed like a thousand crossed out instances of hedge words like "generally" and "might". Outside of academics, Eric also went further than I know many advisors do. He encouraged me to look after my health, knowing (correctly) that the work would follow—and be better—if I did. And probably most importantly, Eric never gave up on me when I was more than ready to give up on myself, all the way up to the very end.

Second, I'm grateful to my consultant, Mohamed Ramedhan. This dissertation would not have been possible without his knowledge and expertise, as well as his patience. I learned how to do field work and elicitation while working with Mohamed, and so he put up with far more awkwardness than any consultant I worked with thereafter. Kindness and patience really describes Mohamed; I recall that after telling him *once* that I was tired and hadn't eaten, Mohamed brought cookies to eat together to *every* elicitation session. So, thank you, Mohamed, for teaching me about your language. I haven't seen you for a

long time, but I hope you're well.

Third, I want to thank my committee, all of whom provided advice and feedback throughout my career. I especially want to thank Sharon Rose, whose feedback has always been incredibly detailed and incisive, and whose wealth of knowledge about African languages has been incredibly helpful throughout the years. I would also like to take this moment to thank faculty at UCSD Linguistics who were *not* on my committee, especially Leon Bergen, Grant Goodall, Andy Kehler, Will Styler, and Eva Wittenberg. No one on this list actually worked with me on Mushunguli, but all of them provided mentorship and advice in other aspects of my graduate career. I also want to thank the amazing staff at UCSD Linguistics, whose knowledge and expertise got me out of more than one tricky situation. Thank you!

Speaking of Africanists, I need to thank Dave Odden. Dave was my first academic mentor, the person who introduced me to field work, and the person who put me on the path to graduate school. This is literal: if Dave hadn't voluntold me to take field methods with him as an undergraduate, I would never have started working on Mushunguli in the first place! Dave was tremendously encouraging to me as an undergraduate, reading what I jokingly referred to as novel-length emails about phonology and returning novel-length replies. I also want to specifically thank Dave for providing me with all of his transcriptions from his own sessions with Mohamed; this helped to bolster some parts of the dissertation. So, thank you for that, and thank you for all of the guidance you've provided over the years.

I would also like to thank Michal Temkin Martinez and Holman Tse, both of whom conducted independent research on Mushunguli, and both of whom provided useful insights and help over the years. In particular, Michal gave me access to the Boise Language Project dictionary long before it was actually published. This was incredibly helpful for verifying and quantifying some of my claims, and I appreciate it. I also want to thank

the other members of the OSU Mushunguli Working Group, especially Betsy Pillion and Tyler Williams. You guys were great friends to me, and I miss you.

To conclude the portion regarding colleagues and mentors, I also want to thank quite a few of my teachers throughout my K-12 education. As I alluded to at the beginning of this section, my educational life was difficult pretty much from kindergarten all the way through high school. However, there was a subset of teachers who always encouraged me to get back up and keep trying when I failed, and who were patient through all of my stumbles. So, to Nancy Stewart, Julie Fryman, John Deliman, Ed Zunic, Greg Varner, and Max Hamman, thank you for believing in me, when I had so much trouble believing in myself.

Thank you to my friends at UCSD. All of the graduate students are amazing, and it is so bizarre to think that I am not going to be seeing you anymore. To my surviving cohort—Andres Aguilar, Eric Meinhardt—you were there from the beginning, and I'm so proud of you. Also, I'm really sorry that I started avoiding you, but it's only because you are both so fascinating to talk to that if I didn't isolate myself, I'd never finish anything. I also want to thank Anna Mai, who watched weird art horror movies with me, and who taught me more about LaTeX and Python than anyone else. And thank you especially to Adam McCollum, now at Rutgers. I looked forward to talking to Adam every morning, and my conversations with him made me exceedingly grateful to not be seriously working on vowel harmony. Finishing out these last few months without my ranting buddy has been so weird, though at least we have continued to rant through texts.

I also want to give a shout out to my conference friends: Danny Hieber, Charlie O'Hara, Caitlin Smith, Jesse Zymet, and Jeff Lamontagne (and Holman Tse, who was acknowledged above). I hope I'll see you again soon.

Finally, thanks to my friends from back home–Nathan, Jess, Brett, Angie, Danielle, Alexa. Nobody believes me when I tell them I'm still in touch with all of my high school

friends, but clearly those people never played D&D together.

Having gotten through friends, I finally turn to family. First, I want to thank my parents and brother, Ian. I love you, and I wouldn't have gotten through this (or my life!) without your love and support.

I also *really* have to thank my husband, Evan. I met Evan while I was still in community college, long before I ever decided to study linguistics. So, truly, Evan has been here every step of the way. I genuinely do not think it would have been possible for me to finish graduate school without him, what with him bringing me food, picking up after me, telling me terrible, terrible jokes, and just being there when I needed him. Evan, darling, I know this time has been almost as stressful for you as it has been for me, I know I haven't always been the partner you needed me to be, but I love you so much, and I'm so grateful, and I swear I promise I will try to remember to throw wrappers away in the future. I'm looking forward to whatever comes next, and I'm so happy you'll be with me for it.

To conclude, I would like to do something unorthodox. Throughout my graduate career, there were a number of chance encounters with people who I may have never spoken to again, moments that they may not remember at all. Sometimes those moments were pivotal, but sometimes they came from people who simply said the exact right thing at the right time, that broke through the impostor syndrome, or embarrassment, or self-defeating thoughts that have plagued me throughout my career. I think it is likely that everyone who finishes has these moments, but it's rarely the case that those individuals are actually acknowledged. However, as I am an anxious person who ruminates, I feel the need to thank some of the individuals who gave me something positive to ruminate about. So, thank you to Michael Diercks, Adam Jardine, Brian Joseph, Darya Kavitskaya, Mike Marlo, Alan Prince, Bruce Tesar, and Anne-Michelle Tessier. I'm not going to go into detail about what each of these moments were, but should one of you read this, feel free to ask.

Chapter 2, in part is currently being prepared for submission for publication of the material. Hout, Katherine. The dissertation author was the primary investigator and author of this material.

Chapters 2 and 6, in part contain material that previously appeared in Africa's Endangered Languages: Documentary and Theoretical Approaches. Hout, Katherine, Oxford University Press, 2017. The dissertation author was the primary investigator and author of this paper.

Chapters 3, 4, and 6, in part contain material that previously appeared in San Diego Linguistics Papers 6. Hout, Katherine. The dissertation author was the primary investigator and author of this paper.

Chapters 3-6, in part also contain coauthored material currently being prepared for submission for publication of the material. Hout, Katherine; Baković, Eric. The dissertation author was the primary investigator and author of this material.

Chapters 4 and 5, in part contain material that previously appeared in Proceedings of the 2018 Annual Meeting on Phonology. Hout, Katherine, Cascadilla Proceedings Project, 2019. The dissertation author was the primary investigator and author of this paper.

VITA

2009	Associate of Arts, Columbus State Community College
2012	B. A. in Linguistics *summa cum laude* with Research Distinction, The Ohio State University
2015	M. A. in Linguistics, University of California San Diego
2020	Ph. D in Linguistics, University of California San Diego

PUBLICATIONS

Katherine Hout, “The disambiguating effect of phonological exceptions in grammar”, Proceedings of the 2018 Annual Meeting on Phonology, 2019.

Katherine Hout, “Dominance-as-markedness: evidence from Bari”, Studies in African Linguistics, 48.2, 206-224, 2019

Katherine Hout, “Exceptions to Hiatus Resolution in Mushunguli (Somali Chizigula)”, Africa’s Endangered Languages: Documentary and Theoretical Approaches, 2017.

ABSTRACT OF THE DISSERTATION

Conspiratorial Exceptionality: A Case Study of Mushunguli

by

Katherine Hout

Doctor of Philosophy in Linguistics

University of California San Diego, 2020

Professor Eric Baković, Chair

Cross-linguistically, there are a variety of attested non-phonological conditions on phonotactics and alternations. The extreme end of this are exceptions, whose idiosyncratic behaviors are unattributable to any morphological, morphosyntactic, or semantic class. Studying exceptions is challenging because they require a substantive language description to even be identified, and because their identification can *appear* to subvert descriptive and analytical generalizations about the grammar. As a result, exceptionality research often focuses on major world languages, and there is considerable contention surrounding whether exceptions are wholly phonological, extragrammatical, or something in between.

This dissertation addresses these gaps via an in-depth case study of segmental phonology in Mushunguli (Somali Chizigula, Kizigua; Narrow Bantu, G.31), an endangered, under-described language spoken by members of the Somali Bantu diaspora. This case study, drawn from original fieldwork conducted in 2011-2012, includes a description and analysis of the hiatus resolution and onset structure conspiracies of Mushunguli, the former of which exhibits what appear to be four operations (asymmetric coalescence, glide formation, secondary articulation, and elision) in identical morphosyntactic contexts.

Situated within this discussion are three exceptional patterns, each representing a separate typological instantiation of exceptional blocking: a set of high vowel-initial stems that exclusively block coalescence, but not other applicable repairs (*simple blocking*); a prefix and a verb root which unexpectedly undergo otherwise unattested palatalization in lieu of elision (*walljumping*); and a set of roots that exceptionally block *all* forms of hiatus resolution (*total non-participation*). Adopting lexically-indexed constraints in Stratal Optimality Theory as a means of capturing these patterns reveals complex interdependencies between exceptions and regular forms in Mushunguli, with the form and behavior of one exception crucially determining the forms and behaviors of other exceptional and regular patterns. This suggests that exceptions are lexical but *not* extragrammatical, instead playing an important role in the grammar as reifying and reinforcing agents.

The study concludes by examining alternative representational analyses of the Mushunguli exceptions, including whole-segment absolute neutralization and underspecification. While these approaches are sometimes capable of capturing the exceptional patterns, they ultimately struggle to unify or situate them with respect to the grammar as a whole.

1 Introduction

A foundational assumption of modern linguistics is that languages tend towards generality. However, an empirical truth is that natural languages afford exceptions–lexical items that idiosyncratically fail to conform to broader generalizations about a language's grammar. Acknowledging the existence of exceptions can potentially contravene otherwise elegant descriptions, and modeling exceptions typically requires some degree of stipulation or significant modification of a theory's core machinery or assumptions.

One common "solution" to handling exceptions is simply not to handle them at all–that is, to treat them as noise, performance errors, or extragrammatical. However, within phonology, this viewpoint has historically been relatively unpopular for two reasons. First, it is generally accepted in phonology that phonological generalizations do not necessarily hold equally across all parts of the lexicon, and that factors besides strictly local segmental (or prosodic) structure can and often do condition phonological alternations. These phenomena–which fall under many names, including domain restrictions, lexical class conditioning, and morpheme-specific phonology–are so pervasive that it can be difficult for a researcher to safely dismiss any pattern as noise.

A second reason why it is difficult for phonologists to dismiss potential exceptions out of hand is that it has been observed that even those lexical items which appear to have no easily identifiable extra-phonological conditioning factor nevertheless may share structural characteristics or pattern together. Evidence for this fact comes from both descriptive/formal and experimental sources. In the former case, one notable body of evidence comes from studies of loanword adaptation strategies have demonstrated that borrowings form independent groups or strata within the lexicon; this is usually described as a core-periphery structure, whereby more nativized forms (which typically, but does not always, correspond to separate periods of borrowing) are subject to more phonological well-formedness restrictions than less nativized ones (Hsu & Jesney 2018; Ito & Mester 1995a, 1999a; Paradis & Lebel 1994; Davidson & Noyer 1997; Saciuk 1969, *inter alia*).[1]

In the latter case, there is a growing body of evidence from corpus studies, nonce word experiments, and artificial language learning studies suggesting that human beings are aware of and can generalize from irregular patterns within the lexicon, even those which are again not conditioned by any easily identifiable phonology-external factor. An example comes from Tagalog, where a corpus study conducted by Zuraw (2000) revealed that stems beginning with labial segments are considerably more likely to undergo an exceptional and apparently unproductive pattern of nasal substitution than either dental- or velar-initial ones. This result is supported by nonce word experiments which have demonstrated that Tagalog speakers are aware of these distributions and can extend them to novel forms (Zuraw 2000, 2010). Similar results (both corpus-based and experimental, involving both humans and computational learners) have been observed in a number of other languages, including Russian (Gouskova & Becker 2013; Becker & Gouskova 2016), Spanish (Albright et al. 2001), English (Albright & Hayes 2002, 2003, 2006), Hungarian (Hayes & Londe 2006; Hayes et al. 2009), Turkish (Becker 2006, 2009), Hebrew (Becker 2009), Dutch (Ernestus & Baayen 2003), and Malagasy (Zymet 2018).

The conclusion of this substantive body of evidence is that exceptions, while lexical, nevertheless exhibit phonological patterning that is accessible to speakers. However, there is little to no consensus as to how–and to what extent–these two factors can be reconciled with respect to formalism. The central debate has been the appropriateness

[1] See especially Hsu & Jesney (2018) for an extensive typological survey of cross-linguistic loanword adaptation patterns.

of abstract phonological representations versus lexical diacritics to mark exceptions, with secondary debates pertaining to the relative merits of one class of (usually diacritic) approaches over another. The widespread adoption of Optimality Theory (Prince & Smolensky 1993/2004) and subsequent related proposals has changed the shape of this debate, but has provided no substantive consensus on any of these questions.

1.1 Exceptions affect the grammar

Obscured by this often contentious debate is an important prediction introduced by the adoption of OT itself: several theories of exceptionality in Optimality Theory, both representational and diacritic, explicitly predict that the presence of exceptions within the lexicon can affect the grammar itself. This fact is important in light of a common criticism levied by proponents of representational approaches against diacritic theories (e.g. cophonologies and indexed constraints), which is that these theories' ability to directly reference the lexicon means that they treat exceptions as "extragrammatical" in a way that representational theories (such as underspecification) do not (Inkelas et al. 1997; Inkelas & Zoll 2007).

In reality, while there are of course important differences between diacritic versus representational approaches to exceptionality, extragrammaticality is not one of them. Provided we assume that there are some kind of formal or methodological restriction(s) on the use of exception marking (which is generally the case, even if the exact shape of those restrictions has not yet been fully determined), the formal principles upon which OT is built requires that exceptions be subject to all constraints and rankings required by regular forms, save for those which are rendered inapplicable by the exception marker itself. A logical extension of this is that any ranking(s) required to account for the behavior of an exception will also necessarily apply to *all* forms in the grammar, exceptional or regular. This results in two sets of related predictions: first, the forms and behaviors of exceptions are bound by the grammar itself. Second, the presence and form of a given exceptional pattern will in turn predict (or rule out) both possible exceptional *and regular* patterns.

That this latter prediction–that exceptions can form interdependent relationships with both exceptions and regular forms–has gone mostly unnoticed (though see below) is unsurprising, as the type of case study needed to observe it is difficult to find. Testing the dependency prediction requires a language with a set of regular phenomena that are already in some way interdependent. An ideal context would be a conspiracy (Kisseberth 1970a, 2011; Kenstowicz & Kisseberth 1977), which is a case where a language exhibits multiple context-specific patterns of alternation that ultimately serve to prevent the same illicit structure from surfacing. The language would also need to have at least one set of exceptions situated within that conspiracy, and to test for the presence of implicational or interdependent relationships between exceptions, more than one set will be needed. This also means that the case study in question will need to be robust enough to reliably distinguish exceptions from regular forms, and there needs to exist a formal treatment of both the conspiracy itself and the exceptions within it.

In the extant exceptionality literature, cases like this are exceedingly rare. Some early work in exceptionality did attempt to address multiple potentially related exceptional patterns in the same language, and/or attempted to situate exceptions with respect to multiple interdependent processes; notable examples come from work on Yine (Kisseberth 1970b), Maltese Arabic (Brame 1972), Seri (Marlett & Stemberger 1983), and Nupe (Hyman 1970, 1973; Harms 1973).[2] However, much more common in the recent

[2]Some of these cases have been reopened, but not necessarily with the intent of investigating interdependencies. This is most true for Yine, which has had at least two modern re-openings (Pater 2010; Zimmermann 2013), each in support of a separate formal proposal. Neither of these contributions take into consideration important patterns identified by other researchers (Matteson 1965; Lin 1997; Hanson 2010) that have bearing on the analysis of the exceptional (or "morpheme-specific") patterns exhibited by the language, as well as the formal proposals themselves. I discuss

literature have been contributions focusing on several shallow examinations of (often) unrelated exceptions in (often) unrelated languages (Pater 2010; Finley 2010; Inkelas & Zoll 2007; Inkelas et al. 1997; Bermúdez-Otero 2012; Hsu & Jesney 2018; Kraska-Szlenk 1997, 1999; Mahanta 2008, 2012, *inter alia*), or a more robust examination of a single exception in a single language, which may or may not be supported by additional shallow cases (Anttila 2002; Archangeli & Pulleyblank 2015; Becker 2006; Compton & Dresher 2011; Gouskova 2012; Gouskova & Becker 2013; Hout 2016, 2017; Hyman 1970; Kabak & Vogel 2011; Paradis & Lebel 1994; Pater 2000; Yearley 1995; Zuraw 2000, *inter alia*). These types of studies are suitable for testing other aspects of formal theories, but are not generally suitable to test whether the prediction of implicational relationships between exceptions and regular forms is a valid one.

That said, the notion that exceptions can influence or illuminate aspects of the grammar has occasionally surfaced in the learnability and loanword adaptation literatures, which both tend to look at larger and more diverse datasets. Recent research by Zymet (2018, 2019) (building off of earlier work by Moore-Cantwell & Pater 2016 and Tanaka 2017) found that the introduction of lexically-specific constraints—a diacritic form of exception marking—into a standard MaxEnt learning model could have deleterious effects on the ability for the model to accurately learn frequency distributions of regular and lexically-specific patterns. Unmodified MaxEnt models treat both lexically-specific and general constraints as regular effects; however, exceptional forms are subject to both sets, while regular forms are only subject to the general set. Thus, an unmodified learner can eventually discard the general constraints in favor of the lexically-specific ones. This results in a learner that can perfectly match the frequencies in training data (= learn the lexicon), but which cannot extend them to novel forms (= fail to learn the grammar), instead preferring a much higher rate of irregularity.

Yine briefly in §7 as a potential avenue for future research.

Note that Zymet's solution to this was not to ban indexed constraints from the grammar, but rather to modify the learner such that general constraints were preferred as an explanatory mechanism over lexically-specific ones. It is also worth pointing out that Zymet's research compares the effects of indexed versus general constraints on the learning of statistical generalizations. Crucially, this work does not test whether the forms and behaviors of exceptions predict the forms and behaviors of other exceptions and regular forms in the same language; rather, it demonstrates that marks which are too powerful can potentially result in the overextension of an irregular pattern.

Something closer to what we are looking for comes from the loanword adaptation literature. As mentioned above, numerous studies of loanword adaptation have observed that borrowings often result in a stratified lexicon, and that these strata can form implicational relationships in the form of a *core-periphery structure* (Ito & Mester 1995a,b). More native (or "nativized") lexical items form the core, which is subject to the largest number of well-formedness (i.e. markedness) restrictions. The stratum immediately peripheral to the core is subject to a proper subset of these markedness violations, and this continues hierarchically all the way out to the most peripheral (= least nativized) stratum.

There have been multiple formal proposals to account for this, but the one I will focus on is stratal faithfulness, as it is in this literature specifically that the ability for some forms of exception marking to influence the ranking of regular constraints was first noticed. To account for the stratified lexicon of Japanese, Ito & Mester 1995a (and subsequent work) propose that the grammar is comprised of general well-formedness (= markedness) constraints, but both general and specific forms of faithfulness constraints, the latter of which correspond to individual lexical strata. These stratal faithfulness constraints only apply to (= can only be violated by) lexical items belonging to the associated stratum. In this proposal, Ito & Mester also observed that the interleaving of these stratally-sensitive faithfulness constraints among the otherwise unranked general markedness constraints resulted in a fully determined ranking by transitivity.

To illustrate, we will take a simple example from Japanese (specifically, from Ito & Mester 1999a). Japanese syllable structure is subject to numerous restrictions; two such restrictions are that codas can contain only place-linked or underlyingly placeless consonants (including the nasal glide N), and non-geminate [p] is banned in onsets. The former structure violates the markedness constraint CODACOND, while the latter violates NO-P. Neither of these constraints are ever violated by native forms, and they do not generally make conflicting decisions. Thus, absent further evidence, the default assumption would be that these constraints are both undominated, and unranked with respect to one another.

However, evidence for a ranking can be found if we look outside the core native vocabulary into loanword strata. It turns out that *all* surface forms in Japanese obey CODACOND, including recent borrowings, as evidenced by words such as English 'bed' /bɛd/ being reanalyzed as [bedd**o**]. The same is not true for NO-P; while native forms and very early borrowings from Chinese (referred to here as Sino-Japanese) seem to uniformly obey it, more recent borrowings can violate it, surfacing with unrepaired singleton *p*.

For example, consider the minimal pair *haN* 'group' (Sino-Japanese) and *paN* 'bread' (from Portuguese or French). Following Ito & Mester (1999a), while these words are assumed to have the same underlying representation /paN/, they are assigned to separate strata (Sino-Japanese and Assimilated Foreign, respectively). This means that they are subject to separate faithfulness constraints; for our purposes, these are FAITH^{SJ}, corresponding to the Sino-Japanese stratum, and FAITH^{AF}, corresponding to the foreign stratum.

Recall that even recent borrowings obey CODACOND; thus, we assume that it must be undominated (for illustrative purposes, we will assume that it dominates all forms of FAITH). To account for the fact that French /pan/ surfaces faithfully as [pan], FAITH^{AF}

must dominate NO-P (1). To account for the fact that Sino-Japanese /pan/ surfaces as [han], FAITH^{SJ} is crucially dominated by NO-P (2).

(1) /paN^{AF}/ → [paN] (adapted from Ito & Mester 1999a)

paN^{AF}	CODACOND	FAITH^{AF}	NO-P	FAITH^{SJ}
a. ☞ paN			*	
b. haN		*!		

(2) /paN^{SJ}/ → [haN] (adapted from Ito & Mester 1999a)

paN^{SJ}	CODACOND	FAITH^{AF}	NO-P	FAITH^{SJ}
a. paN			*!	
b. ☞ haN				*

Notice that as we progress through the tableaux, a ranking emerges between CODACOND and NO-P. However, this ranking is transitive; CODACOND and NO-P can never conflict with one another due to the fact that they apply to independent structures (codas and onsets). Ito & Mester (1995a:190) remark on this transitivity directly, stating that "...it remains to be seen whether this type of ranking argumentation is valid", but do not comment further upon it in subsequent work.

To my knowledge, this prediction has not been substantially explored. This is perhaps unsurprising with respect to Japanese; many of the repairs exhibited (or blocked) by these loanwords are at best loosely-related to one another, making it difficult to definitively tell whether newly-determined rankings are actually important, or simply a byproduct of the model. Moreover, there are additional problems with using loanword adaptation studies—at least as they are typically conducted—as a primary testing ground for the dependency prediction. As seen in the above example, a common tendency in the loanword adaptation literature is to represent lexicon stratification as repairs from foreign *underlying representations* to a nativized surface forms. A second, related tendency is the comparison between what are assumed to be productive repairs conditioned by morphology and what may be diachronic changes between borrowed forms and their surface outputs.

For example, it has been claimed that in Japanese, native vocabulary and foreign loans repair illicit consonant clusters in different ways, with native forms preferring consonant deletion and non-native forms preferring vowel insertion; e.g. English *cream* (= /**kɹ**iːm/) surfaces as [**kɯ.r**iː.mɯ], while Japanese /kak + rɯ/ 'write-NONPAST' surfaces as [ka.k_ɯ] (Smith 2006). However, to treat these as comparable repairs, we must assume that Japanese speakers have stored the unrepaired English input in their lexicon, and that there are no differences between markedness restrictions applying to root-internal sequences and morphophonological alternations. Neither of these assumptions are damning on their own: Japanese speakers may have enough familiarity with English to have plausibly stored a non-native underlying representation, and root-internal well-formedness restrictions are often mirrored by alternations (the duplication problem (Kisseberth 1970a)). However, it seems less likely that native speakers have internalized the underlying representation of *haɴ* 'group' as anything other than *haɴ*, and it is similarly implausible that a native speaker with no exposure to English whatsoever (save for loanwords) would have stored the underlying representation of 'cream' or 'bed' as anything besides something closer to their Japanese pronunciations.

It is also well-known that morphologically complex forms are often *not* always subject to the same restrictions as root-internal structures. For example, in many Bantu languages, there is a cooccurence restriction on the distribution of mid and high vowels in (verbal) roots that extends to derivational extensions (Clements 1991; Hyman 1999, 2003); namely, mid vowels cannot generally appear with high vowels and vice versa.[3] This is seen in Shona, where the applied form of *pera* 'end' surfaces as *perera*, while the

[3] This co-occurrence restriction appears in Mushunguli as well, but is fairly robustly observed in nouns as well as verbs.

applied form of *ipa* 'be evil' is *ipira* (Fortune 1955; Beckman 1997). However, inflectional prefixes concatenated with these stems are not bound by this cooccurrence restriction, e.g. *kuɓereka* 'to carry on the back'. Here, we see that phonotactic restrictions applicable to roots and stems do *not* apply uniformly to morphologically-complex forms. A similar sort of pattern can be observed in ciNsenga (Bantu; N.41) hiatus resolution (Simango & Kadenge 2014): in this language, vowel-vowel sequences occurring at prefix-prefix junctures uniformly undergo either glide formation or elision of V1; however, for verbs, hiatus at a prefix-root juncture remains unresolved, resulting in surface forms with partially-resolved and partially-unresolved hiatus, e.g. /s**i** + **u** + ka + n**i** + **o**n + a/ → [s**u**kan**i.o**na] 'you (sg) will not see me' (root underlined, hiatus contexts in bold).

Neither of these are reasons to dismiss loanword adaptation studies as a source of evidence for implicational relationships between exceptions; the core-periphery structure and lexical stratification is well-attested in multiple languages. However, with respect to the interdependency prediction, these issues introduce the possibility that any result more interesting than the transitively-determined rankings from Ito & Mester 1995a can potentially be attributable to an overlooked domain restriction or an incorrect assumption regarding an underlying representation.

1.2 The case of Mushunguli

This dissertation aims to explore this under-discussed prediction by presenting a novel, in-depth case study of the hiatus resolution conspiracy in Mushunguli (Somali Chizigula, Kizigua; Narrow Bantu, G.311), an endangered, under-described language spoken by members of the Somali Bantu diaspora. As discussed in §2 and §3, Mushunguli has a complicated system of hiatus resolution, exhibiting (descriptively) four separate hiatus repair strategies in similar morphosyntactic contexts. Mushunguli is also unusual in the context of the exceptionality literature in that it exhibits three separate exceptions to hiatus resolution: a set of stems that block one repair (coalescence) but not others; a set of morphemes that exceptionally undergo palatalization (a repair that is otherwise banned in the language); and a set of stems that fail to participate in hiatus resolution whatsoever.

These factors make Mushunguli an ideal test case for the dependency predictions. The fact that the language exhibits complex conspiratorial processes in the same set of morphosyntactic contexts means that we already predict that they will interact; that is, the analysis of even a regular alternation within the conspiracy will have consequences for the analysis of others. Moreover, that there are multiple exceptional patterns exhibiting *separate* behaviors situated within the conspiracy means that that we can test for interdependencies between both the exceptions themselves and the exceptions and regular forms.

The Mushunguli case study also avoids many of the issues seen in the loanword adaptation literature. The Mushunguli exceptions all involve morphophonological alternations; this means that there is no sense here that we are potentially comparing metaphorical apples to oranges. Moreover, while identifying the source of exceptional behavior is challenging for an under-described language with no historical written tradition, comparison with earlier descriptions of the language's sister dialect (Tanzanian Chizigula) suggest that the non-coalescing and palatalization exceptions are native forms that behave irregularly.[4] Only the third set contains loanwords from identifiable non-Bantu sources, and even this set contains at least one native form. As such, it is not necessary to assume that the underlying representations of words correspond with either their historical forms (if native) or their foreign underlying representations (if identifiably loanwords).

Because of these factors, and again, because all of the repairs in question are di-

[4]As discussed briefly in §6, comparison of Mushunguli to Tanzanian Chizigula indicates that the exceptional patterns in the former (blocking of coalescence, palatalization) are the majority patterns of the latter. There is insufficient evidence to determine the directionality of the change, however; that is, it is unclear which of the two dialects are more conservative.

rectly related to each other, Mushunguli is a more compelling test case for not only the existence of implicational relationships between sets of exceptions, but also the prediction of dependency relationships between exceptions and regular forms. As will be seen in §4 and §5, this prediction is supported: using lexically-indexed constraints (Pater 2000, 2010), I demonstrate that the presence of these exceptions within the lexicon reveals deeper underlying structure that would otherwise be impossible to discern. Rather than potentially undermining the hiatus resolution conspiracy, the exceptions in fact support it by revealing otherwise non-identifiable rankings and providing avenues by which ranking paradoxes can be solved. The patterns of Mushunguli indeed suggest that exceptions, while lexical, are not actually extragrammatical–rather, they play a reifying and reflective role within the grammar.

The Mushunguli case study is a significant contribution to the exceptionality literature from three typological perspectives, as well. From a phenomenological perspective, while segmental processes are not an atypical topic in work on exceptionality, exceptions to hiatus resolution specifically are under-discussed relative to other phenomena (though see Roberts-Kohno (1998); Mahanta (2008); Marlett & Stemberger (1983); Inkelas & Zoll (2007)), and I am unaware of any other case studies of exceptions to hiatus resolution that go to the same level of depth as this dissertation. From an areal perspective, while African languages are well-represented in studies regarding domain and positional restrictions on vowel harmony, tone, and reduplication (Baković 2000; Beckman 1997; Clements 1991; Archangeli & Pulleyblank 2015; Downing 2005, 1997; Sande 2019, *inter alia*), formal treatments of *exceptions* in African languages are relatively rare (though see Hyman 1970; Finley 2010; Hout 2019b; Roberts-Kohno 1998; Kraska-Szlenk 1997, 1999), especially when compared to contributions from major world languages, especially European ones. Finally, from the standpoint of typology of exceptions themselves, Mushunguli exhibits at least one rarely-attested form of exceptionality; this is the palatalization exception, which exhibits an alternative repair (referred to here as a *walljumping* pattern). Walljumping patterns are nearly unattested outside of the loanword adaptation literature, which again often involves the comparison of root-internal phonotactic repairs with morphophonological alternations; the only extant case that I am aware of that involves morphophonology comes from a brief discussion of Turkish hiatus resolution by Inkelas & Zoll (2007).

Finally, this case study is a significant contribution to the general Bantu literature as well. While there now exists a community-oriented online dictionary of this language (Dayley et al. 2020), peer-reviewed linguistic description and formal analysis of any aspect of this language is sparse. Note that this is not because Mushunguli lacks interesting data; in addition to the hiatus resolution patterns and their exceptions (Hout 2017, 2019a), the language has both typologically rare segments (Temkin Martinez & Rosenbaum 2017) and phonological contrasts (Tse 2015a,b). The only other aspect of the language with extant description and analysis is work by Barlew (2013, 2016) on locative semantics.[5] Again, Mushunguli is an endangered language, and the situation of its speakers is not improving, as evidenced by the fact that between 2006 and 2019, the count of speakers recorded by Ethnologue (Eberhard et al. 2019) has dropped from 23,000 to 20,000. As such, a secondary goal of this dissertation is to improve both breadth and depth of extant description on this language's morphology and phonology. While the description and analysis contained in this dissertation is far from complete, the grammar sketch presented in §2-4 is crafted to be as complete as is possible without being overly speculative, and I have taken care to both include a diverse set of examples and to avoid repeating examples as much as is feasible.

[5]Non-peer reviewed but publicly-available descriptions and analyses include Hout 2012 (syllable structure and hiatus resolution), Hout 2016 (exceptionality), Pillion 2013 (tone, though see §2.2.5, this dissertation), Hout 2011; Williams 2012 (noun class and agreement), and Tse 2014 (retroflexion and sociophonetics).

1.3 Defining exception

Before moving to the case study in question, I would like to make it clear what I mean when I say "exception" or "exceptionality". The term *exception* (or *exceptional*) is more commonly used colloquially than technically, either referring to a simple counterfactual statement, or to noise in the data that a given researcher may or may not intend to address. However, there are a set of empirical phenomena under the umbrella of *exceptionality*, which are also sometimes referred to as *morpheme-specific*, *class-conditioned*, *morphologically-conditioned* or *lexically-specific* phonology (Pater 2000, 2010; Shih & Inkelas 2016; Shih 2018). While the phenomena described by these terms have some shared characteristics, they are quite diverse, and so before moving to the structure of the dissertation and the case study in question, I feel it is important to make crystal clear why I choose to classify the idiosyncratic patterns in Mushunguli as *exceptions*, with the aim of making it easier for other researchers to situate these patterns with respect to other phenomena attested throughout the literature, especially researchers interested in typology.

"(Phonological) exception" as applied to Mushunguli refers to a **small**, **arbitrary** lexical class that exhibits (more or less) **categorical** phonological behavior that is *idiosyncratic* relative to other similar lexical items that appear in the same morphophonological domain and/or morphosyntactic contexts.

"Small" corresponds to the notion of quantitative weakness identified by Moravcsik (2011); if one was to list out the entire lexicon of a given language, the set of exceptions is presumed to be comprised of fewer members than the set of non-exceptions. Note that this distinguishes the Musuhunguli exceptions from some cases of morpheme-specific phonology where a regular pattern is difficult to identify, such as Finnish allomorphy (Anttila 2002; Pater 2010) or Yine vowel deletion (Kisseberth 1970b; Lin 1997; Pater 2010; Zimmermann 2013), as well as many studies of loanword adaptation.

"Arbitrary" corresponds to the notion of qualitative weakness identified by Moravcsik (2011). The generalization is that exceptions can be distinguished from other cases of extra-phonological conditioning, such as domain or positional restrictions, in that they are exceptional in only one way; their idiosyncratic behavior is not attributable to (or representative of) any well-defined morphological, morphosyntactic, or semantic class. For example, in Modern Hebrew, stress placement is predictable in verbs, but is contrastive (= unpredictable) in nouns (Becker 2003). While it is the case that this distinction is attributable to non-phonological factors, distinctions like these are again not exceptional in the sense intended here, as the predictable nature of stress placement in verbs and the unpredictable nature of stress placement in nouns are both characteristic of the morphosyntactic classes they belong to.

A related, but slightly more controversial notion, from Finley 2010, is that exceptions should have "no phonological explanation" for their behavior. This means that the behavior of the purported exceptions should not be representative of morphemes of that particular phonological shape. For example, many types of vowel harmony have vowels that regularly fail to participate. An example is Akan (Clements 1981; Dolphyne 1988), in which the [-ATR] low vowel /a/ has been described as failing to participate as expected in otherwise productive bidirectional ATR harmony. However, this idiosyncratic behavior is not attributable to any specific morpheme, but rather is a property of /a/ itself.

Both of these criteria are in service of a goal of capturing important linguistic generalizations–because the term "exception" has a secondary meaning of "counterfactual," there is a sense that applying it to a morphemes exhibiting idiosyncratic behavior that has some well-supported linguistic explanation is missing something important that speakers may be aware of. For Mushunguli, while it is the case that there is sometimes some morphosyntactic overlap, and that there are sometimes shared phonological characteristics among a group of exceptions, none of these behaviors are genuinely representative of any easily definable phonological, morphological, morphosyntactic, or semantic class, and cannot be generalized to any such class.

This distinguishes Mushunguli from other Bantu languages with similar patterns that *do* have these types of non-arbitrary explanations for their behavior. The aforementioned example of ciNsenga is one of them. As stated above, ciNsenga bans hiatus resolution at prefix + stem junctures, but this is only (and apparently uniformly) true for verbs. The example given was *sukani.ona* 'you (sg) will not see me', which features the verb *ona* 'see'. As discussed in §5.2, Mushunguli has verbs that also block all forms of hiatus resolution. Interestingly, *ona* is one of them, as evidenced by the example /**si** + **on** + a/ → [s**i**.**oː**na] 'I saw'.

However, there is an important dissimilarity: as will be discussed in §2.2.4, word-internal hiatus resolution in Mushunguli has no morphosyntactic restrictions, including at prefix + stem junctions (e.g. /si + omal + a/ → [soːmaːla] 'I dished up ugali'). In ciNsenga, the behavior of *ona* is generalizable to the class of verbs, and thus it would a mistake to view it as exceptional. In Mushunguli, the behavior of *ona* is idiosyncratic with respect to that of other *o*-initial stems and the class of verbs; as such, it is exceptional. This difference in behavior suggests that there is a genuine typological distinction between different types of "morpheme-specific" phonology. Because of this, it is possible that behaviors that are truly "exceptional" in the sense that I am adopting here may stem from different pathways and/or need to be accounted for in different ways.[6]

Categorical means that the behaviors discussed here were robustly and repeatedly attested in elicitation. However, I acknowledge that because the majority of data presented in this dissertation stems from elicitation with a single speaker, it is possible that intra-speaker variation could occur in other conversational contexts, or that inter-speaker variation exists. That said, I do address some examples of both free variation and optionality in §2, and these examples also mainly came from elicitation with the same spea-

[6]Determining this will require a substantive typological survey of patterns of morpheme-specific phonology and related phenomena, which is outside of the scope of this dissertation.

ker. Moreover, the categoricity of the exceptional patterns of interest in §4-§5 has been corroborated across multiple sources, including a second consultant, consultation with the Boise Language Project's Online Chizigula-English dictionary, comparison with a lexicon of Mushunguli's sister dialect (Tanzanian Chizigula), and/or personal communication with other researchers.

Finally, "idiosyncratic" as a criterion refers to the fact that the exceptions exhibit phonological behavior that is unexpected relative to *phonologically-similar* lexical items. Note that it does not matter whether or not exceptions fall into phonologically-patterned groupings; this is predicted by many theories of exceptionality, including the primary theory adopted in this dissertation (lexically-indexed constraints), as well as the representational theories explored in §6 (whole-segment absolute neutralization and prespecification). Note also and critically that this idiosyncrasy is entirely dependent on synchronic phonological factors: as I have already alluded to in my earlier discussion of the loanword adaptation literature, and as I will state clearly here: for the purposes of *this dissertation*, I do not consider historical source of exceptional behavior as being important when it comes to explaining the behavior of these exceptions in the context of the synchronic grammar. While I do occasionally reference these facts for the purpose of fleshing out the discussion or providing some context, the formalism I am adopting treats all exceptions the same way, regardless of source.

1.4 Structure of the dissertation

In §2, I present a grammar sketch of Mushunguli. In §2.1, I describe the background of the language and its speakers. In the subsequent sections, I describe the phonology (§2.2) and morphology (§2.3) of Mushunguli, and discuss points of future research.

In §3, I present a baseline analysis of the hiatus resolution conspiracy in Optimality Theory. This analysis concludes in an interesting observation and a puzzle. The observation is that some patterns that look like vowel deletion, and which are treated descriptively as such in §2.2.4, are more parsimoniously analyzed as coalescence. The puzzle is that another pattern that looks like vowel deletion cannot be easily or parsimoniously analyzed this way (or any other way), leaving the status of vowel deletion in the language uncertain.

This analysis lays the groundwork for the discussion of the exceptional patterns introduced in §4 and §6. In §4.1, I elaborate upon the functioning of and typological predictions made by the theory of indexed constraints adopted by this dissertation, and establish some analytical biases I will adopt going forward. In §4.2, I present a set of stems that exceptionally block coalescence; here, I show that indexed constraints are capable of capturing and explaining critical phonological generalizations about them. I also show that the existence of the non-coalescing stems clarifies the uncertain status of deletion in this language; namely, under the theoretical assumptions required by lexical indexation, the existence of these exceptions means that vowel deletion is not uncertain, it is *impossible*.

The presence of these exceptions thus opens a pathway for an analysis of the puzzle introduced at the end of §3, and provides insight into the functioning of the grammar. As I discuss in §5, these patterns are better analyzed derivationally as a two-step process of hiatus resolution at level of the word and a postlexical onset repair. I justify this based on the structure of the verb in Mushunguli. The chapter builds an analysis of onset structure in teh language, and then uses Stratal OT to account for the remaining non-exceptional patterns. In §5.1.3 I introduce a case of two morphemes that exhibit a typologically unusual pattern of exceptionality which I term *walljumping*: they exceptionally block the regular repair, but undergo an alternative that is otherwise unattested in the language. I show that the form and behavior of these exceptions follows naturally from the requirements placed upon the grammar by the non-coalescing stems; moreover, their behavior follows from and reinforces the derivational analysis of Mushunguli onset structure repair. Both observations are further reinforced by a third set of exceptions presented in §5.2, which I refer to as total non-participants, due to the fact that they uniformly fail to participate in hiatus resolution at all. Again, I demonstrate that the form and behavior of these exceptions conforms to the predictions and restrictions made by the non-coalescing stems, and reinforces other aspects of Mushunguli morphophonology.

In §6, I examine alternative representational approaches to exceptionality in Mushunguli, with a special emphasis on the non-coalescing stems. In §6.1, I introduce a rule-based analysis of the non-coalescing stems that relies on absolute neutralization of whole segments. I show that this analysis is worse at linking phonological structure with exceptional phonological behavior than the indexed constraint analysis. I also show that while absolute neutralization of segments is a possible solution in Stratal OT (under very narrow circumstances), it is not a possible solution *for Mushunguli*. In §6.2, I sketch an analysis based in prespecification (Inkelas et al. 1997), which is a principled form of underspecification. I show that while this approach is capable of capturing the non-coalescing stems from §4.2, it struggles to capture some broader generalizations about vowel length in the language, and especially struggles with the set of exceptions introduced in §5.2.

In §7, I summarize the analysis presented throughout the dissertation, as well as the alternative accounts. I also point out avenues for future research, both from descriptive and typological perspectives.

2 Mushunguli Phonology and Morphology

To motivate the analysis of exceptions in Mushunguli, it is necessary to have a firm grounding in the descriptive facts of the language. This chapter serves as an introduction to Mushunguli. The chapter is laid out as follows: §2.1 will discuss the background of the language and its speakers, as well as provide some history of the work. §2.2 will provide a high-level overview of the phonology of the language, including phonemic inventory, syllable structure and repairs, and some aspects of prosody. §2.3 will cover morphology relevant to the cases at hand. Finally, §2.2.4 will delve more in depth into a specific subcase in the morphophonology, which is the hiatus resolution conspiracy. The purpose of this final section is to establish the groundwork for the following analyses of exceptions to this conspiracy.

A secondary purpose of this chapter is to ensure that some description of this language's phonology and morphology exists in a single place. As discussed in the introduction, there are very few published works on the language, and no (peer-reviewed) published descriptions of the morphophonology (save for small sketches in Hout 2017 and Hout 2019a). However, while I have taken care to be as complete as possible, this chapter should not be viewed as a substitute for a proper grammar; as will be seen, there are a few significant gaps in this description, particularly with respect to tone, verbal extensions, and some aspects of TAM morphology. Note too that the majority of these data were those used in the development of Hout 2012; in re-assessing this earlier work, some transcription errors were found. In any case where this description contradicts the previous one, the description in the dissertation should take precedence.

2.1 Background

Mushunguli, also known as Somali Chizigula, Chizigua, or Kizigua (G.311, ISO [xma]), is an endangered, under-described Somali Bantu language.

Mushunguli and its speakers originate in Tanzania, where a partially-intelligible sister dialect, Tanzanian Chizigula, is still spoken. The speakers of the Somali dialect descend from Zigula people imported as slaves into Somalia in the nineteenth century, who later escaped slavery and resettled in the lower Jubba River valley (Van Lehman & Eno 2003; Eno & Eno 2007). Together with escaped slaves from other ethnic groups (primarily Yao and Makhuwa), they formed what became known as the Gosha community, and so Zigula people are sometimes also referred to as Gosha. However, unlike many other Bantu ethnic groups affected by the Arab slave trade, the Zigula people retained both their Bantu identity and their language.

Both the terms "Somali Bantu" and "Mushunguli" itself are complicated. *Somali Bantu* refers collectively to any group with Bantu heritage who settled in the Jubba River valley. However, not all Somali Bantu who still identify as Bantu ethnically actually speak a Bantu language; many primarily speak Maay, a Cushitic language. Similarly, not all Somali Bantu who still speak Bantu languages identify as Bantu ethnically, as is the case for the Mwiini. Again, the Zigula people retained both their ethnic identity and their language, and Mushunguli is the dominant Bantu language spoken in the region (Chanoff 2002; Van Lehman & Eno 2003).

The origins of the term “Mushunguli” are murky, but it has been claimed to be a Somali reanalysis of the Chizigula word *mzigula* ‘person’ (Chanoff 2002), and corresponds to both the people and the language itself. The word carries negative connotations such as “slave” and “worker,” and so for some speakers (particularly younger ones), this name is dispreferred over the endonym Chizigula or the Swahili exonym Kizigua (Michal Temkin Martinez p.c). Regardless, I adopt it here for two reasons. The first and more important is that this was the word my primary consultant consistently used, despite being aware of other options. The second is more pragmatic: using “Mushunguli” as opposed to “Kizigua” or “Chizigula” allows for easier distinction between the Somali and Tanzanian dialects of the language, the latter of which is better described.

Members of the Gosha community were displaced a second time due to the Somali Civil War in the 1990s. Some pockets of Mushunguli speakers remain in Somalia (mainly in Mogadishu); however, many more have spent some time in the refugee camps in Kenya, Uganda, and Tanzania. While some have resettled in these countries, a large number of Somali Bantu–including Zigula people–were eventually resettled in multiple cities in the United States as a part of the larger Somali diaspora. Major Somali Bantu communities are located in Columbus, OH, San Diego, Boise, Buffalo, Twin Cities, and Pittsburgh (Nurse 2010).

This language was up until recently classified by Ethnologue as “vigorous,” with 23,000 speakers. However, it has been recently reclassified as “threatened,” and the number of speakers has dropped to 20,000 (Eberhard et al. 2019). This reclassification is consistent with my experiences working with speakers of this language. In Somalia, Mushunguli had extensive contact with Somali and Maay, but remained distinctly Bantu due to the group’s strong ethnic identity. However, the splintering of the community due to resettlement and the inclusion of the Somali Bantu within the larger Somali diaspora has greatly affected the transmission of the language to younger speakers, as well as the language itself. For example, while my primary consultant, Mohamed Ramedhan, reported that he had been teaching the language to his children, he noted that his younger children spoke it less well than his older children, and had trouble communicating with their grandmother, who spoke no English. Personal communication with another member of the community in San Diego, Hamadi Jumale, supported this, in particular noting the loss of vocabulary and knowledge specific to life and culture in Somalia (especially ethnobotanical and ethnomedical terms). These observations were corroborated by members of the Boise Research Group, who worked with younger speakers than my primary consultant, and the dictionary they have developed indicates that there has been increasing adoption of Swahili and English loanwords, likely due to the influence of the refugee camp system as well as resettlement in English-speaking countries such as the United States. Given the decreasing proficiency of younger speakers and the displacement and splintering of the community, it is unclear to me whether the language will persist beyond the next couple of generations; if it does survive, my assumption is that it will be substantively different than the relatively stable language it had been before the Somali Civil War.

The majority of the data presented in this dissertation come from an elicited corpus of over 2500 tokens collected from a native speaker in 2011 and 2012. My primary consultant, Mohamed Ramedhan, was born in Mogambo and had resettled in Columbus, OH. Mohamed was in his late 50s or early 60s at the time these data were collected. In addition to Mushunguli, he is fluent in Somali and Maay, as well as proficient in Arabic, Swahili and some English. Elicitation was primarily conducted in English, but Swahili and Arabic were occasionally used to navigate linguistic barriers. A few supplementary examples are taken (with permission) from David Odden's data, who also worked with Mohamed. A second set of supplementary examples are taken from the Online Chizigula-English dictionary recently completed by the Boise Language Project (Dayley et al. 2020). As it would be cumbersome to fully cite external examples every time they occur, I will be using a superscript on the gloss when they are introduced: examples from Odden are marked with a superscript ^{D}, while examples from Boise are marked with a superscript ^{B}.

2.2 Phonetics and phonology

This section is a descriptive account of aspects of Mushunguli phonology that are not of primary concern to the case studies in the following two chapters; that is, everything except hiatus resolution, which is given a formal treatment in §2.2.4. In this section, I will be using a narrower phonetic transcription for consonants than I will be adopting throughout the rest of the dissertation. As such, I will also be using this section to establish the broader transcription method I use to render examples during analysis.

2.2.1 Consonants

Mushunguli has a large phonetic consonant inventory, and due to the heavy influx of borrowings and complications introduced by morphophonology, it is difficult to assess which sounds are definitively contrastive in some cases. Personal communication between myself and other researchers working on Mushunguli suggests that there is a considerable amount of inter-speaker variation, so here I provide analysis of my primary consultant's idiolect, noting controversy as it arises.

Mushunguli has around 33 phonetic consonants, illustrated below.

Before I discuss the distribution of consonants by manner, I would like to make clear some general distributional facts about consonants in the language. Unless otherwise noted in the forthcoming description, there are relatively few truly robust restrictions between consonants and vowels. It is also quite difficult to identify true minimal pair contrasts between consonants of the same manner or place, so my evidence for contrastiveness largely comes from (lack of) distributional restrictions. This is made more

	Bilabial	Labiodental	Dental	Retroflex	Palatal	Velar	Uvular	Glottal
Oral Stop	p ɓ		t ɗ	ᶑ	tʃ ʄ	k ɠ	q	
Prenas. Stop	ᵐ̊p ᵐb			ᶯ̊ʈ ᶯɖ	ᶮɟ	ᵑ̊k ᵑg		
Nasal	m		n		ɲ	ŋ		
Fricative		f v	s z		ʃ	x ɣ		h
Trill			r					
Approx	w				j			
Lat. appr.			l					

Figure 2.1: Consonants of Mushunguli

complicated due to the morphological complexity of Mushunguli (see below) and the heavy influence of other languages, including (but probably not limited to) Arabic, Somali, Maay, Italian, Swahili, and English. A secondary source of complication stemming from the sociolinguistic situation of Mushunguli is the possible early influence of Bantu languages spoken by other ethnic groups affected by the Arab slave trade (see Tse 2014, 2015a,b for more discussion on this point).

There is also an important distinction to be made in this description between *word*-initial/internal and *root*-initial/internal. In general, I will be giving examples of word-initial and word-internal segments. That said, certain consonants are considerably more likely to occur root-internally than root-initially (and vice versa). As discussed in more detail in §2.3, Mushunguli is an agglutinating language, and much of the inflectional morphology comes in the form of prefixes, with *very* few lexical items appearing regularly in unprefixed contexts. This affects the distribution of consonants; for example, some segments—namely prenasalized stops—are robustly attested *word*-initially, but not *root*-initially. This can most easily be seen by comparing nominal roots occurring in multiple classes, e.g. [ᵐbuga] ‘rabbit’ (in noun classes 9/10), but [kabuga] ‘bunny’ (in noun class 12). Similarly, voiced segments (especially sonorants and prenasalized stops) are more common root-internally than voiceless ones, but *word*-medial voiceless consonants are robustly attested due to the fact that they are much more likely to occur *root*-initially. That said, I have no examples in my data whatsoever of a word ending in a consonant.

Segments with secondary articulations (prenasalization, labialization, and velarization) are very common both word-initially and root-internally; labialization/velarization is also common root-initially (see discussions regarding stops, nasals, and approximants, below). True consonant clusters are uncommon in my data, and exclusively occur in borrowings. They are somewhat more common in the BLP dictionary (Dayley et al. 2020), though still restricted to borrowings and still the minority of cases.[1]

2.2.1.1 Stops and affricates

Mushunguli has roughly 17 phonetic stops and affricates, including both pulmonic and implosive stops, as well as voiced and voiceless prenasalized stops. As I stated above, it is difficult to find true minimal pair contrasts between stops; however, some stops do have predictable distributions based on syllable position and following vowels.

Voiceless oral stops are unaspirated in all contexts. Voiced oral stops are generally implosive in Mushunguli; this is most apparent word-initially.[2] Word-internally, implosivity is weaker, but provided the stop is released, there is still a long lag VOT with prevoicing. In very fast speech, intervocalic voiced stops are often lenited to fricatives or glides. This can override otherwise robust phonotactic restrictions against certain sequences, e.g. [m̩gosi m̩butu**wu**tu] ‘fat man’ (which features a typically banned [wu] sequence).

There do not seem to be any strong distributional restrictions on voiced and voiceless oral stops by syllable position, other than the tendency for implosivity to be reduced. However, as stated above, some stops have predictable distributions based on the follo-

[1]In the forthcoming description, I avoid using examples from the BLP dictionary. While their corpus is considerably larger than my own, the dictionary entries are rendered orthographically using a modified form of the standard Swahili orthography, with much phonetic detail omitted. As I have very little access to their recordings and did not work with their speakers, I am more comfortable making judgments based on data collected by myself or others who worked with my primary consultant.

[2]Entries in Dayley et al. 2020 suggest that there is some variation between pulmonic and non-pulmonic stops in loanwords; for example, it lists both *diksheneri* and *dhiksheneri* for ‘dictionary’ (where <dh> indicates [ɗ]). My consultant didn’t exhibit this type of variation.

wing vowel. This is most robustly observed with the coronal stops; Mushunguli has both dental and retroflex coronal stops.[3] For the voiced oral stops, there is no strong evidence to suggest that there is a contrast; retroflex stops occur before back vowels /o/ and /u/ (e.g. [ᶑuᶑu] 'bug'; [m̩ᶑoᶑo] 'small (cl 1)'), while alveolar stops occur elsewhere ([ɗege] 'bird'; [idi] 'two (cl 9)'). There do not appear to be any cases of non-prenasalized [ʈ] (e.g. [tʃitu^{ŋ}gulu] 'onion'), but prenasalized coronal stops are generally retroflex in all contexts.

Across descriptions, the status of the coronal stops is controversial. Tse (2014, 2015a,b) claims that there exist cases of [^{n}d] in loans (e.g. [^{n}doni] 'boat'). However, where these loans exist in the data I have collected, they are definitely retroflex (e.g. [^{ɳ}ɖoni] 'boat'). Michal Temkin Martinez (p.c.) has suggested to me that retroflexes are actually [tr] and [dr] clusters (which is how they are generally written orthographically). It is the case that retroflex stops have an impressionistically rhotic quality, especially word-initially ones. However, I agree with Tse—who worked in part with publicly-available data collected from my primary consultant—that these are single segments. Again, consonant clusters are rare to the point of potentially being unattested outside of borrowings from non-Bantu sources, and most words containing retroflex stops in the corpus I have collected do not stem from likely non-Bantu sources (Arabic, English, Italian, or Somali), e.g. [$^{ɳ̥}$ʈ^{h}e^{m}bo] 'elephant'; [$^{ɳ̥}$ʈ^{h}o^{ɳ}ɖo] 'star'; [ᶑole] 'finger'). There is also no evidence in the data I have collected to suggest that retroflex [ᶑ] and [ɗ] are distinct in a way that would suggest that one is actually two segments; both are implosive, and roughly the same length (Figure 2.2). However, given that retroflexion does not seem to be contrastive in this language, and given that my understanding is that the Boise Language Project consultants are both younger and more proficient in English than my primary consultant, it may be the case that they have reanalyzed or are in the process of reanalyzing the

[3] While I am using a narrower transcription for this section, for readability's sake, I transcribe dentals without a diacritic.

retroflexes as clusters.[4]

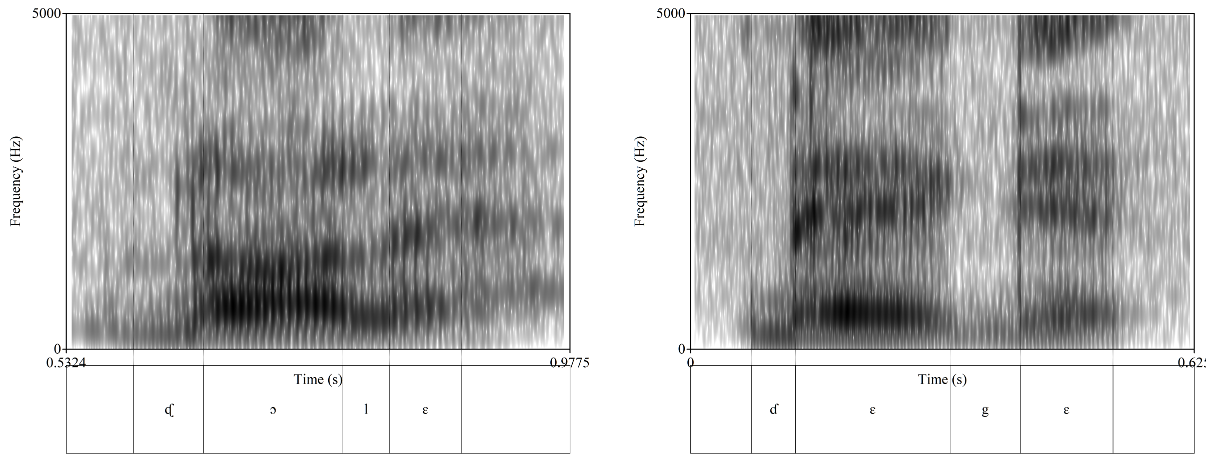

(a) [ɖɔlɛ] ‘finger’

(b) [dɛgɛ] ‘bird’

Figure 2.2: Comparison of retroflex (a) and non-retroflex (b) voiced coronal stops (with narrowly transcribed vowels)

There is considerably less controversy outside of the coronal stops. The voiceless counterpart of the palatal stop [ʄ] was realized as an affricate [tʃ] for my speaker. This segment is much more common before front vowels (especially [i]), but because it can occur in morphologically complex contexts preceding any vowel (e.g. *tʃ-aŋgu* ‘my (cl 7)’; *wa-tʃ-ogoh-eð-a* ‘they scared us’), I hesitate to call this a genuine distributional *restriction*.

/k/ and /$^{ŋ̊}$k/ are sometimes realized as [q] and [$^{N̥}$q], respectively, before back vowels (e.g. *m̩kono* ‘arm’ is more narrowly transcribed as [m̩qono]). However, there is otherwise no evidence of an existing contrast between these sounds in my data, and this pattern is impressionistically weaker and less robust than the aforementioned controversy surrounding the coronal stops.

Mushunguli also has both voiced and voiceless prenasalized for all places of articulation, though I have no examples of a voiceless counterpart of [^{ɲ}ɟ] (i.e. no *[$^{ɲ̊}$tʃ]).[5]

[4]The only BLP recording I have access to does have an example of word-medial [^{n}d] before [e]; however it unfortunately does not have any examples of coronal stops before back vowels.

[5]Dave Odden’s data has an example transcribed [nchíndu] ‘red colobus’ agreeing in class 9/10–

The nasal portion of voiced prenasalized stops is usually audible and agrees with the oral portion in place. The nasal portion of word-initial voiceless prenasalized stops is typically completely voiceless and is sometimes inaudible (especially in isolation and word-initially); in these cases, the primary audible acoustic cue is strong aspiration. A phonetic study by Temkin Martinez & Rosenbaum (2017) has verified, however, that nasal airflow does persist throughout the closure of voiceless prenasalized stops. Word-internal voiceless prenasalized stops generally have an at least partially-audible nasal portion which, again, agrees with the place of the following consonant. The aspiration is also sometimes reduced, though it is never lost entirely.

Demonstrably root-initial prenasalized stops are much rarer than word-initial or root-internal ones; the majority of word-initial prenasalized stops occur in classes 9/10, which are marked by the concatenation of a nasal prefix (/ɲ-/). Cases of truly root-initial prenasalized stops (e.g. *-ᶯ̊ʈu* 'person') are rare and do not ever actually occur word-initially. The distribution of voiceless prenasalized stops is also slightly more restricted than their voiced counterparts. All voiced prenasalized stops can occur root-internally or word-initially. However, for the voiceless prenasalized stops, only root- or word-internal [ᵑ̊kʰ] is robustly attested, with [ᶯ̊ʈʰ] occurring rarely.

Examples of all stops in word-initial and word-internal contexts are given in Table 2.1.

Because retroflexion, implosivity, and aspiration are predictable, non-contrastive, and do not bear on the relevant patterns in the analysis of Mushunguli hiatus resolution, they will not be reflected in transcriptions outside of this section. This is largely for the sake of readability. For a similar reason, I will not be superscripting the nasal portion of prenasalized stops after this section, though voicelessness and syllabicity will still be diacritically marked. These choices are summarized below in Table 2.2.

this orthographic transcription likely corresponds to [ᶮ̊tʃíᶯɖu], but I don't have a recording of this word to verify.

Table 2.1: Word-initial and word-medial stops in Mushunguli

	Word-initial		Word-internal	
[p]	pi:ka	'cook!'	tope	'mud'
[$^{\text{m̥}}$p$^{\text{h}}$]	$^{\text{m̥}}$p$^{\text{h}}$ula	'nose'	Unattested	
[ɓ]	ɓu$^{\text{ɳ}}$ɖiki	'gun'	kuhisaba	'to count'
[$^{\text{m}}$b]	$^{\text{m}}$buguni	'ostrich'	utu$^{\text{m}}$bo	'intestine'
[t]	taɠi	'egg'	$^{\text{ŋ}}$goto	'sheep'
[$^{\text{n̥}}$t̪$^{\text{h}}$]	$^{\text{n̥}}$t̪$^{\text{h}}$a$^{\text{ŋ}}$gulu	'basket'	ba$^{\text{n̥}}$t̪$^{\text{h}}$i	'door'
[ɗ]	ɗi:g$^{\text{w}}$a	'cl 5 fell'	we:ɗi	'good (cl 3)'
[ᶑ]	ᶑuᶑu ɗam̩ti	'caterpillar'	kuᶑumu:la	'to chop'
[$^{\text{ɳ}}$ɖ]	$^{\text{ɳ}}$ɖoni	'boat'	maqoɳɖe	'fists, claws'
[tʃ]	tʃirevu	'chest'	k$^{\text{w}}$erek$^{\text{w}}$etʃe	'francolin'
[ʄ]	ʄe$^{\text{m}}$be	'hoe'	luɟe$^{\text{n}}$do	'chameleon'
[$^{\text{ɲ}}$ɟ]	$^{\text{ɲ}}$ɟina	'louse'	ma$^{\text{ɲ}}$ɟano	'yellow'
[k]	kazana	'baby'	ɲoka	'snake'
[$^{\text{ŋ̊}}$k$^{\text{h}}$]	$^{\text{ŋ̊}}$k$^{\text{h}}$a$^{\text{ɳ}}$ɖe	'food'	k$^{\text{w}}$i:$^{\text{ŋ̊}}$k$^{\text{h}}$a	'to give'
[q]	quwi	'turtle'	m̩qono	'arm'
[$^{\text{ŋ̊}}$q$^{\text{h}}$]	$^{\text{ŋ̊}}$q$^{\text{h}}$oɠa	'eagle'	vibo$^{\text{ŋ̊}}$q$^{\text{h}}$o	'hippos'
[ɠ]	ɠasa	'hand'	$^{\text{n}}$duɠu	'sibling'
[$^{\text{ŋ}}$g]	$^{\text{ŋ}}$guku	'chicken'	m̩na$^{\text{ŋ}}$go	'doorway'

Table 2.2: Transcription conventions for stops

Segment	**Transcription**
[ɓ, ᶑ/ɗ, ʄ, ɠ]	<b, d, ɟ, g>
[$^{\text{m}}$b, $^{\text{ɳ}}$ɖ, $^{\text{ɲ}}$ɟ, $^{\text{ŋ}}$g]	<mb, nd, nɟ, ŋg>
[p, t, tʃ, k/q]	<p, t, tʃ, k>
[$^{\text{m̥}}$p$^{\text{h}}$, $^{\text{n̥}}$t̪$^{\text{h}}$, $^{\text{ŋ̊}}$k$^{\text{h}}$/$^{\text{ŋ̊}}$q$^{\text{h}}$]	<m̥p, n̥t, ŋ̊k>

2.2.1.2 Fricatives

Mushunguli has nine fricatives: [f, v, s, z, ʃ, ð, x, ɣ, h]. Voiced fricatives are often weakly voiced relative to voiced stops. As with stops, the coronal fricatives tend to be dental; I have not observed the same tendency for retroflexion before back vowels that occurs with the coronal stops. Note that there are no prenasalized fricatives; the only case of a non-syllabic nasal preceding a fricative that I have observed is the word [hanʃi] 'paper.'

Some of these fricatives have a restricted distribution. In general, the most robustly-attested fricatives in both word-initial and word-internal positions are [v], [s], [z∼ð], and [h]. Word-internal [f] is common, but almost always occurs in a morphologically complex context (i.e. root-initial [f]).

The velar fricatives [x] and [ɣ] are very rare; I only have examples of word-initial [x], which exclusively occur in loanwords (e.g. *xamisi* 'Thursday'), while all cases of [ɣ] in the corpus surface variably with [ɡ] (especially [ɡ^{w}]). I also have collected relatively few examples of [ʃ], though it does occur word-internally with some frequency.

Finally, there is free variation between [z] and [ð], for example *m̩teza∼m̩teða* 'peanut.' There are cases of /z/ that do not exhibit this variation, but they are rare, e.g. *kukazi^{ŋ}ga* 'to fry'. I have never observed an analogous pattern of variation between [s] and [θ].[6]

Examples of fricatives are given in Table 2.3.

Table 2.3: Word-initial and word-internal fricatives in Mushunguli

	Word-initial		Word-internal	
[f]	fu^{ŋ}go	'civet'	folofota	'lung'
[v]	vu^{ɳ}ɖe	'cloud'	m̩ɓavi	'thief'
[s]	suwi	'leopard'	$^{m̥}$p^{h}asa	'twin banana'
[z∼ð]	ziso	'eye'	koː$^{ŋ̊}$k^{h}eða	'to suckle'
[ʃ]	ʃawaka	'bednet'	k^{w}aːʃa	'to burn'
[x]	xamisi	'Thursday'	Unattested	
[ɣ]	Unattested		haraɣ^{w}e	'bean sp.'
[h]	heːtʃitabu	'to the book'	lalahi	'fish'

I do not adopt any orthographic shortcuts for fricatives in forthcoming transcripti-

[6] Sharon Rose (p.c.) suggests that this may be due to influence from either Arabic or Somali, noting that MSA ð has become [z] in some dialects, and that Somali has no [z], but does have [ð] as an intervocalic variant of /d/; the latter is true for Maay as well (Paster 2006). I think these are both reasonable hypotheses, but I do not have enough carefully-collected data illustrating when the variation does and does not occur to verify them. Michal Temkin Martinez (p.c.) has told me that her consultant(s) did not exhibit this variation, but it is robustly attested in data from my speaker.

ons, save that I will be transcribing *z*∼*ð* as <z>.

2.2.1.3 Nasals

Mushunguli has four nasals, [m, n, ɲ, ŋ]. I have observed no apparent distributional restrictions on these nasals with respect to following vowels, save for [m], which only occurs rarely before [u]. That said, the distribution of the nasals is uneven. [ŋ] is much rarer than the other nasals in all positions; it is far more common to surface as the nasal portion of a prenasalized velar stop, or as a syllabic nasal generated by copula reduction (see below).

Examples of non-syllabic nasals are given in Table 2.4.

Table 2.4: Word-initial and word-internal non-syllabic nasals in Mushunguli

	Word-initial		Word-internal	
[m]	mavuha	‘bones’	ɲimi	‘tongues’
	meːno	‘teeth’	m̩nomo	‘mouth’
	muɲu	‘salt’	mame	‘mother’
[n]	naːɟa	‘I am eating’	m̩qono	‘arm’
	nimoːto	‘it’s a fire’	nene	‘nine’
	noːgeːra	‘I am going far’	k^{w}inula	‘to take a pot off of the fire’
[ɲ]	ɲama	‘meat’	k^{w}ivaɲa	‘to hear together’
	ɲu^{m}ba	‘house’	weːɲu	‘your (pl) (cl 3)
	ɲi^{ŋ}gi	‘many’	naniwasiːɲe	‘I will look at them’[D]
[ŋ]	ŋo^{m}be	‘cow’	kuŋala	‘(be) white’
	ŋoŋo^{ɳ}ɖo	‘weaver bird’[D]		

Mushunguli also has cases of morphologically-derived syllabic nasals, most commonly /m̩/, which is the surface allomorph of the class 1 and class 3 noun class prefixes, as well as the second person plural subject and object prefix. The copula /ni-/ can also be reduced to a syllabic [n̩], and this reduced copula optionally harmonizes with following segments (e.g. [wata] ‘duck’, [ŋ̍wata] ‘it’s a duck’[D]; [$^{m̥}$p^{h}a^{ŋ}ga] ‘machete’, [m̩pa^{ŋ}ga] ‘it’s a machete’).

Nasals are the only segments that may potentially occur in codas in native forms, and only as a result of concatenation of a prefix, such as the augment, before a syllabic nasal prefix, e.g. [m̩ti] ‘tree,’ but [u:mti] ‘the tree’. Note that voiceless stops following syllabic nasals (whether part of a coda or on their own) are not aspirated.

Except for very rare exceptional cases, there are no cases of non-syllabic nasals preceding any segment besides a vowel. In morphophonological contexts where this could potentially occur, such as classes 9 and 10 (see §2.3.1), no nasal surfaces, e.g. /ɲ-suwi/ → [suwi] ‘leopard.’ Syllabic nasals can and do regularly surface before all segments, save for [l]; *no* nasals, syllabic or otherwise, can precede [l], as evidenced by the singular/plural pair *m̩neᵑge/mileᵑge* ‘moon/moons’.

To distinguish syllabic and non-syllabic nasals, and especially to distinguish syllabic nasals from prenasalized stops, I will mark syllabicity diacritically, e.g. <m̩, n̩>. This allows for distinctions between a voiceless prenasalized labial stop <m̥p> (e.g. ***m̥p**aŋga* ‘machete’), a voiced prenasalized stop <mb> (e.g. ***mb**ala* ‘antelope sp’), and a syllabic nasal preceding non-prenasalized labial stops <m̩p, m̩b> (e.g. *m̩pira* ‘ball’; *m̩bavi* ‘thief’).

2.2.1.4 Liquids

The liquids are one of the few sets of consonants that show strong distributional restrictions or allophonic relationships. Mushunguli has two liquids: the lateral [l] and the trill [r], the latter of which is often tapped in fast speech. Neither sound is common word- or root-initially, though both do occur ([lalahi] ‘fish’; [roti] ‘bread’; [si-laga:la] ‘I fell’, [tʃi-revu] ‘chin’). In root-internal contexts, however, these sounds are in complementary distribution: [r] only occurs following front vowels, while [l] can be preceded by back vowels and [a].

There are some rare words that do not conform to this generalization, such as *moroti* ‘millet’ and *ᵑgurumo* ‘pig’, the latter of which variably surfaces as [ᵑgulumo]. I do

Table 2.5: Distribution of [l] and [r]

[l]		[r]	
^{m}bala	‘bushbuck’	m̩pira	‘ball’
m̩tab^{ɣ}ali	‘hawk’	^{m}bigiri	‘giraffe sp.’
$^{n̪}$t̪^{h}a^{ŋ}gulu	‘basket’	b^{ɣ}irib^{ɣ}iri	‘spider’
kulu	‘big’	$^{m̥}$p^{h}era	‘warthog’
^{ŋ}gola	‘knife’	k^{w}ereka	‘to give birth’
$^{ŋ̊}$k^{h}olo^{ŋ}go	‘lake’	m̩vere	‘woman’

not have any cases of root-internal [l] being similarly preceded by a front vowel, though root-initial [l] can be if that front vowel belongs to an agreement prefix, e.g. *milomo* ‘mouths’. As stated above, root-initial [l] also cannot be preceded by a nasal; if such a context arises, nasal assimilation occurs, e.g. [m̩**n**omo] ‘mouth’. Note that there are roots that genuinely begin with /n/, e.g. [m̩nu^{ŋ}gu]/[minu^{ŋ}gu] ‘god/gods’. I have no examples of root-initial [r] following a nasal, syllabic or otherwise.

Despite the apparent allophonic relationship between [l] and [r], because I have no evidence of a morphophonological alternation, I will be transcribing them as individual segments, even in underlying representations.

2.2.1.5 Approximants

Mushunguli has two glides, /j/ and /w/. Tautosyllabic glide-vowel sequences that are homorganic for place and height appear to be entirely forbidden (i.e. no *ji* or *wu*). *Heterosyllabic* sequences *i.j* and *u.w* are tolerated (e.g. *ju**w**e* ‘stone,’ *ihi**j**e* ‘bad’).

The labiovelar glide /w/ can also follow consonants; however, when this occurs, it is generally realized as a secondary articulation and not as a full glide. This secondary articulation is predictably labial if the preceding consonant is not labial, e.g. [m̩s^{w}a] ‘mosquito’; otherwise, it is velar, e.g. [m^{ɣ}aːna] ‘child.’ As I discuss in §5.1.2, I generally assume that root-internal secondary articulations (as in ‘mosquito’) are underlying, while secondary articulations at edges (as in ‘child,’ which is a morphologically complex form

/mu + ana/) stem from alternations. Regardless, in forthcoming transcriptions, I render all “labial” secondary articulations as <ʷ>, unless otherwise indicated. Palatal secondary articulations are generally disallowed (see §3.5, but cf. §5.1.3); in the exceedingly rare cases where they do arise, it is always at a morpheme edge (e.g. [kogofʲa] ‘to frighten’).

There is an optional process of intervocalic glide deletion that mainly affects underlying glides. This occurs much more frequently with /w/ than /j/, especially in /a + a/ contexts. Glide deletion results in unresolved hiatus, which has varying results in terms of vowel length (see below).

Glide deletion can sometimes also apply in phrase-level (postlexical) contexts. For example, /mi + teza # i + etu/ ‘our peanuts’ can surface as either [miteza jeːtu] or [miteza eːtu].[7] This does not apply within words; for example, the 2SG subject prefix /u-/ is often in hiatus due to its concatenation with TAM morphology, but in cases where glide formation applies, glide deletion will not. Thus, /na + u + oger + e/ ‘you will swim’ can only surface as [nawoːgeːre], never *[na.oːgeːre].

2.2.2 Vowels

Mushunguli has a five vowel system, /a, i, u, e, o/. All non-low vowels are typically pronounced rather lax, and this is particularly true of the mid vowels, which are more narrowly transcribed as [ɛ] and [ɔ]. I adopt the <a, i, u, e, o> convention here for orthographic convenience. /a/ is phonetically central and unround, but it can be influenced by the place of a following consonant, becoming more [æ]-like or more [ɑ]-like, though never reaching anything resembling a canonical pronunciation of either vowel.

It is much easier to find true and near-minimal pairs for vowels than it is for con-

[7] In all cases that I have collected with an underlying sequence like this (/...a # i + V/), even if the glide does not surface, there is an audible distinction between the [a] and the following vowel; that is, postlexical coalescence from /a # i/ does not apply here.

sonants. Examples of (near-) minimal pair contrasts are given in Table 2.6.

Table 2.6: Vowel (near-) minimal pairs

m̩ti	'tree'	$^{ŋ̊}$k^{h}u^{ɳ}ɖe	'cow peas'
m̩to	'river'	$^{ŋ̊}$k^{h}a^{ɳ}ɖe	'food'
boko	'banana'	$^{ŋ̊}$k^{h}o^{ɳ}ɖe	'farm'
buku	'book'	$^{ŋ̊}$k^{h}o^{ɳ}ɖo	'war'
m̩bago	'forest'	$^{ŋ̊}$k^{h}ola	'snail'
$^{m̥}$bogo	'buffalo; wildebeest'	$^{ŋ̊}$k^{h}ala	'crab'
$^{m̥}$p^{h}era	'warthog sp.'	ɲama	'meat'
$^{m̥}$p^{h}ula	'nose'	ɲimi	'tongues'
m̩pira	'ball'	-itika	'answer'
$^{m̥}$piri	'adder'	-ituka	'be startled'
me:zi	'months'	-edi	'good'
mazi	'water'	-idi	'two (cl 10)'

2.2.2.1 Frontness and roundness

All front vowels in Mushunguli are unround, and all back vowels are round. Front and back vowels often pattern together in roots, but there are plenty of examples that contravene this generalization. Some of these are likely loans, e.g. [kartoni] 'box' (from Italian *cartoni* or English *carton*), [tʃiko^{m}be] 'cup' (from Maay *koombe*). However, others are clearly Bantu in origin, e.g. [^{ɳ}ɖevu] 'beard'). Moreover, unlike the height cooccurrence restriction discussed below, verbal extensions do not agree with root vowels for frontness/roundness (e.g. [kuzi:sa] 'to ask a lot', cf. [u:za] 'ask!'; more examples in Table 2.7).

Given the relatively small size of the corpus, I hesitate to call this a *restriction* so much as a *tendency*.[8]

[8] Looking just at polysyllabic noun roots containing ≥2 non-central vowels (n = 132), ∼58% (77) agree for frontness/roundness.

Table 2.7: Cooccurrence of front and back vowels

ɓoⁿɖe	'valley'	m̩gosi	'man'
ⁿɖevu	'beard'	ɠutʷi	'ear'
tʃisigino	'elbow'	vuⁿɖe	'cloud'ᴰ
ɠome	'nutshell'	kogohe:za	'to scare'
ᵐbuli	'word'	su:ᵐbika	'I piled things up'
makoⁿɖe	'fists; claws'	suzi:sa	'I asked a lot'
suwi	'leopard'	kʷiguta	'to be satiated'
m̩sukule	'zombie'	ko:ᵑgeza	'to add'

2.2.2.2 Height and height harmony

There are three tiers of height in Mushunguli, evidenced primarily by the fact that high, mid, and low vowels behave differently in hiatus contexts (see §2.2.4). There is also a general cooccurrence restriction between high and mid vowels in roots; this is more apparent in verbs than in nouns, but the tendency for height matching in nouns is higher than the tendency for back/round matching.[9] Low vowels are unrestricted in their distribution and so may pattern with either freely.

Like many other Bantu languages, Mushunguli exhibits limited height harmony in extensional suffixes containing front vowels. These suffixes surface as [e] when the final root vowel is mid; otherwise, they surface as [i]. Examples are given in (3).[10]

(3) Height harmony

[9] Looking again at the same set of 132 polysyllabic noun roots with ≥ 2 non-low vowels, ~71% (94) conform to the height restriction. This is lower than a similar count of the BLP dictionary (Dayley et al. 2020:11/06/2019), which only looked at polysyllabic nouns beginning with <n> (excluding proper names and periphrastic compounds). This set was chosen because the organization of the dictionary at the time made it difficult to divide nouns from other parts of speech and especially to distinguish the first vowel of a root from the first vowel of a word (i.e. a prefix vowel). <n>-initial entries include all overtly-marked nouns belonging to class 9/10 (save for those beginning with a labial nasal or prenasalized labial stop) and naturally excludes verbs. There were 210 such entries, and height matching was observed in ~79% (166) of these.

[10] I collected relatively few examples of verbs with extensions, so many of these examples are taken from Odden's transcriptions. These transcriptions were rendered orthographically (similar to the orthography adopted by the BLP); I have translated the orthography into IPA, but I do not have these recordings, so there may be errors.

a. tʃi-l**o**ᵑg-**e**ːs-a ‘we spoke a lot/too much’[D]
b. kʷ-**i**ːv-**i**ːs-a ‘to hear a lot/too much’
c. wa-his**ab**-**i**ːsa ‘they counted a lot/too much’[D]

Suffixes containing /a/, such as the reciprocal /-an/ and the final vowel/indicative mood marker /-a/, do not participate in height harmony. The /-e/ allomorph of the final vowel also does not alternate in any context. Both are illustrated in (4).

(4) Failure of harmony to apply to /a/ and /e/ in stems

a. tʃ-**o**ːᵑ**g**e**z**-**a**ːn-**a** ‘we are adding each other’[D]
b. tʃ-**o**ːᵑ**g**e**z**-**ez**-**an**-**a** ‘we are adding for e.o.’[D]
c. kʷ-**a**ːⁿɖ**am**-**ir**-**a**ːn-**a** ‘to lean back together’
d. na-n-**it**-**i**ːs-**e** ‘I will go a lot/too much’
f. na-ð**a**ːm-**e** ‘he will be heavy’[D]

It is not possible to determine whether low vowels actually *block* harmony or simply do not participate; this is because in my corpus, there are no examples of verb roots where a mid vowel precedes a low vowel, and derivational extensions containing /a/ (which do not alternate) are ordered after those containing non-low vowels (which do).

2.2.2.3 Length

Vowel length is not contrastive, but surface long vowels do arise in some contexts. Most forms of hiatus resolution regularly result in a long vowel due to compensatory lengthening. In elicited speech, this is very easy to detect, while in naturalistic speech, it is less pronounced but still measurable. The only case where this is not true is the resolution of sequences of identical vowels, which always results in a single *short* vowel, e.g. [sitiːka] ‘I answered’ (*[siːtiːka]). This is general, not exceptional, and will be discussed in more detail in §3.4.

A few examples of hiatus resolution resulting (or failing to result) in compensatory lengthening are given in (5); justification for the underlying forms, as well as many more examples, will appear in §2.3-§2.2.4 and all chapters thereafter.

(5) Compensatory lengthening

a.	/ka + itik + a/	k**eː**tiːka	'(s)he answered'
b.	/ku + itik + a/	k^{w}**iː**tiːka	'to answer'
c.	/a + a + ereker + a/	**eː**rekeːra	'(s)he is floating'
d.	/i + asam + a/	j**aː**saːma	'it (cl 9) opened its mouth wide'
e.	/si + omol + a/	soːmoːla	'I dished up ugali'

There is a general restriction against the final syllable of an utterance being long, and so compensatory lengthening will fail to apply if word-internal hiatus occurs at the juncture of the penultimate and final syllables of an utterance. This is a difficult context to elicit; however, this tendency can be seen with the proximal form of 'that', which is a morphologically complex form /AUG + DEM + o/ that always results in hiatus, e.g. /buku # i + di + o/ → [buku iːɟo] 'that (prox) book' (*[buku iːɟoː]).[11]

The lengthening of the penultimate syllable here is due to the fact that this phrase was elicited in isolation: the penultimate syllable of any utterance is also lengthened, as illustrated by the examples in (6). This is usually, but not always, associated with a high pitch target; while penultimate lengthening is quite consistent, tone and intonation are not.

(6) Penultimate lengthening

a.	sidum**uː**la	'I chopped'
b.	sikemoː m̩p**iː**ra	'I threw the ball'
c.	sikemeː mipira k^{w}aːmaj**oː**^{ɳ}ɖa	'I threw the balls to the baboons'

Note that nouns elicited in isolation only variably exhibited this lengthening (a possible effect of list intonation), but as the above examples illustrate, nouns can be lengthened in this context. I do not have any examples of multi-clausal sentences in the data I collected; there may be more contexts in which this type of lengthening can occur.

Because distinguishing between unresolved hiatus and long vowels is of critical importance to this dissertation, long vowels are transcribed here with a <ː> (e.g. [aː]), regardless of source, while vowels in hiatus will be transcribed as <V.V>.

[11]The alternation between /di/ and /ɟ/ here is discussed in §5.1.3.

2.2.3 Syllable types

There are four permissible syllable types in *native* forms in Mushunguli: /CV/, /V/, /N̩/, and (rarely) /(C)VN/.

Table 2.8: Syllable types in native forms (non-exceptional)

/CV/	weː.nuː.ka	'they stood up'
	mˠaː.na	child
	ᵑgo.to	sheep
	jʷeː.di	'good (cl 1)'
/V/	i.zi	'voice'
	iː.ta	'go!'
	a.ma.jo.ⁿɖa	'the baboons'
	bu.ku. iː.ɟo	'that (prox) book'
	si.ᵑ̊kʰa. im.bʷa. ma.vu.ha	'I gave the dogs bones'
/N̩/	m̩.bˠa	'dog'
	m̩.pa.ᵑga	'it's a machete'
	m̩.ku.lu	'big (cl 1/3)'
	ŋ̍.wa.ta	'it's a duck'
/(C)VN/	uːm.to	'the river'
	iːm.bˠa	'the dogs'
	sim.to.aː.ni	'I beat you pl.'

Of these, only /CV/ syllables are unrestricted. The initial consonant of these syllables may be singly-articulated or may involve a secondary articulation, most commonly prenasalization, labialization, or velarization.

/V/-type syllables (= vowel hiatus) are also broadly or optionally permissible in some contexts, but outside of exceptional cases, these generally only occur at the left edge of a word or utterance. Some V-initial noun roots have no overt noun class prefix (class 5), and surface with an onsetless syllable, as exemplified by *izi* 'voice' in Table 2.8. Similarly, imperative verbs are usually unmarked, meaning that V-initial verbs in the imperative will surface with a /V/ syllable (exemplified by *iːta* 'go!'). Finally, as discussed in §2.3, there are many agreement prefixes of the shape /V-/ or /VCV-/ that can (or obligatorily) occur at the beginning of words (exemplified by *amajo*^{n}*ɖa* ‘the baboons’).

When /V/-type syllables occur utterance-initially, the resulting onsetless syllable is universally tolerated. However, a *word*-initial /V/-type syllable is not guaranteed to surface if it is not also in utterance-initial position. Hiatus is only universally tolerated between words in the context /$V_{[+high]}$ # V/, iff the second vowel is not identical to the first. For /a # V/ contexts, hiatus is optionally repaired. This occurs more often in fast speech, and results in a long vowel, e.g. [sikemoː m̩p**iː**ra] ‘I threw the ball’ (from /si + kem + a # u + mu + pira/). For sequences of identical vowels, long or short vowels can obtain. It is unclear whether the former should be viewed as a long vowel (i.e. a /CVː/ syllable) or two vowels in hiatus (a /CV/ syllable followed by a /V/ syllable). I leave the question open for now.[12]

Syllabic nasals, as discussed in §2.2.1.3, are allowed word-initially. These are often reduced in fast speech. /(C)VN/ syllables are only permited when a V-final prefix precedes what would otherwise be expected to surface as a syllabic nasal. Prefixes that can be affected by this include the 3SG (= class 1) and 2PL object prefixes, which are always preceded by a subject agreement and/or TAM markers. It also includes noun class prefixes for classes 1 and 3 when preceded by the augment (pre-prefix), as well as the class 9/10 noun prefixes in the same context *iff* they precede a monomoraic root.

From an analytical standpoint, it is somewhat unclear whether the nasals in /(C)VN/ sequences should truly be viewed as nasal codas, or as syllabic nasals. I have no examples of native or borrowed words ending in any consonant word-finally. Moreover, the syllabic nasal prefixes are a part of the small subset of morphemes that consistently carry tone. While these prefixes are definitely shorter in word-internal contexts than word-initially, in many tokens, pitch drops or rises noticeably during the nasal. Similarly, while in some tokens the vowel is as long as the nasal, in others the

[12] Given that hiatus is only optionally resolved between words, treating this as two vowels in hiatus is more consistent with the analysis presented in §3.4.

vowel is quite short and weak compared to the nasal. However, even non-syllabic onset nasals (especially [m]) are often as long as or longer than short vowels in this language. Due to inconsistency across tokens, I leave this question open for future research.

There are also marginal syllable types and distributions in Mushunguli that only occur in borrowings or identifiable exceptions. With respect to syllable types, consonant clusters not involving nasals can appear in borrowings (e.g. *kartoni* 'box'[D], *kaskazini* 'north'[B], *skulu*[B]).[13] Again, I have no examples of word-final consonants.

With respect to distributions, as the introduction of this dissertation as well as some prior discussion has already alluded to, it is possible to get /V/-type syllables in exceptional and optional contexts. Some examples are given in (7).

(7) Marginal V-type syllables

a.	ka.u.juːsa	'(s)he revived/transplanted'
b.	ku.o.neː.sa	'to see a lot/too much'
c.	tʃa.i	'tea'
d.	waː.o	'their (cl 2)'
e.	si.wa.ka.uː.la	'I searched for 3PL'
f.	si.i.hisaːba ([siːhisaːba])	'I counted them (cl 4)'
g.	kutoː.a	'to beat'
h.	kʷeː.iː.ga	'to pretend'
i.	m̩ɲa.u∼m̩.ɲa.wu	'cat'
j.	amajonda∼ama.oⁿɖa	'the baboons'
k.	iː.o∼iːjo	those (cl4)/that (cl9) (prox.)

Examples (7a,b) feature unresolved hiatus at a prefix-stem boundary; specifically a non-coalescing stem (discussed in §4.2) and a total non-participant (discussed §5.2). These exceptions conform to the criteria established in §1.3 and are of primary interest to this dissertation.

Examples (7c-e) are cases of root-internal hiatus. These types of words are uncommon, and tend to occur in loans or high-frequency words. I do not analyze cases of root-internal phonotactic violations in this dissertation, but it is a point of future research.

[13] Note that I do not have any examples of these in accessible recordings, save for [kaskazini] (from a BLP recording). In this latter case, there is no evidence of a repair of this syllable.

Examples (7f-h) feature potentially unresolved hiatus in possibly exceptional or regular contexts; that is, it is unclear whether this behavior can be described as *arbitrary* or not. Example (7e) features an object prefix as V2 in a hiatus context (/si + i/); these cases can result in surface long vowels (cf. /si + hisab + a/ → *sihisa:ba* 'I counted'). This is a behavior specific to object prefixes; phonologically similar but morphologically dissimilar underlying contexts, such as an /i/-final subject prefix concatenated with an /i/-initial stem, surface with a short vowel, e.g. /si + ita^{ŋ}g + a/ → [sita:^{ŋ}ga].

This behavior could be exceptional; it could also be attributable to a number of other factors, such as morphosyntactic position, prosodic domain, discourse contexts, or the retention of lexical tone. Because I have very few examples of object marking of this type in this context, I leave the question open for future research.

Examples (7g,h) are the only cases I have observed of mid vowels in a potential V1 position. In the first case, which features the verb *toa* 'beat', hiatus is not resolved between the root edge and the final vowel. The second case involves the reflexive prefix /e-/, hiatus resolution *does* apply in the context in which /e/ is V2 (else we would expect *[ku.e.i:ga]), but not where it is V1. This behavior is observed in the imperative as well, e.g. [e.i:ga] 'pretend!'. It may be the case that this behavior is generalizable to mid vowels in V1 position (in which case it is not exceptional), or it may be the case that these morphemes exceptionally have strong right edges. Barring further evidence, I again leave this question open.

Finally, examples (7i-k) feature examples of optional intervocalic glide deletion mentioned in §2.2.1.5. This is variable and is more likely to occur in fast speech.

2.2.4 Hiatus resolution

As alluded to throughout the preceding discussion, the operations of morphology in Mushunguli result in a large number of underlying contexts that, if unrepaired, would result in a surface onsetless syllable, known as *hiatus*. Underlying hiatus can occur in almost any morphophonological context, save possibly suffix-suffix boundaries. However, the only position in which hiatus is uniformly tolerated is word-initially, and then only if the word itself is utterance-initial.[14] Examples of *regular* underlying hiatus contexts (and their surface results) are given in Table 2.9.[15]

Table 2.9: Underlying hiatus contexts

Word-initial	/**i** + puluk + a/	**i**pulu:ka	'it (cl 9) flew'
Prefix-prefix	/k**u** + **e** + jag + a/	k^{w}**e:**ja:ga	'to scratch oneself'
Prefix-root	/m**u** + **o**jo/	m**o:**jo	'heart'
Root-extension	/ku + d**i** + **i**s + a/	kud**i:**sa	'to eat a lot/too much'
Root-FV	/si + d**i** + **a**/	si**:ɟa**	'I ate'
Between words (optional)	/si + to + **a** # **i**zi + ɲ + b^{w}a/	sito**e:** zimb^{w}a	'I beat the dogs'

As evidenced by Table 2.9, underlying hiatus at morpheme boundaries is generally not tolerated. Mushunguli is typologically interesting in that it exhibits an unusually diverse set of hiatus-related alternations that are largely blind to morphosyntactic position or semantic class; that is, the same set of hiatus repairs apply across the board. While it is not uncommon for languages to have more than one hiatus repair, it is the case that most languages usually have maximally two strategies in any given morphophonological context (Rosenthall 1997; Casali 1996, 1997).[16]

Mushunguli can be described as having four classes of repairs that apply in eight different morphophonological contexts. These are asymmetric height coalescence, which merges a low V1 and high V2 into a single mid vowel corresponding to the place of V2

[14]Hiatus within roots is sometimes tolerated (e.g. *m̩ti wa:.o* 'their tree'), but this is rare.

[15]Justification for proposed underlying representations in Table 2.9 is given in §2.3. Note also that the root-FV example featuring the verb *eat* (/di-/) is one of two morphemes exhibiting an exceptional repair (§5.1.3). V-final roots are rare, so I was forced to use the exceptional example in this case.

[16]See §3 for additional discussion on this point.

(/a + i/ → [eː]); glide formation, which changes a high V1 preceding a non-identical V2 into a corresponding glide (/i + V/ → [jVː], /u + V/ → [wVː]); secondary articulation, which changes a postconsonantal high round vowel preceding a non-round vowel into a labial (if the consonant is not labial) or velar (if it is) secondary articulation (/$C_{[\text{-lab}]}$u + $V_{[\text{-rd}]}$/ → [C^wVː], /$C_{[+\text{lab}]}$u + $V_{[\text{-rd}]}$/ → [$C^ɣ$Vː]), and elision, which applies to all of sequences of identical vowels, low vowels preceding mid vowels, postconsonantal high round vowels preceding round vowels, and all postconsonantal prevocalic high front vowels.

Examples of each are given below in (8); note that every case here involves a subject prefix attached directly to a non-exceptional V-initial verb stem.

(8) Hiatus resolution

a.	w**a** + **i**va	w**eː**va	'they heard'	**Coalescence**
b.	k**a** + **o**mboka	k**oː**mboːka	'(s)he went far'	**Elision (Low V_1)**
c.	s**i** + **i**taŋga	s**i**taːŋga	'I called'	**Elision ($V_1 = V_2$)**
d.	**i** + **a**sama	**jaː**saːma	'it (cl 9) gaped'	**Glide Formation**
e.	k**u** + **i**va	k^w**iː**va	'you heard'	**Labialization**
f.	m**u** + **i**va	m^ɣ**iː**va	'you pl. heard'	**Velarization**
g.	s**i** + **a**sama	s**aː**saːma	'I gaped'	**Elision (/Ci + V/)**
h.	k**u** + **o**gera	k**oː**geːra	'you swam'	**Elision (/Cu + o/)**

I leave analysis to §5.1; for now, I will only be presenting examples and descriptions of each hiatus resolution strategy.

2.2.4.1 Asymmetric coalescence

Coalescence (sometimes also referred to as contraction or fusion) is a process that merges a sequence of two adjacent vowels into a single vowel that retains some qualities of each input vowel. Mushunguli exhibits what has been described as *asymmetric height* coalescence (Casali 1996), where a low + high sequence merges to form a single mid vowel corresponding to the place of the second (high) vowel. The output of coalescence in this case has thus retained the non-high characteristic of the input low vowel and the place characteristic of the input high vowel. Like most hiatus resolution processes in

Mushunguli, asymmetric height coalescence always results in compensatory lengthening.

Examples of asymmetric height coalescence are given in (9).

(9) Examples of coalescence

a.	/k**a** + **i**v + is + a/	k**eː**viːsa	'(s)he heard a lot'
b.	/m**a** + **i**no/	m**eː**no	'teeth'
c.	/h**a** + **i** + tʃi + tabu/	h**e**ːtʃitabu	'to the book'
d.	/ni + **a** + **u** + to + a/	n**oː**toːa	'I am hitting it (cl 3)'
e.	/w**a** + **u**mbal + a/	w**oː**mbaːla	'they are piled up'

Coalescence is one of two repairs that applies fairly regularly across words. The result is again generally lengthened. Some examples are given in (10).

(10) Examples of coalescence across word boundaries

a.	sito**eː** zimbʷa 'I beat the dogs'	
b.	ᵑgʷeneː kuŋala	'white crocodile' (*lit* 'a crocodile is white')
c.	sitʃimaliːsoː kucheːma	'I stopped the singing'
d.	sidumuloːmti	'I chopped the tree'

Note that this is optional; hiatus can and does go unresolved in this context, e.g. *soːgoheːza idijoːnda* 'I scared the baboons' (from /si + ogoh + ez + a # idi + jonda/). Note also that if a triplicate sequence arises due to the presence of word-internal hiatus at the beginning of the second word (i.e. /V_1 # $V_2 + V_3$...), only hiatus between V_2 and V_3 will be resolved, resulting in hiatus between words, e.g. *sivuna usaːzi* 'I broke the bed' (from /si + vun + a # u + usazi).

2.2.4.2 Glide formation and secondary articulations

Prevocalic high vowels (/i,u/) become glides ([j,w]) in /V + V/ contexts, provided the second vowel is not identical. The results of glide formation also exhibit compensatory lengthening, though there is sometimes variability in its realization. In most cases, the following vowel is simply long, but in some cases, the glide is impressionistically longer than the following vowel. This variability occurs in both fast and slow speech, and does not appear to be lexically-specific.[17]

Examples of glide formation are given in (11).

(11) Glide formation

a.	/**u** + edi/	→	**w**e:di	'good (cl 3)'
b.	/**u** + a + hem + a/	→	**w**a:he:ma	'you are breathing'
c.	/**u** + oger + a/	→	**w**o:ge:ra	'it (cl 3) swam'
d.	/**u** + itaŋg + a/	→	**w**i:ta:ŋga	'it (cl 3) called'
e.	/**i** + aŋgu/	→	**j**a:ŋgu	'my (cl 9)'
f.	/**i** + ose/	→	**j**o:se	'all/the whole (cl 4; cl 9)'
g.	/**i** + uz + a/	→	**j**u:za	'it (cl 9) asked'
h.	/**i** + ereker + a/	→	**j**e:reke:ra	'it (cl 9) floated'

Prevocalic, postconsonantal high round vowels exhibit a comparable pattern of secondary articulation. This secondary articulation takes the form of labialization on the preceding consonant if the consonant is not labial; if it is, then the articulation is velar. This can only apply if V2 is unround; otherwise, no glide or vowel surfaces. Compensatory lengthening applies again in this context.

Examples are given in (12)

(12) Secondary articulations

a.	/k**u** + itaŋg + a/	→	**k**^{w}i:ta:ŋga	'to call'
b.	/m**u** + ana j**u** + edi/	→	**m**^{w}a:na **j**^{w}e:di	'good (cl 1) child'
c.	/k**u** + e + jag + a/	→	**k**^{w}e:ja:ga	'to scratch oneself'
d.	/lu + ajo/	→	l^{w}a:jo	'foot'

While there are only a limited set of contexts available at morpheme boundaries (/ju-/, /mu-/, /ku-/, /lu-/), other types of labialized and velarized consonants show up in roots. These segments do not seem to cause lengthening of preceding or following syllables.

(13) Root-internal labialized and velarized consonants

[17] The fact that Mushunguli exhibits compensatory lengthening at all is somewhat typologically unusual because the language does not have contrastive vowel length (Hayes 1989; Odden & Odden 1999), cf. Lin (1997); Hanson (2010); Hyman (1985).

a. bʷiribʷiri 'spider'
b. m̩bʷa 'dog'
c. m̩tʷi 'head'
d. n̥tʷiga 'giraffe sp.'
e. n̥tundʷi 'fruit sp.'
f. m̩sʷa 'termite'
g. n̥tulʷa 'tomatillo'
h. kʷerekʷetʃe 'francolin'
i. ŋgʷena 'crocodile'
j. kuŋgʷi 'eagle sp'

Neither glide formation nor labialization/velarization can apply across word boundaries; hiatus is uniformly unresolved in these cases.

2.2.4.3 Elision and simplification

Mushunguli also has a number of underlying contexts that result in no surface exponent of the initial vowel.

Low vowels preceding mid vowels result in a surface long mid vowel (14).

(14) /a + e/ → [eː], /a + o/ → [oː]

a.	/a + **a** + **o**mbok + a/	**oː**mboːka	'you are going far'
b.	/k**a** + **e**res + a/	k**eː**reːsa	'she gave birth'
c.	/w**a** + **e** + jag + a/	w**eː**jaːga	'they scratched themselves'
d.	/ka + tumbiri k**a** + **o**se/	katumbiri k**oː**se	'the whole vervet'

All postconsonantal high round vowels preceding /o/ similarly result in a long mid vowel (15)

(15) No labialization/velarization before [o]

a.	/k**u** + omal + a/	→	k**oː**maːla	'to dish up ugali'
b.	/m**u** + oger + a/	→	m**oː**geːra	'you (pl) swam'
c.	/k**u** + ombok + a/	→	k**oː**mboːka	'you went far'
d.	/mu + gosi j**u** + ose/	→	m̩gosi j**oː**se	'the whole man'

All postconsonantal high vowels in V1 position surface with no exponent of the vowel; instead, the vowel is lost and V2 is lengthened (16).

(16) No palatalization

a.	/tʃ**i** + asa/	→	**tʃa:**sa	'we divorced'
b.	/si + d**i** + aza/	→	si**da:**za	'I lost it (cl 5)'
c.	/m**i** + ezi/	→	**me:**zi	'months'
d.	/z**i** + etu/	→	**ze:**tu	'our (cl 10)'
e.	/s**i** + u + siɲ + a/	→	**su:**si:ɲa	'I looked at it (cl 3)'

I have no examples in the data I have collected of any of these hiatus repairs applying between words.

Finally, sequences of identical vowels are simplified to a single *short* vowel (17).[18]

(17) *Simplification*

a.	/s**i** + **i**taŋg + a/	→	s**i**ta:ŋga	'I called'
b.	/w**a** + **a**mbiz + a/	→	w**a**mbi:za	'they helped'
c.	/k**u** + **u**mbal + a/	→	k**u**mba:la	'to be piled up'

Unlike the preceding three repairs, simplification can apply between words, e.g. *sito:a katumbiri* 'I beat the vervet' (from /si + to + a a + katumbiri/). This is again optional.

2.2.4.4 Mid vowels

As I discussed briefly in §2.2.3, there are rare cases of mid vowels in V1 position. These include the reflexive prefix /e-/ and the root /-to/ 'beat'. These are the only examples in my corpus of mid vowels in V1 position.

The reflexive prefix will participate in hiatus resolution as V2, e.g. *kʷe:bulu:ga* 'to smear self'. However, if it is V1 (or V2 in a triplicate), hiatus remains unresolved at the right edge of the prefix, e.g. *e.i:va* 'hear yourself!', *kʷe:.i:va* 'to hear oneself'. The latter can only ever arise in contexts where it is V1, in which case it always surfaces with unresolved hiatus, e.g. *kuto:.a* 'to beat'.

2.2.4.5 Summary of hiatus resolution outputs

A summary of hiatus resolution outputs in /V + V/ and /CV + V/ contexts are given in Tables 2.10 and 2.11. Note that entries in *italics* are hypothesized outputs; I do not

[18] Note that the output can be lengthened if it falls in the penultimate syllable of an utterance.

have examples of mid vowels as V1 in most phonological contexts.

Table 2.10: Hiatus resolution outcomes (regular) for /V + V/ sequences

V1↓ V2→	**i**	**u**	**e**	**o**	**a**
i	i	juː	jeː	joː	jaː
u	wiː	u	weː	woː	waː
e	e.i	*e.u*	*e.e*	e.o	e.a
o	*o.i*	*o.u*	o.e	*o.o*	o.a
a	eː	oː	eː	oː	a

Table 2.11: Hiatus resolution outcomes (regular) for /CV + V/ sequences

CV1↓ V2→	**i**	**u**	**e**	**o**	**a**
Ci	Ci	Cuː	Ceː	Coː	Caː
Cu	C^wiː	Cu	C^weː	Coː	C^waː
Ca	Ceː	Coː	Ceː	Coː	Ca

2.2.5 Tone

Tanzanian Chizigula has been described as having both lexical and phrasal tone, including a distinction between toneless and high-toned roots (Kenstowicz 1989; Kenstowicz & Kisseberth 1990; Kisseberth 1992; Kisseberth & Cassimjee n.d.). However, in working with my consultant, I observed very little evidence of anything besides some lingering lexical tones on affixes and some roots. I collected no examples of noun or verb roots that are distinguished only by tone, and attempts to elicit tone minimal pairs in contexts where it could make a crucial distinction (such as distinguishing the class 4 and class 9 subject prefixes, both of which are /i-/) resulted in irregularity, both within a single elicitation session and across sessions. This particular difficulty in eliciting tone with the primary consultant of the OSU Mushunguli research group was corroborated by other members, including Dave Odden and Betsy Pillion (p.c.), both of whom specifically focused on tone.[19] Researchers working with other speakers have also corroborated the

[19]Odden p.c. 03/08/18: “I had a real question as to whether there is a lexical contrast in roots...there were some roots that seemed to ”act differently”, but he was very inconsistent.”

observation that tone is at best moribund (Michal Temkin Martinez, p.c.).

Non- or marginally-tonal Bantu languages are rare but not unattested (e.g. Swahili, Chitumbuka (Downing 2006), Chimwiini (Kisseberth & Abasheikh 2011)). Speakers of Mushunguli have had considerable contact with speakers of non-tonal languages, including Swahili, English, and Maay-Maay (Somali); this is a possible source of the loss of tone.

Because of the irregularity in my data, and because there is no compelling evidence that any of the exceptions of import to the following chapters are affected by some residual tonal distinction, I do not transcribe tone in this dissertation. This is a considerable gap in this literature that needs to be addressed, but given the difficulties I experienced in eliciting tone consistently, I am hesitant to even speculate. I direct interested parties to work by Pillion (2013), but note however that Pillion has suggested to me (p.c.) that the results in that paper may not be generalizable.

2.3 Morphology

The morphology of Mushunguli, like many Bantu languages, is agglutinating. Both nouns and verbs are morphologically complex; outside of particles and (some) conjunctions, monomorphemic words in Mushunguli are very rare. This section (and the examples in this dissertation more broadly) largely focus on agreement marking on nouns and verbs.[20]

2.3.1 Noun class and agreement

Like most Bantu languages, Mushunguli has a complex system of noun class agreement. Mushunguli's system is especially complicated, due both to the large number

[20]Note that from this point on, I adopt the broader transcriptions discussed throughout §2.2.

of borrowed words, the total loss of class 13, and the loss of unique class agreement prefixes for classes 11 and 14.

2.3.1.1 Noun class prefixes

Every noun belongs to one of fourteen noun classes, conventionally numbered 1-10, 11, 12, 14, and 15. Classes 1-10 are grouped roughly into singular-plural pairs; lexical items belonging to the remaining three classes will take a plural counterpart from one of the plural (= even-numbered) classes. Class membership is indicated by a class prefix (or lack thereof) on the noun root and agreement behavior with verbs and modifiers.

A summary of singular-plural pairings is given in Table 2.12. Note that I have not included phonologically predictable allomorphs of noun class prefixes (see Table 2.13). Note as well that while some noun classes have unique noun class prefixes (11, 14) or unpredictable allomorphs (5, 9, 10), none of these cases correspond to unique agreement or concord prefixes.

Examples of phonologically predictable allomorphs of the noun class prefixes are given in (2.13). Note that I have not included every possible predictable allomorph of the class 9/10 prefixes.

A few notes on the noun class prefixes. First, some nouns that agree in class 9 lack a noun class prefix. This occurs in both phonologically predictable and unpredictable environments. The former is exemplified by *suwi* 'leopard' in Table (2.13): non-syllabic nasals are banned before fricatives, and so the prefix does not surface.[21] In other cases (especially loans), there simply is no prefix, e.g. *baraza* 'veranda(s)'. In these cases, class membership is determinable by looking at modifier agreement (see below).

Second, I have described the class 1/3 prefix here as /mu-/ (consistent with other Bantu descriptions). However, it very rarely surfaces this way; the only examples I have

[21] Syllabic nasals such as those created by copula reduction are allowed, e.g. *n̩suwi* 'it's a leopard'[D].

Table 2.12: Singular-plural pairings

Class	Sg	Pl	Examples	
1-2	mu-	wa-	m̩vere, wavere	‘woman’
1-4	mu-	mi-	m̩ɲawu, miɲawu	‘cat’
3-4	mu-	mi-	m̩ti, miti	‘tree’
3-10	mu-	ɲ-	mo:jo, ɲojo	‘heart’
3	mu-		muɲu	‘salt’
5-6	Ø-	ma-	juwe, majuwe	‘stone’
5-6	z-	ma-	zino, me:no	‘tooth’
6		ma-	mazi	‘water’
7-8	tʃi-	vi-	tʃirevu, virevu	‘chin’
9-10	ɲ-	ɲ-	ɲoka, ɲoka	‘snake’
9-10	Ø	Ø	baraza, baraza	‘veranda’
11-11	lu-	lu-	luɟendo, luɟendo	‘chameleon’
11-10	lu-	ɲ-	lʷajo, ɲajo	‘foot’
12-2	ka-	wa-	kahuɟi, wahuɟi	‘hawk’
14	u-		utʃiza	‘darkness’
14-4	u-	mi-	ulosi, milosi	‘language’
14-10	u-	ɲ-	utumbo, ɲutumbo	‘intestine’
15	ku-		kutʃema	‘singing’

Table 2.13: Phonologically predictable allomorphs of the noun class prefixes

Class 1	**m̩**gosi	‘man’	**Class 2**	**wa**gosi	‘men’
	mʷa:na	‘child’		**w**ana	‘children’
Class 3	**m̩**neŋge	‘moon’	**Class 4**	**mi**leŋge	‘moons’
	mʷe:zi	‘month’		**me**:zi	‘months’
Class 5	kuwi	‘turtle’	**Class 6**	**ma**kuwi	‘turtles’
	ziso	‘eye’		me:so	‘eyes’
Class 7	**tʃi**fulo	‘bubble’	**Class 8**	**vi**fulo	‘bubbles’
	tʃisuse	‘scorpion’		**vi**suse	‘scorpions’
Class 9	**m̥**puku	‘rat’	**Class 10**	**m̥**puku	‘rats’
	ŋgʷena	‘crocodile’		**ŋ**gʷena	‘crocodiles’
	ɲutʃi	‘bee’		**ɲ**utʃi	‘bees’
	suwi	‘leopard’		suwi	‘leopards’
	m̩bʷa	‘dog’		**m̩**bʷa	‘dogs’
Class 11	**lu**ɟe	‘vein’	**Class 12**	**ka**buga	‘bunny’
	lʷi:mbo	‘song’^{D}		**ka**zana	‘baby’
Class 14	**u**losi	‘language’	**Class 15**	**ku**muli:ka	‘lightning’
	utʃi	‘honey’		**k**ʷa:ʃa	‘burning’

observed are *muʃuŋguli* ‘Mushunguli’ and *muɲu* ‘salt’. It is much more typical for these prefixes to be reduced to [m̩] before consonants; this is true for other prefixes of the same shape, including the adjective concord prefixes and the (human) 2PL subject and object agreement prefixes.

Third, the class 12 prefix can be used productively as a diminutive marker. If this occurs, the class 12 prefix replaces the expected noun class prefix and the resulting noun will agree with class 12, e.g. *kagosi* ‘small man,’ *kagosi kado:do* ‘very small man’ (c.f. *m̩gosi m̩do:do* ‘small man’). Diminutives of animals generally refer to babies, e.g. *kabuga* ‘bunny,’ (*mbuga* ‘rabbit’). They can also refer to smaller animals of a specific archetype; for example, the root *-tumbiri* corresponds to monkey. Its class 9 form *n̥thumbiri* refers to a colobus monkey, while its 12 form refers to a vervet. I collected no examples of roots that exclusively belong to class 12 save for perhaps *-zana* ‘baby’ (cf. *-ana* ‘child’); the plural of this root agrees in class 2 as well.

2.3.1.2 NP-internal agreement

NP-internal modifiers are marked with concord prefixes corresponding to the class of the head noun.[22] Similarly, verbs are marked with agreement prefixes corresponding to the class of the subject and (sometimes) object. I will be focusing on NP-internal agreement in this section; I discuss subject and object marking in the following section.

There are two classes of noun concord prefixes in Mushunguli. These are adjective concord prefixes, which apply mainly to cardinal numerals and “true” adjectives (as opposed to stative verbs, which are more common, and which take subject agreement prefixes); and demonstrative concord prefixes, which apply to demonstratives, interrogative pronouns, possessive pronouns, and quantifiers.[23] While these form distinct sets, there

[22]Conjoined noun phrases are more complicated than this; see Williams (2012).

[23]Note that some modifiers are always unmarked, e.g. *usebu* ‘new,’ *kejasi* ‘some,’ *gani* ‘which’. Some of these are loans; for example, ‘new’ most likely originates from Somali *cusub*.

is no evidence suggesting that some of these markers should be considered clitics and others affixes; in terms of their phonological behavior, they behave more or less identically, with the notable exception of the class 5 demonstrative prefix /di-/ (discussed in detail in §5.1.3).

Adjective agreement prefixes are typically identical in shape to the noun class prefixes, as seen below in Table 2.14, using *kulu* 'big' to illustrate.

Table 2.14: Adjective agreement prefixes (with examples)

Class 1 (mu-)	m̩gosi **m̩**ku:lu	'big man'	**Class 2 (wa-)**	wagosi **wa**ku:lu	'big men'
Class 3 (mu-)	m̩ti **m̩**ku:lu	'big tree'	**Class 4 (mi-)**	miti **mi**ku:lu	'big trees'
Class 5 (Ø-)	jonda ku:lu	'big baboon'	**Class 6 (ma-)**	majonda **ma**ku:lu	'big baboons'
Class 7 (tʃi-)	tʃirevu **tʃi**ku:lu	'big elbow'	**Class 8 (vi-)**	virevu **vi**ku:lu	'big elbows'
Class 9 (ɲ-)	simba **ŋ̥**ku:lu	'big lion'	**Class 10 (ɲ-)**	simba **ŋ̥**ku:lu	'big lions'
Class 12 (ka-)	katumbiri **ka**ku:lu	'big vervet'	**Class 15 (ku-)**	kumulika **ku**ku:lu	'big lightning'

Note the lack of class 11 and 14 agreement prefixes. All nouns belonging to these classes take identical agreement patterns to class 3.

Class 5, which is typically unmarked, can irregularly surface with *di-* (*jonda dikulu* 'big baboon'); I have very few examples of this, so it may be an accident. Similarly, class 10 can sometimes surface with *zi-*, e.g. (*simba zin̥titu* 'black lions').[24]

Demonstrative concord prefixes are a distinct set; while for some classes the demonstrative agreement prefix is again identical to the noun class prefix, for others a different prefix is used. This is illustrated below in Table 2.15, using *no* 'this/these'. Note that /-no/ is a monosyllabic root, and so these examples all have penultimate lengthening (i.e. the vowel length should not be assumed to be part of the prefix).

Definiteness is marked with a pre-prefix before the noun class marker, referred to here as the augment. This marker is a vowel that matches the vowel of the demonstrative prefix. Notice in Table 2.16 that for classes 5 and 10, an overt noun class marker surfaces on the noun.

[24]Many color terms in Mushunguli are stative verbs and take subject agreement prefixes, but 'black' is not one of them, as evidenced in this example by the presence of the prenasalized stop).

Table 2.15: Examples of demonstrative agreement prefixes

Class 1 (ju-)	m̩vere **ju:**no	'this woman'	**Class 2 (wa-)**	wavere **wa:**no	'these women'
Class 3 (u-)	m̩teza **u:**no	'this peanut'	**Class 4 (i-)**	miteza **i:**no	'these peanuts'
Class 5 (di-)	buku **di:**no	'this book'	**Class 6 (ma-)**	mabuku **ja:**no	'these books'
Class 7 (tʃi-)	tʃituŋgulu **tʃi:**no	'this onion'	**Class 8 (vi-)**	vituŋgulu **vi:**no	'these onions'
Class 9 (i-)	ŋuluwe **i:**no	'this pig'	**Class 10 (zi-)**	ŋuluwe **zi:**no	'these pigs'
Class 12 (ka-)	kazana **ka:**no	'this baby'	**Class 14 (u-)**	ulosi **u:**no	'this language'
Class 15 (ku-)	kutoa **ku:**no	'this beating'			

Table 2.16: Examples of the augment (pre-prefix) in definite constructions

Class 1	**u**m^{w}a:na	'the child'	**Class 2**	**a**wana	'the children'
Class 3	**u**mt^{w}i	'the head'	**Class 4**	**i**mit^{w}i	'the heads'
Class 5	**idi**jonda	'the baboon'	**Class 6**	**a**majonda	'the baboons'
Class 7	**i**tʃiboŋko	'the hippo'	**Class 8**	**i**viboŋko	'the hippos'
Class 9	**i**mb^{w}a	'the dog'	**Class 10**	**izi**mb^{w}a	'the dogs'
Class 11	**u**luɟendo	'the chameleon'	**Class 12**	**a**kalogo	'the duiker'
Class 14	**u**ta	'the bow'	**Class 15**	**u**k^{w}a:ʃa	'the burning'

Table 2.17 summarizes all regular agreement markers.

Table 2.17: Agreement markers

Class	Augment	Adjective concord	Demonstrative concord
1	u-	mu-	ju-
2	a-	wa-	wa-
3	u-	mu-	u-
4	i-	mi-	i-
5	idi-	Ø- (di-)	di-
6	a-	ma-	ja-
7	i-	tʃi-	tʃi-
8	i-	vi-	vi-
9	i-	ɲ-	i-
10	izi-	ɲ- (ziɲ-)	zi-
11 (3)	u-	mu-	u-
12	a-	ka-	ka-
14 (3)	u-	mu-	u-
15	u-	ku-	ku-

2.3.1.3 Irregularities in the noun class system

There are a number of idiosyncracies within the noun class system, either with respect to class assignment or agreement.

One factor is animacy. There is some evidence for animacy effects in Mushunguli. As discussed in more detail by Williams (2012), verbs taking conjoined noun phrases disagreeing for class as their subjects resolve this conflict by agreeing with the plural of the more animate noun, with classes 1 and 2 (referring to humans) taking precedence over all others. For example, *m̩vere yuno na ŋguluwe ino* ***wagʷa*** 'this woman and this cow fell', takes the class 2 plural marker on the verb, not the class 10 one (**zigʷa*).[25] However, it is possible for nouns not belonging to classes 1/2 to act as subjects of transitive verbs (though bizarre scenarios, like walking trees, were semantically odd for my consultant), as well as experiential and stative verbs. Similarly, my consultant had no problems producing possessive constructions in which the possessor was either an animate or inanimate noun, e.g. *miti maɟani jaːke* 'the trees' leaves', *idijonda boko ɟaːke* 'the baboon's banana'.

That said, there are some non-human nouns referring to animate subjects that exhibit class 1 agreement patterns in the singular, but class 4 agreement patterns in the plural. Examples are given in (18), using the root *-ɲawu* 'cat'.

(18) Animacy of 'cat'

a.	m̩ɲawu yʷeːdi	'good cat'	(*m̩ɲawu weːdi)
b.	m̩ɲawu kaːɲʷwa	'a cat drank'	(*m̩ɲawu uːɲʷwa)
c.	miɲawu jeːdi	'good cats'	
d.	miɲawu iːɲʷa 'cats drank'		

Other examples nouns exhibiting this behavior include *m̩ɲula* 'leech', *m̩saŋgo* 'parasitic worm', *muzi* 'bird sp.', and *m̩gulo* 'coucal sp'.[26]

[25]Gender conflict resolution is complicated in Mushunguli; see Williams 2012 for a much more complete account.

[26]Note that there are class 3 nouns referring to living things that do not show this pattern of agreement, e.g. *m̩sʷa* ***uːɲʷa*** 'a termite drank', cf. *m̩ɲawu* ***kaːɲʷa*** 'a cat drank'.

Similarly, there are some nouns referring to humans that are marked with the class 9/10 prefix in the singular and plural. These nouns nevertheless take class 1/2 agreement, as in (19), using *-dugu* ‘sibling; cousin’.

(19) Humans in class 9/10

a. ndugu m̩ku:lu ‘big sibling’ (*ndugu ŋ̥kulu)
b. ndugu waku:lu ‘big siblings’ (*ndugu ŋ̥kulu, *ndugu ziŋ̥kulu)

Another source of irregular agreement is the historical reduction of the noun class system; noun class 13 (the plural of class 12) is gone entirely, and classes 11 and 14 do not have unique agreement patterns. As I stated above, class 12 nouns now appear to invariably have class 2 plural forms (and are marked with the class 2 prefix /wa-/); in many cases, this distinguishes them semantically from their non-diminutive counterparts, e.g. *wabuga* ‘bunnies’, but *mbuga* ‘rabbits’.[27]

(20) Class 12/2

a. kazana kaku:lu ‘big baby’
b. wazana wa:kulu ‘big babies’
c. kalogo ke:di ‘good duiker’
d. walogo we:di ‘good duikers’

The loss of class 11 has resulted in the reassignment of many class 11 nouns into several other classes, some of which have retained their original prefix (*lu-*) in a fossilized state. Class 11 nouns marked in this way take class 3 agreement patterns in the singular; inanimate class 11 nouns have class 10 plural forms and agreement, while animate class 11 nouns retain their class 11 prefix in the plural, but take class 2 agreement patterns.

(21) Class 11/10 and 11/2

[27] I have no example of non-living diminutives in my corpus, so it is unclear whether something like (hypothetical) $^{?}$*kati* ‘small tree’ would have a class 2 plural. Given that there are no examples of inanimate class 2-agreeing nouns in my corpus, my guess is that it would not.

a.	l^wimbo m̩tali	'long song'[D]
b.	ɲimbo ɲiŋgi	'many songs'[D]
c.	l^wajo weːdi	'good foot'
e.	ɲajo zeːdi	'good feet'
f.	luɟendo m̩kulu	'big chameleon'
g.	luɟendo waʃano	'five chameleons'

Finally, class 14 generally contains abstract or mass nouns (e.g. *utʃiza* 'darkness'); most mass nouns have no plural form, and take class 3 agreement. However, class 14 also contains some nouns that do have plurals; some of these appear to be taken from class 11 (e.g. *ulimi* 'tongue'), while others are countable abstract nouns (e.g. *ulosi* 'language'). Most countable class 14 nouns take class 10 plurals and agreement patterns.[28]

(22) Class 14, 14/10

Class 14 (uncountable)

a.	utʃiza weːdi	'good darkness'[D]
b.	utʃiza m̩kulu	'big darkness'[D]

14/10

c.	uta weːdi	'good bow'
d.	uta m̩titu	'black bow'
e.	ɲuta zeːdi	'good bows'
f.	ɲuta ndodo	'small bows'

2.3.2 Verbal morphology

Verbal morphology is several orders of magnitude more complicated than that of nouns. Verbs minimally consist of a subject prefix, root, and final vowel (typically /-a/, though /-e/ is used for the subjunctive mood and /-i/ sometimes also arises in loans). However, verbs are typically more morphologically complex than this, and may contain one or more additional overt tense-aspect prefixes (depending on the tense or aspect), a negative prefix, as well as any number of derivational or inflectional extensions. Object prefixes may also be present depending on discourse and morphosyntactic factors.

[28]The plural of *ulosi* 'language' is *milosi*, but I do not have any examples in the corpus of this noun being modified.

A full description of the intricacies of verb formation in Mushunguli is well outside the scope of this dissertation; for the remainder of this section, I will simply outline aspects of verbal morphology most typically referenced in relevant examples.

2.3.2.1 Tense, aspect, and mood marking

Mushunguli generally marks tense directly on the verb with a tense-marking prefix. Aspect and mood, if marked, is typically suffixal (including the final vowel, as mentioned above).

The examples in this dissertation primarily focus on the following TAM combinations: simple past (the unmarked aorist (Dayley et al. 2020)), present progressive, and future subjunctive. The infinitive and imperative forms of verbs are also regularly referenced. The markers are summarized in Table 2.18.[29]

Table 2.18: Tense markers

Infinitive	ku-	ku-ROOT-a	**ku**pi:ka	'to cook'
Imperative	Ø	Ø-ROOT-a	pi:ka	'cook!'
Present	a-	SUBJ-a-ROOT-a	n**a:**pi:ka	'I am cooking'
Future	na-	na-SUBJ-ROOT-e	**na**nipi:k**e**	'I will cook'
Simple Past	Ø	SUBJ-ROOT-a	**si**pi:ka	'I cooked'

2.3.2.2 Subject marking

Subjecthood is marked directly on the verb via an agreement prefix, which usually occurs at the extreme left edge of a verb (save for the future subjunctive, whereby the non-present prefix /na-/ precedes it, and negative constructions, in which case the negative prefix /ha-/ precedes it). This prefix agrees with the noun class of the subject of the clause

[29] There are obviously far more constructions than this, but these are the only constructions for which I have recorded full paradigms using the same verbs. Future work on Mushunguli needs to address this gap.

(if non-human), or person (if human). The latter set of markers also varies by tense (see below).

Subjects do not have to be overtly stated if they were previously mentioned in the discourse context, e.g. *simba iːɟa* 'a lion ate' is always licit, and *iːɟa* 'it (cl 9) ate' is licit provided that the subject was previously mentioned.

For non-human subjects, there are no differences in the forms of subject prefixes across tenses.[30] Examples of these are given below in (2.19), using the intransitive verb *kugʷa* 'to fall'.

Table 2.19: Non-human subject marking

Class 3	u-	**u**ːgʷa	**Class 4**	i-	**i**ːgʷa
Class 5	di-	**di**ːgʷa	**Class 6**	ja-	**ja**ːgʷa
Class 7	tʃi-	**tʃi**ːgʷa	**Class 8**	vi-	**vi**ːgʷa
Class 9	i-	**i**ːgʷa	**Class 10**	zi-	**zi**ːgʷa
Class 12	ka-	**ka**ːgʷa	**Class 14**	u-	**u**ːgʷa
Class 15	ku-	**ku**ːgʷa			

Human subject prefixes encode both person and number; there is no masculine or feminine gender marking.[31] The 1-3SG prefixes also have allomorphs corresponding to past and non-past tenses.

The set of subject prefixes used for present and future constructions (= non-past) are illustrated below in (2.20), using the present form of the verb *hema* 'breathe' and the future form of the verb *gula* 'buy'.

In Table 2.20, I have included both the phonetic realizations of the prefixes as well as proposed underlying forms. This is because affix ordering in Mushunguli makes it impossible to elicit a present paradigm where the subject prefixes are not in a hiatus context; it is similarly impossible to elicit *any* non-past paradigm where the 2SG and 3SG

[30]There are phonologically predictable allomorphs of these prefixes that occur when they precede or follow a vowel.

[31]Non-human subjects that take class 1 or 2 agreement patterns also take at least the 3SG and/or 3PL prefixes. I did not elicit examples of 1-2 SG/PL agreement with non-human subjects (e.g. a cat talking about itself, or addressing a leech).

Table 2.20: Human subject marking (non-past)

	Singular		Plural	
Future				
1	/ni-/	na**ni**gu:le	tʃi-	na**tʃi**gu:le
2	/u-/ ([o-])	no:gu:le	/mu-/ ([m-])	na**m**gu:le
3	/a-/ ([Ø])	nagu:le	wa-	na**wa**gu:le
Present				
1	/ni-/ ([n-])	**n**a:he:ma	/tʃi-/ ([tʃ-])	**tʃ**a:he:ma
2	/u-/ ([w-])	**w**a:he:ma	/mu-/ ([m^w-])	**m^w**a:he:ma
3	/a-/ ([Ø])	**a**he:ma	/wa-/ ([w-])	**wa**he:ma

subject prefixes are not in a hiatus context. Hiatus resolution was discussed in §2.2.4, but I will summarize the patterns here quickly, to explain the proposed underlying forms.

The underlying forms of the 1SG/PL prefixes are justified by the fact that they surface with a high vowel in the future paradigm. In the present paradigm, the fact that no high vowel surfaces, but there is lengthening of the following vowel, is consistent with the general behavior of postconsonantal high front vowels (§3.5).

The underlying form of the 2SG prefix is justified by the fact that [o:] is the expected surface result of /a + o/ hiatus resolution (§3.3), and [wa:] is the expected surface result of /u + a/ hiatus resolution (§3.2). The form of the 2PL prefix is justified by the fact that it behaves identically to the 1/3 noun class and adjective concord prefixes; the fact that the proposed /u/ surfaces as a secondary articulation when it is V1 in a hiatus context is also consistent generally with the behavior of postconsonantal /u/ (§3.5).

The form of the 3PL prefix is justified by the fact that it surfaces as [wa-] before a C-initial root. The fact that the result of /wa-a/ in the present paradigm is [wa] is also consistent with the behavior of sequences of identical vowels (3.4). This is also true for the proposed 3SG prefix; however, it is always in an identical vowel context, so in principle, 3SG could be unmarked. I propose that there is a prefix due to the fact that in very careful speech (i.e. illustrative), my consultant would often deliberately fail to resolve hiatus; in these cases, he would produce a long vowel for this context (e.g. [aːheːma]).

Singular forms of verbs in the past tense have different subject prefixes. Unlike the future and present paradigms, this context does not result in hiatus unless the root itself begins with a vowel (or if there is an object prefix). Examples are given below, using the past tense form of the verb *ɟa* 'eat.'[32]

Table 2.21: Human subject marking (past)

	Singular		Plural	
1	si-	**siː**ɟa	tʃi-	**tʃiː**ɟa
2	ku-	**kuː**ɟa	mu-	**m̩**ɟa
3	ka-	**kaː**ɟa	wa-	**waː**ɟa

2.3.2.3 Object marking

Objects can be marked on the verb directly using a prefix (this is not always obligatory; see below), which occurs at the left edge of the verb stem. For non-human objects, object prefixes are identical *in shape* to the subject prefixes; though note that I was never able to elicit a verb with a class 15 object prefix, so it is unclear what the shape of this prefix would be, if it exists.

Table 2.22: Non-human object marking

Class 3	u-	s**uː**leːta	**Class 4**	i-	si.**i**leːta
Class 5	di-	si**di**leːta	**Class 6**	ja-	si**ja**leːta
Class 7	tʃi-	si**tʃi**leːta	**Class 8**	vi-	si**vi**leːta
Class 9	i-	si.**i**leːta	**Class 10**	zi-	si**zi**leːta
Class 12	ka-	si**ka**leːta	**Class 14**	u-	s**uː**leːta

For human objects, the set of object prefixes largely mirrors that of the general class of subject prefixes, though note that the 2SG prefix matches the past tense, and the 3SG prefix mirrors the class 1 adjective concord prefix. This is illustrated in Table 2.23, using the present forms of the verb *toa* 'beat'.

[32]The root for eat is /di-/; see §5.1.3.

Table 2.23: Human object marking

	Singular			Plural		
1	ni-	ka**ni**to:a	'he beat me'	tʃi-	wa**tʃi**to:a	'they beat us'
2	ku-	si**ku**to:a	'I beat you'	mu-...-ni	si**m**toa:**ni**	'I beat you (pl)'
3	mu-	si**m**to:a	'I beat him/her'	wa-	si**wa**to:a	'I beat them'

Whether or not an object prefix will appear (and whether or not it *must* appear) is dependent on morphosyntactic and discourse factors. Examples are given in (23). Note that this description is certainly incomplete; my corpus does not contain very many examples of object prefixes.

(23) Distribution of object prefixes

a.	sito:a m̩go:si	'I beat a man'	No OP
b.	si**m**to:a **u**mgo:si	'I beat the man'	OP + AUG
c.	si**m**to:a jeje	'I beat him'	OP + pronoun
d.	si**m**to:a	'I beat him'	OP only
e.	*simto:a m̩go:si	'I beat a man'	OP only
f.	*sito:a jeje	'I beat him'	Pronoun, no OP

If a lexical NP object is contained within the same sentence as the verb, then the presence of an object prefix is determined by definiteness. There is no object prefix if the object is indefinite (23a), and the use of an object prefix in this context appears to be ungrammatical (23e). If the object is definite (marked by the augment), use of an object prefix is obligatory (23b). If the NP object is an overt pronoun in the same clause, then the object prefix is again obligatory (23c,f). Finally, if the NP has been previously mentioned in the discourse, it is possible to simply use an object prefix and not state the lexical NP object.

2.3.2.4 Verbal extensions

Mushunguli has several verbal extension suffixes that can be concatenated with verb roots. More than one suffix can be used at once. A list of the extensions which may appear in examples within this dissertation (which is certainly incomplete), their Bantu conventional glosses, and their rough meanings, are given below in Table 2.24. As was mentioned in 2.2.2, verbal extensions containing front vowels alternate for [± high], with [e] surfacing after mid vowels, and [i] surfacing after both high and low vowels.

Table 2.24: Verbal extensions

-iz	causative	'make *x*'
-ir[33]	applicative	'*x* for (benefactive)'; 'use to *x*'
-ik	intransitive	'be *x*-able'
-is	habitual	'*x* a lot or too much'
-an	reciprocal	'*x* for each other'
-aɲ	aggregate	'*x* together'
-w[34]	passive	
-ire	perfective	

These suffixes vary considerably in how productive they are, and there are many cases of unproductive or lexicalized readings, especially when more than one is used at once. For example, the unmodified verb *kʷa:sa* means 'to leave' (as in 'go away') or 'to divorce,' but the applicative (benefactive) form of the verb, *kʷa:si:za* means 'to leave for; bequeath'. A more extreme example is the verb *kʷa:ndamira:na*, which is the benefactive reciprocal form of the verb 'to lean,' and which means 'to lean back-to-back.'

Investigating the nature of extensions in this language is a worthwhile subject for future research, but is outside the scope of this dissertation.

This chapter, in part, contains material that previously appeared in *Africa's Endangered Languages: Documentary and Theoretical Approaches*. Hout, Katherine, Oxford University Press, 2017. This chapter, in part, is also currently being prepared for submission. Hout, Katherine. The dissertation author was the primary investigator and author of all of this material.

[33] This suffix usually surfaces as [-iz] after sibilants.

[34] This suffix has an allomorph *-(i)gʷ*. I do not have enough passive data to determine whether this allomorph has a predictable distribution.

3 Analysis of Hiatus Resolution

This chapter presents an analysis of regular hiatus resolution strategies in Mushunguli. This analysis will ultimately be the baseline from which the discussion of the exceptions in §4 and §6 will be built; as we will see there, the behavior of these exceptions will both follow from and reinforce the analysis of the regular aspects of the language.

Recall from §2.2.4 that Mushunguli, descriptively, has four hiatus repairs with eight surface outcomes. Examples of each are repeated below in (24); note that every case here involves a subject prefix attached directly to a non-exceptional V-initial verb stem.

(24) Hiatus resolution

a.	w**a** + **i**va	w**eː**va	'they heard'	**Coalescence**
b.	k**a** + **o**mboka	k**oː**mboːka	'(s)he went far'	**Elision (Low V_1)**
c.	s**i** + **i**taŋga	s**i**taːŋga	'I called'	**Elision ($V_1 = V_2$)**
d.	**i** + **a**sama	**jaː**saːma	'it (cl 9) gaped'	**Glide Formation**
e.	k**u** + **i**va	kʷ**iː**va	'you heard'	**Labialization**
f.	m**u** + **i**va	mˠ**iː**va	'you pl. heard'	**Velarization**
g.	si + **a**sama	s**aː**saːma	'I gaped'	**Elision (/Ci + V/)**
h.	k**u** + **o**gera	k**oː**geːra	'you swam'	**Elision (/Cu + o/)**

All eight of these patterns are robustly attested at prefix + prefix and prefix + root boundaries in verbs. They are also repeatedly and regularly observed in other morphosyntactic contexts, including prefix + prefix boundaries in locatives, prefix + root and prefix + prefix boundaries in nouns, and prefix + prefix and prefix + root boundaries of noun modifiers. Examples of these are given in (25).

(25) Hiatus resolution (non-verbs)

a.	w**a** + **i**ŋgi	w**eː**ŋgi	'many (cl 2)'	**Coalescence**
b.	**ja** + **o**se	**j** **oː**se	'all (cl 6)'	**Elision (Low V_1)**
c.	z**i** + **i**hije	z**i**hiːje	'bad (cl 10)'	**Elision ($V_1 = V_2$)**
d.	**i** + etu	**jeː**tu	'our (cl 4)'	**Glide Formation**
e.	k**u** + **e**di	k**ʷeː**di	'good (cl 15)'	**Labialization**
f.	m**u** + **e** + zi	m**ˠeː**zi	'month'	**Velarization**
g.	tʃ**i** + **a**ni	tʃ**aː**ni	'which? (cl 7)'	**Elision (/Ci + V/)**
h.	u + k**u** + **o**	uː**ko**	'that (prox) (cl 15)'	**Elision (/Cu + o/)**

It is not atypical for Bantu languages to have more than one repair in the same context; in particular, glide formation with one of deletion or coalescence is typologically common (Casali 1996). However, it is unusual to see a language with all of coalescence, glide formation, secondary articulation, and elision in the same set of morphosyntactic contexts. The most famous examples of these come from the Nguni subgroup of Southern Bantu languages (Sibanda 2009), especially Xhosa (McLaren 1955; Aoki 1974; Casali 2011), which has attested patterns similar to that of Mushunguli. As will be seen, this surface diversity makes Mushunguli a challenging language to analyze relative to other Bantu languages that have been described as having three or more repairs. This is because these other cases either exhibit repairs that, while descriptively distinct, are formally similar enough to one another to be given similar formal treatments, or because these other cases exhibit diverse repairs in *separate* morphosyntactic contexts.

An example of the former comes from Nambya; in nominal constructions prevocalic low vowels are elided, while prevocalic high vowels are either glided in /V + V/ contexts or form secondary articulations (labialization or palatalization) in unmarked /CV + V/ contexts (in marked contexts, they are also elided) (Kadenge 2013). Similar patterns are also attested in Karanga (Mudzingwa & Kadenge 2011) and ciNsenga (Simango & Kadenge 2014).[1]

[1]This type of description and analysis is controversial, due to the fact that secondary articulations apply to the same vowels as glide formation, and have a glide-like surface result. There is considerable controversy in Bantu literature regarding whether such surface patterns should be considered glide formation (= consonant clusters) or secondary articulations (see Odden 2015 for

Rule-based frameworks have no difficulty in treating each of these repairs as distinct. However, when these types of patterns are analyzed in an OT framework, we see that there emerges an interesting difference between description and analysis; namely, what appear to be many distinct repairs can typically be reduced to one or two.[2] This is illustrated by three examples from Kadenge (2013).

(26) Hiatus resolution in Nambya nominals (from Kadenge 2013)

a.	/ì-á^{ŋ}gù/	já^{ŋ}gù	'my (cl 9)'	Glide formation
b.	/lì-ópò/	l^{j}ópò	'eyelids'	Secondary articulation
c.	/ʧì-òtó/	ʧòtó	'fire-place'	Elision

Descriptively, there do appear to be three distinct "repairs" here. However, crucially, competing repairs such as epenthesis and coalescence do not occur in this context. As a result, these repairs are similar enough to be collapsed into one, and in fact the proposed analysis of these three examples does exactly this. Hiatus resolution in Nambya nominals violate one or more of three faithfulness constraints: MAX(Rt), which penalizes deleting a segment (but not a mora); MAX(μ), which penalizes deleting a mora (but not a segment); and MAX[coronal], which penalizes deleting a coronal feature from a segment. Notably, each of these constraints is a MAX-IO constraint; i.e. a constraint penalizing deletion. Each of the three repairs (as analyzed by Kadenge) violates a subset of these constraints, resulting in an implicational hierarchy of deletion.[3]

(27) Hierarchy of deletion for Nambya nominals

a summary and references).

[2]This is similar, but not identical, to a formal and typological claim made by Rosenthall (1997), which is that languages exhibit at most a bipartite distribution of prevocalic vowels; that is, a language can have distinct operations applying to high and low V1s, but these operations are assumed to be non-overlapping. However, OT itself does not *require* these repairs to be non-overlapping.

[3]Kadenge does not consider faithfulness violations incurred by a number of important candidates, including glide formation with compensatory lengthening, the formation of secondary articulations (with or without compensatory lengthening), or coalescence.

		MAX[cor]	MAX(Rt)	MAX(μ)
Elision	/tʃì-òtó/ → tʃòtó	*	*	*
Sec. Art	/lì-ópò/ → lʲópò		*	*
Glide formation	/ì-áⁿgù/ → jáⁿgù			*

As will be discussed in §3.3 and §3.4, Mushunguli has a similar implicational relationship between faithfulness violations incurred by a subset of repairs that look distinct. In that section, I argue that formally these are *not* different repairs. Similarly, I would argue that this formal analysis of Nambya does not actually support the idea that Nambya has three repairs in the same context: it has one repair (deletion) with context-specific surface results.

Similarly, Nambya has been described as having all of elision, glide formation, secondary articulation, coalescence, and epenthesis (Kadenge 2013); critically, however, these repairs do not all occur in the same morphosyntactic contexts. For example, epenthesis is only observed in verbs, while coalescence only applies post-lexically. This means that formal analysis of these other repairs can crucially rely on constraints that refer to morphosyntactic contexts or semantic classes; while no analysis is actually given, for example, it is possible to abstract away from the problem of epenthesis applying in verbs but not in nouns by assuming that there is version of DEP specific to nouns, which dominates the hiatus resolution-driving markedness constraint. This type of analysis has been proposed to account for the failure of hiatus resolution to apply at prefix + root boundaries in verb in ciNsenga, as discussed in the introduction (Simango & Kadenge 2014).

Unlike these other languages, the analysis of Mushunguli will not be able to make reference to morphosyntactic context or semantic class to explain the four classes of repairs. The consequence of this is that any analysis required for one pattern, such as coalescence, will also necessarily apply to *all* patterns.

The analysis to come will demonstrate that Mushunguli only actually has two classes of hiatus repairs, with contextually-sensitive surface results. In §3.3 and §3.4, it will be demonstrated that the only available analysis for examples (8a-c), including those described as elision, is actually coalescence. In §3.5, it will be demonstrated that neither one-step secondary articulations nor vowel deletion are applicable analyses for examples (8e-h). The surface forms of all of these examples will ultimately be demonstrated to be the output of the repair of derived glides–that is, these examples are all subject to glide formation, just as (8d) is.

3.1 Starting assumptions

In this analysis, I assume a binary set of height and place features for vowels. I also assume full specification of all vowels save for /a/, which is unspecified for any place feature.

(28) Height and place specifications for vowels

/a/	[-high, +low, 0back, 0round]
/i/	[+high, -low, -back, -round]
/e/	[-high, -low, -back, -round]
/u/	[+high, -low, +back, +round]
/o/	[-high, -low, +back, +round]

The markedness constraint driving hiatus resolution that I will adopt is ONSET:

(29) ONSET: $*[_\sigma V$ (McCarthy & Prince 1993)

Syllables must have onsets.

The majority of cases of interest to this analysis concern heterosyllabic, adjacent vowels. However, as noted by McCarthy & Prince (1993:33), ONSET is technically violated by word-initial onsetless syllables, which are not uncommon in this language. We can potentially avoid this by proposing the following ANCHORING constraint (following McCarthy & Prince 1995).

(30) L-ANCHOR(Stem,PrWd): the leftmost segment of the stem has a correspondent at the left edge of the PrWd

For this to work, "stem" needs to be defined non-recursively; that is, the generalization for Mushunguli is exclusively that the leftmost segment of any input word surfaces as the leftmost segment of the output word. However, it is not clear that this constraint is active in Mushunguli, at least prior to phrase-level (= postlexical) phonology (where hiatus is variably tolerated). The example discussed by McCarthy & Prince (1993) is Axininca Campa, which has productive word-internal epenthesis. In Mushunguli, the only possible hiatus repairs are those which require the interactions between two input segments. There are no cases of epenthesis occurring word-internally, word-initially, or between words, so I assume that the constraint DEP-C, which penalizes epenthesis, is undominated.

(31) DEP-C: every consonant in the output has a correspondent in the input

We will also, over the course of this section and working through §4.2, eventually see that MAX-V, which penalizes vowel deletion, is also undominated.

(32) MAX-V: every vowel in the input has a correspondent in the output

Regardless, given that we have not yet actually established that either DEP-C or MAX-V is undominated (i.e. by examining a word-internal hiatus context), and given that there are some hiatus resolution processes that *look* like vowel deletion, we will assume for now that L-ANCHOR(Stem,PrWd) dominates ONSET, which will allow examples such as /i + hem + a/ 'it (cl 9) breathed' to surface as [ihe:ma].

(33) Protection of word-initial onsetless syllables

/i + hem + a/	L-ANCHOR	ONSET	DEP-C	MAX-V
a. ☞ .i.he.ma		*		
b. Ci.he.ma	*!		*	
c. he.ma	*!			*

That said, unless otherwise noted, I am not going to be marking word-initial ONSET violations, as, again, all of the cases of interest are word-internal. For the sake of time and space, we will also make the additional *a priori* assumption that constraints penalizing deletion or gliding of root vowels in V_2 position (following Casali 1997) are undominated, as these patterns never occur in exceptional or regular contexts.

3.2 Glide formation in /V + V/ contexts

Prevocalic high vowels /i/ and /u/ become corresponding glides [j] and [w] at morpheme boundaries, except for when the following vowel is identical (§3.4), or the vowel is preceded by a tautosyllabic consonant (§3.5).

Examples of glide formation are given below in (34).[4]

(34) Glide formation

a.	/**u** + edi/	→	**w**eːdi	'good (cl 3)'
b.	/**u** + a + hem + a/	→	**w**aːheːma	'you are breathing'
c.	/**u** + oger + a/	→	**w**oːgeːra	'it (cl 3) swam'
d.	/**u** + itaŋg + a/	→	**w**iːtaːŋga	'it (cl 3) called'
e.	/**i** + aŋgu/	→	**j**aːŋgu	'my (cl 9)'
f.	/**i** + ose/	→	**j**oːse	'all/the whole (cl 4; cl 9)'
g.	/**i** + uz + a/	→	**j**uːza	'it (cl 9) asked'
h.	/**i** + ereker + a/	→	**j**eːrekeːra	'it (cl 9) floated'

Recall from §2.2.4 that glide formation results in compensatory lengthening. This characteristic of Mushunguli distinguishes it from other recent analyses of glide formation in several Bantu languages, including Nambya (Mudzingwa & Kadenge 2011; Kadenge 2013), ciNsenga (Simango & Kadenge 2014), Karanga (Mudzingwa & Kadenge 2011), and Shona (Mudzingwa 2010). All of these analyses assume that the change from an underlying vowel to a corresponding glide violates a faithfulness constraint MAX-IO(μ) (defined below), which forces the associated segmental melody to be associated with the

[4]Examples (34c-d) are marginal, as truly animate class 3 nouns take parasitic class 1 agreement.

syllable onset in the output this which forces the reassociation of segmental material to the onset position of a syllable.

(35) MAX-IO(μ) (MAX-μ): A mora in the input must have a correspondent in the output. (Rosenthall 1997:146)

Because the above languages do not exhibit compensatory lengthening as a result of glide formation (or any other process), glide formation can be accounted for by having ONSET and a constraint penalizing long vowels *V: (Rosenthall 1997:147) dominate Max-μ.

(36) NO-LONG-VOWELS (*V:): assign a violation for any vocalic melody in the output that is associated to more than one mora (Rosenthall 1997:146)

This is illustrated in (37), using an abbreviated tableau from the ciNsenga analysis proposed by Simango & Kadenge (2014).

(37) Glide formation without CL (adapted from Simango & Kadenge 2014)

$i_1 + a_2wo$	ONSET	*V:	MAX-μ
a. $i_1.a_2wo$	*!		
b. $j_1a{:}_2wo$		*!	
c. ☞ j_1a_2wo			*

Two things to note here: first, this analysis assumes that all vowels are underlyingly moraic. I will be following this assumption as well. Second, notice that this analysis seems to assume that there is no faithfulness violation incurred by glide formation—or if there is one, it is not important enough to specify. What this tableau is actually suggesting is that glide formation is assumed to be *default*, and what is being penalized by the ranking *V: » MAX-μ is whether or not *compensatory lengthening* obtains.

That said, the fact that the proposed analysis of ciNsenga lacks a faithfulness violation incurred by glide formation is formally suspect. If vowels are underlyingly moraic, then glide formation eliminates the input association between mora and vowel–this *is* a faithfulness violation. I assume that this must be a notational shortcut, and while Simango & Kadenge (2014) do not justify it in the paper, I will do so for the analysis of Mushunguli.

To penalize glide formation, I will adopt the constraint MAX-ASSOCIATION (Keer 1999).[5]

(38) MAX-ASSOCIATION (MAX-ASC): If τ_1 is a mora in the input and it is associated to ζ_1 and $\tau_1 \Re \tau_2$, and $\zeta_1 \Re \zeta_2$, then τ_2 is associated to some ζ_2. (Keer 1999:47)

The change of a high vowel to a glide violates MAX-ASSOCIATION because glide formation–unlike coalescence–fully dissociates segmental content from a mora. I assume that high vowels and corresponding glides are otherwise featurally identical (Rosenthall 1997), so this is the minimal change necessary to resolve hiatus in these cases. Note as well that MAX-ASSOCIATION as defined above is also violated by segment or mora deletion, as these processes inherently eliminate an association line.

In Mushunguli, hiatus resolution does result in compensatory lengthening (except for when the two vowels in hiatus are identical; see §3.4). As such, we will need to assume that the opposite ranking than the one proposed for ciNsenga, {ONSET, MAX-μ} » *V: holds in this language. In order for glide formation to obtain, MAX-ASSOCIATION must be dominated by ONSET. To allow for a lengthened result, MAX-μ must also dominate *V:. To prevent deletion, MAX-V must also dominate MAX-ASSOCIATION. This is illustrated in (39).

(39) Glide formation with CL (Mushunguli)

[5] This is a modified version of NO-SPREAD (McCarthy 1997), which penalizes any cases in which there is a change in association between moras and their segmental melodies. MAX-ASSOCIATION as formulated only penalizes mora de-linking, not segmental re-association. I use this constraint because NO-SPREAD would penalize both mora de-linking and the reassociation of moras with other segmental content (see Keer 1999).

$i_1 + a_2sam + a$	ONSET	MAX-μ	MAX-ASC	*V:
a. $i_1.a_2sama$	*!			
b. ☞ $j_1a\!:_2sama$			*	*
c. j_1a_2sama		*!	*	

Notice in (39) that MAX-ASSOCIATION on its own does not actually play a role in ruling out either loser candidates. It would be possible to remove MAX-ASSOCIATION from this ranking and still obtain the same surface result, as illustrated in (40).

(40) Glide formation with CL (no MAX-ASC)

$i_1 + a_2sam + a$	ONSET	MAX-μ	*V:
a. $i_1.a_2sama$	*!		
b. ☞ $j_1a\!:_2sama$			*
c. j_1a_2sama		*!	

Here, we see that ONSET must dominate *V: to prevent blocking of hiatus resolution in favor of *not* making a long vowel. Similarly, we see that MAX-μ must dominate *V: to allow for a lengthened result. This is, of course, the Mushunguli equivalent of the ranking adopted by Simango & Kadenge (2014) for ciNsenga.

The weakness of MAX-ASSOCIATION can be illustrated further by considering deletion candidates, which violate MAX-V. Again, MAX-ASSOCIATION is violated by any form of deletion that also eliminates an association line; this includes mora and segment deletion. This means that any violation of MAX-V or MAX-μ entails violation of MAX-ASSOCIATION. As illustrated in the comparative tableau in (41), this combined with the previous observations means that MAX-ASSOCIATION does not actually make any crucial decisions.

(41) No crucial decisions made by MAX-ASSOCIATION

/i + a/	ONSET	MAX-V	MAX-μ	MAX-ASC	*Vː
☞ jaː				1	1
a. ∼ i.a	W			L	L
b. ∼ ja			W	=	L
c. ∼ Øaː		W		=	=
d. ∼ Øa		W	W	=	L

Here, we see that MAX-ASSOCIATION does not distinguish between any of the three loser candidates that resolve hiatus; these candidates are all ruled out by either MAX-V or MAX-μ (or both). The one candidate for which MAX-ASSOCIATION has the potential to make a decision is the losing fully-faithful candidate [i.a], where ONSET » MAX-ASC allows for glide formation with a lengthened result. However, its decision-making power is shared in this one case with *Vː, which is crucially dominated by MAX-μ to allow for a lengthened result. So, while it is the case that glide formation violates MAX-ASSOCIATION (insofar as there must be a faithfulness violation associated with glide formation), that violation is not important. As such, in the forthcoming tableau and discussion, I will generally be excluding MAX-ASSOCIATION.

At this stage, we can also begin definitively ruling out conceivable alternative hiatus resolution strategies, i.e. vowel deletion and consonant epenthesis. Both MAX-V and DEP-C (defined again below) will need to dominate *Vː.

(42) MAX-V: every vowel in the input has a correspondent in the output

(43) DEP-C: every consonant in the output has a correspondent in the input

Note that MAX-V is violated only by deleting the entirety of an underlying vowel's *segmental* content. It does not penalize mora deletion, and violation of neither MAX-V *nor* MAX-μ entails violation of the other. This distinction will prove important to accounting for both the fact that Mushunguli only has derived long vowels, and the analysis of what happens to sequences of identical vowels.

I have illustrated all of these strategies in tableau (44), including both lengthened and non-lengthened candidates. In subsequent tableaux, I will only include lengthened candidates, and other simplifications will be adopted, as discussed below.[6]

(44) Glide formation vs. other applicable strategies

	$u_1 + i_2 v + a$	DEP-C	MAX-V	MAX-μ	ONSET	*V:
a.	$u_1.i_2.va$				*!	
b. ☞	$w_1 i{:}_2.va$					*
c.	$w_1 i_2.va$			*!		
d.	$u_1.Ci_2.va$	*!				
e.	$Ø_1.i_2.va$		*(!)	*(!)		
f.	$Ø_1.i{:}_2.va$		*!			*

In (44), we see that the glide formation candidate (b) is optimal due to the low ranking of *V:. Its non-lengthened competitor in (c) is ruled out due to the higher rank of MAX-μ, and all other candidates (unresolved hiatus, deletion with and without lengthening, and epenthesis) are ruled out for similar reasons.

Because vowel deletion is a potential way to analyze some of the examples from (24) and (25), I will include MAX-V and candidates that violate it in subsequent tableaux and discussion, but I will henceforth just assume that DEP-C is undominated due to the complete unattestedness of consonant insertion, real or apparent.

The Hasse diagram in figure 3.1 summarizes the important rankings established for hiatus resolution thus far. Notice that at this point, MAX-V is unranked with respect to the other constraints; this is because glide formation with compensatory lengthening only violates *V:, so it harmonically bounds any candidate that violates more than one

[6] Three other tableau conventions are illustrated here that will be adopted throughout this analysis (unless otherwise noted): (a) subscript numerals to distinguish input segments (subscript $_P$ and $_S$ are also used when the distinction between prefix and stem is important); (b) a period <.> to denote a syllable boundary; and (c) a null sign <Ø> to denote a deleted segment. I follow the correspondence theory convention of assuming that deleted material is entirely deleted.

constraint. We will discuss glide formation in postconsonantal contexts in §3.5, and for now will turn our attention to coalescence.

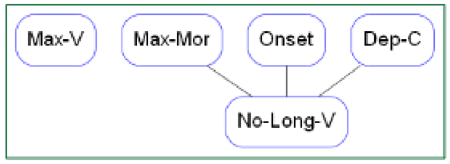

Figure 3.1: Ranking established for glide formation

3.3 Coalescence of non-identical vowels

Coalescence applies when a low vowel precedes a high vowel. The result is a single lengthened mid vowel corresponding to the place feature of V_2: front (and unrounded) *e* if /i/; back (and rounded) *o* if /u/. For reason that will become clearer shortly, I will refer to this pattern as high coalescence; examples are given in (45).

(45) Examples of high coalescence

a.	/k**a**+**i**v+is+a/	k**e:**vi:sa	'(s)he heard a lot'
b.	/m**a**+**i**no/	m**e:**no	'teeth'
c.	/h**a**+**i**+tʃi+tabu/	h**e:**tʃitabu	'to the book'
d.	/ni+**a**+**u**+to+a/	n**o:**to:a	'I am hitting it (cl 3)'
e.	/w**a**+**u**mbal+a/	w**o:**mba:la	'they are piled up'

Coalescence is formally represented here as many-to-one correspondence, in violation of the faithfulness constraint UNIFORMITY, following McCarthy & Prince (1995) (see also Pater 1999, but cf. Keer 1999).

(46) UNIFORMITY: no element in the output has multiple correspondents in the input.

Coalescence also incurs several identity violations. The raising of the first, outer prefix vowel from low to mid violates IDENT(low).

(47) IDENT(low): assign a violation to any output segment whose value for [±low] differs from that of its input correspondent

The lowering of the second, inner stem vowel from high to mid also violates IDENT(high).

(48) IDENT(high): assign a violation to any output segment whose value for [±high] differs from that of its input correspondent

Because UNIFORMITY, IDENT(low), and IDENT(high) are all violated by the winning candidates of high coalescence, and because none are violated by alternative candidates we have thus far considered (faithful hiatus, vowel deletion), all three constraints must be ranked below ONSET and MAX-V. This is illustrated in (49).[7]

(49) Coalescence of a low + high sequence

	$ka_1 + i_2v\text{-}a$	ONSET	MAX-V	ID(hi)	ID(lo)	UNIFORMITY
a.	$ka_1.i_2.va$	*!				
b. ☞	$ke\!:_{1,2}va$			*	*	*
c.	$k\emptyset_1 i\!:_2va$		*!			

This ranking has an important formal consequence. Consider again the examples of /low + mid/ hiatus resolution, from 2.2.4.

(50) Loss of /a/ before non-high vowels

a.	/k**a** + **o**mbok + a/	**ko:**mbo:ka	'(s)he went far'
b.	/k**a** + **e**res + a/	k**e:**re:sa	'she gave birth'
c.	/w**a** + **e** + jag + a/	w**e:**ja:ga	'they scratched themselves'
d.	/wa + tumbiri w**a** + **o**se/	watumbiri w**o:**se	'all vervets'

This repair was previously described as elision. However, coalescence of a low + mid sequence is a formal possibility we must consider: this candidate does not violate IDENT(high), but it *does* violate IDENT(low) and UNIFORMITY. This is a proper subset of the constraints violated by high coalescence; as a result, the ranking just established for what was purposefully referred to as "high" coalescence now excludes the possibility of

[7] I assume, consistent with the definition of MAX-ASSOCIATION, that coalescence of two vowels into a single vowel retains the association between moras and vowels; i.e. coalescence does not violate MAX-ASSOCIATION.

deletion in this context. What appears to be deletion is in fact another form of coalescence, which we will refer to as mid coalescence.

(51) Mid coalescence

	$wa_1 + o_2se$	ONSET	MAX-V	ID(hi)	ID(lo)	UNIFORMITY
a.	$wa_1.o_2se$	*!				
b. ☞	$wo\!:_{1,2}se$				*	*
c.	$w\emptyset_1 o\!:_2se$		*!			

Table 3.1 summarizes the faithfulness constraint violations, in addition to UNIFORMITY, incurred by the winning candidates of mid and high coalescence.[8]

Table 3.1: Faithfulness violations incurred by winning coalescence candidates

Input	Output	Violations	
$/a_1 + i_2/$	$e\!:_{1,2}$	V_1: IDENT(low) V_2: IDENT(high)	high coalescence
$/a_1 + u_2/$	$o\!:_{1,2}$	V_1: IDENT(low) V_2: IDENT(high)	high coalescence
$/a_1 + e_2/$	$e\!:_{1,2}$	V_1: IDENT(low) V_2: none	mid coalescence
$/a_1 + u_2/$	$o\!:_{1,2}$	V_1: IDENT(low) V_2: none	mid coalescence

There are three observations of note here. The first is, again, that the winning mid coalescence candidates violate a proper subset of the constraints violated by the winning high coalescence candidates. The second is that only in cases of high coalescence does the second vowel (which usually corresponds to the initial vowel of the stem) incur any violations at all. The third is that the only constraint violated by the second vowel in cases of high coalescence is IDENT(high). All of these facts will become important in the discussion of the exceptions in 4.

[8]As will be discussed in §5.1.1, high and mid coalescence both also violate IDENT(V-Place), due to the fact that the placeless prefix /a/ becomes associated with a place feature. IDENT(V-Place) violations are more important for the discussion of the exceptions in §4, so I abstract away from this for now.

Before moving on, there are some further aspects of this ranking that need to be explored. First, under this ranking, glide formation is guaranteed in high + low sequences (e.g. /i + a/). Both coalescence and glide formation violate *Vː (because compensatory lengthening applies). However, glide formation does not share faithfulness violations with coalescence, as shown by the tableau in (52).[9] As a result, glide formation is the preferred strategy in this context.

(52) Glide formation

$i_1 + a_2$ŋgu	Ons	Id(hi)	Id(lo)	Uniformity	*Vː
a. $i_1.a_2$.ŋgu	*!				
b. ☞ j_1aː$_2$.ŋgu					*
c. eː$_{1,2}$.ŋgu		*(!)	*(!)	*(!)	*

To ensure coalescence in low + high sequences, we need to prevent gliding of low prefix vowels. Gliding of low vowels is a very rare strategy for resolving hiatus (see e.g. Casali 2011 §2.4, note 6), a typological fact that is beyond the scope of this work. There are two results of potential concern; first the potential for low /a/ to become some kind of low (or otherwise non-high) glide, a result which I presume to be ruled out by an undominated constraint similar in effect to the HiGlide constraint discussed by Casali (1997:516): 'Glides must be [+high].' Like consonant epenthesis and Dep-C, I henceforth ignore all low-gliding candidates and the HiGlide constraint.

The second and greater concern is the possibility of low /a/ becoming high *j* or *w*. This low-gliding candidate must lose both to high coalescence. All three hypothetical candidates violate Ident(low), and both low gliding and high coalescence violate Ident(high). However, coalescence violates Uniformity, while glide formation does not. This means that all things being equal, gliding is preferred to coalescence.

[9] Uniformity dominates *Vː here (and in all subsequent tableaux where both constraints are illustrated) because as the former is defined in (46), it could in principle be violated by merging two moras as well. If *Vː dominated Uniformity, then a short output candidate *[keva] would be chosen for all coalescence candidates.

(53) Low gliding wins over coalescence

ka_1-i_2v-a	ONS	ID(HI)	ID(LO)	UNIFORMITY	*V:
a. $ka_1.i_2$.va	*!				
b. ☹ $ke{:}_{1,2}$va		*	*	*!	*
c. ☞ $kw_1i{:}_2$va		*	*		*

We can mitigate this somewhat by re-introducing MAX-ASSOCIATION; provided it dominates UNIFORMITY, we will not choose to form glides in this context.

(54) Coalescence wins over low gliding

ka_1-i_2v-a	ONS	ID(HI)	ID(LO)	MAX-ASC	UNIFORMITY	*V:
a. $ka_1.i_2$.va	*!					
b. ☞ $ke{:}_{1,2}$va		*	*		*	*
c. $kw_1i{:}_2$va		*	*	*!		*

However, this will prove insufficient: if we had a situation whereby somehow coalescence was blocked, this analysis would predict that the low gliding candidate will beat the candidate with unresolved hiatus. As will be discussed in §4.2, there are cases where unresolved hiatus beats out all other alternative repairs; as such, we will want to be able to independently rule out low gliding. To do this, we will introduce an undominated local conjunction of IDENT(high) and IDENT(low).

(55) **IDENT(high)&IDENT(low)**: assign a violation to any output segment whose value for [±high] and [±low] differs from that of its input correspondent

This conjunction is violated by raising a low vowel all the way to a high vowel or glide (and vice versa). I follow Łubowicz (2005) in assuming that a local conjunction is only violated if the conjuncts share a locus of violation. Low gliding violates IDENT(high)&IDENT(low) because the locus of violation of both conjuncts is the initial low vowel ($/a_1/ \rightarrow |w_1|$). Regular high coalescence does not violate this conjunction, however, because as we saw in Table 3.1, violations of IDENT(high) and IDENT(low) do not share a locus of violation; IDENT(low) is violated by V_1 and IDENT(high) is violated by V_2.

(56) Ensuring coalescence

ka_1-i_2v-a	ONS	ID(hi)&ID(lo)	ID(HI)	ID(LO)	UNIF
a. $ka_1.i_2$.va	*!				
b. ☞ $ke{:}_{1,2}$va			*	*	*
c. $kw_1i{:}_2$va		*!	*	*	

The following Hasse diagram summarizes some important aspects of the constraint ranking so far. This ranking is consistent with, but further embellishes and articulates, the ranking that was established at the end of §3.2. MAX-V now has a ranking (it is undominated) along with ONSET; they are also unranked with respect to each other. These dominate the three constraints penalized by coalescence. Finally, UNIFORMITY dominates *V: (rendered here as No-Long), to allow for a lengthened result.

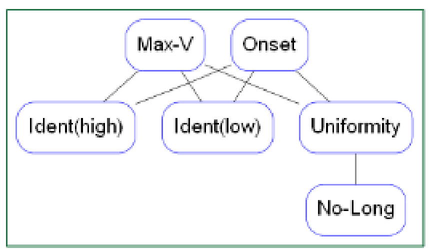

Figure 3.2: Ranking established for glide formation and coalescence

3.4 Simplification and vowel length

So far, we have only dealt with hiatus resolution contexts that result in compensatory lengthening. However, as discussed in §2.2.4, sequences of identical vowels are simplified to a single *short* vowel (unless the resultant vowel is in the penultimate syllable, in which case it is lengthened by penultimate lengthening). Examples of simplification are repeated in (57).

(57) *Simplification*

a.	/s**i** + **i**taŋg + a/	→	s**i**taːŋga	'I called'
b.	/w**a** + **a**mbiz + a/	→	w**a**mbiːza	'they helped'
c.	/k**u** + **u**mbal + a/	→	k**u**mbaːla	'to be piled up'

This distinction between simplification and the other hiatus resolution strategies suggests that it is due to an operation distinct from the operations responsible for glide formation and coalescence, the obvious candidate being deletion of one or the other of the two identical vowels. It could then be claimed that the coalescence operations (as well as glide formation) target only the vowel melody, preserving its timing unit that is consequently filled by lengthening, while the simplification *qua* deletion operation targets both the vowel melody and its timing unit, thus precluding lengthening.

There are two reasons why such an approach will not work within the larger analysis of Mushunguli. One is the form and behavior of some exceptions to hiatus resolution, discussed in considerably more detail in §4.2. The other, which I will address now, is simply the behavior of the regular grammar.

First, abstracting away from the issues of lengthened versus non-lengthened results, we have already determined that MAX-V must dominate all faithfulness constraints violated by high coalescence; this is why the resolution of /low + mid/ sequences is now viewed as coalescence, not deletion. Merging two identical vowels violates no other faithfulness constraints besides UNIFORMITY; as such, coalescence is predicted by our analysis to be optimal over vowel deletion. To distinguish this result from high and mid coalescence, I refer to it as *identity coalescence*.

(58) Identity coalescence (abstracting away from lengthening)

	$ka_1 + a_2sam + a$	ONSET	ID(hi)&ID(lo)	MAX-V	UNIF
a.	$ka_1.a_2.sa.ma$	*!			
b. ☞	$ka_{1,2}.sa.ma$				*
c.	$kØ_1aː_2.sa.ma$			*!	
d.	$kw_1aː_2.sa.ma$		*!		

For sequences of two identical *high* vowels, we will need to find a way to block tautosyllabic glide-vowel sequences that are homorganic for height and place (*ji* and *wu*). Fortunately, there is already independent motivation for this: there are no derived or word-internal surface *ji* or *wu* sequences in this language.[10] We can thus assume that there is yet another undominated markedness constraint that penalizes homorganic glide-vowel sequences, which we will call $*G_iV_i$ (cf. Baković 2006).

(59) $*G_iV_i$: assign a violation for any surface tautosyllabic glide-vowel sequence that is homorganic for height and place

The operation of this constraint is illustrated in the tableau in (60).

(60) Identity coalescence, not glide formation

	$i_1 + i_2$taŋg + a	ONSET	$*G_iV_i$	UNIFORMITY
a.	i_1.i_2.ta.ŋga	*!		
b. ☞	$i_{1,2}$.ta.ŋga			*
c.	j_1iː$_2$.ta.ŋga		*!	

We have established that coalescence is optimal over deletion and glide formation, but we have not yet established how high and mid coalescence result in compensatory lengthening while simplification *qua* identity coalescence does not. This fact can be linked with another generalization about the language, which is that there is no *contrastive* vowel length in this language, but surface long vowels nevertheless obtain.

Richness of the Base (Prince & Smolensky 1993/2004:205, 225) requires that our grammar must be able to account for all hypothetical inputs. Currently, our analysis predicts that potential underlying long vowels should surface faithfully (61).

(61) *Faithfulness to underlying vowel length*

[10] Save for two possible exceptions, see §6.1.1.

i:	MAX-μ	*V:
a. ☞ i:		*
b. i	*!	

We cannot safely assume that potential underlying long vowels would simply surface faithfully, nor can we safely abstract away from this problem. Words borrowed from languages that have either contrastive or non-contrastive vowel length seem to be uniformly reanalyzed with short vowels, e.g. *sukulu*∼*skulu* 'school' (English: [sku:l]); *bakeri* 'cup' (Maay: [bakeeri] (Paster 2006)). Proto-Bantu is also typically reconstructed with a length distinction. While neither of these facts necessarily motivate a synchronic analysis of vowel shortening, taken in tandem with the asymmetric behavior of identity coalescence, a more substantive analysis is warranted.

To this end, we will capitalize on an important distinction between the set of inputs that result in surface long vowels and the set of inputs that does not. Compensatory lengthening only applies when the input vowel sequence is non-identical, meaning that the surface long vowel is associated with two moras from *non-identical* segmental melodies. All inputs that do not result in a surface long vowel (besides underlyingly short vowels) are either a sequence of two identical vowels, or two moras attached to a single segment; in both cases, the moras were (or are) associated with *identical* melodies. That the language makes a distinction between these two cases suggests that the link between melody and mora is stronger than may be typically assumed; that is, segmental content is in some sense 'visible' to moras, and in Mushunguli, the preference is to have non-matching moras. As such, I will propose a constraint, *MATCH-μ, which penalizes this distinction.

(62) *MATCH-μ: assign a violation for each segment associated with two moras in the output if both of those moras are associated with identical segmental melodies in the input

*MATCH-μ is violated both by long output vowels associated with two featurally identical input correspondents (= identity coalescence), as well as long output vowels associated with a single input correspondent (= underlying long vowels).[11] This constraint captures the intuition that compensatory lengthening is restricted to cases with non-identical inputs. It also captures the descriptive fact that although vowel length is non-contrastive, compensatory lengthening nevertheless obtains.[12]

Ranking *MATCH-μ above MAX-μ will result in underlying long vowels being shortened.

(63) *Shortening of underlying long vowels*

iː	*MATCH-μ	MAX-μ	*Vː
a. iː	*!		*
b. ☞ i		*	

For simplification contexts, adding this constraint to our ranking (along with $*G_iV_i$) ensures that identity coalescence with mora pruning is optimal, even in cases where glide formation or deletion could in principle apply. This is illustrated in (64).

(64) *Identity coalescence with mora pruning*

[11] I assume that penultimate lengthening is a postlexical process that violates a constraint DEP-μ. As *MATCH-μ is currently defined, a long vowel created by mora insertion would not violate it, as an inserted mora has no segmental melody in the input.

[12] This is comparable to patterns of compensatory lengthening exhibited in Yine. Yine has no phonemic geminate consonants or long vowels (Lin 1997; Hanson 2010, cf. Matteson 1965). Despite this, CL triggered by consonant deletion obtains. Unlike Mushunguli, lengthening in Yine is only triggered by repair of word-internal geminates, but this is attributable to the fact that *only* geminates are repaired in Yine; disagreeing consonant clusters are tolerated, regardless of sonority.

	/$i_1 + i_2$taŋg + a/	ONS	*G_iV_i	*MATCH-μ	MAX-V	MAX-μ	UNIF	*V:
a.	$i_1.i_2$taŋga	*!						
b.	$j_1i_2$taŋga		*!			*		
c.	j_1i:$_2$taŋga		*(!)	*(!)				*
d.	$Ø_1i_2$taŋga				*!	*!		
e.	$Ø_1$i:taŋga			*!	*(!)			*
f. ☞	$i_{1,2}$taŋga					*	*	
g.	i:$_{1,2}$.ta.ŋga			*!			*	*

Notice that in (64), the ranking *MATCH-μ » MAX-μ ensures that lengthened candidates will *always* lose when the two input vowels are identical. Conversely, the ranking MAX-μ » *V: ensures that shortened candidates will *always* lose for other forms of coalescence, as illustrated in (65).[13]

(65) *High coalescence always results in lengthening*

	ka_1-i_2v-a	ONS	*MATCH-μ	IDENT(hi)	MAX-μ	*V:
a.	$ka_1.i_2$.va	*!				
b.	$ke_{1,2}$va			*	*!	
c. ☞	ke:$_{1,2}$va			*		*

The Hasse diagram in Fig 3.3 schematizes the ranking established for mora pruning. Here, ONSET and *MATCH-μ ("No-Match-Mora") dominate MAX-μ, which in turn dominates *V:.

With simplification *qua* identity coalescence now accounted for, we are largely going to abstract away from vowel length for the remainder of the analysis of both regular and exceptional forms in this language, as all of the exceptional patterns of primary interest do not have behaviors that are attributable to vowel length.

[13]The only situation that contravenes this generalization is the fact that hiatus resolution of non-identical vowel sequences in utterance-final position do not exhibit lengthening. I will abstract away from the details of prosody in this language and simply assume there is an undominated constraint *V:$_{Pwd}$] that penalizes word-final long vowels.

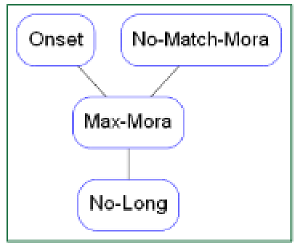

Figure 3.3: Ranking for simplification and mora pruning

As a final note before moving on to the last set of hiatus resolution contexts, consider that there is a hierarchy of costliness with respect to the three types of coalescence. As was illustrated in Table 3.1, mid coalescence violates a proper subset of the constraints violated by high coalescence; now we see that identity coalescence in turn violates a proper subset of the constraints violated by mid coalescence. Furthermore, of the three, high coalescence is the only type for which one of the constraints — IDENT(high) — is violated solely by a change to the second vowel (which typically is the initial vowel of the stem). For the others, IDENT(low) is borne solely by V_1, and violation of UNIFORMITY is shared between the two vowels. This observation, summarized in (66), will be relevant to the analysis of the exceptions in §4.

(66) Hierarchy of costliness (adapted from Hout 2019a)

	IDENT(high)	IDENT(low)	UNIF
High (a_1+i_2 → $e_{1,2}$)	V_2	V_1	V_1,V_2
Mid (a_1+e_2 → $e_{1,2}$)		V_1	V_1,V_2
Identity (a_1+a_2 → $a_{1,2}$)			V_1,V_2

An interesting consequence of this analysis thus far is that it suggests that the mere existence of high coalescence excludes the possibility of a deletion analysis for any other cases of potential coalescence (at least within the same morphosyntactic context). That is to say, absent some form of domain or lexical specification, coalescence and deletion are now predicted to be mutually exclusive repairs. *However*, it is critical to note that these facts are insufficient to determine the ranking between ONSET and MAX-V; that is, the fact that deletion is an impossible solution in coalescence contexts does not entail that it is impossible in *every* hiatus context.

3.5 Postconsonantal prevocalic high vowels

As previously established, glide formation applies to non-identical prevocalic high vowels. In situations where no consonant precedes the initial high vowel, the surface result is always a full glide, as was seen in (34). However, as was seen in (8), /Ci + V/ and /Cu + V/ contexts do not have the same result. Rather, either a secondary articulation occurs, or no exponent of the vowel surfaces.

We will start with the simpler case of /Cu + $V_{[\text{-rd}]}$/ contexts. Here, a reflex of the underlying /u/ surfaces as a secondary articulation on the preceding consonant. Examples are given in (67):

(67) Secondary articulations

a.	/k**u** + itaŋg + a/	→	**k**^{w}i:ta:ŋga	'to call'
b.	/m**u** + ana j**u** + edi/	→	**m**^{w}a:na **j**^{w}e:di	'good (cl 1) child'
c.	/k**u** + e + jag + a/	→	**k**^{w}e:ja:ga	'to scratch oneself'

We have already established that Mushunguli avoids complex onsets, repairing them via secondary articulation or consonant deletion (as was seen in §2.2.1). This suggests that the constraint *COMPLEX is active in the grammar.

(68) *COMPLEX: No more than one C or V may associate to any syllable position node.
(Prince & Smolensky 1993/2004:96)

Because we are assuming that glide formation results in the reassociation of a vowel melody to the onset of a syllable, in /CV + V/ contexts this would create a complex onset. As such, as long as *COMPLEX is ranked at least as high as ONSET, we predict full glides to be penalized. To distinguish these candidates from candidates with secondary articulations, we will also introduce two constraints: IDENT(back)–which penalizes palatalization–and IDENT(round)–which penalizes labialization.[14]

(69) IDENT(back): assign a violation to any output segment whose value for [±back] differs from that of its input correspondent

(70) IDENT(round): assign a violation to any output segment whose value for [±round] differs from that of its input correspondent

For the sake of space, we will combine these into a single cover constraint IDENT(V-Place).

(71) IDENT(V-Place): assign a violation to any output segment whose value for [±back] or [±round] differs from that of its input correspondent

IDENT(V-Place) here is intended as a cover constraint for both IDENT(back) and IDENT(round); it is *not* a local conjunction. As it is currently defined, IDENT(V-Place) is violated by changing, adding, or losing a vowel place feature. For example, an input-output pair /$\mathrm{i}_1+\mathrm{u}_2$/ → [$\mathrm{u}_{1,2}$] violates IDENT(V-Place) twice, penalizing the change in [-back] and [-round] incurred by /i_1/. An input-output pair /$\mathrm{i}_1+\mathrm{u}_2$/ → [$\mathrm{ɯ}_{1,2}$] also violates it twice, penalizing the change from [-back] → [+back] incurred by /i_1/ and the change from [+round] → [-round] incurred by /u_2/. Because central vowels are assumed to be placeless, an input-output pair /$\mathrm{i}_1+\mathrm{u}_2$/ → [$\mathrm{ɨ}_{1,2}$] would violate it three times, due to the loss of both place features and the change from [+round] to [-round].

For the practical purposes of this analysis, we will generally only be dealing with candidates that violate IDENT(V-Place) once. These include all attested forms of (non-identity) coalescence (because /a/ is placeless, and V_2 donates its place feature), as well as the re-association of a vowel place feature to a consonant that lacks that feature. For

[14]These are not the only constraints violated by a segment with a secondary articulation; however, they are the only unique faithfulness violation incurred by the creation of a labial or palatal secondary articulation. We will return to the details of onset structure in §5.1.1.

example, hypothetical /kw/ → [k^w] violates IDENT(V-Place) due to the addition of a [+round] feature, as does hypothetical /k^w/ → [kw] or [k] (due to the loss of a [+round] feature).

Labialization can be ensured over glide formation or vowel deletion by ranking *COMPLEX over IDENT(V-Place). This is illustrated below in (72).[15]

(72) *Labialization*

	$ku_1 + i_2v + a$	ONSET	MAX-V	*COMPLEX	ID(V-Pl)
a.	$ku_1.i_2.va$	*!			
b. ☞	$k^w{}_1 iː_2.va$				*
c.	$kw_1 iː_2.va$			*!	
d.	$kØ_1 iː_2.va$		*!		

However, labialization is not general. For cases of /Cu + o/, no glide or secondary articulation ever surfaces, as illustrated in (73).[16]

(73) *Failure of glide formation in /Cu + o/ contexts*

a.	/k**u** + omal + a/	→	k**oː**maːla	'to dish up ugali'
b.	/m**u** + oger + a/	→	m**oː**geːra	'you (pl) swam'
c.	/k**u** + ombok + a/	→	k**oː**mboːka	'you went far'
d.	/mu + gosi j**u** + ose/	→	m̩gosi j**oː**se	'the whole man'

A similar pattern emerges in all /Ci + V/ contexts; here, we might expect that an analogous secondary articulation (e.g. palatalization) could occur. However, this is not the case: in these contexts no glide or secondary articulation *ever* surfaces. This applies regularly to all prefix + root and prefix + prefix boundaries regardless of lexical class, as illustrated in (74):

(74) Distribution of /Ci + V/

[15] I assume also that the tendency for labialized segments to be realized with velarization is a late (postlexical) reinterpretation of labialization.

[16] /$Cu_1 + u_2$/ outputs [$u_{1,2}$], due to *G_iV_i and *MATCH-μ. The same is true for /$Ci_1 + i_2$/.

a.	/tʃ**i** + asa/	→	**tʃa:**sa	'we divorced'	Subject Prefix
b.	/si + d**i** + aza/	→	si**da:**za	'I lost it (cl 5)'	Object Prefix
c.	/m**i** + ezi/	→	**me:**zi	'months'	Noun Class Prefix
d.	/z**i** + etu/	→	**ze:**tu	'our (cl 10)'	Demonstrative
e.	/s**i** + u + siɲ + a/	→	**su:**si:ɲa	'I looked at it (cl 3)'	Subj + Obj

At first glance, it seems that we should be able to analyze all of these cases as vowel deletion. For the /Cu + o/ examples in (73), we would assume that a specific markedness constraint penalizing labialized segments before round vowels (e.g. $*C^wO$) dominates MAX-V; this will block labialization in favor of deletion. Similarly, for /Ci + V/ contexts, we would assume that a general constraint against palatalized segments (e.g. NOPAL) dominates MAX-V, which again would block the application of palatalization.

However, this interpretation of the data is already precluded by the current analysis. Even if we assume ONSET dominates MAX-V (which is required by a deletion analysis), the fact that the coalescence-related faithfulness constraints must be dominated by MAX-V means that if we block glide formation, we predict that a coalescence candidate will win. This is illustrated by the tableau in (75), using a /Ci + a/ context.

(75) *Coalescence of a high + low sequence*

si + di_1 + a_2z + a	ONS	*CPLX	NOPAL	MAX-V	ID(hi)	ID(V-Pl)	UNIF
a. si.di_1.a_2.za	*!						
b. si.$d^j{}_1$a:$_2$.za			*!			*	
c. si.dj_1a:$_2$.za		*!					
d. ☹ si.$dØ_1$a:$_2$.za				*!			
e. ☞ si.de:$_{1,2}$.za					*	*	*

This result reinforces the observation that coalescence and vowel deletion are mutually exclusive repairs in the same morphosyntactic context, at least for a language like Mushunguli. Unfortunately, this leaves us unable to account for the patterns exhibited by /$CV_{[+hi]}$ + V/ sequences. This is a genuine ranking paradox: the ranking that allows us to have labialization in /Cu + $V_{[-rd]}$/ contexts (76) is incompatible with the ranking that will allow for vowel deletion in /Ci + V/ and /Cu + o/ contexts (77).

(76) Ranking for labialization: {ONSET, MAX-V, *COMPLEX} » IDENT(V-Place)

/Cu + a/	ONSET	MAX-V	*COMPLEX	IDENT(V-Place)
☞ C^waː				1
a. ~ Cu.a	W			L
b. ~ CØaː		W		L
c. ~ Cwaː			W	L

(77) Ranking for vowel deletion: {ONSET, *COMPLEX, IDENT(V-Place)} » MAX-V

/Ci + a/	ONSET	*COMPLEX	IDENT(V-Place)	MAX-V
☞ CØaː				1
a. ~ Ci.a	W			L
b. ~ Cjaː		W		L
c. ~ C^jaː			W	L

Fortunately, this puzzle will prove to be solvable: in §4.2, we will see that the presence of an exceptional pattern will clarify the ranking between MAX-V and ONSET; namely, MAX-V must dominate ONSET. This will exclude the possibility of vowel deletion in *any* context, a fact which will have implications both for the missing portion of our analysis, as well as the analysis of other exceptions in the language (§5.1). With vowel deletion ruled out, the only option remaining to us will be glide formation. These derived glides will then be deleted or changed to secondary articulations at a later derivational level (Kiparsky 2015). This analysis will in turn be supported by the behavior of a second set of exceptions, whose behavior is consistent both with the adoption of multiple levels and the wholesale exclusion of vowel deletion as a hiatus resolution strategy in Mushunguli.

This chapter, in part, contains material that has previously appeared in San Diego Linguistics Papers 6. Hout, Katherine, 2016. It also, in part, contains coauthored material currently being prepared for submission for publication. Hout, Katherine; Baković, Eric. The dissertation author was the primary investigator and author of all of this material.

4 Exceptions in Conspiracies

The preceding chapter demonstrated that hiatus is regularly and productively repaired at morpheme boundaries. These repairs form a conspiracy: multiple context-sensitive repairs are attested, all serving the same goal of eliminating heterosyllabic sequences of adjacent vowels. However, we were left with a paradox: the ranking that seemed to be required to account for patterns of coalescence and $/_{\sigma}V_{[+\text{high}]}+V/$ contexts could not account for glide formation in postconsonantal contexts.

In this chapter, we will also see that these repairs are not exceptionless–there are three sets of exceptions, each representing a separate typological instantiation of exceptional blocking: *simple blocking*, in which only one repair is prevented from applying (§4.2); *walljumping*, in which the regular repair is blocked and an alternative applies (§5.1); and *total non-participation*, in which all applicable repairs are blocked (§5.2). Examples of each are given in Table 4.1; note the superscripts on the underlying representations of the exceptions in question, which denote members of each set.

Table 4.1: Three forms of exceptional blocking in Mushunguli

Type	UR	Expected	Actual	Gloss
Simple Blocking	/ka + it^{L1} + a/	*[keːta]	[ka.iːta]	'(s)he went'
	/ku + it^{L1} + a/	[kʷiːta]	[kʷiːta]	'to go'
Walljumping	/di^{L2} + aŋgu/	*[daːŋgu]	[ɟaːŋgu]	'my (cl 5)'
	/ku + di^{L2} + an + a/	*[kudaːna]	[kuɟaːna]	'to eat together'
Total Non-Participation	/ka + on^{L3} + a/	*[koːna]	[ka.oːna]	'(s)he saw'
	/si + on^{L3} + a/	*[soːna]	[si.oːna]	'I saw'

In this and the following chapter, I will demonstrate that analysis of these exceptions using lexically-indexed constraints (Pater 2000, 2010) reveals that they do not undermine the hiatus resolution conspiracy. Rather, they act as both reifying and reinforcing agents in the grammar, both obeying it and imposing their own restrictions upon it.

In §4.1, I discuss in detail the formal properties of lexical indexation (Pater 2000, 2010), with the goal of showing how these properties make cross-linguistic and language-internal predictions about the typology of blocking. In §4.2, I introduce the simple blocking set of exceptional stems, which resist high coalescence, but not glide formation or identity coalescence. I demonstrate that the introduction of an indexed constraint into the grammar to handle this pattern both captures structural generalizations about the set, and also restricts possible analyses of both other exceptions and regular forms.

The most important of these restrictions is the elimination of vowel deletion as a solution to hiatus resolution; this has direct implications for the analysis of postconsonantal glide formation. In §5.1, we will see that with vowel deletion eliminated, the only option remaining to us is to adopt lexical levels in the form of Stratal OT (Bermudez-Otero 2011; Kiparsky 2015). This result is reinforced by the behavior of the second, walljumping, set of exceptions, which exceptionally undergo palatalization, an otherwise unattested repair in the language. These exceptions exhibit behavior consistent with the adoption of levels, and are used to explore the implications of introducing indexed constraints into Stratal OT.

Finally, in §5.2, I present the third set of total non-participants in hiatus resolution, and demonstrate that their forms and behaviors are similarly consistent with the form and behavior of the exceptions discussed in the preceding sections. This analytical unity provides evidence that exceptions are not only governed by higher-level aspects of the grammar, they reinforce aspects of a complex conspiracy, and clarify otherwise ambiguous aspects of the grammar. These observations suggest that while exceptions are lexical, this is not equivalent to them being extragrammatical.

4.1 Lexical indexation

There are a myriad of formal mechanisms to capture and explain aspects of lexical exceptions. The theory I will adopt in this dissertation is the version of constraint indexation proposed by Pater (2000, 2010). This theory allows exceptional morphemes to be linked to special clones of otherwise universal constraints, which may be ranked differently than regular (= unindexed) instances of those constraints. Going forward, unless otherwise noted, any exceptional morphemes and the constraints that they are indexed to will be notationally distinguished by a superscript, i.e. CON^{L}. Following the example from Table 4.1, a number is also used to differentiate sets of exceptions and indexed constraints, e.g. CONA^{L1}, CONB^{L2}.

All forms of constraint indexation assume that morphemes can be indexed to a clone of an existing constraint in the grammar, which is (typically, but not always) ranked somewhere above its counterpart. Content-wise, an indexed constraint is evaluated in the same way as its unindexed counterpart, but its scope is limited to the morphemes associated with it; unindexed morphemes and morphemes indexed to other constraints are invisible to these constraints. For example, IDENT(high) and its indexed variant IDENT(high)^{L} are both violated by a change to the [±high] feature of a segment, but IDENT(high)^{L} is only violated if that segment belongs to an indexed morpheme. This effectively subdivides the lexicon into "exceptional" and "non-exceptional" components; however, all components of the lexicon, exceptional or regular, are subject to the same constraint rankings modulo the indexed constraints themselves. This is distinct from lexically-specific rankings (cophonologies), which assume both an (arbitrarily) subdivided lexicon and the existence of a separate grammar for each subdivision.[1]

[1]I am not going to address seriously the ongoing debate regarding whether indexed constraints

4.1.1 Characteristics specific to this form of lexical indexation

Pater's theory has some aspects that distinguish it from other similar proposals. The most often-discussed factor is that it allows for indexed markedness constraints; prior proposals in this vein generally only allow lexically-specific faithfulness (e.g. Ito & Mester 1995a, 1999a, 2001; Fukazawa 1999; Alderete 2001), rejecting indexed markedness on the grounds that it is too powerful. However, Pater argues that indexed markedness allows for the straightforward capture of exceptional (non-)triggering patterns, and there are indeed cases of exceptionality that appear to be driven by a preference to avoid some especially marked structure, making indexed markedness more descriptively attractive as well (see e.g. Pater's discussion of exceptional blocking by markedness in Finnish). As will be discussed in §5.2, indexed markedness also makes it possible to straightforwardly capture the behavior of total non-participants without appeal to multiple indexed constraints.

The two other aspects of this theory that are relevant to the case presented in this dissertation are locality and inconsistency detection. I will discuss each briefly, and then move on to typological predictions made by this theory.

4.1.1.1 Locality

The most critical restriction on this theory is the locality condition, stated below in (78).

---

or cophonologies are better at capturing , especially given that as newer exceptionality formalisms adopt weighted constraints, the formal distinction between these two classes of theories has become increasingly muddy. Some authors do refer to diacritic uses of constraint weighting as "cophonology" (e.g. Sande 2019); however, these authors are generally either using "cophonology" to mean a non-arbitrary division of the lexicon (which is what Inkelas et al. (1997) and Inkelas & Zoll (2007) argue cophonologies should primarily be used for), or are comparing some form of scaled constraint *weighting* to ranked indexed constraints, without seriously considering how the formal implementation of lexically-specific constraints into a constraint weighting system would differ substantively from the formal implementation of lexically-specific rankings.

(78) Locality: lexically-indexed constraints "apply if and only if the locus of violation contains some portion of the indexed morpheme" (Pater 2010:133)

The concept 'locus of violation' is familiar from McCarthy (2003, 2007), and refers to any portion of a form that minimally meets the structural description of the constraint. The locality condition ensures that a lexically-indexed constraint is only violated by each portion of a form that minimally meets the constraint's structural description and that includes an exponent of the indexed morpheme. This can be illustrated by considering a hypothetical language that, unlike Mushunguli, *truly* obeys ONSET. We will call this language Toy. The grammar of Toy handles word-initial onset violations by deleting the initial vowel, requiring a ranking ONSET » MAX. Thus, a form like /apaka/ will surface as [Øpaka].

(79) Initial deletion

/apaka/	ONSET	MAX
a. .a.pa.ka	*!	
b. ☞ Øpa.ka		*

Let us say that this language has a root, /abaka^{L}/ which exceptionally blocks deletion. This can be captured indexing this root to a clone of MAX, MAX^{L}. Ranking the clone above ONSET means the root will be allowed to surface faithfully, as seen in (80).

(80) Deletion blocked

/abaka^{L}/	MAX^{L}	ONSET	MAX
a. ☞ .a.ba.ka		*	
b. Øba.ka	*!		*

Candidate (80b) is ruled out by the high ranking of MAX^{L}. However, this ranking will *not* block hiatus resolution in a morphologically complex context. For example, if a V-final prefix, e.g. /ki-/, is concatenated with this root (/ki + abaka^{L}), hiatus will still be resolved by deleting the *prefix* vowel. This is illustrated in (81).

(81) Onset violations only tolerated by the root

/ki + abaka^{L}/	MAX^{L}	ONSET	MAX
a. ki.a.ba.ka		*!	
b. ☞ kØa.ba.ka			*
c. ki.Øba.ka	*!		*

Both the application of deletion and the form deletion takes in this case is attributable to the locality condition: MAX^{L} is only violated by deleting segments in the indexed morpheme. As such, the ranking MAX^{L} » ONSET cannot uniformly override the ranking ONSET » MAX; thus, fully faithful candidate [ki.a.ba.ka] will lose. However, provided that MAX^{L} is never crucially dominated by another constraint enforcing deletion, no segment in /abaka^{L}/ can *ever* be deleted. Thus, candidate [ki.Øba.ka] will also lose. However, the prefix has no index to protect it, so it cannot violate MAX^{L}. In this context, its vowel be deleted to satisfy ONSET.

As discussed in §4.2.1 and §4.2.2, when locality is applied to a real case, it both restricts the typological predictions of lexical indexation, as well as guarantees a link between phonological structure and exceptional behavior.

4.1.1.2 Inconsistency resolution and constraint demotion

Pater's theory includes a skeletal learnability proposal, built off of the inconsistency detection property of Recursive Constraint Demotion (Tesar & Smolensky 1998, 2000); this is intended to restrain the power of a theory of indexed constraints. My aim here is not to build substantively on this learnability proposal (which Pater notes is incomplete), but rather to explain how its assumptions apply to evaluating typological predictions made by lexically-indexed constraints, and to establish some analytical biases I will be adopting in generating an analysis.

This theory assumes that indexed constraints are cloned from existing constraints in the grammar in response to a learner identifying an inconsistent pattern in the lexicon.

More specifically, if a learner, in the process of discovering a grammar, encounters a situation in which there is no constraint ranking that can account for all inputs (e.g. by finding an exception) a constraint is cloned and re-ranked. This means that while indexed constraints are assumed to be generated from universal constraints, they themselves are not universal. This adds a complication in evaluating typological predictions of this model: if an indexed constraint could never make a crucial decision with respect to some candidate, then this theory would assume that the learner would never generate it in the first place.

The purpose of this chapter (and this dissertation) is not to build or improve upon a learnability proposal, but to build an analysis of Mushunguli exceptions, and use this analysis to explore the dependencies between exceptions and regular forms in a single language. To do this effectively, I feel it is important to outline some analytical biases that I adopt in shaping the analysis to come.

4.1.2 Analytical biases adopted by this dissertation

The operation of RCD does not natively provide much guidance in terms of how an analysis should be shaped. However, there are two analytical sub-biases that I will adopt going forward, both of which are discussed briefly by Pater, but not taken as requirements. These are the following: (a) minimize indexed constraints; and (b) lexically index the smallest set of forms. Both of these serve a general assumption made by this analysis that indexed constraints should be reserved exclusively for exceptions, not other forms of morpheme-specific phonology.

4.1.2.1 Minimizing indexed constraints

I follow the assumption that fewer indices are better, and that phonological (or morphological) generalizations should supercede arbitrary indices. One way to do this is to collapse phonological generalizations made by the indexed constraints in the same way that they are collapsed for unindexed ones; that is, to "re-use" indexed constraints. For example, a simplified view of a conspiracy is a case where the satisfaction of one markedness constraint requires the violation of different faithfulness constraints, depending on context; this produces related alternations. As will be demonstrated by the case of Mushunguli, having multiple exceptions within conspiracies like this results in dependency relationships: exceptions impose requirements both on which other constraints can be indexed, and on the ranking(s) that comprise the grammar itself.

However, the inverse case is also possible: a language can have contexts where violation of one faithfulness constraint can satisfy multiple unrelated markedness constraints, producing unrelated alternations. If this language has exceptions to both alternations in this case, an unbiased learner (or analysis) would treat cloning the faithfulness constraint once and ranking it above both markedness constraints as no better or worse than cloning it twice, once for each exception. The bias towards fewer indexed constraints means that we expect the grammar to either only choose the first option, or to eventually collapse multiple clones of the same constraint into a single one, provided that doing so does not create new ranking conflicts. This theory would predict that exceptions that can "recycle" indexed constraints in this way should be more permissible in the grammar than those which require the cloning of additional ones, especially the cloning of the same constraint twice.

4.1.2.2 Indexing the smallest set of forms

A second way to minimize indexed constraints is to guarantee, insofar as it is possible, that indexed constraints are limited to exceptions alone. This means collapsing together as much as possible any other forms of domain or morpheme specification, and/or relying on other forms of stringency relationships to handle cases of regular opacity.

Doing this would of course require a way to reliably separate exceptionality from similar phenomena. This would be tricky, but a possible avenue would be to formalize the definitional criteria from §1.3 such that they rise at least to the level of a methodological restriction on the adoption of an indexed constraint. This is similar to the informal proposals made by Inkelas et al. (1997); Inkelas & Zoll (2007) with respect to cophonology. These proposals reserve cophonology for cases of non-arbitrary morpheme-specific phonology (e.g. a noun vs. verb distinction); however, as discussed in §6 (especially §6.2), this is largely because these authors are in favor of a representational approach to exceptionality, not because they obviate the need for an exceptionality formalism at all.[2]

Formalizing the definitional criteria would require filling a substantial extant gap in the exceptionality literature, which is a typological survey of cross-linguistic attested examples of both exceptionality and other forms of morpheme-specific phonology. An important next step in this research program is conducting just such a survey, but it is outside the scope of this dissertation. For now, we will still use the definitional criteria as an informal mechanism to index the smallest set of forms, as the lexicon of Mushunguli makes this a much easier problem than some other languages. In the Dayley et al. (2020) BLP dictionary, roughly 11% of verbal entries (including derived forms) contain V-initial roots. Of these forms, $<$1% block hiatus resolution. As such, blocking hiatus resolution is a non-representative behavior of V-initial verb stems.

As another example of this informal restriction applied to Mushunguli, consider the paradoxical behavior of results of postconsonantal glide formation. This behavior does not conform to the criterion of arbitrariness, because this behavior is representative of *all* morphemes of the shape /$CV_{[+hi]}$-/. Conversely, the behavior of the extremely small number of palatalization exceptions (a subset of /$CV_{[+hi]}$-/ morphemes) contravenes this larger generalization. Thus, the former problem will not be solved with lexically-indexed

[2] As discussed in §6.2, the adoption of a representational analysis does not solve any problems introduced by a lexical indexation analysis of Mushunguli, and introduces new ones on its own.

constraints (§5.1), while the latter problem will be (§5.1.3).

In tandem with the general preference to reduce the number of indexed constraints, this results in fewer indexed constraints that have more power. It also restricts the analytical space, which makes it easier to see both how indexed constraints can affect the grammar, as well as how they can be affected by it. For example, as will be discussed in §5.1.3, when indexed constraints are used in tandem with a theory relying on a more richly-structured lexicon, there emerge ranking conditions under which indexed constraints can become inactive.

4.1.3 Typology of blocking

To provide context for the analysis of Mushunguli, I am going to briefly outline the way that given a set of constraints and a set of regular and indexed inputs, lexical indexation can make both language-internal predictions regarding possible and impossible forms of exceptionality and regularity, as well as cross-linguistic typological predictions about the behavior of exceptions more generally.

An early observation regarding the typology of exceptions in a synchronic context comes from SPE (Chomsky & Halle 1968:374):

> In fact, not infrequently an individual lexical item is exceptional in that it alone fails to undergo a given phonological rule or, alternatively, in that it is subject to some phonological rule.

To handle this observation, Chomsky & Halle introduced diacritics that allow for words to be marked as [± rule]. However, this was shown rather quickly to be inadequate both formally and descriptively in the discussion of Yine by Kisseberth (1970b). Kisseberth argued that a characterization (and concomitant formalism) of exceptions as chiefly [± rule] cannot account for attested patterns in Yine (see §7.2.2 for a summary of the Yine facts) and many other languages. This is because under the SPE formalism, it is not possible to mark a form as [-rule] in order to prevent it from serving as the context of a rule without *also* marking it as something that cannot undergo the rule itself. This is unfortunately exactly the pattern exemplified by Yine: some suffixes fail to trigger one process (deletion), but they still undergo it. Kisseberth thus proposed a formal mechanism—and by extension, a typology–the following way: exceptions are lexical items that are marked as "either undergoing a rule or not, and as either serving as the context for a rule or not" (Kisseberth 1970b:57).

The shift away from rule-based formalism spearheaded by the widespread adoption of Optimality Theory (Prince & Smolensky 1993/2004) also changed the way that exceptions were discussed, with a greater emphasis on exceptions in terms of the interplay of markedness and faithfulness. Exceptional "blocking" has been the primary focus of discussion of exceptions, with a general emphasis on associating exceptional behavior with faithfulness. This association with faithfulness is explicit in proposals concerning mechanisms such as indexed faithfulness constraints, (McCarthy & Prince 1995; Fukazawa 1999; Ito & Mester 1995a, 1999a, 2001; Alderete 2001), but is also implicit in most representational analyses, as these proposals typically assume a greater degree of faithfulness to an underlying representation as being the driver of exceptional behavior.

Perhaps due to the emphasis on cases of blocking, Kisseberth's typological distinctions are still widely cited and have remained largely unchanged over the last fifty years. Pater (2010) re-frames Kisseberth's categories as *blocking* and *application*. Finley (2010) also re-frames these categories to (non-)undergoer and (non-)trigger in her typological investigation of exceptions to vowel harmony, a phenomenon for which these terms are most appropriate. The former set of terms is useful in discussing the fact that there are multiple ways to block an alternation, but I have been borrowing the latter because they can be treated as atomistic and thus allow for closer pinpointing of what aspects of the (non)-alternations under examination are exceptional.

An important assumption made by lexical indexation is that indexed constraints are assumed to be a part of the same grammar as regular forms. An under-discussed consequence of this is that the ranking required to generate regular patterns in the language will necessarily apply to exceptions (except where indices subvert it). The converse is also true: any ranking required in order to generate an exceptional pattern will necessarily apply to all other forms, related or unrelated. That is to say, absent a justifiable reason to introduce some other form of additional specification (such as belonging to independent lexical strata), the presence and form of an exception can determine the presence and form of regular patterns. Moreover, because exceptions themselves follow from regular patterns in the grammar, the presence and form of an existing exception can *also*, in principle, determine the presence and form of *other* exceptions. These observations are summarized below.

(82)Effects of exceptions on the grammar

a. The presence of one exception predicts (or rules out) the forms of regular patterns in the same grammar.
b. The presence of one exception predicts (or rules out) the possibility of other exceptions in the same grammar.

These implications can be illustrated more concretely using a toy example of hiatus resolution. Imagine a language, which we will call Sanuma[3], that productively deletes the first vowel in a sequence of two adjacent heterosyllabic vowels at morpheme boundaries; for example, the sequence /bi + adi/ surfaces regularly as [badi]. The general schema for a process is M » F, which we will assume for this case is ONSET » MAX-V.

Let us assume that Sanuma also has a homophonous prefix /bi-^{L}/ that exceptionally resists deletion; that is, /bi^{L} + adi/ surfaces as [bi.adi]. This exception can be captured under lexical indexation by cloning MAX-V and ranking the clone (MAX-V^{L}) over ONSET,

[3]'Toy,' in Mushunguli.

which generates deletion in regular cases (83) and unrepaired hiatus in exceptional cases (84).

(83) Regular deletion

	$bi_1 + a_2di$	Max-V^{L}	Onset	Max-V
a.	$bi_1.a.di_2$		*!	
b. ☞	$b\text{Ø}_1a_2.di$			*

(84) Exceptional blocking of deletion

	$bi_1{}^{L} + a_2di$	Max-V^{L}	Onset	Max-V
a. ☞	$bi_1{}^{L}.a_2.di$		*	
b.	$b\text{Ø}_1{}^{L}a_2.di$	*!		*

However, as we have already established, hiatus resolution is more properly viewed as a conspiracy in many languages. This means that other applicable constraints can potentially be ranked *below* our indexed constraint, generating alternative patterns that do not apply in this context. For our hypothetical scenario, we will only consider one such alternative solution: consonant epenthesis, which violates Dep-C. We will assume that in our current example language, epenthesis never occurs, and thus the minimal ranking to successfully generate the regular deletion pattern is the partial order {Onset, Dep-C} » Max-V (schematically: {M,C} » F).

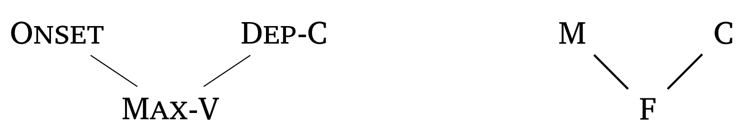

Figure 4.1: Schematic for a partial order

If there are no examples of productive epenthesis anywhere in the language, then provided that Onset (M) and Dep-C (C) both dominate Max-V (F), the ranking between M and C will be indeterminate. This is unproblematic; however, if we reintroduce our exceptional prefix /bi^{L}-/ and its linked indexed constraint, this indeterminate ranking cannot persist. The cloning and promotion of the lower-ranked faithfulness constraint requires one of the higher-ranked constraints be promoted above the other; otherwise, in exceptional cases, no winner can be chosen. This is illustrated below in Figure 4.2:

bi$_1{}^L$ + a$_2$di	MAX-V^L	ONS	DEP-C	MAX-V
a. ? bi$_1{}^L$.a$_2$.di		*		
b. bØ$_1{}^L$a$_2$.di	*!			*
c. ? bi$_1{}^L$.Ca$_2$.di			*	

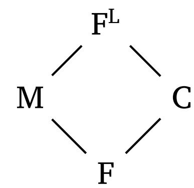

Figure 4.2: Schematic and tableau for impermissible indeterminate ranking

For this hypothetical example, promotion of either M or C to undominated status results in two possible languages. If we promote C (DEP-C), the result will be unrepaired hiatus, the same as (84). This scenario, which we will refer to as *simple blocking*, is illustrated in Figure 4.3.

bi$_1{}^L$ + a$_2$di	MAX-V^L	DEP-C	ONS	MAX-V
a. ☞ bi$_1{}^L$.a$_2$.di			*	
b. bØ$_1{}^L$a$_2$.di	*!			*
c. bi$_1{}^L$.Ca$_2$.di		*!		

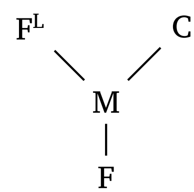

Figure 4.3: Schematic and tableau for simple blocking ranking

The promotion of M, which in this case is ONSET, however, results in a perhaps typologically surprising form of blocking: an undominated M must be satisfied by any means necessary, and if C is not undominated, a means exists. In this case, when deletion is blocked, epenthesis occurs as an *alternative*. This scenario, which will we refer to as *walljumping*, is illustrated in Figure 4.4.[4]

The scenario I am illustrating here is somewhat different than the possibility of so-called "alternative" (Pater 2004) or "divergent" (Hsu & Jesney 2018) repairs someti-

[4]"Walljumping" is used here to evoke the sense that when the typical pathway to satisfaction is metaphorically "walled off," the grammar finds an alternative and perhaps unexpected path past said wall.

	$bi_1{}^L + a_2di$	MAX-V^L	ONS	DEP-C	MAX-V
a.	$bi_1{}^L.a_2.di$		*!		
b.	$bØ_1{}^La_2.di$	*!			*
c. ☞	$bi_1{}^L.Ca_2.di$			*	

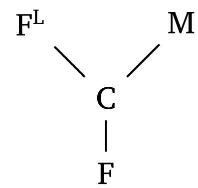

Figure 4.4: Schematic and tableau for walljumping ranking

mes discussed in studies of loanword adaptation (and these are rare even within that literature). More often than not, the identification of such repairs is really a comparison between a productive alternation that applies to native forms with word-internal phonotactic repairs from (presumed) foreign underlying representations. The walljumping scenario assumes that both the productive and unproductive patterns are alternations, with identical underlying representations.

It is also worth noting that where these patterns have been discussed in theories relying on constraint ranking, the proposed schematic for walljumping is $\{F_B{}^L, M\}$ » F_A » F_B. In proposals that rely on constraint weighting, such as lexically-scaled weights (Hsu & Jesney 2018), walljumping patterns are still assumed to be driven by the conflict of two faithfulness constraints with different scaling factors and a markedness constraint that is weighted more heavily than their combined weights. The toy example I have developed here demonstrates that the interactions of two faithfulness constraints with a markedness constraint is *a* ranking schematic that can generate walljumping behavior under certain circumstances. However, as will be demonstrated by the real case study of Mushunguli in §5.1, the actual ranking conditions that can lead to walljumping behavior are more diverse than this. For Mushunguli, the schematic $\{F^L, M_A\}$ » M_B » F generates this type of pattern; this is why I refer to the intermediate constraints as C.

This example shows that although indexed constraints are only violable by indexed morphemes, the fact that they require a fully-determined ranking to function means that the mere presence of an exception can have an impact on the entire grammar. This forced determination has two consequences: first, the presence of one type of exception will necessarily predict the possibility or impossibility of others. For example, if we have a simple blocking exception that requires the regular ranking C » M » F, we predict that we will not see an equivalent walljumping exception in the same context. This can be seen in (85), which features a third exceptional prefix /mi^{L2}/, indexed to a copy of ONSET.[5] This prefix is only subject to ONSET^{L2}, which means that even if it is ranked above DEP-C, a deletion candidate will still win.

(85) Walljumping impossible

$mi_1^L + a_2di$	MAX-V^L	ONSET^{L2}	DEP-C	ONSET	MAX-V
a. $mi_1.a_2.di$		*!		*	
b. ☞ $mØ_1a_2.di$					*
c. $mi_1.Ca_2.di$			*!		

This is happening because MAX-V^L cannot "see" an input that is not indexed to it; thus, deletion wins because it is preferred by the regular ranking. This means that technically, this learner would never generate an indexed constraint at all, since there is no inconsistency here to detect. The only scenario in which it is possible for a walljumping pattern to emerge is if our third exceptional prefix was indexed to *both* MAX-V^L and NO-HIATUS^{L2}, and both dominate DEP-C.[6]

(86) Walljumping impossible

$mi_1^{L1,2} + a_2di$	ONSET^{L2}	MAX-V^{L1}	DEP-C	ONSET	MAX-V
a. $mi_1.a_2.di$	*!			*	
b. $mØ_1a_2.di$		*!			*
c. ☞ $mi_1.Ca_2.di$			*		

[5]ONSET is arbitrarily chosen for this example; the same is true for indexed DEP-C.
[6]Here, indexing the prefix to a copy of DEP-C will not have the same effect; instead, it will reinforce the simple blocking pattern.

I am unaware of any walljumping pattern in a real language that has been analyzed with multiple indices like this. For the three types of blocking exceptions discussed for Mushunguli, it will never be necessary to index more than one constraint per morpheme, regardless of how many repairs are exceptionally blocked or conditioned.

The second prediction is that while regular lexical items are not subject to indexed *constraints*, they *are* subject to *rankings* determined by the introduction of indexed constraints into the grammar. This means that the existence of an exception can determine the possibility or impossibility of other *regular* patterns. For example, in the simple blocking scenario discussed above, the existence of a prefix that failed to resolve hiatus required DEP-C to dominate ONSET; as such, for this language, no regular epenthesis candidate can win.

Unlike the previous prediction, this really is all-or-nothing: simple blocking in this configuration is effectively extreme faithfulness, so a grammar with a simple blocking exception to some process is mutually exclusive with one that also exhibits a walljumping exception *to that same process*. If by some mechanism (such as constraint re-ranking between lexical strata) that ranking were to be undone, then a walljumping pattern *will* emerge.

In the next section, we will examine the first case study, the Mushunguli non-coalescing stems, which will demonstrate the way that indexed constraints determine rankings, and the consequences these have for the analysis of both regular and exceptional patterns.

4.2 Non-coalescing stems

Our first case study from Mushunguli is a set of high vowel-initial stems that exceptionally fail to undergo high coalescence (i.e. the alternation /a + i/ → [e] fails), leaving hiatus unresolved. Critically, these stems' behavior is not exceptional with respect to other forms of hiatus resolution in the language: they regularly trigger glide formation and undergo identity coalescence normally. This is illustrated in (87), using a derived form of one such exceptional stem, *-itisa* 'go a lot' compared to a regular stem *-ivisa* 'hear a lot'.[7]

(87) Regular vs. exceptional high-vowel stems

	Regular (*-ivisa* 'hear a lot')	**Exceptional** (*-itisa* 'go a lot')	
1sg past (/si-/)	s**i**vi:sa	s**i**ti:sa	**Simplification**
Infinitive (/ku-/)	kʷ**iː**vi:sa	kʷ**iː**ti:sa	**Glide Formation**
3sg past (/ka-/)	k**eː**vi:sa	k**a.i**ti:sa	**High Coalescence**

As seen in (87), exceptional stems resolve hiatus normally in glide formation and simplification contexts; that is, they behave as if they are regular vowel-initial stems. However, these same stems exceptionally fail to resolve hiatus in coalescence contexts, behaving instead as if they were *consonant*-initial stems. This pattern is further illustrated in (88) below, where the behavior of a consonant-initial, regular vowel-initial, and exceptional vowel-initial stem are compared.

(88) Comparison of exceptional stems with regular V- and C-initial stems

	C-initial (*-lima* 'breathe')	**Exceptional V-initial** (*-ita* 'go')	**Regular V-initial** (*-iva* 'hear')
1sg. past (/si-/)	s**iliː**ma	s**iː**ta	s**iː**va
Infinitive (/ku-/)	k**uliː**ma	kʷ**iː**ta	kʷ**iː**va
3sg past (/ka-/)	k**aliː**ma	k**a.iː**ta	k**eː**va

Here, the dashed-line box groups together the situations where regular and exceptional vowel-initial stems behave similarly, and the solid-line box groups together situations where the exceptional vowel-initial and consonant-initial stems behave similarly.

[7] The verbal extension *-is* is used here to move the initial stem vowel out of the penultimate syllable, which makes it possible to see that simplification has applied (recall from §2.2.2 that the penultimate syllable of any verb is lengthened, and from §3.4 that the resolution of hiatus between sequences of identical vowels is the only repair that does not result in compensatory lengthening).

The majority of these exceptional stems are verbs, the unprefixed imperative forms of which unambiguously begin with a high vowel. In addition, a few nouns belonging to class 5/6 (*ini* 'liver,' *izi* 'voice,' *ivu* 'ash'), as well as the native word for 'two' (*idi*) behave similarly, though this is only diagnosable by the failure of coalescence to apply at a prefix-root boundary; unfortunately, appropriate and unambiguous morphophonological contexts for glide formation or simplification do not exist for these lexical items.[8]

An exhaustive list of lexical items exhibiting this form of exceptionality is given below in (89), using the 3rd singular past *ka-* to illustrate the exceptional behavior in the case of the verbs, and the class 2 demonstrative marker to illustrate the behavior of 'two.'

(89) Exceptional stems

a.	ka.iːta	'go'	b.	ka.iːha	'be bad, angry'
c.	ka.iːmba	'sing'	d.	ka.iːra	'cry'
e.	ka.uːsa	'take out'	f.	ka.uːmba	'mold (shape)'
g.	ka.uguːla	'lament'	h.	ka.uːŋga	'want, like'
i.	ka.uːja	'come back'	j.	ka.uːza	'ask'
k.	ka.uguːða	'care for someone sick'	l.	ka.ujuːsa	'revive, transplant'
m.	ma.ini	'livers' (sg: *ini*)	n.	ma.izi	'voices' (sg: *izi*)
o.	ma.ivu	'ash'	p.	wa.idi	'two (cl 2)

The choice to treat these as exceptions follows from the definition established in §1.3. As demonstrated by the near-minimal pair above in (87), the non-coalescing stems exhibit phonological behavior that is distinct from phonologically similar forms. There are no impressionistic or measurable phonetic distinctions between the non-coalescing and coalescing sets in terms of vowel quality, length, or tone. The only reason we do not consider *ita* and *iva* to be true minimal pairs is because of this difference in behavior. The non-coalescing set is also not morphologically or lexically distinct; lexical items exhibiting this behavior belong to multiple semantic classes. Finally, while vowel-initial stems are

[8]For *-idi*, it is possible to generate a simplification context for class 8 (/vi + idi/); however, no phrasal contexts in which this sequence was outside of the penultimate syllable (which is subject to lengthening independent of hiatus resolution) were elicited. A glide formation context is truly impossible; locative forms of nouns do not take locative agreement markers, and while class 15 nouns have the prefix /ku-/, my consultant treated them as uncountable and thus was unable to reliably generate forms modified by 'two.' See §2.3.1 for illustrations of the relevant morphemes.

less common than consonant-initial ones (roughly 8% of the verbal lexicon), the non-coalescing set makes up less than 1% of these.

The behavior of these exceptions is also strikingly systematic and regular, which allows for the identification of two critical generalizations about the set. These are stated below in (90).

(90) *Critical generalizations*

a. The non-coalescing stems all begin with high vowels.

b. The non-coalescing stems are exceptional *only* with respect to coalescence.

These generalizations are linked to the stems' phonological behavior and shape, so it is important that our analysis captures them; recall from §1.1 that a criticism of theories like lexical indexation are that they potentially treat structural factors as accidental or unimportant (and thus treat exceptions as ungrammatical). For this particular case, this is a result we especially want to avoid, as it can be demonstrated without appeal to any form of exceptionality theory that the generalizations from (90) are in fact intrinsically linked.

To see this, consider in more detail the actual changes involved in each of the hiatus resolution strategies. Each strategy is schematically illustrated (using front vowels) in (91), where the subscript 'P' indicates the first, outer prefix vowel and the subscript 'S' indicates the second, inner stem vowel.

(91) Results of hiatus resolution strategies

Strategy	*Schematic illustration*	*Change to V_P*	*Change to V_S*
Glide Formation	/i_P + a_S/ → $j_P a\!:_S$	glided	**no change**
Simplification	/i_P + i_S/ → $i_{P,S}$	no change	**no change**
Mid coalescence	/a_P + e_S/ → $e\!:_{P,S}$	raised (and fronted)	**no change**
High coalescence	/a_P + i_S/ → $e\!:_{P,S}$	raised (and fronted)	**lowered**

As these schematic representations make clear, glide formation affects only the first (prefix) vowel; similarly, mid coalescence also only affects the prefix vowel, raising (and fronting) it. Because we must treat simplification as vacuous coalescence, it affects no vowel at all. Of all of the hiatus resolution strategies, only high coalescence has a material effect on the initial vowel of the stem. “Material effect” here means a change in a contrastive property of the vowel; changes in vowel quality, syllable attachment (i.e., changing a vowel to a glide), and deletion are thus all considered material changes.[9]

At this point, without yet introducing any indexed constraints, it makes sense that *only* high coalescence is blocked in the case of the exceptional stems: stems that begin with high vowels are the only stems expected to result in any sort of contrastive change to the stem-initial vowel, and only when they are expected to undergo coalescence. This suggests a degree of interconnection between the generalizations from (90) that will ideally be captured in their analysis. These can be revised to form the unified generalization below:

(92) Unified generalization

A member of the set of non-coalescing stems may only block hiatus resolution if the stem would be materially affected by the otherwise normally expected hiatus resolution strategy.

Any reasonable analysis will capture the individual generalizations from (90), but a truly explanatory analysis will capture the unified generalization in (92) and thus provide the best account for the observed behavior of this set of exceptional stems. The analysis I propose in the following section accomplishes the explanatory result using indexed constraints. Specifically, I propose that a copy of the faithfulness constraint IDENT(high) is lexically indexed to these exceptional stems and ranked above the markedness constraint

[9]Because length is not contrastive in Mushunguli, compensatory lengthening does not constitute a material change.

driving hiatus resolution, ONSET. Only high coalescence can be blocked by this lexically-indexed copy of IDENT(high), because only coalescence with stem-initial high vowels risks violating this constraint.

4.2.1 Application of locality to the non-coalescing stems

To motivate the choice of IDENT(high) as the best candidate for indexation, consider again the near-minimal pair of forms /ku_P-i_Sv-a/ 'to hear' and /ku_P-$i_S t^L$-a/ 'to go'. The optimal resolution of hiatus between the prefix vowel /u/ and the stem vowel /i/ is to produce *wi:* via glide formation in both cases.[10] Only MAX-ASC is violated by each of these mappings, and this violation is borne solely by the prefix vowel's mapping from /u/ to *w*. Because of this, a lexically-indexed instance of this constraint, MAX-ASC^L, would fail to distinguish this near-minimal pair, capturing the fact that they both behave regularly in this respect. Indeed, there is no faithfulness constraint at all that could be indexed *to a stem* that will exclusively prevent a *prefix* vowel from undergoing glide formation. This follows from the locality condition, and is a good result for Mushunguli, as there are no exceptional roots that exclusively block glide formation in prefixes.

Next, consider the near-minimal pair of forms /si_P-i_Sv-a/ and /si_P-$i_S t^L$-a/. The optimal resolution of hiatus between the prefix vowel /i_P/ with the stem vowel /i_S/ is identity coalescence to produce *i* in both cases. Only UNIFORMITY is violated by this output, and this violation is shared by the mapping of both the prefix vowel and the stem vowel from /$i_P + i_S$/ to $i_{P,S}$. A lexically-indexed instance of this constraint, UNIFORMITY^L, would thus be able to distinguish this near-minimal pair. In order to capture the fact that they both behave regularly, then, UNIFORMITY^L must be ranked at least as low as UNIFORMITY.

Finally, consider the near-minimal pair of forms /ka_P-i_Sv-a/ and /ka_P-$i_S t^L$-a/. The resolution of hiatus between the prefix vowel /a_P/ with the stem vowel /i_S/ via high coa-

[10]The surface output of post-consonantal glide formation will be discussed in §5.1.

lescence to produce $e_{P,S}$ is optimal in the former, regular case, but not in the latter, exceptional case. A shared violation of IDENT(low) is borne solely by the prefix vowel's mapping from /a_P/ to $e_{P,S}$ — recall the hierarchy of costliness from (66) — and thus lexically-indexed IDENT(low)^L would fail to distinguish this near-minimal pair.

A violation of UNIFORMITY is shared by the mapping of both the prefix vowel and the stem vowel from /a_P-i_S/ to $e_{P,S}$; lexically-indexed UNIFORMITY^L would thus be able to distinguish this near-minimal pair, but as we know from the case of identity coalescence above, UNIFORMITY^L must be ranked at least as low as its non-indexed counterpart. Moreover, UNIFORMITY^L would only capture the generalization that coalescence is blocked; it would not capture the generalization that it is exclusively high-vowel initial roots that block it.

Having considered and discarded all possible relevant candidates for indexation, we finally come to the only possible analysis available to the non-coalescing stems: IDENT(high). This violation is uniquely borne by the stem vowel's mapping from /i_S/ to $e_{P,S}$. Lexically-indexed IDENT(high)^L is thus able to distinguish this near-minimal pair. Ranking it above ONSET can be effectively used to capture the fact that they behave differently. This is shown by the contrast between the two tableaux below:

(93) High coalescence allowed with regular (non-indexed) stems

ka_P-i_Sv-a	IDENT(hi)^{L1}	ONS	IDENT(hi)	IDENT(lo)	UNIF
a. $ka_P.i_S$ː.va		*!			
b. ☞ $ke_{P,S}$ːva			*	*	*

(94) High coalescence blocked with exceptional (indexed) stems

ka_P-$i_S t^{L1}$-a	IDENT(hi)^{L1}	ONS	IDENT(hi)	IDENT(lo)	UNIF
a. ☞ $ka_P.i_S$ː.ta		*			
b. $ke_{P,S}$ː.ta	*!		*	*	*

Critically, the fact that the optimal result here is tolerance of hiatus as opposed to triggering some alternative resolution strategy must be due to the subordination of ONSET to other constraints that could in principle be violated to satisfy it. Most of these constraints were already demonstrated to be undominated in §2.2.4. However, one was not: MAX-V, which must dominate ONSET in order to account for the fact that deletion isn't available as a back-up strategy in the face of high coalescence failure, as illustrated in (95).

(95) Accidental deletion due to ONSET » MAX-V

ka_P-$i_S t^{L1}$-a	ID(hi)^{L1}	ONS	MAX-V	ID(hi)	UNIF
a. ☹ $ka_P.i_S$ː.ta		*!			
b. $ke_{P,S}$ː.ta	*!			*	*
c. ☞ $kØ_P$i.ta			*		

This means that MAX-V must dominate ONSET, ruling out vowel deletion entirely as a hiatus resolution strategy in Mushunguli.

(96) Deletion blocked; hiatus preserved

ka_P-$i_S t^{L1}$-a	ID(hi)^{L1}	MAX-V	ONS	ID(hi)	UNIF
a. ☞ $ka_P.i_S$ː.ta			*		
b. $ke_{P,S}$ː.ta	*!			*	*
c. $kØ_P$i.ta		*!			

To reiterate, the non-coalescing exceptions specifically rule out vowel deletion as a solution to hiatus resolution in Mushunguli. This means that anything that looks like deletion in the language *cannot* be analyzed as such.

The constraint ranking established so far for regular and exceptional patterns of coalescence in Mushunguli is given in Figure 4.5.

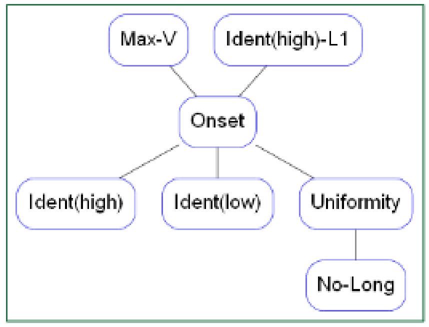

Figure 4.5: Ranking for hiatus resolution and exceptions to coalescence

4.2.2 Guaranteeing the unified generalization

The proposed analysis guarantees the unified generalization governing exceptional stem behavior, restated in (97).

(97) A member of the set of non-coalescing stems may only block hiatus resolution if the stem would be materially affected by the otherwise normally expected hiatus resolution strategy.

Given the regular pattern of hiatus resolution in Mushunguli and the constraint ranking responsible for it, lexically-indexed IDENT(high)^{L} can only have a discernible effect given indexed stems that begin with high vowels, and only when those indexed stems are in coalescence contexts. That is, the form and behavior of the exceptional stems *follows* from the general patterns of the language.

First, because IDENT(high)^{L} is only (relevantly) violated by changes from [+high] to [–high] in indexed stem vowels, only indexed stems beginning with [+high] vowels are predicted to behave exceptionally. This is because, as already noted, only high coalescence incurs a violation of IDENT(high). This means that even if there happened to be an indexed stem beginning with a mid vowel, mid coalescence could not be blocked and so the stem could not realistically be considered exceptional. This is shown in (98).

(98) No blocking of lexically-indexed mid-vowel stems

	$\text{ka}_\text{P}\text{-o}_\text{S}\text{mal}^\text{L}\text{-a}$	IDENT(hi)^L	ONS	IDENT(hi)	IDENT(lo)	UNIF
a.	$\text{ka}_\text{P}.\text{o}_\text{S}.\text{maː.la}$		*!			
b. ☞	$\text{koː}_{\text{P,S}}.\text{maː.la}$				*	*

Second, the fact that high coalescence is the only one of the regular hiatus resolution strategies that can be blocked by IDENT(high)^L is guaranteed by the locality condition in (78). Any constraints indexed to an exceptional stem are only relevant if the locus of violation includes a part of the stem itself, and high coalescence is the only hiatus resolution strategy for which a locus of violation is part of the stem.

To see this, we need to consider the only other productive hiatus resolution process in this language, which is glide formation. So far, we have largely been ignoring MAX-ASSOCIATION (redefined below) in our discussion due to the fact that it evaluates only losers, is not crucially dominated by any constraints (save for ONSET), and does not crucially dominate any constraints itself, but rather shares decision-making power with MAX-μ, which *does* crucially dominate *Vː to allow for lengthened outputs.

(99) MAX-ASSOCIATION (MAX-ASC): If τ_1 is a mora in the input and it is associated to ζ_1 and $\tau_1 \Re \tau_2$, and $\zeta_1 \Re \zeta_2$, then τ_2 is associated to some ζ_2. (Keer 1999:47)

Evaluating MAX-ASSOCIATION is slightly tricky. Deletion of segments or deletion of moras entails the loss of an association line, so violation of either MAX-V or MAX-μ entails violation of MAX-ASSOCIATION. Non-identity coalescence does not violate MAX-ASC since the output retains both input correspondents; thus, the mora's association with its input segment is not lost.[11] Glide formation *does* violate MAX-ASC, however, because glide formation de-links the input mora from its underlying segmental melody. Because glide formation does obtain in this language, ONSET must dominate MAX-ASSOCIATION.

Now, imagine that a stem is indexed to a clone of MAX-ASC, MAX-ASC^L, and that MAX-ASC^L outranks ONSET. Glide formation would not be blocked here, because the de-

[11]Identity coalescence does violate MAX-ASC because the mora is deleted due *MATCH-μ.

linked mora was associated to the prefix vowel–that is, the locus of violation is entirely within the prefix.

(100) No blocking of glide formation by MAX-ASC^{L}

$ku_P + i_S t^L + a$	MAX-ASC^{L}	ONSET	MAX-ASC
a. $ku_P.i\!:_S.ta$		*!	
b. ☞ $kw_P i\!:_S.ta$			*

Of course, if MAX-ASC^{L} was indexed to the prefix, then the locus of violation of MAX-ASC^{L} would include a part of the indexed morpheme and glide formation would be blocked — but this would be true for every instance of this prefix, completely independent of its attachment to any particular stem. This kind of pattern is attested in Kimatuumbi, for example, where one plural prefix /ki-/ fails to undergo glide formation (Odden 1996).[12]

Now suppose there is an instance of IDENT(low) indexed to some stems, IDENT(low)^{L}, outranking ONSET. Coalescence cannot be blocked because although it does involve the stem vowel, the stem vowel itself does not change from [+low] to [–low]. This result is shown in (101) with an example of mid coalescence.

(101) No blocking of coalescence by IDENT(low)^{L}

ka_P-$o_S mal^L$-a	IDENT(low)^{L}	ONSET	IDENT(low)	UNIF
a. $ka_P.o_S.ma\!:.la$		*!		
b. ☞ $ko\!:_{P,S}.ma\!:.la$			*	*

Here again, if the morpheme that IDENT(low)^{L} is indexed to were the prefix, then the locus of violation of this constraint would include a part of the indexed morpheme and mid coalescence would be blocked — but this would be true for every instance of this prefix, independent of its attachment to any particular stem (and it would be true of high coalescence as well). I am unaware of the existence of a non-coalescing prefix in the world's languages, but under our current assumptions it is not an implausible pattern.

[12]At least two other lexical items in Kimatuumbi exhibit the same resistance to glide formation, including a derived verb stem -*angali*- and a verb focus affix -*ti*- (ibid:68).

Finally, suppose that there is an instance of bottom-ranked UNIFORMITY indexed to the exceptional stems, UNIFORMITY^{L}. I have already demonstrated in §4.2.1 that this constraint must be ranked at least as low as its non-indexed counterpart in Mushunguli, but I now consider the predictions for (hypothetical) exceptional stems indexed to this constraint if it were to be ranked above ONSET.

(102) Typology of exceptions linked to UNIFORMITY^{L}

if the stem-initial vowel is...	then high coalescence is...	and mid coalescence is...	and identity coalescence is...
high	blocked	*n/a*	blocked
mid	*n/a*	blocked	blocked
low	*n/a*	*n/a*	blocked

Again, I am unaware of any examples of this type of pattern in the world's languages, but it would require a language which has multiple hiatus repairs (including unambiguous coalescence), and a mixed set of exceptional forms that exclusively fail to undergo coalescence.[13] Regardless, the bias we are currently assuming would prefer IDENT(high)^{L1} over UNIFORMITY for Mushunguli, as this captures a phonological generalization that would otherwise be missed if UNIFORMITY were indexed.

Moreover, if we were to discover evidence suggesting that this is an undesirable prediction, there are options to avoid it. We could, for example, follow the proposal

[13] There are exceptions in Seri (Hokan) to patterns *described* both as coalescence and deletion (Marlett & Stemberger 1983). Like Mushunguli, the applicability of "coalescence" is restricted to specific segments; in Seri, it only applies when a prefix ending in /o/ is attached to a stem beginning in /o/. Unlike Mushunguli, this form of "coalescence" does not actually involve featural preservation, and so is probably better described as a mutation. For example, /jo+otx/ (<yo-otx>) surfaces as *jatx* (<y-atx>) 'arise (distal)'. Some roots block this, e.g. /jo+o:sx/ (<yo-o:sx>) 'sprinkle (distal)', which surfaces as *joo:sx* <yo-o:sx> rather than expected **jasx*. As such, the existence of the Seri pattern does not support an analysis of indexed UNIFORMITY, and no extant analysis treats it that way (Martlett & Stemberger rely on empty timing slots, while Gouskova (2012) proposes indexed ALIGNMENT).

by Keer (1999) that UNIFORMITY is not a part of CON; rather, all forms of non-identity coalescence violate the family of IDENT constraints, making the coalescence of identical elements would be cost-free. Not only would this make UNIFORMITY unavailable for lexical indexation, it is consistent with the low-ranked status of UNIFORMITY in the constraint ranking that has thus far been established for Mushunguli.[14]

The nonexistence of UNIFORMITY would also mean that deletion of one of two adjacent identical vowels is harmonically bounded by identity coalescence, which predicts that deletion can never be the optimal strategy for hiatus resolution between identical vowels. The surface manifestation of deletion and identity coalescence is presumably the same, so no markedness constraint can distinguish them; identity coalescence violates no faithfulness constraints, and deletion violates MAX-V. I am not aware of any typological evidence bearing on this prediction, but it is difficult to imagine that any such evidence could be forthcoming given the presumed sameness of the two surface manifestations.

An alternative way to undermine the predictions in (102) is to strengthen the locality condition from (78) such that the locus of violation of an indexed constraint must be entirely contained within the indexed morpheme. UNIFORMITY is unique among the faithfulness constraints in this analysis in that it crucially references both the vowel in a non-indexed prefix and the vowel of a potentially-indexed stem. This strengthened locality condition would thus not allow UNIFORMITY^{L} to have any effect here, but it would also have the potentially undesirable consequence of preventing markedness constraints that crucially reference multiple segments from having any effect when lexically indexed as well.

This restriction would also need to be carefully formalized, as there are a number

[14]Currently, the only work that UNIFORMITY is doing is penalizing mora merger over mora deletion as a solution for simplification. If we wanted to remove UNIFORMITY, we would need to stipulate that moras can never merge, only delete. Alternatively, if we found an alternative formulation of simplification that treated merger of non-identical moras as a faithfulness violation (just as we do for segments), then we could safely remove UNIFORMITY from CON.

of cases that would be left ambiguous; for example, the ranking for Mushunguli compensatory lengthening is *MATCH-μ » MAX-μ » *V:. Indexing *V: and allowing it to dominate ONSET cannot block hiatus resolution at all, so the unified generalization is not undermined. However, if *V:^{L} is ranked above MAX-μ, it *can* block compensatory lengthening; this is because the locus of violation of *V: in both glide formation and coalescence contexts includes the stem vowel, either wholly (in the former case) or partially (in the latter).

(103) Glide formation without CL due to *V:^{L}

	$u_P + i_S t^L + is + a$	*V:^{L}	ONSET	MAX-μ	*V:
a.	$u_P.i_S$.ti.sa		*!		
b.	$w_P i\!:_S$.ti.sa	*!			*
c. ☞	$w_P i_S$.ti.sa			*	

(104) Coalescence without compensatory lengthening due to *V:^{L}

	$ka_P + i_S t^L + is + a$	*V:^{L}	ONSET	IDENT(hi)	MAX-μ	*V:
a.	$ka_P.i_S$.ti.sa		*!			
b.	$ke\!:_{P,S}$.ti.sa	*!		*		*
c. ☞	$ke_{P,S}$.ti.sa			*	*	

The problem with this scenario is that in the former, glide formation case, the locus of violation of *V:^{L} is wholly the stem vowel, as it only has a single correspondent. However, in the latter, coalescence case, the locus of violation of *V:^{L} is a multiply-corresponded output. If the more restricted form of the locality condition were to hold here, it is unclear whether *V: or any constraint like it should be allowed to be indexed.

In sum, the proposal that the exceptional stems of Mushunguli are indexed to IDENT(high)^{L}, ranked above ONSET, ensures the unified generalization: namely, the fact that those exceptional stems begin with high vowels and the fact that they are only exceptional with respect to (high) coalescence.

4.3 Summary

In this chapter, we have seen that lexically-indexed constraints make explicit typological and formal predictions regarding the behaviors of exceptions in conspiracies. The most important of these going forward was that the indexed constraints conditioned by the presence of an exception can force determination of otherwise unspecified rankings, meaning that the form and behavior of one exception can predict (or rules out) the form and behavior of both other exceptional and regular patterns.

An example of this was attested in the analysis of the non-coalescing stems. Here, we saw that the introduction of IDENT(high)^{L} required a strict ranking between ONSET and MAX-V, clarifying the uncertain status of deletion as a word-internal hiatus repair in Mushunguli. MAX-V now dominates ONSET, and an important consequence of this is that all cases of apparent word-internal vowel deletion must be analyzed as something else. This was already demonstrated to be the case for coalescence; in the next chapter, we will explore the consequence of this prediction on resolving the ranking paradox encountered in §3.5.

Finally, we also explored how constraint indexation can capture and unify phonological generalizations about the form and behavior of non-coalescing stems. This result will be important to the discussion of representational theories in §6, where we will see that these generalizations are not necessarily guaranteed by adopting a representational approach to exceptionality without sacrificing other generalizations captured by the lexical indexation analysis of Mushunguli.

This chapter, in part, contains material that has previously appeared in San Diego Linguistics Papers 6. Hout, Katherine, 2016. It also contains material that appeared in Proceedings of the 2018 Annual Meeting on Phonology. Hout, Katherine, Cascadilla Proceedings Project, 2019. This chapter, in part also contains coauthored material currently being prepared for submission for publication of the material. Hout, Katherine; Baković,

Eric. The dissertation author was the primary investigator and author of this material.

5 The Clarifying Role of Exceptions

In §2.2.4, we established a ranking to deal with the majority of hiatus resolution outcomes in Mushunguli; this in turn informed our discussion of the non-coalescing exceptions. However, we left one set of patterns unexplained, which was the analysis of the resolution of /Ci + V/ and /Cu + o/ sequences. The examples from (73) and (74) are repeated below:

(105) Failure of glide formation in /Cu + o/ contexts (repeated)

a.	/k**u** + omal + a/	→	k**oː**maːla	'to dish up ugali'
b.	/m**u** + oger + a/	→	m**oː**geːra	'you (pl) swam'
c.	/k**u** + ombok + a/	→	k**oː**mboːka	'you went far'
d.	/mu + gosi j**u** + ose/	→	m̩gosi joːse	'the whole man'

(106) Distribution of /Ci + V/ (repeated)

a.	/tʃ**i** + asa/	→	**tʃaː**sa	'we divorced'	Subject Prefix
b.	/si + d**i** + aza/	→	si**daː**za	'I lost it (cl 5)'	Object Prefix
c.	/m**i** + ezi/	→	**meː**zi	'months'	Noun Class Prefix
d.	/z**i** + etu/	→	**zeː**tu	'our (cl 10)'	Demonstrative Prefix

The problem posed by these two patterns is that both appear to require deletion. However, we have now had multiple forms of evidence converge to suggest that deletion of vowels is entirely unavailable as a solution to hiatus resolution in Mushunguli. In §3.3 and §3.4, it was demonstrated that there is a sub-superset relationship between IDENT(high) and all other constraints implicated in coalescence. We also saw that MAX-V must therefore dominate IDENT(high) to allow high coalescence to occur, which meant that other superficial cases of deletion (/a + $V_{[mid]}$/ → $V_{[mid]}$ and /V_i + V_i/ → [V_i]) must be treated as coalescence as well. This left us with an indeterminate ranking of MAX-V and ONSET.

In §4.2, we saw that the introduction of an indexed form of IDENT(high)^{L} made it necessary to determine the ranking of MAX-V and ONSET in order to choose a winner in exceptional cases. MAX-V was shown to crucially dominate ONSET, as otherwise we predicted deletion to occur as an alternative repair when coalescence was blocked. Following our predictions from (82), the existence of the non-coalescing exceptions now definitively rules out the possibility of deletion in regular or exceptional cases.

The question, then, is what to do about the examples from (105) and (106). Recall from §3.5 that we were left with a ranking paradox: the ranking required for secondary articulations ({ONSET, MAX-V, *COMPLEX} » IDENT(V-Place)) was mutually exclusive with the one for vowel deletion ({ONSET, *COMPLEX, IDENT(V-Place)} » MAX-V). We have ruled out the latter ranking, but promoting MAX-V in the former one will now leave us with secondary articulations across the board.

(107) Incorrect palatalization

/Ci + a/	MAX-V	ONSET	*COMPLEX	IDENT(V-Place)
a. Ci.a		*!		
b. CØaː	*!			
c. Cjaː			*!	
d. ☞ Cʲaː				*

If we reverse the ranking of *COMPLEX and IDENT(V-Place), we will instead form full glides (this is equally true if we allow our specific markedness constraint, NOPAL, to dominate *COMPLEX).

(108) Glide formation

/Ci + a/	Max-V	Onset	Ident(V-Place)	*Complex
a. Ci.a		*!		
b. |CØaː	*!			
c. ☞ Cjaː				*
d. Cʲaː			*!	

Neither of these is the correct outcome for /Ci + a/, but we appear to be stuck. Either we uniformly allow for secondary articulations to be formed, or we uniformly allow for glides to be formed. We can potentially make a decision by determining the actual ranking for onset structure in Mushunguli (which is exactly what we will do in §5.1.1); either way, we will need to find a way to repair the unwanted outcomes.[1]

5.1 Onset structure and *walljumping* palatalization

The solution we will adopt for the latter problem is Stratal OT (Kiparsky 2000, 2015; Bermudez-Otero 2011). I assume that there are two lexical levels (or strata), hereafter referred to as "stem" (for the base level) and "word" (for the second level), as well as a postlexical level. Stratal OT assumes that each level contains a parallel constraint ranking, which can (and generally will) differ across them, with the output of each level feeding into the next. Each level is still evaluated in parallel, and Richness of the Base still applies; furthermore, it is impossible for any input to skip a level or progress through them in a different order.

Analyses relying on multiple levels based in and outside of OT have been proposed before for many other morphologically-complex languages (see Kiparsky 2015 for a summary), including several other Bantu languages, such as Kikamba (Roberts-Kohno 1998), some Nguni languages (Xhosa and Zulu) (Sibanda 2009), and Kimatuumbi (Odden 1996). In general, there is a very strong tradition of assuming multiple lexical levels in the Bantu

[1]We will ultimately see that the only option is glide formation, and that this fact is reinforced by the behavior of a set of exceptions that exceptionally undergo palatalization (§5.1.3).

literature. The structure of the verb, as reconstructed by Meeussen (1967), is hierarchical. This involves the base, which contains the root (or radical) and verbal extensions, the stem (the base + the final vowel), and the word (or verb), which includes inflectional prefixes. This is illustrated in Figure 5.1

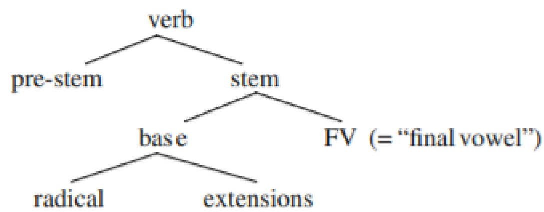

Figure 5.1: Traditional structure of the Bantu verb (taken from Hyman 2009), citing Meeussen 1967

More typical terms for this organization are the derivational stem (= the base), the inflectional stem (= the stem), the macrostem (= the stem + object prefixes), and the inflected verb (the macrostem + inflectional prefixes) (Downing 1997). Evidence for these different levels comes from the fact that different morphological and morphonological processes in a given language will apply to these different constituents differently. My corpus of Mushunguli has this type of evidence for the inflectional stem and the inflected verb; it has slightly weaker evidence for the derivational stem, and does not have direct evidence for the macrostem.

Evidence for the derivational stem comes from the fact that, as was discussed in §2.2.2.2, Mushunguli exhibits height harmony in roots and extensions. Examples are repeated below in 109, with additional examples added.

(109) Height harmony (repeated)

a.	tʃi-l**o**^{ŋ}g-**e**:s-a	'we spoke a lot/too much'^{D}
b.	k^{w}-**i**:v-**i**:s-a	'to hear a lot/too much'
c.	wa-his**a**b-**i**:sa	'they counted a lot/too much'^{D}
d.	si-w-o:-goh-e:s-a	'I frightened them a lot/too much'
e.	na-n-it-i:s-e	'I will go a lot/too much'
f.	na-n-iv-i:s-e	'I will hear a lot/too much'

As seen here, harmony applies to roots and extensions, but not to inflectional prefixes. This indicates that at least some form of stem is distinct from the inflected verb. However, this is not strong evidence for the derivational stem on its own; while it is the case that (109e,f) contain a final vowel that does not agree with the root vowels in height, it is also the case that alternating derivational extensions are probably best analyzed as containing underlying high vowels, due to the fact that high vowels surface in extensions following both root-final low and high vowels, while mid vowels only surface after mid vowels.

(110) High vowels in extensions

a. k^{w}-**aː**^{n}ɖ**a**m-**i**r-**aː**n-**a** ‘to lean back together’
b. k^{w}-as-iz-a ‘to bequeath’
c. k-umbik-is-a ‘to pile up a lot/too much’

A second form of evidence for the inflectional stem comes from reduplication. Reduplicated verbs in Mushunguli contain the final vowel, but not inflectional prefixes, as illustrated in (111).

(111) Reduplication

a. ku-toatoa ‘to hit a little’[D]
b. ku-siɲasiɲa ‘to look a little’[D]
c. ku-labuɲalabuɲa ‘to chew a little’[D]
d. a-gonagona ‘(s)he is sleeping a little’[D]
e. na-pulukapuluka ‘I am flying a little’[D]

I unfortunately do not have examples of an extended verb that is reduplication in the data available to me, but the fact that the final vowel is included in reduplication predicts that they would be duplicated as well.

This behavior will also justify an important detail for the analysis Mushunguli, which is that noun and modifier roots will be treated as an inflectional stem as well. This is because inflectional prefixes do not repeat when these words are reduplicated.

(112) Noun reduplication

a. m̩dodododo ‘small (cl1)’
b. m̥piri ‘adder’[D]
c. m̥piripiri ‘pepper’

The question now is how this applies to lexical levels in the analysis of Mushunguli onset structure and hiatus resolution. It is the case that the morphological structure of Mushunguli makes it difficult to test this; in general, roots are of the shape -(C)VC-, while extensions are of the shape -VC. This means that it is difficult to find examples of complex onsets *or* hiatus at these boundaries, and the most common cases -di- ‘eat’ and -to- ‘beat’ are definitely and possibly exceptional, respectively. However, the fact that prefixes are not reduplicated and do not participate in harmony is sufficient evidence to separate the prefixes into a second level, even if we must assume that hiatus resolution applies generally.[2]

Moreover, as discussed in §2.2.4, there is also clear evidence for a distinct phrase (= postlexical) level, given that there is optional coalescence across word boundaries (e.g. /simba i+no/ → [simbeːno]~[simba iːno] ‘this lion’), but glide formation never occurs (e.g. /tʃi+tuŋgulu i+tʃi+o/ → [tʃituŋgulu iːtʃo], *[tʃituŋgulʷiːtʃo] ‘that (prox) onion’). “Optional” is important here. Word-internally, hiatus is repaired categorically, and exceptions to it are similarly categorical. Hiatus resolution across word boundaries, however, does not appear to be obligatory. This indicates that while hiatus resolution does apply between words, it does not do so equally, meaning that we have access to a postlexical level whereby we can deal with the output of glide formation.

The analysis to come will have some similarities to the analysis proposed by Sibanda (2009) for the Nguni family of languages. Notably, these languages have an important set of similarities to Mushunguli: they have five-vowel invetories, high coalescence, glide formation, and a number of context-specific patterns of secondary articulations and

[2]The eventual analysis of secondary articulations and exceptional palatalization in the case of ‘eat’ will rely on an undominated indexed version of MAX-C; thus, even if glide formation applies at both the level of the derivational stem and the inflected word, the resulting derived glide will be protected all the way to the postlexical level.

(descriptive) elision. Crucially, like Mushunguli, these all apply in similar morphosyntactic contexts. Sibanda argues on these grounds that it is necessary to adopt Stratal OT to account for them, though he abstracts away from any independent justification for multiple levels, instead treating the phonology as cyclic and post-cyclic.

For Mushunguli, the analysis of exceptions to hiatus resolution will support this type of analysis. The behavior of the non-coalescing stems crucially rules out deletion as a general pattern of hiatus resolution, and has revealed the possibility of treating secondary articulations as postlexical. The behavior of the palatalization exceptions, in turn, supports this distinction in levels: as discussed in §5.1.3, it will not be possible to analyze the exceptional palatalization pattern in one step. The fact that both a prefix and a verb root exceptionally undergoes palatalization is further evidence for the application of an onset structure repair at the postlexical level, as postlexical phonology is often assumed to be allowed to apply within fully derived words.

Because hiatus resolution appears to apply to both the stem and word level, for the purposes of this discussion we are going to conflate the two levels with respect discussion onset cluster repairs. Specifically, we will be looking at how the grammar deals with consonant clusters and secondary articulations, both derived and underlying. The following discussion and analysis will necessarily be more abstract than any preceding; unlike hiatus resolution, the set of contexts that can productively result in a consonant cluster are considerably rarer and less diverse (indeed, glide formation is by far the most productive way to potentially form a consonant cluster or secondary articulation). This is due to two factors: first, that words are minimally bisyllabic means that most morphemes contain a moraic segment, which for Mushunguli is nearly always a vowel. Second, the underlying structure of Mushunguli morphemes makes it much more difficult to bring consonants together than vowels; all suffixes are /-V(C)/, nearly all prefixes are /(C)V-/, and (native) roots are similarly nearly all /(C)VC/. However, the forthcoming analysis still needs to satisfy Richness of the Base. In order to do so, I will be considering both inputs that we have direct evidence for, as well as a set of hypothetical abstract inputs that cannot be allowed to surface (conforming to the generalizations just outlined).

Once the ranking for onset structure is determined, it will be combined with the ranking already generated for hiatus resolution, and crucial interactions will be discussed. In §5.1.2, I will build an analysis of the word level, and discuss the implications of re-ranking constraints. In §5.1.3, I will discuss a set of morphemes that exceptionally undergo palatalization. Similar to the non-coalescing stems from §4.2, the behavior of these exceptions definitively rule out certain possible grammars; namely, they reinforce the fact that Mushunguli creates full glides, not secondary articulations, at the stem level.

5.1.1 Word-level onset cluster repairs

To build our analysis of onset structure in Mushunguli, there are several important generalizations regarding syllable structure and phonotactics that need to be accounted for. Some of these were discussed in §2, while others have gone unelaborated upon. These are summarized below.

(113) Onset structure generalizations

a. The structure of the syllable in Mushunguli is /(C)V/; there are no (unexceptional) surface consonant clusters, and codas are disallowed.

b. Consonant clusters arising from concatenation of agreement prefixes with C-initial roots are repaired either via deletion of one of the two consonants or the formation of a secondary articulation.

c. Some secondary articulations (prenasalized stops and labialized segments) also occur within roots.

d. Labial segments can occur before round vowels, but *labialized* segments cannot.

e. Palatal segments are attested (e.g. /j/, /ɟ/), but (unexceptional) *palatalized* segments are not.

f. In the exceptional contexts in which palatalized segments do occur, it is always at the edge of a morpheme.

Our word-level ranking will need to account for these generalizations, either by selecting the optimum at this level, or by selecting an appropriate input to the postlexical level. In turn, the postlexical level will need to select the correct surface outputs, *without* accidentally (re-)creating impermissible ones. I will also assume, standardly, that the best possible analysis will be the simplest one that retains all of our important generalizations. With the addition of lexical strata, this will take the form of minimizing constraint reranking, as well as our prior goals of minimizing indexed constraints.

Because we know that surface consonant clusters do not generally arise in Mushunguli, complex onsets and codas must generally be impermissible.[3] We will restrict our discussion to the former. A complex onset (whether underlying or intermediately derived) violates the markedness constraint *COMPLEX:

(114) *COMPLEX: No more than one C or V may associate to any syllable position node. (Prince & Smolensky 1993/2004:96)

Word-level alternations and root-internal phonotactics both provide insight into how the grammar deals with underlying consonant clusters. The noun class 9/10 prefix generally surfaces as [ɲ] before V-initial roots, prenasalization before stop-initial roots, and otherwise does not surface at all.[4] Some examples are given in (115); note that for this and

[3]There are cases, particularly in the BLP dictionary, of consonant clusters occurring in loanwords from European languages, such as ***skulu***$_{B}$ 'school,' ***dikʃeneri***$_{B}$ 'dictionary,' ***kartoni***$_{B}$ 'box,' etc. I don't have an example of the latter two, but for my speaker, the former was pronounced as *sukulu*, suggesting that he has reanalyzed it as conforming to CVCV structure.

[4]Two caveats: first, some nouns that agree in class 9 do not have a noun class prefix, or have

the following set of examples, I return to a narrower transcription to make clear cases of prenasalization.

(115) Behavior of class 9/10 prefix

a.	/ɲ + oka/	→	ɲoka	'snake'
b.	/ɲ + buga/	→	ᵐbuga	'rabbit' (c.f *kabuga* 'bunny')
c.	/ɲ + puku/	→	m̥pʰuku	'rat'
d.	/ɲ + suwi/	→	suwi	'leopard'
e.	/ɲ + ŋoᵐbe/	→	ŋoᵐbe	'cow'

The evidence from class 9/10 tells us that consonant deletion and the formation of secondary (nasal) articulations are licit operations in Mushunguli. This is supported further by the appearance of productive labialization as a result of glide formation (e.g. *kʷiːva*), and the failure of expected labialization or palatalization to apply in restricted contexts.

Outside of some loanword patterns, root-internal phonotactics mirror patterns of alternation; we see both prenasalized stops and labialized consonants, but we do not see labialized consonants before round vowels, palatalized consonants, prenasalized fricatives, and so on. Some examples are given in (116).

(116) Root-internal secondary articulations

a.	m̩**tʷ**i	'head'
b.	m̩**sʷ**a	'louse'
c.	ŋ̊kʰa**ⁿd**e	'food'
d.	katu**ᵐb**iri	'vervet monkey'

That phonotactics mirrors patterns of alternation suggests that the same ranking should apply to both.

We will begin by building our analysis around (non-syllabic) nasals, as nasal prefixes productively undergo both prenasalization and deletion. As discussed in §3.5, I

the wrong noun class prefix, e.g. *baraza* 'veranda,' *luɟendo* 'chameleon.' I assume that these are morphologically suppletive, not phonological exceptions. Second, this prefix sometimes surfaces as a syllabic nasal, e.g. m̩bʷa 'dog'. Recall from §2.2.3 that words in Mushunguli appear to be minimally bimoraic (which for Mushunguli, is equivalent to bisyllabic). All forms like this I am aware of are monosyllabic, so I assume the syllabic nasal here is a manifestation of a minimal word requirement.

assume that forming any kind of a secondary articulation (or simply having one to begin with) violates a generalized markedness constraint *SECART.

(117) *SECART: assign a violation to any output segment with multiple articulations

Like vocalic secondary articulations, nasal ones are not wholly forbidden–prenasalized stops are allowed, but not prenasalized fricatives or sonorants. I will assume, following Padgett (1994), that there is a constraint that penalizes nasal articulations associated with continuants.

(118) $*^{n}$S: if [+nas], then [-cont]

Forming or *eliminating* a nasal secondary articulation also violates a faithfulness constraint IDENT(nasal).

(119) IDENT(nasal): assign a violation to any output segment whose value for [±nasal] differs from that of its input correspondent

Finally, deleting any consonant (nasal or glide) violates MAX-C.

(120) MAX-C: assign a violation for every input consonant without an output correspondent

I assume that underlying root-internal prenasalized stops surface faithfully. The ranking IDENT(nasal) » *SECART protects them from being turned into regular stops, while the ranking MAX-C » *SECART prevents them from being deleted.[5] This is illustrated in (121), using the word *bambo* 'marabou stork.'

[5]Two notes: first, deletion of a prenasalized stop would also create hiatus, so ONSET is also violated, but its violation is less important here than MAX-C. Second, it isn't actually *necessary* to assume that prenasalized stops surface faithfully. Surface evidence only tells us that word-internal prenasalized stops are allowed, while prenasalized anything else is not. Evidence from alternation, in turn, only tells us that there is productive prenasalization and deletion. As such, it is entirely safe to switch the ranking of *SECART and IDENT(nasal), which would uniformly neutralize word-internal prenasalized segments to their non-prenasalized forms. If this were the case, then all word-internal prenasalized stops surface from underlying nasal+stop sequences. I arbitrarily choose to treat prenasalized stops as underlying.

(121) Preservation of underlying prenasalization

/ba^{m}bo/	MAX-C	IDENT(nasal)	*SECART
a. ☞ |ba^{m}bo|			*
b. |ba.bo|		*!	
c. |ba.o|	*!		

Our current ranking is {MAX-C, IDENT(nasal)} » *SECART. To account for productive prenasalization in /N + stop/ contexts, we must add *COMPLEX to the ranking. *COMPLEX must dominate both IDENT(nasal) and *SECART, as in (122), using the example of /ɲ + buga/ → [^{m}buga] 'rabbit'.[6]

(122) Prenasalization across morpheme boundaries

/ɲ + buga/	*COMPLEX	MAX-C	IDENT(nas)	*SECART
a. |.mbu.ga|	*!			
b. ☞ |^{m}bu.ga|			*	*
c. |Øbu.ga		*!		

In (122), we see that MAX-C must also dominate IDENT(nasal) (and *SECART) to avoid deleting the nasal prefix, but the ranking of *COMPLEX and MAX-C has not yet been determined. Note that this ranking will also now choose a prenasalized stop in the case of a (hypothetical) underlying nasal + stop *sequence*.

(123) Word-internal prenasalization

/bambo/	*COMPLEX	MAX-C	IDENT(nas)	*SECART
a. |ba.mbo|	*!			
b. ☞ |ba^{m}bo|			*	*
c. |baØbo~bamØo|		*!		

Our ranking is now {*COMPLEX, MAX-C} » {IDENT(nasal), *SECART}. This ranking currently predicts that prenasalization occurs across the board, but we know that

[6]In this and all subsequent discussion, we are abstracting away from nasal place assimilation.

prenasalized non-stops never surface. We have evidence from the behavior of class 9/10 (which feature nasal stop prefixes) that nasals are deleted before fricatives; thus, *COMPLEX must dominate MAX-C. *ⁿS must also dominate MAX-C, to prevent prenasalization. This is illustrated in (124), using the example of /ɲ + suwi/ → [suwi] 'leopard'.

(124) Nasal prefix deletion

/ɲ + suwi/	*COMPLEX	*ⁿS	MAX-C	IDENT(nas)	*SECART
a. |.nsu.wi|	*!				
b. |ⁿsuwi|		*!		*	*
c. ☞ |Øsu.wi|			*		

As a final note, this ranking does not treat (hypothetical) underlying prenasalized non-stops the same way as it does potential derived ones. Instead, they lose their prenasalization (125), while underlying /N + non-stop/ sequences undergo deletion of one of the segments (126).

(125) De-nasalization of word-internal underlying ⁿs

/baⁿsa/	*COMPLEX	*ⁿS	MAX-C	IDENT(nas)	*SECART
a. |ba.ⁿsa|		*!			*
b. ☞ |ba.sa|				*	
c. |baØ.a			*!		
d. |ba.nsa|	*!			*	

(126) Word-internal consonant deletion

/bansa/	*COMPLEX	*ⁿS	MAX-C	ID(nas)	*SECART
a. |ba.nsa|	*!				
b. |ba.ⁿsa|		*!		*	*
c. ☞ |banØa|~|baØsa|			*		

The ranking for nasal repairs is given in Fig 5.2. As we turn our attention to vocalic inputs, the important rankings going forward are *COMPLEX » *SECART, *COMPLEX »

MAX-C, and MAX-C » *SECART. This ranking (*COMPLEX » MAX-C » *SECART) is fully determined.

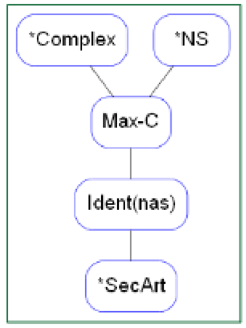

Figure 5.2: Ranking established for repair of clusters involving nasals

There are three categories of concern: underlying secondary articulations, (hypothetical) underlying /CG/ sequences, and underlying /$CV_{[+hi]}$ + V/ contexts. The only licit *surface* secondary articulation is labialization before unround vowels. In all other cases, the underlying vocalic portion needs to be neutralized entirely. As discussed in §3.5, I assume that forming or removing a vocalic secondary articulation (= labialization or palatalization) violates the cover constraint IDENT(V-Place).

(127) IDENT(V-Place): assign a violation to any output segment whose value for [±back] or [±round] differs from that of its input correspondent

As illustrated in (116), labialized consonants do occur before unround vowels in roots. Similar to how we assumed that underlying prenasalized stops can surface faithfully, we will assume that an underlyingly labialized consonant also surfaces faithfully. For this to be possible, IDENT(V-Place) must dominate *SECART, as illustrated in (128). I have also included *COMPLEX as well, to illustrate a harmonically bounded candidate.

(128) Preservation of underlying labialization

/C^wa/	IDENT(V-Place)	*COMPLEX	*SECART
a. ☞ |C^wa|			*
b. |Ca|	*!		
c. |Cwa|	*(!)	*(!)	

In the tableau in (128), we see that the fully faithful candidate |C^wa| wins due to the low ranking of *SECART. The second candidate, in which the secondary articulation is simply lost, loses because IDENT(V-Place) dominates *SECART. The third candidate is fission, which splits the labialized segment into a CG onset. This candidate is harmonically bounded by the second, because fission of this segment both changes a place feature AND creates a complex onset.

We have already established that constraints penalizing epenthesis (§3.1), forming glides from low vowels (§3.3), and vowel deletion (§4.2) are all undominated. As IDENT(V-Place) is defined, it is violated once by any instance of (non-identity) coalescence. Thus, ranking IDENT(V-Place) above ONSET will result in hiatus being unresolved in coalescence contexts. This is illustrated in (129), using a deletion candidate for illustrative purposes.

(129) Accidental blocking of coalescence

/$ka_1 + i_2v + a$/	MAX-V	IDENT(V-Place)	ONSET
a. ☞ |$ka_1.i_2$.va|			*
b. ☹ |$ke{:}_{1,2}$.va|		*!	
c. |$k\text{Ø}_1 i{:}_2$va|	*!		

Because of this, ONSET must dominate IDENT(V-Place), just as it does all other constraints involved in coalescence.

(130) Coalescence permitted

/$ka_1 + i_2v + a$/	MAX-V	ONSET	IDENT(V-Place)
a. |$ka_1.i_2.va$|		*!	
b. ☺ |$ke\!ː_{1,2}.va$|			*
c. |$k\varnothing_1 i\!ː_2 va$|	*!		

Importantly, that IDENT(V-Place) is violated by coalescence candidates does *not* undermine the unified generalization from (97). The locus of violation of IDENT(V-Place) is the prefix (due to the change of placeless to having a place feature), so indexation IDENT(V-Place) (or either of the two constraints it technically represents) cannot block it. Like UNIFORMITY, however, IDENT(V-Place) could be indexed to *prefixes* to block coalescence.[7]

IDENT(V-Place) also must dominate *COMPLEX in order to avoid coalescing two non-identical high vowels. Because *COMPLEX has already been shown to dominate *SECART (122), the optimal candidate for an underlying /$CV_{[+high]} + V$/ context is glide formation; that is, we have now ruled out one-step labialization in a context where labialization is allowed to (and eventually will) surface. This means that for *every* /$(C)V_{[+high]} + V$/ context (save for /(C)i + i/ and /(C)u + u/), the only option is glide formation, not secondary articulations.

(131) Glide formation, not coalescence or secondary articulation

/$ku_1 + i_2v + a$/	ONSET	IDENT(V-Pl)	*COMPLEX	*SECART
a. |$ku_1.i_2.va$|	*!			
b. ☞ |$kw_1 i\!ː_2.va$|			*	*
c. |$k^w{}_1 i\!ː_2.va$|		*!		*
d. |$ki\!ː_{1,2}.va$|		**!		

Recall from (124) that *COMPLEX dominates MAX-C; thus, by transitivity, IDENT(V-

[7]It could also be indexed to underlying /C^jV/ and /$C^wV_{[+rd]}$/ in order to preserve the secondary articulations, as well as /Cu-/ prefixes to block labialization at the word level (§5.1.3). In the latter case, it would not block stem-level glide formation; only MAX-ASC can be indexed to exclusively block glide formation and secondary articulations in this language.

Place) does as well. This ranking will thus select a deletion candidate for *any* input consonant-glide sequence (/CGV/), regardless of whether the hypothetical resultant secondary articulation would otherwise be licit in the language. This is illustrated by the tableau in (132).

(132) Stem-level general glide deletion

/Cwa/	IDENT(V-Pl)	*COMPLEX	MAX-C	*SECART
a. |Cwa|		*!		
b. ☞ |CØa|			*	
c. |C^{w}a|	*!			*

The ultimate consequence for the stem level is that we must assume that all word-internal labialized consonants are underlying, and not derived from an underlying /Cw/ or /Cu/ sequence. This assumption will not harm the analysis of onset structure in Mushunguli, and there is some evidence that it is reasonable on other grounds. There exist (rare) minimal pairs of roots that distinguish labialized consonants from non-labialized ones (e.g. *m̩ti* 'tree' vs *m̩t^{w}i* 'head'), though some are only possible to see if we allow morphologically complex forms to serve as minimal pairs (e.g. *tʃiːg^{w}a* 'we fell' vs *tʃiga* 'leg'). Labialized consonants in roots do not appear to cause compensatory lengthening, though given that roots tend to be either mono- or bimoraic in Mushunguli, it is difficult to get a labialized consonant outside of a final position (which is obligatorily shortened) or penultimate position (which is obligatorily lengthened, at least in phrases).[8] Finally, there is a difference in distributions between glides and labialized segments, as well as derived and underlying labialized segments. Root-internal instances of /w/ surface before [o], but labialized consonants do not. Similarly, labialized glides ([j^w]) only surface word-initially, and only in derived contexts.[9]

[8]I do have one example from a compound, which exhibits penultimate lengthening, but not compensatory lengthening: *m̩b^{w}ambaːgo* 'wild dog' (*lit* 'forest dog').

[9]That [j^w] surfaces at all further supports the assumption that labialization is a two-step pro-

There are two other outputs we need to consider: labialized consonants before *round* vowels, and palatalized consonants. Neither output is attested in roots or as the result of the operation of (regular) morphophonology; thus, we need to ensure that all relevant contexts do not select these outputs at any level. We have already ruled out the possibility of both outputs being derived at the stem level from underlying /CGV/ contexts. However, our current ranking predicts that *all* underlying secondary articulations surface faithfully, as in (133).[10]

(133) /C^wo/ → *|C^wo|

/C^wo/	IDENT(V-Place)	*COMPLEX	*SECART
a. ☞ |C^wo|			*
b. |Cwo|	*!	*	
c. ☹ |Co|	*!		

To prevent this result (as well as palatalization), we will assume as we did in §3.5 that IDENT(V-Place) is dominated by two context-sensitive markedness constraints: NOPAL, which penalizes palatalized segments generally, and *C^wO, which penalizes labialized segments before round vowels.

(134) **NOPAL**: no secondary palatal articulations

(135) ***C^wO**: no labialized segments before round vowels

cess. We must assume that underlying /j^w/ is neutralized by a high-ranked markedness constraint penalizing certain place of articulation combinations (e.g. *[-back, +round] or something similar, which by Richness of the Base must be active anyway to explain the lack of surface front round vowels). This means that independent of the ranking of IDENT(V-Place) and ONSET, we expect glide formation in the context /ju+V/. Subordination of this constraint to IDENT(V-Place) at the word level will allow |jw| to surface as [j^w].

[10]Note that a candidate that simply deletes the entire consonant would both violate MAX-C (which is relatively low-ranked, but does dominate *SECART) and ONSET (which dominates all of these constraints either crucially or by transitivity). Depending on where exactly this hypothetical underlying input falls in the word, as well, this could potentially violate L-ANCHOR.

The distinctions “palatalized” and “labialized” are critical for these constraints because as was discussed in §2.2.1, Mushunguli has both palatal and labial consonants, with no distributional restrictions. Palatal consonants can precede front vowels (*ɲimi* ‘tongues’), and labial consonants can precede round ones (*omola* ‘dish up ugali’). The only significant restriction is that homorganic glide-vowel sequences *ji* and *wu* are prohibited, which we accounted for in §3.4.

As illustrated by the tableaux in (136) and (137), ranking these constraints above IDENT(V-Place) will result in the neutralization of these hypothetical inputs to a singly-articulated consonant.

(136) Loss of palatalization

/C^jV/	NOPAL	IDENT(V-Pl)	*COMPLEX	*SECART
a. |C^jV|	*!			*
b. ☞ |CV|		*		
c. |CjV|		*	*!	

(137) Loss of labialization before round vowels

/C^wo/	*C^wO	IDENT(V-Pl)	*COMPLEX	*SECART
a. |C^wo|	*!			*
b. ☞ |Co|		*		
c. |Cwo|		*	*!	

To summarize, we now have a full ranking for onsets in this language. The ranking is responsible for the deletion of glides for all input /CG/ sequences, the preservation of labialization for all input /$C^wV_{[-rd]}$/ sequences, and the elimination of secondary articulations for /$C^wV_{[+rd]}$/ sequences as well as all cases of /C^j/. All other input consonant clusters undergo deletion of the first consonant, save for nasal + stop sequences, which instead become prenasalized stops.

Our primary interest is not onset structure alone, but the results of glide formation in postconsonantal contexts. To determine this full ranking, we will need to "zipper" together our current ranking with the hiatus resolution ranking already determined in §4.2.[11]

We have already determined much of this ranking. MAX-V conflicts with all other hiatus resolution constraints, including ONSET and IDENT(V-Place). It dominates *COMPLEX in order to choose glide formation in /CV$_{[+high]+V}$ contexts over vowel deletion. However, MAX-V has already been determined to be undominated due to the non-coalescing stems, so the relative important of this ranking is small. We have also already seen that ONSET crucially dominates IDENT(V-Place) (130) and *COMPLEX (131). ONSET also dominates MAX-C and *SECART, by transitivity.

IDENT(high) does not conflict with any constraints except *COMPLEX, which it must dominate in order to prevent coalescence for the underlying context /Ci + e/.[12]

(138) Labialization ruled out for glide formation contexts

	/vi$_1$ + e$_2$tu/	ONSET	IDENT(V-Pl)	IDENT(hi)	*COMPLEX
a.	|vi$_1$.e$_2$tu|	*!			
b. ☞	|vj$_1$eː$_2$tu|				*
c.	|v$^j{}_1$eː$_2$tu|		*!		
d.	ve$_{1,2}$ːtu			*!	

Finally, we must consider our specific markedness constraints. Because NOPAL and *C^{w}O dominate IDENT(V-Place), they must be ranked at least as high as ONSET. Therefore, while IDENT(V-Place) makes a crucial decision for unmarked cases of labialization (131), for these more specific contexts, *C^{w}O and NOPAL rule out secondary articulations. The

[11]To account for the behavior of -di- 'eat', I assume that the hiatus resolution ranking applies at the stem level as well. As with root-internal consonant clusters, root-internal cases of hiatus are presumably handled by exception features.

[12]The surface output of e.g. /vi + etu/ is [veːtu]; this could, in principle, be coalescence. However, that coalescence does *not* obtain in this context is crucial, due to the behavior of two exceptions, discussed in §5.1.3.

result for glide formation is shown in the tableau in (139), using a /Cu + o/ context to illustrate.

(139) Labialization ruled out for glide formation contexts

	$/\text{ku}_1 + \text{oma}_2\text{l} + \text{a}/$	ONSET	*C^{w}O	IDENT(V-Pl)	*COMPLEX
a.	$\|\text{ku}_1.\text{o}_2.\text{ma.la}\|$	*!			
b. ☞	$\|\text{kw}_1\text{oː}_2.\text{ma.la}\|$				*
c.	$\|\text{k}^w{}_1\text{oː}_2.\text{ma.la}\|$		*!	*	

With this, we have established a word-level ranking that will always select candidates with full glides when the input is a prevocalic high vowel. This is illustrated by the tableau in (140), which features a /Ci + a/ context. Note that I elected to rank NOPAL as high as MAX-V, but in actuality both NOPAL and *C^{w}O are unranked with respect to both MAX-V and ONSET.

(140) Glide formation always wins

	/Ci + a/	MAX-V	NOPAL	ONS	ID(V-Pl)	ID(hi)	*CPLX	*SECART
a.	$\|\text{Ci}_1.\text{a}_2\|$			*!				
b. ☞	$\|\text{Cj}_1\text{aː}_2\|$						*	
c.	$\|\text{C}^j{}_1\text{aː}_2\|$		*!		*			*
d.	$\|\text{Ceː}_{1,2}\|$				*(!)	*(!)		
e.	$\|\text{CØ}_1\text{aː}_2\|$	*!						

The important rankings for onset structure and hiatus resolution at the word level are schematized in the Hasse diagram in Figure 5.3.

Before moving to the postlexical level, where we will be exclusively focused on repairing complex onsets, I want to return briefly to some important intricacies that may be otherwise difficult to glean from this violation tableau and Hasse diagram.

The ranking NOPAL » IDENT(V-Place) » *COMPLEX » MAX-C is fully determined; however, some of these rankings are transitive. While both NOPAL and IDENT(V-Place)

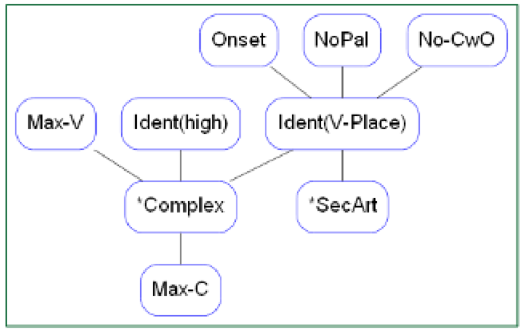

Figure 5.3: Word-level rankings for onset structure and hiatus resolution

dominate *COMPLEX, only the latter ranking is *crucial*; there are no cases in which the ranking NOPAL » *COMPLEX is the sole determinant of an optimum.

This is easier to see using comparative tableaux. There are three contexts in which a palatalization candidate *could* win: underlying palatalization, underlying /CjV/ sequences, and underlying /Ci + V/ sequences. First, we see in (141) that the crucial role played by NOPAL at the stem level is to prohibit underlying, root-internal palatalization from surfacing faithfully.

(141) Underlying palatalization

/C^ja/	NOPAL	IDENT(V-Pl)	*COMPLEX
☞ |Ca|		1	
a. ~ |C^ja|	W	L	
b. ~ |Cja|		=	W

Candidate (141a) is ruled out by the ranking NOPAL » IDENT(V-Place); while (b), which is fission, is harmonically bounded. Thus, there is no conflict between NOPAL and *COMPLEX.

For underlying /CjV/, we see in (142) that, again, there are no conflicts between these two constraints.

(142) Underlying /Cja/

/Cja/	NoPal	Ident(V-Pl)	*Complex	Max-C
☞ |CØa|				1
a. ~ |C^{j}a|	W	W		L
b. ~ |Cja|			W	L

Here, NoPal and *Complex are in agreement; both must dominate Max-C in order to prevent glides from surfacing, but there is no fixed ranking between them. The only context in which NoPal and *Complex can conflict are cases of /Ci + a/ hiatus, as seen in (143).

(143) /Ci + a/ hiatus resolution

/si + az + a/	NoPal	Onset	Ident(V-Pl)	*Complex
☞ |sjaːza|				1
a. ~ |si.a.za|		W		L
b. ~ |s^{j}aːza|	W		W	L

In (143), we see that if we were to omit NoPal from this ranking, we would still expect palatalization to lose in favor of glide formation. Notice as well that the violation profile of candidate (142a) has the same structure as the violation profile of candidate (143b); that is, NoPal and Ident(V-Place) share decision-making power. This indicates that the ranking NoPal » Max-C is also transitive.

To summarize, there are no cases that provide a strong ranking argument between NoPal and *Complex for this language. The only case in which they conflict is hiatus resolution, in which case Ident(V-Place) can be said to be subordinate. Moreover, there *are* strong ranking arguments for Ident(V-Place) » *Complex, as well as NoPal » Ident(V-Place); so, the ranking between NoPal » *Complex is transitive.

5.1.2 The postlexical level

The ranking at the word level produced the following intermediate outputs for hiatus resolution and related processes.

(144) Word-level outputs

a. Coalescence (with compensatory lengthening) of all /a + V/ (§3.3)

i. /ka + iv + a/ → |keːva| (49)

ii. /ka + ombok + a/ → |koːmboka| (51)

b. Coalescence (*without* compensatory lengthening) of all /V_i + V_i/ sequences and shortening of (hypothetical) underlying long vowels (§3.4)

i. /ka + asam + a/ → |kasama| (64)

ii. /iː/ → |i| (63)

c. Glide formation (with compensatory lengthening) of (non-identical) underlying /$V_{[+\text{high}]}$ + V/ (§3.2) and /$CV_{[+\text{high}]}$ + V/ (§5.1.1) sequences.

i. /i + asam + a/ → |jaːsama| (40), (44)

ii. /si + asam + a/ → |sjaːsama| (131), (140)

d. Preservation of underlying labialization before unround vowels

i. /C^wa/ → |C^wa| (128)

e. Elimination of all (hypothetical) underlying palatalization and labialization before round vowels

i. /C^ja/ → |Ca| (136)

ii. /C^wo/ → |Co| (137)

f. Deletion of all (hypothetical) glides in all underlying /CGV/ contexts.

i. /Cwa/ → |CØa| (132)

g. Exceptional blocking of high coalescence (§4.2)

i. /ka + it + a/ → |ka.ita| (94)

Moving to the postlexical level, most of these inputs have already been neutralized to an unmarked form. The primary inputs of concern are those from (144d) and a subset of those from (144c); that is, the post-consonantal glides.[13] We know that the ultimate results of glide formation mirror those of root and word-level phonotactic restrictions pertaining to glides and secondary articulations in the grammar; that is, the ranking at the postlexical level does not *introduce* new (non-exceptional) patterns; it only reinforces the ones that already exist.

The ranking that allowed us to preserve underlying labialization (144d) was IDENT(V-Place) » *SECART. This ranking obtains through all levels, as in (145).

(145) Preservation of underlying labialization at postlexical level

$\|gut^wi\|$	IDENT(V-Place)	*SECART
a. ☞ gut^wi		*
b. gut^wi	*!	

However, the unmodified word-level ranking will not allow for appropriate secondary articulations to be formed from the intermediate outputs of *glide formation*. At the word level, IDENT(V-Place) dominates *COMPLEX, which blocks the formation of secondary articulations, and *COMPLEX dominates MAX-C, which prefers deletion. As a result, if we feed a |CG| onset into our current ranking, we will simply delete the glide, just as we did at the word level. This is illustrated by the tableau in (146), which features the intermediate output of /ku + iv + a/. Note that in *this* tableau, we are assuming no constraint re-ranking.

[13] To account for the fact that glide formation does not apply at the postlexical level, I assume that the constraint penalizing glide formation, MAX-ASC is promoted to undominated status. Accounting for optional coalescence will presumably rely on some kind of variable ranking of UNIFORMITY, such as Stochastic OT (Boersma & Hayes 2001). I will leave this question for future research.

(146) Lexical ranking produces incorrect word-level glide deletion

|kwiː.va|	IDENT(V-Place)	*COMPLEX	MAX-C	*SECART
a. kwiː.va		*!		
b. ☹ kʷiː.va	*!			*
c. ☞ kØiː.va			*	

For the example illustrated in (146), we know that the output should be labialization and not deletion; these constraints need to be re-ranked. What is and is not a permissible re-ranking is a non-trivial question, however. The fact that lexical strata are assumed to output *regular* phonology does not change the fact that they are fundamentally cophonologies, and as such are subject to the same problems of unconstrained re-ranking (= cophonological explosion) (Anttila 2002; Inkelas et al. 1997; Inkelas & Zoll 2007).

Kiparsky (2015) hypothesizes that when moving from one lexical level to the next, the only possible re-ranking is to promote one or more constraints to undominated status. However, this restriction won't work for Mushunguli. If we promote *COMPLEX, we predict that it is always better to delete a glide than it is to ever form a secondary articulation, even in cases where we expect one. This is illustrated in (147).

(147) Accidental glide deletion

|kwiː.va|	*COMPLEX	IDENT(V-Place)	MAX-C	*SECART
a. kwiːva	*!			
b. ☹ kʷiːva		*!		*
c. ☞ kØiːva			*	

If we promote MAX-C, then the resultant low ranking of *COMPLEX means that we will simply leave glides unchanged, even in contexts where they do not surface. This is illustrated in (148):

(148) Incorrect glide preservation

|Cjaː|	MAX-C	IDENT(V-Place)	*COMPLEX	*SECART
a. ☞ Cjaː			*	
b. C^jaː		*!		*
c. ☹ CØa	*!			

This means we must at least promote both *COMPLEX and MAX-C above IDENT(V-Place) to get labialization in appropriate contexts.[14]

(149) Labialization permitted

|kwiː.va|	*COMPLEX	MAX-C	IDENT(V-Place)	*SECART
a. kwiːva	*!			
b. ☞ k^wiːva			*	*
c. kØiːva		*!		

Sadly, this ranking will not produce the right results for input |CjV| and |Cwo|; rather, we now predict that secondary articulations will be allowed across the board.

(150) Incorrect palatalization

|sjaːsama|	*COMPLEX	MAX-C	IDENT(V-Place)	*SECART
a. sjaːsama	*!			
b. ☞ s^jaːsama			*	*
c. sØaːsama		*!		

Adding our specific markedness constraints (in this case, NOPAL) will block palatalization; however, for deletion to obtain, *COMPLEX must dominate MAX-C.

(151) Correct deletion of derived |j|

|sjaːsama|	NOPAL	*COMPLEX	MAX-C	IDENT(V-Pl)	*SECART
a. sjaːsama		*!			
b. s^jaːsama	*!			*	*
c. ☞ sØaːsama			*		

[14] IDENT(V-Place) still must dominate *SECART to allow underlying labialization to surface faithfully, e.g. /gut^wi/ → |gut^wi| → [gut^wi]

Mushunguli thus presents a challenge for Kiparsky's hypothesis; two constraints do need to be promoted, but one of them crucially dominates the other.

Now, recall from the end of §5.1.1 the word level had a fully-determined ranking of four constraints: NOPAL » IDENT(V-Place) » *COMPLEX » MAX-C. However, the full determination of the ranking masks the fact that there are several intricate relationships between these four constraints; in particular, while NOPAL is the top-ranked constraint, it is arguably the least important of the four. This can be more easily seen if we examine the six pairwise rankings that these four constraints form and the candidates that they make crucial decisions about; this is summarized in Table 5.1.

Table 5.1: Crucial decisions made by pairwise rankings.

Ranking	**Crucial Decisions**
NOPAL » IDENT(V-Place)	/C^jV/ → |CV| (*|C^jV|)
NOPAL » *COMPLEX	none
NOPAL » MAX-C	none
IDENT(V-Place) » *COMPLEX	/Cu + V/ → |CwVː| (*|C^wVː|)
IDENT(V-Place) » MAX-C	/CwV/ → |CØV| (*|C^wV|)
*COMPLEX » MAX-C	/N + $C_{[+cont]}$V/ → |ØC| (*|NCV|)

In Table 5.1, we see that there are crucial, independent decisions made by four of the pairwise rankings; this is what leads to the full determination of the ranking. However, we also see that of the four constraints, top-ranked NOPAL only makes one crucial decision, which is the elimination of hypothetical underlyingly palatalized segments. For all other contexts, IDENT(V-Place) shares decision-making power, playing a back-up role in ruling out palatalization candidates in favor of glide formation and deletion.

By comparison, IDENT(V-Place) is quite important. Its ranking with respect to both *COMPLEX and MAX-C is crucial for preventing one-step secondary articulations and eliminating underlying glides, respectively. It also crucially dominates *SECART to preserve underlying labialization, and shares decision-making power with ONSET to prevent coalescence of two high vowels.

The promotion of *COMPLEX » MAX-C at the postlexical level, however, changes this. At the word level, deletion of underlying glides was preferable to forming a secondary articulation. At the postlexical level, this is reversed: labialization (outside of marked contexts) is now preferable to deleting a glide, which allows a word like (underlying) /ku + asam + a/ to eventually surface as [kʷaːsaːma]. As was seen in (149), this requires the subordination of IDENT(V-Place) to MAX-C.

The consequence of this is that NOPAL now has the opportunity to make crucial decisions, albeit with a more restricted set of candidates. There are no postlexical contexts that provide new /Ci + V/ inputs, and due to the fact that glide formation does not obtain between words, hypothetical /i # V/ inputs can be blocked by assuming that MAX-ASSOCIATION is promoted to an undominated status. Similarly, already palatalized segments have been repaired; thus, we are only considering |CjVː| inputs. We know that in all regular cases, these surface as [CØVː], with no lingering featural exponent of the underlying vowel. The comparative tableau in (152) illustrates the full effect of the re-ranking of *COMPLEX » MAX-C on these inputs.

(152) Postlexical input |CjVː|

|djaːsama|	*COMPLEX	NOPAL	MAX-C	IDENT(V-Place)
☞ dØaːsama			1	
a. ~ djaːsama	W		L	
b. ~ dʲaːsama		W	L	W

Here, we see two effects of re-ranking. First, because IDENT(V-Place) has lost any decision-making power for handling secondary articulations, the ranking NOPAL » MAX-C has become crucial to prevent a palatalized output from surfacing. Second, we see that the ranking of NOPAL and *COMPLEX has become indeterminate. The generalization to be made here is that the grammar of Mushunguli has nothing to say about the relative goodness or badness of palatalization versus retention of a complex onset. At the word level, palatalized inputs and underlying CG inputs are both neutralized, either by changing the place features (= eliminating the secondary articulation) or deleting the underlying consonant entirely. Forming a complex onset is preferred over palatalization, but this is because creating vocalic secondary articulations is generally avoided at the level of the word, not because palatalized segments are demonstrably worse. At the postlexical level, palatalization *is* avoided because it is palatalization, but since all input |CGV:| sequences are repaired, there is no way to determine which is worse.

That a ranking does not have anything particular to say about certain candidates relative to others is not problematic, provided that there truly is no evidence from the language at all to create a conflict. However, as will be seen in the following subsection, there actually is evidence from exceptions that forces the determination of *COMPLEX and NOPAL. This evidence provides insight into the behavior of indexed constraints within a leveled grammar; it will also further support the central claim that patterns of exceptionality are dependent on and reinforce larger generalizations about the grammar.

The Hasse diagram in Fig 5.4 schematizes the important aspects of the ranking established for the postlexical level thus far.

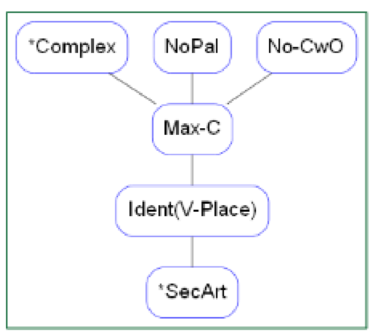

Figure 5.4: Ranking for postlexical repair of complex onsets

5.1.3 *Walljumping* palatalization

As was discussed in §2.3, noun class 5 has a more or less uniform set of agreement prefixes, all of which are /di-/. The only "exception" (in a morphological sense) is that like noun class 10, the class 5 prefix itself has a zero allomorph; however, it still surfaces as [di-] when the augment is attached (e.g. *idiboko* 'the banana', c.f. *boko* 'banana').

Because true adjectives are unmarked in class 5, we will focus on the three main classes of agreement prefixes: subject, object, and demonstrative. All of these prefixes surface faithfully before consonants, as illustrated in (153).

(153) Class 5 agreement prefixes

a.	/**di** + puluk + a/	**di**pulu:ka	'it (cl 5) flew'	Subject
b.	/si + **di** + pik + a/	si**di**pi:ka	'I cooked it (cl 5)'	Object
c.	/ɟula # **di** + no/	ɟula **di:**no	'this (cl 5) frog'	Demonstrative

The examples in (154), which feature prevocalic contexts, reveal that one of these prefixes is not like the others.

(154) Exceptionally-behaving demonstrative prefix

a.	/di + asam + a/	da:sa:ma	'it (cl 5) gaped'	Subject
b.	/si + di + az + a/	sida:za	'I lost it (cl 5)'	Object
c.	/ɟula # **di** + aŋgu/	ɟula **ɟ**a:ŋgu	'my (cl 5) frog'	Demonstrative

While the subject and object prefixes surface without an exponent of the underlying vowel, as expected, the demonstrative prefix instead surfaces as a palatal stop in hiatus contexts.[15] Importantly, this behavior is not representative of the morphological class of demonstrative prefixes; all others of the shape /Ci-/ undergo glide deletion. It is also not representative of the class of coronals; while there is no /ti-/, we have already seen that /ni-/, /si-/, and /zi-/ all surface unpalatalized in all relevant contexts.

[15]While a careful phonetic study is warranted, there are no measurable or impressionistic differences between the output of exceptional palatalization and underlying palatal stops that surface faithfully. The only detectable difference is that implosivity is weaker in intervocalic contexts, but this is true of voiced stops generally, including underlying palatal stops (e.g. the [ʄ] in *ɟula* is more strongly pre-voiced than the one in *maɟula*).

(155) Demonstrative and coronal prefixes

/vi-/	vituŋgulu **veːtu**	'our (cl 8) onions'	
/zi-/	ŋ̊kunde **zaːŋgu**	'my (cl 10) beans'	**Morphological Class**
/tʃi-/	tʃiŋ̊ko **tʃaːke**	'his/her (cl 7) elbow	
/ni-/	ka**naː**za	'(s)he lost me'	
/si-/	**soː**geːra	'I swam'	**Alveolars**
/zi-/	**zaː**ɟa	'they (cl 10) are eating'	

Interestingly, this alternation is not exclusive to the class 5 demonstrative prefix. It can also be observed in the verb 'eat'. Before non-high vowels, it surfaces as *ɟ* (including citation form *kuːɟa*). However, in contexts where an *i*-initial verbal extension immediately follows the root, it instead surfaces as [d]. This is illustrated in (156).

(156) Palatalization and identity coalescence in 'eat'

a. siː**ɟ**a 'I ate'
b. naniː**ɟ**e 'I will eat'
c. ku**ɟ**aːna 'to eat together'
d. si**d**iːsa 'I ate a lot'
e. nani**d**iːse 'I will eat a lot'
f. ku**d**iːra[B] 'to eat for'

Due to lexical gaps, it is impossible to generate a full paradigm of /di + V/ for both 'eat' and the demonstrative prefix; to my knowledge, there are no /i/-initial modifiers that take demonstrative concord prefixes, and there are no C-initial extensions.[16] However, a parsimonious analysis would assume that both the prefix and the root for 'eat' are underlyingly /di/. They behave identically before /e/ and /a/, and the surface form of 'eat' before /i/-initial extensions is consistent with our analysis of hiatus resolution, which prefers identity coalescence over glide formation. Thus, we will be treating these two morphemes as exceptions, and the same type of exception, at that.

This is a walljumping pattern: the expected repair (glide deletion) has been blocked, while an alternative (palatalization) has applied. However, this is complicated, because the surface results of postconsonantal glide formation are arrived at via derivation.

[16]There is the possibility of the passive suffix /-w/; however, *di* takes the V-initial passive allomorph *-igʷ* (*digʷa*).

This introduces two wrinkles: first, it is unclear how indexed constraints should function across lexical levels, and second, it is unclear whether these are exceptions to word-level glide formation (one-step palatalization), or exceptions to postlexical glide deletion (two-step palatalization).

With regard to the first wrinkle, there is no consensus in the literature about the relative compatibility of different approaches to morpheme-specific phonology. However, proponents of Stratal OT have raised weak (Kiparsky 2015) or strong (Bermúdez-Otero 2012) objections against lexically-indexed constraints. Kiparsky (2015) grants that there are "genuine" lexical exceptions, but does not elaborate on what should be done with them. Instead, he suggests that Stratal OT is the correct solution for "opacity," pointing at lexically-indexed constraints as one of several unnecessary piecemeal solutions.

Bermúdez-Otero (2012) argues much more fervently against indexed constraints, stating that they lack empirical content and suggesting that abstract representational specifications are the way to handle difficult cases. First, I object to the notion that a purely formal device like indexed constraints needs to have "empirical content"; formalisms *model* empirical content, they do not innately have any themselves. Second, Stratal OT is itself a formal device loosely supported by some general assumptions about cognition; this is equally true of indexed constraints, which have been implemented successfully into multiple learning models that capture speaker knowledge about phonotactic frequency distributions and propensities for some lexical items to undergo or not undergo alternations (Moore-Cantwell & Pater 2016; Zymet 2018, 2019). Indeed, Zymet's work in particular has successfully adopted indexed constraints as a solution to lexically-specific variation, indicating that they *can* be active at a postlexical level. Third, as I discuss in §6.1.2 and §6.2, some representational alternatives for the Mushunguli case are wholly incompatible with Stratal OT, while others are, at best, capable of capturing some patterns, but not others.

Either way, both of these arguments seem to be predicated on the notion that indexed constraints and lexical levels are incompatible formal devices. But this hasn't been demonstrated to be the case, and indeed Mushunguli seems to be a counterexample to this. Stratal OT is used here as a solution to regular opacity that is supported by morphology and phonology, while indexed constraints are used to capture arbitrary lexically-specific patterns. The only question is exactly when indexed constraints are introduced into a leveled grammar; however, the answer for this case study is fairly straightforward. Lexical indexation assumes that constraints are cloned as a response to the detection of a surface-level inconsistency between expected outputs. As we have seen, the constraint then becomes part of the grammar, and can affect it. Stratal OT assumes that constraints are promoted across levels, not created. While we are obviously relaxing the latter assumption, the more restrictive hypothesis is that regardless of when an indexed constraint becomes "active," it must exist at all levels of the grammar. This makes the walljumping exceptions in Mushunguli an interesting case for examining the way that indexed constraints interact with a level-ordered grammar.

This brings us to the second wrinkle: the question is whether we must treat palatalization as a one-step process–which would seem to undermine the derivational nature of glide formation in this language–or a two-step process, which would support it. Interestingly, only the latter case is a possible solution for Mushunguli, which means that the exceptions themselves reinforce the derivational nature of glide formation in this language. I will discuss this first, and then consider the alternative case.

5.1.3.1 Mushunguli: two-step walljumping palatalization

Recall from the end of §5.1.2 that we were left with the following ranking of the four onset structure constraints: {*COMPLEX, NOPAL} » MAX-C » IDENT(V-Place). Because IDENT(V-Place) no longer makes important decisions with respect to |CGV:| inputs, we will eliminate it from consideration: this leaves us with an indeterminate ranking between two top-level constraints that must dominate some subordinate constraint; that is, our ranking {M, C} » F (from §4.1.3). Again, this ranking is unproblematic provided that there are no exceptions to glide deletion. However, we *do* have an exception to glide deletion, which leaves us with no choice but to index MAX-C (MAX-C^{L2}). As we saw with the hypothetical skeletal ranking, simply promoting MAX-C^{L2} above our indeterminate ranking will result in a grammar that cannot make a decision between preserving the entire glide or simply its place feature. This is illustrated in (157).

(157) Indeterminate glide protection

|dj^{L2}aːŋgu|	MAX-C^{L2}	*COMPLEX	NOPAL	MAX-C
a. ? djaː.ŋgu		*		
b. dØaː.ŋgu	*!			*
c. ? dʲaː.ŋgu			*	

Just as was the case with MAX-V and ONSET for the non-coalescing stems, the indeterminate ranking between *COMPLEX and NOPAL cannot persist once we introduce MAX-C^{L2}. The introduction of IDENT(high)^{L1} resulted in the disambiguation from {M, C} » F to the simple blocking schema C » M » F; that is, deletion was worse than leaving hiatus unresolved. In this case, *COMPLEX, the markedness constraint driving the repair (M) is what must be promoted, resulting in the *walljumping* schema M » C » F; that is, leaving a complex onset unrepaired is worse than palatalization. This is illustrated in (158):

(158) Walljumping palatalization

|dj^{L2}aŋgu|	MAX-C^{L2}	*COMPLEX	NOPAL	MAX-C
a. djaː.ŋgu		*!		
b. dØaː.ŋgu	*!			*
c. ☞ dʲaː.ŋgu			*	

That this exception's behavior is derived via the interactions of lexical levels and indices raises important questions regarding the notion that indexed constraints can make phonological generalizations. In §4.2.2, we saw that the indexed constraint used to capture the behavior of the non-coalescing stems (IDENT(high)^{L1}) was capable of unifying generalizations about the form and behavior of these exceptions. However, MAX-C^{L2} has no ability to adjudicate over either morpheme at the level it is introduced; rather, it takes advantage of intermediate derived structure.

There is evidence to support this, however. While we have been treating the stem and word level as having the same ranking, it is the case that if glide formation applies to the root-final vowel of the root, the glide will be created at an earlier lexical level than the one created for the class 5 demonstrative concord prefix; this glide will need to be protected from becoming a secondary articulation. MAX-C^{L} can do this for us; essentially, what the grammar seems to be doing is recycling the same indexed constraint for multiple problems.

Moreover, it is the case that there is more than one type of phonological generalization. There is nothing representationally special about (synchronic) /d/ that would suggest it is more amenable to (synchronic) palatalization than other segments. While there is a palatal counterpart, this is also true of /n/, yet no version of /ni-/ palatalizes as an alternative to glide formation. Moreover, there is at least one other word that appears to undergo palatalization, which is *kogoːfʲa* 'to frighten'. Given the disparity in representations, it seems like the generalization is not one that is merely representational.[17]

As such, the phonological generalizations to be made here have to do with the operation of the grammar itself. First, that these morphemes surface palatalized underscores the derivational nature of glide formation in this language; vocalic secondary articulations only happen in two steps. Second, this underscores the generalization that consonant

[17] I do not have examples of this word in a derived context, so it isn't possible to determine what its behavior would be before a derivational extension. However, /f/ and /d/ do not form a natural class, and whether 'frighten' is underlyingly /ogofi/ or /ogofj/, either will be protected if it is indexed to MAX-C^{L2}

clusters are especially marked in this language, moreso than unresolved hiatus. Hiatus is tolerated regularly at the left edge of a word, variably across word boundaries, and exceptionally both word- and root-internally. Consonant clusters never surface except for rare cases of root-internal clusters in loanwords, and those which surface tend to be exactly the kind that presumably could not form licit secondary articulations in this (or perhaps any other) language, e.g. *kartoni* 'cardboard box'. Forms like this can also be indexed to MAX-C^{L2}; the fact that they result in surface clusters can be attributed to the fact that it simply isn't possible for them to form secondary articulations.

The ranking for onset structure repairs at the postlexical level is schematized in the Hasse diagram in Figure 5.5.

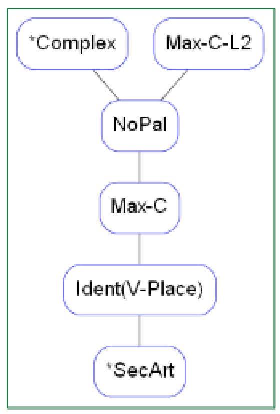

Figure 5.5: Postlexical-level ranking for onset structure (including MAX-C^{L2})

5.1.3.2 Not Mushunguli: one-step walljumping palatalization

The analysis we adopted for the walljumping exceptions implies that despite the fact that these look like exceptions to word-level glide formation, they are actually exceptions to postlexical glide *deletion*. That is, what we have here are genuine lexical exceptions, but the exceptionality pertains to an operation that can only occur at the postlexical level. Standard assumptions born from Lexical Phonology (Kiparsky 1982b), which is the precursor to Stratal OT, are that exceptionality (especially categorical exceptionality) is a property of lexical levels. This raises the question as to whether an alternative, one-step palatalization, can apply. While this would not obviate the need for postlexical onset cluster repairs, it *would* at least remove the potentially uncomfortable assumption that categorical exceptionality is applying at a postlexical level. However, the answer to this question for Mushunguli is no; while it is possible to generate an analysis that performs walljumping palatalization in one step, the output of that ranking is inconsistent with attested forms in Mushunguli.

To see this, we construct an alternative version of Mushunguli, built from the same set of word-level inputs and the same constraint set as the analysis of Mushunguli (including the non-coalescing stems). The only way to achieve word-level one-step palatalization in lieu of glide formation is to index *COMPLEX; we will call this constraint *COMPLEX^{E} to distinguish it from indexed constraints used in the real analysis. The addition of this indexed constraint linked to exceptional outputs /di + ^{E}aŋgu/ and /di^{E} + etu/ produces the same set of intermediate inputs to the word level, except that it also outputs exceptionally palatalized forms. It also results in a nearly identical word-level ranking; crucially, the relationships between the onset structure constraints are the same in both scenarios.[18]

For brevity's sake, we are going to significantly truncate discussion of the word level. The violation tableau in (159) illustrates that this grammar outputs exceptional palatalization (as opposed to glide formation, deletion, or coalescence) in one step; note that solid and dashed lines here illustrate the fully-determined ranking, not only the ranking determined by this example.

[18]The difference is that for the real word-level analysis of Mushunguli, IDENT(high) has no crucial or interesting relationship with either NOPAL or IDENT(V-Place). The former dominates IDENT(high) by transitivity, while the latter is unranked with respect to IDENT(high). For the hypothetical case of one-step palatalization, the introduction of *COMPLEX^{E} forces determination: IDENT(high) crucially dominates NOPAL (and IDENT(V-Place) by transitivity) in order to prevent coalescence of /i + e/ when glide formation is blocked in exceptional contexts.

(159) Alternative Mushunguli: word-level one-step palatalization

/di^E + etu/	*CPLX^E	MAX-V	ONS	ID(hi)	NOPAL	ID(V-Pl)	*CPLX
a. di.etu			*!				
b. dje:tu	*!						*
c. ☞ d^{j}e:tu					*	*	
d. dØe:tu		*!					
e. de$_{1,2}$:tu				*!			

Glide formation is still chosen in regular contexts (160).

(160) Alternative Mushunguli: word-level glide formation

/vi + etu/	MAX-V	ONS	ID(hi)	NOPAL	ID(V-Pl)	*CPLX	*SECART
a. vi.etu		*!					
b. ☞ vje:tu						*	
c. v^{j}e:tu				*!	*		*
d. vØe:tu	*!						
e. ve$_{1,2}$:tu			*!				

Note that the output of this grammar is that any inputs to which *COMPLEX^E is indexed are now palatalized. However, no form of *COMPLEX makes important decisions regarding inputs with secondary articulations: the only candidate it could rule out is the fission candidate [dj], which is harmonically bounded due to the fact that splitting |d^j| into [dj] violates both *COMPLEX (and *COMPLEX^E) and IDENT(V-Place). With its decision-making power eliminated, the metaphorical protection provided by *COMPLEX^E is nullified, and the exceptions are now formally equivalent to regular palatalized inputs.

With our intermediate outputs in hand, we move to the postlexical level. For a one-step analysis to be possible, the exceptions must surface with palatalization intact. This leads to two possible grammars, each interesting in its own right, but neither of which is actually Mushunguli.

Recall from (145) above that the ranking IDENT(V-Place) » *SECART still must hold at the postlexical level in order to prevent neutralizing underlying labialized inputs.

There is a deeper generalization here: for *any* grammar to allow a secondary articulation to surface, all markedness constraints penalizing that secondary articulation must be subordinated to all faithfulness constraints that protect it. Thus, any grammar that wants to preserve palatalization at the postlexical level must subordinate NoPal to Ident(V-Place) (or rather, promote Ident(V-Place) above NoPal). This prevents input |dʲ| from reverting back to [d], as seen in (161). Notice in this tableau that *Complex^{E} no longer makes any crucial decisions, and indeed has no (non-stipulative) relationship with its unindexed counterpart at all.

(161) All scenarios: postlexical preservation of exceptional palatalization

|d^{jE}eːtu|	*Complex^{E}	*Complex	Ident(V-Pl)	NoPal	*SecArt
☞ dʲeːtu				1	1
a. ∼ djeːtu	W	W	W	L	L
b. ∼ deːtu			W	L	L

As seen in (161), the only work being done by *Complex^{E} is the penalization of a harmonically bounded candidate. However, the subordination of NoPal is *crucial.*

With the exceptions preserved, we need to consider the regular outputs of glide formation. The two expected outcomes for word-level inputs |CGVː| are labialization (for |CwVː|) and deletion (for |CjVː|). It turns out that we can have one, but not the other.

If we prioritize labialization (which we will call Scenario 1), then two other rankings must obtain: *Complex » Ident(V-Place), which favors labialization over complex onsets, and Max-C » Ident(V-Place), which favors labialization over deletion. This is illustrated for input |kwiːva| in (162).

(162) Scenario 1 outputs postlexical labialization (good)

|kwi:va|	*COMPLEX	MAX-C	IDENT(V-Place)	*SECART
a. kwi:va	*!			
b. kØi:va		*!		
c. ☞ k^wi:va			*	*

Combining these two pairwise rankings with the ranking IDENT(V-Place) » {NOPAL, *SECART} results in the ranking {*COMPLEX, MAX-C} » IDENT(V-Place) » {NOPAL, *SECART}. This will still output palatalization in exceptional cases, because we have maintained the ranking IDENT(V-Place) » NOPAL.

(163) Scenario 1 outputs exceptional palatalization (good)

|d^{jE}e:tu|	*CPLX	*CPLX^E	MAX-C	ID(V-Pl)	NOPAL	*SECART
a. ☞ d^je:tu					*	*
b. de:tu				*!		
c. dje:tu	*(!)	*(!)		*(!)		
d. Øe:tu			*!			

Unfortunately, this will preferentially palatalize *all* |CjV:| inputs. This is illustrated in (164), using a hypothetical input |zje:tu| (from /zi + etu/).

(164) Scenario 1 outputs universal palatalization (bad!)

|zje:tu|	*COMPLEX	MAX-C	IDENT(V-Pl)	NOPAL	*SECART
a. zje:tu	*!				
b. ☹ zØe:tu		*!			
c. ☞ z^je:tu			*	*	*

Note that if a grammar has universal palatalization, then there is no evidence whatsoever to suggest that the exceptions were ever exceptional in the first place. This version of lexical indexation relies on inconsistency detection, but if all /Ci + V/ outputs result in palatalization (one way or another), there is no inconsistency to detect! Thus, the learner would never have generated an indexed constraint at all.

Prioritizing glide deletion (which we will call Scenario 2) is possible under this theory of lexical indexation, but it also does not output Mushunguli. For this scenario, *COMPLEX and IDENT(V-Place) must both dominate MAX-C. This ranking will preferentially delete |j| over palatalization, as in (165).

(165) Scenario 2 outputs palatal glide deletion (good)

|zje:tu|	*COMPLEX	IDENT(V-Pl)	MAX-C	NOPAL	*SECART
a. zje:tu	*!				
b. ☞ zØe:tu			*		
c. z^{j}e:tu		*!		*	*

Again, this grammar outputs palatalization for the exceptional cases.

(166) Scenario 2 outputs exceptional palatalization (good)

|d^{jE}e:tu|	*CPLX	*CPLX^{E}	ID(V-Pl)	MAX-C	NOPAL	*SECART
a. ☞ d^{j}e:tu					*	*
b. de:tu			*!			
c. dje:tu	*(!)	*(!)	*(!)			
d. Øe:tu				*!		

Unfortunately, this ranking also prefers to delete *any* glide in a |CGV:| context, as seen in (167).

(167) Scenario 2 outputs universal glide deletion (bad!)

|kwi:va|	*COMPLEX	IDENT(V-Place)	MAX-C	*SECART
a. kwi:va	*!			
b. ☞ kØi:va			*	
c. ☹ k^{w}i:va		*!		*

The implications of this second scenario are more interesting than the first, because this is *not* a case where the indexed constraint has had no effect. Remember: this is alternative Mushunguli, in which vowel deletion is independently ruled out due to the existence of the non-coalescing exceptions. Regardless, this grammar outputs something that looks like deletion for *all* input /Ci + V/ and /Cu + V/ sequences. The existence of the walljumping exceptions, even in a grammar like this, nevertheless both provides positive evidence against deletion, and also reinforces that forming secondary articulations is a two-step process for any language like Mushunguli. Even though the indexed constraint is rendered inactive at the postlexical level, its effects are still felt in the output of the grammar.

5.2 Total non-participants in hiatus resolution

So far, our discussion of the typology of exceptional blocking has focused on exceptions to a single alternation within a conspiracy. There is a third option: exceptions which fail to participate in a conspiracy *at all*. We will refer to these as *total non-participants* (TNPs for short). We have already seen an example of this with the Yine TNP /-wa/ 'yet/still,' which neither triggered nor underwent vowel deletion.

Mushunguli's instantiation of this typological option is a bit more complicated in character, but not in analysis. There is a set of roots that are total non-participants in the hiatus resolution *conspiracy*, as illustrated by the paradigm for one such root, *ona* 'see' in (168).

(168) Paradigm for total non-participant *ona*

a.	oːna	'see!'	/on + a/
b.	ku.oːna	'to see'	/ku + on + a/
c.	si.oːna	'I saw'	/si + on + a/
d.	i.oːna	'it (cl 9) saw'	/i + on + a/
e.	u.oːna	'it (cl 3) saw'	/u + on + a/[19]
f.	ka.oːna	'(s)he saw'	/ka + on + a/
g.	naː.oːna	'I am seeing'	/ni + a + on + a/
h.	k^{w}iː.oːna	'to see it/them (cl 9/cl 4)'	/ku + i + on + a/
i.	ku.oneːsa	'to see a lot/too much'	/ku + on + is + a/

[19]This form was coerced as part of a paradigm; animate class 3 nouns take class 1 agreement.

Verbs like *ona* never trigger glide formation, nor do they participate in coalescence; rather, they behave identically to C-initial roots, save for the fact that no consonant ever surfaces. Note that this behavior is local to the verb root itself, as illustrated by examples (168g,h). Note also from (168i) that *ona* triggers harmony, meaning that it meets the representational requirements necessary for it to do so.

The majority of TNPs begin with mid or low vowels, and all mid- and low-vowel-initial exceptions to hiatus resolution are TNPs. I have identified no cases of exceptions that fail to trigger glide formation without also blocking coalescence, provided that relevant morphological contexts are applicable to the root in question; that is, the only glide formation-exclusive exceptions are the cases of walljumping palatalization discussed in the preceding section. This is consistent with the hierarchy of costliness from (66) and the generalizations discussed in §4.2.2; our analysis predicts that the only type of exceptions that could exclusively block coalescence are those beginning with high vowels. As such, the form of the TNP exceptions in Mushunguli again concords with other patterns.

However, importantly, unlike the non-coalescing stems, there *are* some TNPs that begin with high vowels, e.g. *iɟariʃa* 'to rent' (<kuidyarisha>[B] 'to rent', <naidyarisha>[B] 'I am renting'). This indicates that vowel height (and quality) is immaterial for these exceptions, suggesting that their failure to participate is not featurally-driven.[20]

This generalization is important, because although the TNP exceptions follow from the general patterns of the grammar, their behavior is unique as it is not exclusively a non-undergoing pattern. Unsurprisingly, neither indexed constraint proposed so far can capture this pattern. IDENT(high)^{L1} can only block high coalescence, and only for stems beginning with high vowels; even if we applied it to the example of 'rent', glide formation would still not be blocked without the indexation of an additional constraint.

[20]Why there is a disparity between initial vowel qualities of these exceptions is unclear. Given that many more of the TNP class are identifiable borrowings, it may be the case that the phonotactics of the donor language(s) play a role in this. It is also possible that high-vowel initial borrowings are more likely to be nativized quickly or grouped with the non-coalescing stems.

Similarly, MAX-C^{L2} is *truly* inappropriate here, as the root-initial vowel in these cases is never expected to alternate with a glide.However, it is possible to capture this pattern relatively straightforwardly using indexed markedness. One such approach is to use an indexed version of ALIGN(L, root, L, σ).

(169) **ALIGN-L(Root^{L3},σ)**: the left edge of an indexed root must align with the left edge of a syllable

ALIGN-L(Root^{L3},σ) can block coalescence and glide formation, as both processes realign the edge of the root as part of a syllable nucleus. Like the other two indexed constraints, this one must also be undominated at every level; this will prevent the resolution of hiatus all the way to the surface. This is illustrated by the tableaux in (170) and (171), illustrating glide formation and coalescence, respectively.

(170) Total non-participation in glide formation

/i + on + a/	ALIGN-L^{L3}	MAX-V	ONSET	MAX-ASC
a. ☞ i.ona			*	
b. joːna	*!			*
c. Øona		*!		

(171) Total non-participation in coalescence

/ka + on + a/	ALIGN-L^{L3}	MAX-V	ONSET	UNIFORMITY
a. ☞ ka.ona			*	
b. koːna	*!			*
c. kØona	*!	*		

The tableau in (172) illustrates that ALIGN-L^{L3} will not prevent glide formation from applying to the prefix in an example like (168h).

(172) Glide formation permitted

	/ku + i + on + a/	ALIGN-L^{L3}	MAX-V	ONSET	MAX-ASC
a.	ku.i.ona			**!	
b. ☞	kwiː.ona			*	*
c.	kujoːna	*!			*
d.	kuØona		*!	*	

ALIGN-L^{L3} is an interesting choice of indexed constraint because it captures a complementary grammatical generalization. As discussed by McCarthy & Prince (1993); Ito & Mester (1999b), ONSET can be interpreted as an Alignment constraint, ALIGN-L(σ,C). ALIGN-L(Root^{L3},σ) is in a certain sense an inverse version of ONSET, albeit one that specifies a prosodic edge. That these exceptions seem to be undermining ONSET entirely may in turn appear to undermine the notion that exceptions are not extragrammatical. However, these exceptions *are* affected by, and reinforce, grammatical dependencies. First, this analysis correctly predicts that there should be no stems that *exclusively* and *arbitrarily* block glide formation; because these exceptions have strong left edges, we predict that any hiatus resolution strategy should be blocked. Second, that they lack unifying featural generalizations and block both coalescence and glide formation was predicted by the analysis of the non-coalescing stems in §4.2. The form and behavior of the non-coalescing stems predicted the form and behavior of other exceptions to coalescence, and the TNP exceptions conform to this prediction. Third, the fact that deletion does not apply in contexts like /i + ona/ reinforces both the inapplicability of deletion as a repair and the tolerance of hiatus at the left edge of the word. Fourth, and finally, these exceptions reinforce the fact that the stem is an independent morphological constituent from the word. Thus, even these exceptions which seem to entirely circumvent the hiatus resolution system in the language, nevertheless still conform to and reinforce larger generalizations.

5.3 Summary

In this chapter, we have completed our grammar fragment of Mushunguli. Our ranking now correctly repairs all hiatus resolution inputs, consonant clusters, as well as underlying and derived secondary articulations. Crucially, this was only possible to do after taking into consideration evidence from a set of simple blocking exceptions to coalescence (§4.2), which ruled out the possibility of vowel deletion as a possible hiatus repair in the language. The fact that only the exceptional stems were capable of determining this ranking indicates that exceptions play an important role in the grammar; by apparently undermining the preference to avoid unrepaired hiatus, the non-coalescing stems actually end up reinforcing a stronger prohibition against deleting vowels.

This had an immediate consequence on the grammar: given that deletion was no longer possible, these exceptions illuminated a path to sovling the ranking paradox introduced at the end of §3.5; namely, that glide formation had a secondary, postlexical derivational step, which we accounted for by adopting Stratal OT to model the postlexical grammar. This also now has reduced the descriptive four repairs in Mushunguli to two: glide formation, and coalescence.

Two other forms of exceptional blocking were also examined and analyzed: walljumping palatalization (§5.1.3), and total non-participants in hiatus resolution (§5.2). The former of these once again strongly reinforced the derivational nature of postconsonantal glide formation in this language, while the latter conformed both to the inapplicability of vowel deletion and the generalization that exclusive blocking of coalescence is predicted to only be possible if it can be tied to a featural quality of one of the input vowels. Taken together, all three of these exceptional patterns serve as reinforcing agents in the grammar; while they are lexical and unproductive, it is incorrect to treat them as extragrammatical or unimportant.

This chapter, in part, contains material that appeared in Proceedings of the 2018

Annual Meeting on Phonology. Hout, Katherine, Cascadilla Proceedings Project, 2019. It also contains coauthored material in preparation for submission for publication. Hout, Katherine and Baković, Eric. The dissertation author was the primary investigator and author of this material.

6 Alternative Representational Analyses of Mushunguli

The analysis presented in §4 and §5 relied on two separate strategies, one for exceptions (lexical indexation), and one for dealing with regular opacity that arose due to the interactions of morphophonological domains (Stratal OT). Both of these are examples of what are commonly referred to as *diacritic* theories (Kiparsky 1968/1982a; Pater 2010), which allow for the phonology to have direct evidence to some aspect of the morphology and/or the lexicon. A substantial class of alternative approaches to exceptionality exists, which posits a strictly "phonological" (that is, representational) treatment of exceptions. Theories of this type generally assume that most (if not all) exceptionality can be explained via appeal to some abstract underlying structure.

There are several common arguments for representational approaches over diacritic ones. The first is that proponents of representational solutions assume that phonologically-exceptional behavior should be handled by the phonology itself. Using a lexical diacritic to mark items as exceptional is effectively off-loading them to the lexicon and thus equivalent to treating them as extragrammatical. Moreover, applying an arbitrary lexical diacritic to a set of forms with shared characteristics can be viewed as treating those shared characteristics as coincidental; this is especially apparent in rule-based frameworks, but is also possible with indexed constraints and rankings if inappropriate constraints are indexed (e.g. indexing UNIFORMITY for the non-coalescing stems). Representational solutions by their very nature emphasize the role of phonological structure in determining phonological behavior; thus, while some authors have argued that diacritics can or should be used for marking exceptions which indeed appear to be completely idiosyncratic (Kisseberth 1970b; Kenstowicz & Kisseberth 1977), others have taken the hard-line stance that morpheme-specific diacritics are inappropriate for analyzing most or all cases of morpheme-specific phonology, including exceptionality (Inkelas et al. 1997; Inkelas & Zoll 2007; Kabak & Vogel 2011).

This latter hard-line stance is motivated by a second criticism, which is that diacritics are overly powerful; this is the problem of "cophonological explosion" as discussed by Inkelas et al. (1997). Here, once a marked constraint or ranking is invoked to capture one potentially interesting but unproductive pattern, there is nothing stopping the grammar from using these tools to explain every unproductive pattern, including apparently trivial ones such as the fact that some roots end with consonants and others end with vowels. This results in a potentially infinite set of cophonologies or indexed constraints (assuming an entirely naive theory).

As has already been demonstrated, lexically-indexed constraints can in fact capture phonological generalizations, and more importantly, they have clear explanatory and predictive benefits. While indexed constraints *do* treat exceptions as lexically-specific, they also tie their behavior directly to either underlying or output phonological structures. This places grammatical restrictions on the extent and form of exceptional behavior in any given language; it also allows exceptions to affect the grammar in subtle ways (as we have seen with the forced promotion of MAX-V). In other words, exception-marking in this way cannot be dismissed as treating exceptions as extragrammatical. In this chapter, I show that while representational solutions are in principle more restrictive than diacritic ones, the difference between the two approaches in practice is not significant enough to tip the scales in favor of representations. Furthermore, representational solutions within OT are in fact inadequately expressive to capture both exceptional and regular patterns in Mushunguli.

In the interest of space, I will primarily be focusing on alternatives to the non-coalescing stems, as of the Mushunguli exceptions, these are the most amenable to a representational analysis. This chapter is divided into two sections, to correspond roughly to two classes of representational solutions: §6.1 covers absolute neutralization of whole segments using both rules (§6.1.1) and Stratal OT (§6.1.2). It will be seen that the rule-based abstract representational analysis of the regular and exceptional patterns of hiatus resolution in Mushunguli, which is a slightly modified version of the analysis presented in Hout (2012, 2017), accounts for the two key generalizations from (90), but critically, it is unable to guarantee the unified statement of patterned exceptionality in (92) and thus to link these two generalizations in the way that the lexically-indexed constraint account does. A similar analysis is proposed for Stratal OT, but will be discarded on the grounds that it is unworkable.

§6.2 explores a more restricted abstract representational approach relying on featural prespecification, following Inkelas (1995); Inkelas et al. (1997). Prespecification is chosen over other forms of underspecification as it critically relies on a principled method of assigning underlying feature values. This analysis suffers from a similar issue as the rule-based abstract representational analysis: it is capable of capturing the critical generalizations, but struggles to unify them.

6.1 Absolute neutralization of whole segments

This section covers two forms of absolute neutralization of abstract segments: a rule based analysis (drawn from Hout (2012, 2017)) and absolute neutralization in Stratal

OT. The former will be demonstrated to be elegant, but insufficiently explanatory; the second will be shown to be impossible.

6.1.1 Absolute neutralization with rules

The Mushunguli non-coalescing stems represent a classic case of underapplication opacity (Kiparsky 1968/1982b; McCarthy 1999; Baković 2011). The absence of *any* obvious motivator for this behavior combined with some distributional facts about Mushunguli makes the non-coalescing stems amenable to elegant analysis via absolute neutralization of an abstract underlying segment.

This analysis has two main claims: first, that there are separate rules for coalescence, simplification, and glide formation. Skeletal forms of these are given in (173).

(173) Phonological rules

a. Coalescence: $V_{[+\text{low, -high}]} + V_{[\alpha\text{place}]} \rightarrow V_{[\text{-high, -low, }\alpha\text{place}]}$

b. Simplification: $V_i + V_i \rightarrow ØV_i$

c. Glide Formation: $V_{[+\text{high}]} + V \rightarrow GV$

The second part of the analysis is that the exceptional non-coalescing stems begin with an abstract consonant distinguishing them from regular stems. To motivate this, recall from §3.4 that there are no root-initial homorganic glide + vowel (*ji*, *wu*) sequences. Other root-initial glide + vowel sequences do exist, as shown in (174).

(174) Root-initial GV sequences

a.	jaga	'scratch (verb)'	e.	wowo	'them'
b.	jega	'shoulder'	f.	wawa	'wing'
c.	jonda	'baboon'	g.	wivi	'chameleon'
d.	juwe	'stone'	h.	weŋgu	'spleen'

As can be seen in (174), (what we assume to be) underived glides are not especially rare as syllable onsets, and appear both word-initially and word-medially. Moreover, derived glides and their exponents (in the form of secondary articulations) are also common. This makes the lack of *ji* and *wu* is quite striking.

Within my corpus, there are only two observed cases of surface *ji* or *wu*, both of which are marginal, optional, and root-internal: *m̩ɲawu* 'cat' and *mbajidi* 'hartebeest'. The general lack of *ji* and *wu* in Mushunguli has been corroborated in personal communication with Michal Temkin Martinez and with the online Chizigula-English dictionary developed by the Boise Language Project, which has a much larger lexicon (>14,000 entries) (Dayley et al. 2020). Similarly, an unpublished lexicon of the Tanzanian dialect of Chizigula collected by Charles Kisseberth and Farida Cassimjee has few if any instances of *ji* or *wu*, with the notable exception of *tʃꜜáːji* 'tea' (in the lexicon: <ch!á:yi>).[1] Neither 'cat' nor 'hartebeest' show up in either source transcribed with a glide (in the case of 'cat') or at all (in the case of 'hartebeest').[2] Moreover, some tokens of these words do not have a detectable glide, and none have as long of glides as unambiguous cases, suggesting that these may be phonetic transitions rather than phonological glides. Either way, the fact that a tiny number of exceptions exist does not contravene the fact that homorganic glide-vowel sequences *ji* and *wu* are generally phonotactically illicit in Mushunguli.[3]

Now recall from (88), repeated in (175), that exceptional stems behave as though they are vowel-initial in glide formation and simplification contexts but as though they

[1]Kisseberth and Cassimjee state that they follow the Kiswahili orthographic convention of omitting glide transitions between vowels, but actually use both conventions in the lexicon itself. This makes the actual status of intervocalic glides in Tanzanian Chizigula uncertain, but suggests that the case of 'tea' probably does actually have a glide in it. Regardless, there is no strong evidence that reflexes of underlying glides in *ji* and *wu* contexts obtain in *Somali* Chizigula, and my consultant produced this word as [tʃa.i].

[2]The Chizigula online dictionary does not contain an entry like *mbajidi*, and does not have an entry for 'hartebeest' at all; the Tanzanian Chizigula lexicon has *ŋkongoni* (in the lexicon: <nkhongoni>). My early transcriptions (as well as Odden's) have *mbajidi* glossed as 'rhinoceros' (which is *mpera*), later corrected to 'hartebeest'. Hartebeest are not native to Somalia, where my primary consultant grew up, so it's hard to tell if this is a lexical idiosyncrasy, a mistranslation, or a neologism.

[3]If we assume that deletion is the default repair for illicit consonants, then 'cat' and 'hartebeest' can both be accounted for by indexing the roots to MAX-C^{L2}.

are consonant-initial in coalescence contexts.

(175) Comparison of exceptional stems with regular V- and C-initial stems

	C-initial	**Exceptional V-initial**	**Regular V-initial**
	(-*lima* 'breathe')	(-*ita* 'go')	(-*iva* 'hear')
1sg. past (/si-/)	s**ili:**ma	s**i:**ta	s**i:**va
Infinitive (/ku-/)	k**uli:**ma	kʷ**i:**ta	kʷ**i:**va
3sg past (/ka-/)	k**ali:**ma	k**a.i:**ta	k**e:**va

A reasonable proposal given these observations is that the reason these stems sometimes behave as though they begin with a consonant is because they *do* begin with a consonant underlyingly; namely, that they begin with the homorganic glide + vowel sequences *ji* and *wu*, which are brought into line with the phonotactics of Mushunguli by a rule deleting the first component of a (stem-initial) homorganic glide + vowel sequence. That is, the exceptions listed in (89) have the abstract underlying representations given in (176).

(176) Abstract lexical representations of non-coalescing stems

a.	/jit/	'go'	i.	/wus/	'take out'
b.	/jini/	'liver'	j.	/wuj/	'come back'
c.	/jih/	'be bad'	k.	/wumb/	'mold'
d.	/jivu/	'ash'	l.	/wuz/	'ask'
e.	/jimb/	'sing'	m.	/wugul/	'lament'
f.	/jizi/	'voice'	n.	/wuguz/	'care for a sick person'
g.	/jir/	'cry'	o.	/wuŋg/	'want'
h.	/jidi/	'two'	p.	/wujus/	'revive'

The exceptional behavior of these stems can then be accounted for via serial rule ordering: homorganic glide deletion is ordered after coalescence but before glide formation and simplification. This allows the correct surface forms for both exceptional and regular stems to be generated, as shown in (177).

(177) Derivations of exceptional and regular stems

	Exceptional (/jit/)		
UR	/**si** + **ji**ta/	/k**u** + **jit** + a/	/k**a** + **jit** + a/
Coalescence	—	—	BLOCKED
HGD	s**ii**ta	k**ui**ta	k**ai**ta
Simplification	sita	—	—
Glide Formation	—	k**wi:**ta	—
Postlexical	si**:**ta	kʷ**i:**ta	k**ai:**ta
Surface	[si**:**ta]	[kʷ**i:**ta]	[k**ai:**ta]

	Regular (/iv/)		
UR	/**si** + **iv** + a/	/k**u** + **iv** + a/	/k**a** + **iv** + a/
Coalescence	—	—	k**e**va
HGD	—	—	—
Simplification	siva	—	—
Glide Formation	—	k**wi:**va	—
Postlexical	si**:**va	kʷ**i:**va	k**e:**va
Surface	[si**:**va]	[kʷ**i:**va]	[k**e:**va]

This analysis is elegant insofar as it effectively takes care of two problems (exceptions, and distributional restrictions on glide-vowel sequences) with one rule. It can also account for both of the key generalizations in (90). First, the fact that the exceptional stems all begin with high vowels is accounted for by the fact that only high vowels can be preceded by homorganic glides. Second, the fact that these stems are only exceptional to coalescence is accounted for by the ordering of coalescence before homorganic glide deletion, and the ordering of a homorganic glide deletion rule after coalescence but before glide formation and simplification is the key to the analysis. This analysis also has some advantages over other forms of analysis relying on abstract segments or specifications: it does not rely on unpronounceable segments, but rather unpronounceable sequences, and motivates an independently-required rule to account for why the sequences *ji* and *wu* do not appear on the surface in Mushunguli. That is to say, while this analysis does rely on whole segments, it is arguably more intuitive for a learner than something like the analysis of Maltese Arabic (Brame 1972), which relies on an abstract segment that does not actually exist in the surface phonetic inventory of the language.

However, there are several problems with this analysis. Ideal analyses in this vein either have historical support (e.g. Nupe (Hyman 1970); Inuit (Compton & Dresher 2011)) or independent synchronic behavior supporting the presence of the abstract segment, such as it surfacing variably (e.g. Kikamba (Roberts-Kohno 1998)). No consonant ever surfaces for the non-coalescing stems, so the latter option is out. Moreover, while a glide-initial analysis may have some historical support, a *homorganic* glide-initial analysis is much less well-supported.

To be clear, all available evidence does point to the source of this exceptional behavior being some kind of incomplete sound change, and the loss of some kind of initial consonant is as likely a reason as any other. While it is not possible to find all of the non-coalescing stems in the Tanzanian Chizigula lexicon (Kisseberth & Cassimjee n.d.), many are recorded there, and behave the same as their Mushunguli equivalents. Charles Kisseberth (p.c.) notes that he generally found the the "exceptional" pattern (= blocking coalescence) to be more common than coalescence in Tanzanian Chizigula. All of this indicates that these are probably native words, not loans, and suggests that this exceptional pattern stems from a sound change that predates the initial displacement into Somalia, rather than another source, such as influence of (non-Bantu) languages spoken in the region.

Investigation into the likely Proto-Bantu sources of V-initial stems suggests that the majority (if not all) originally began with consonants, especially **b* and **j*, the latter of which was probably either *z* or *j* (Odden 2015). It has been proposed that the loss of historical **j*, at least, is a pathway for vowel-initial words (Bostoen & Bastin 2016) This is supported by the reconstructions of regularly-behaving V-initial roots in Mushunguli, taken from the Bantu Lexical Reconstructions 3 database (BLR 3) (Bastin et al. 2002), as seen in (178).[4]

[4]The orthographic representations of the Proto-Bantu roots are taken directly from BLR 3, while the Mushunguli roots are represented as they have been throughout the dissertation. Note that

(178) Proto-Bantu reflexes of regular roots

a.	*jígu	-iv-	‘hear’
b.	*jíngɪd	-iŋgir-	‘enter’
c.	*jìkad	-ikal-	‘sit’
d.	*jàngat	-igat-	‘carry’
e.	*jínk	-ink-	‘give’
f.	*jít	-itaŋg-	‘call’[5]
g.	*jítɪk	-itik-	‘answer’
h.	*jóg	-oger-	‘swim’
i.	*jombʊk	-ombok-	‘cross, pass, go through’
j.	*jonk	-onk(ez)-	‘suck’
k.	*jongɪd	-ongez-	‘add to’
l.	*jàmb	-amb(iz)-	‘help’
m.	*bód	-ol-	‘be rotten’

There is also synchronic evidence of historical lenition from *b > w; for example, the noun class 2 prefix *wa-* is typically reconstructed as **ba-*) (and is indeed *ba-* in many other Bantu languages), and of a general dispreference for glide-initial roots. The Chizigula of Somalia dictionary (Dayley et al. 2020) contains roughly 5875 verbal entries (including derived forms).[6] 9% of these are vowel-initial, both exceptional and regular. Glide-initial entries make up <2% of all verbs (around 1% for *w*-initial and <1% for *j*-initial). The disproportionate number of vowel-initial roots further supports the hypothesis that lenition and loss of historical glides is at least one possible source of vowel-initial stems, exceptional or otherwise. Second, as was mentioned in §2.2, intervocalic glides optionally fail to surface in some cases. Given that root-initial glides typically occur in intervocalic contexts, especially for verbs, the presence of this synchronic pattern provides a possible

the BLR 3 database conflates IPA *j* and *ʑ* as <*j>, which represents an ongoing controversy regarding the historical status of these consonants. BLR 3 also distinguishes high vowels <*ɪ, *ʊ> from super-high vowels <*i, *u>. Finally, note that I have only included reconstructions categorized as belong to Guthrie's Bantu zone G, which is the origin point for Chizigula. Uncategorized reconstructions are not included, as well as any which I haven't been able to locate, of course.

[5] There are reconstructions of this word as <*jítan> in other Bantu zones.

[6] I consider here only verbs due to the organization of the dictionary at the time this research was conducted (late 2018). Nouns are cited including their class prefix; this makes it difficult to reliably diagnose V-initial status in most cases, as it requires cross comparison with singular and plural forms. The citation form of verbs is the infinitive, which means verbs regularly begin with one of <ku>, <ko>, or <kw>.

pathway for *how* the glides were lost.

Turning to the exceptional stems, a slightly different pattern emerges. Some examples are given in (179).

(179) Proto-Bantu reflexes of exceptional stems

a.	*jʊgʊd	-ugul-	‘groan; be ill’
b.	*bʊj	-uj-	‘come back’
c.	*bʊ́ʊdɪ	-uz-	‘ask’
d.	*bʊ́mb	-umb-	‘mould’
e.	*jɪmb	-imb-	‘sing’
f.	*bìdɪ́	-idi	‘two’
g.	*bii	-ih-	‘be bad’
h.	*jibu	ivu	‘ash’
i.	*jʊ́ì	izi	‘voice’
j.	*dàng	-uŋg-	‘like, desire’
k.	*pít	-it-	‘go’

First, the fact that reconstructions of most of the roots are recoverable further supports the notion that these are native words. However, it is worth noting that only five of the examples in (179) have been reconstructed with homorganic glide-vowel sequences (assuming **b* > w). Meanwhile, all but one of the reconstructions of the high-vowel initial roots from (178), at least, *have* been reconstructed with homorganic glide-vowel sequences. While it might be reasonable to hypothesize that there are synchronic underlying glides, it isn't clear that we should assume they are homorganic. If anything, the evidence suggests that **j* → Ø was the pathway that many V-initial roots took, while the non-coalescing stems both have a more mixed origin.

A second problem with this analysis is that it relies on an absolute neutralization of whole segments, which presents challenges with respect to learnability and theoretical expressiveness. An abstract representational analysis of this kind feels intuitive to a trained analyst, but requires a number of assumptions that it is unclear a learner would have. In this case, we are assuming that learners would *notice* the lack of *ji* and *wu*, and thus reconstruct the presence of the glides in response to cases where exceptions block coalescence. Moreover, it is difficult to determine how to formally restrict a theory that allows for the unrestricted positing of abstract segments–that is, there is no reason why the initial segments in this case *must* be homorganic glides; they could just as easily be any other segment with a limited distribution, or a segment that does not exist in the phonetic inventory of Mushunguli at all.

While we can make it an informal requirement that the analysis be well-motivated, it is also the case that learners would also need to develop alternative representational solutions, hopefully just as well-motivated, for the other exceptions in the language. While we could posit an abstract segment to block hiatus resolution for the total non-participating exceptions discussed in §5.2, it will be much more difficult to motivate what that segment should be. This is especially true given that unlike the non-coalescing stems, the TNPs seem to be a mixed bag of native- and non-native forms, with no substantive shared phonological characteristics. Some, like *ona* 'see,' are Bantu (PB: **-bon-*), and their exceptional behavior might share the same pathway **b* > w as other exceptional stems. Others, like *iɟariʃa* 'rent,' are loans (from Maay *ijaar*). Thus, while we can posit an abstract consonant that survives all the way to the end of the derivation, the form of that consonant will have to be entirely *ad hoc*. For the walljumping palatalization cases (which have no easily identifiable historical source at all), there obviously is no way to posit a whole segment that will prevent the derived glide from being deleted.[7] It is unsurprising that the only attempt to analyze this pattern with rules ultimately relied on a rule diacritic, though in this case, one that made reference to syntactic position (Hout 2012).

[7]My best attempts at sourcing these exceptions have resulted in a tentative conclusion that they stem from a shift in permissiveness of CG clusters (or palatalization) that occurred *after* the displacement into Somalia. The transcriptions in the Tanzanian Chizigula lexicon (Kisseberth & Cassimjee n.d.) indicate that [Cj] sequences derived by glide formation are not as uniformly prohibited as in Mushunguli. For example, compare [**ch**-a-lagá:la] '[cl.7] is falling' with [**vy**-a-lagá:la] '[cl.8] are falling' (cf. Mushunguli ***tʃa:****gʷa* and ***va:****gʷa*). Class 5 agreement prefixes are typically transcribed as <dy> prevocalically but <di> otherwise, e.g [**dy**-a-lagá:la] '[cl.5] is falling', [di-lagâ:la] '[cl.5] fell' (cf. Mushunguli ***da:****gʷa* and ***di:****gʷa*); this is also true for demonstratives, e.g. *bulangeti* ***dy****-angu kulu* 'my big blanket'. A similar pattern is observed for 'eat', e.g. [w-a:**dy**a] 'you are eating', [ku**di̱í:**la] 'eat with'. It seems that the avoidance of [Cj] became regularized, even for class 5, save for the class 5 demonstrative prefix and 'eat.'

A second issue with this analysis is that it fails with respect to the stated aim of capturing, explaining, and unifying the critical generalizations from (92). This rule ordering is *ad hoc*, unmotivated by any independent considerations. For example, if the order of glide formation and coalescence were reversed, it would be possible to describe a language exactly like Mushunguli where a set of stems beginning with high vowels exceptionally fail to trigger glide formation, but not coalescence. If the order of simplification and coalescence were reversed, we would predict a set of exceptional stems that only resist simplification, despite the fact that simplification has no material effect on the stem vowel. In the latter case, if the position of glide formation is not also changed, we end up predicting an (in this case) unwanted Duke of York derivation for /(C)u + wu/ contexts: simplification fails to occur and homorganic glide deletion applies, creating an intermediate form |u + u|. This form is now subject to later-applying glide formation, resulting in *[wu]–exactly the sequence we were trying to avoid in the first place. Having the HGD rule thus does not even guarantee that we will actually avoid the sequences it was intended to eliminate. Contrast this with the lexical indexation analysis, in which blocking of high coalescence alone is guaranteed by the fact that it is the only hiatus resolution strategy in the language that materially affects stem vowels. This is the fundamental difference between these two analyses: under the derivational analysis, the critical generalizations are captured, but under indexation, the critical generalizations are both captured and *guaranteed*.

A third issue is that under this analysis, we have lost all dependencies between the exceptional patterns; this is true independent of the *ad hoc* rule ordering issue. Because we can appeal to any phonological structure (motivated or otherwise) to account for these patterns, and because we can write our rules to be wholly descriptive, there are no restrictions placed on the system by introducing any form of exception feature. Under the lexical indexation analysis, the elimination of vowel deletion as an option has ramifications for the analysis of both exceptional and regular forms, and all exceptional patterns concord with these predictions. The rule-based derivational analysis places *no restrictions* on the forms of exceptions within this system, whereas the lexical indexation analysis situates them within the hiatus resolution and onset structure conspiracies.

Finally, it should be noted that there is no principled way to strengthen this particular analysis, e.g. by motivating the necessary rule ordering by demonstrating that e.g. coalescence is a stem- or word-level rule while homorganic glide deletion, simplification, and glide formation are all word-level or postlexical. Hiatus resolution applies at all lexical levels, but only coalescence can apply postlexically, e.g. /ɲ + simba i + no/ → [simba iːno]~[simbeːno] 'this lion'). Crucially, glide formation cannot apply postlexically, e.g. /tʃi-tuŋgulu i-tʃi-o/ → [tʃi.tuŋ.gu.lu.iː.tʃo], *[tʃi.tuŋ.gu.lwiː.tʃo] 'this onion'. As such, it would be misguided to assign coalescence to the lexical level and glide formation to the postlexical level for the purposes of motivating their different orders. The lexical indexation analysis, on the other hand, accounts for our key generalizations and their interconnectedness without making any further assumptions; where levels are adopted, they are done so based on a combination of both phonological and morphological evidence. This problem will become more serious in the following subsection.

6.1.2 Absolute neutralization in Stratal OT

We have discarded the rule-based abstract representational solution on the grounds that it fails to unify generalizations about the non-coalescing stems, and because it fails to capture apparent dependencies between the sets of exceptions. However, some of these problems are intrinsic not to abstract representations, but to the adoption of *ad hoc* rule ordering. Stratal OT reintroduces the possibility of derivation as a formal device for capturing exceptionality (and we have used it for exactly this purpose; the effects of indexed MAX-C are not felt at the lowest level, but rather the intermediate one), even if the levels are independently-motivated. As such, it would be prudent to examine whether we have accidentally reintroduced absolute neutralization as a viable strategy to capture the non-coalescing stems.

The answer is an uncomplicated ‘no’, at least for Mushunguli, and for this type of exception. It is the case that one can, in principle, rely on neutralization of an abstract segment to block an alternation using lexical strata, *if and only if* there is no evidence that alternation applies at any higher level. For example, imagine a language, just like “Sanuma” (from §4.1.3), that fails to delete the vowel of an exceptional prefix. The example there relied on an indexed form of MAX; however, it would be just as possible to assume that the prefix ends in an illicit consonant, which we will arbitrarily decide is **ʔ**. If this is the case, then any deletion candidate will be less harmonic than one that is fully faithful, as seen in (180). Note that in this tableau, we will also assume a constraint *ʔ penalizing this segment, which is crucially dominated by MAX.

(180) Exceptional blocking of deletion (word level)

/biʔ + adi/	ONSET	MAX	***ʔ**
a. ☞ biʔadi			*
b. biØ.adi	*!	*	

If we assume that there is a level at which the abstract consonant is neutralized (which is naively assumed in all cases of absolute neutralization), then provided that deletion cannot occur at this level (i.e. re-ranking MAX so that it dominates ONSET), the consonant can be eliminated safely by also promoting ***ʔ** above MAX, as illustrated in (181).

(181) Exceptional blocking of deletion (postlexical)

|biʔadi|	***ʔ**	MAX	ONSET
a. [biʔadi]	*!		
b. ☞ [biØadi]		*	*

Thus, we see that it is in principle possible to use an abstract segment to block an alternation in Stratal OT. However, it is not actually possible to reinstate the abstract glide solution in Stratal OT for Mushunguli, by, for example, relying on $*G_iV_i$ to play the role of HGD in the ordered rule analysis.[8]

This analysis is precluded because the previously-discussed mismatch between alternations and levels. We have evidence that coalescence *can* occur postlexically (*m̩bʷa iːno∼mbʷeːno* 'this dog'), while glide formation (and any concomitant processes) cannot (*buku iːɟo*, *bukʷiːɟo 'that (prox) book'); more generally, coalescence can apply at every level glide formation can, but not vice versa. The non-coalescing exceptions block high coalescence (of course), but trigger all relevant forms of glide formation. Thus, order to allow the cases where hiatus is repaired to surface correctly, we will need to make the following assumptions: (a) glide formation applies at the stem and word level, but not the postlexical level; (b) the underlying glide is neutralized *before* the postlexical level (so, the word level); and (c) coalescence does *not* apply at the word level. This will allow glides to be formed and then repaired postlexically. However, this also means that at the postlexical level, we now predict the non-coalescing stems to exhibit variable word-internal coalescence, which, of course, they do not.

This is easier to see if we remove glide formation from consideration and just allow the abstract glide to reach the postlexical level. As long as coalescence is a valid option at this level (variable or otherwise), the ranking ONSET » IDENT(high) must obtain. Thus, any ranking that neutralizes the glide will *still* prefer a coalescence candidate. This is illustrated by the tableau in (182), where I use the aforementioned $*C_iG_i$ to drive the

[8]Note that in the current analysis, $*G_iV_i$ is undominated, which prevents simplification candidates from selecting *|ji|. This would have to be adjusted for this alternative; $*G_iV_i$ would need to be dominated by MAX-C, as we do not want to accidentally delete our abstract glides. Because MAX-C is either crucially or transitively dominated by the majority of other constraints in the word-level grammar, the absolute neutralization alternative being sketched here would assume that simplification of high vowel sequences results in glide formation, which is presumably repaired at a later level.

deletion of the abstract glide. Note that for the second candidate, the <Ø> representing the deleted glide is arbitrarily placed due to the convention of representing coalesced vowels as <V$\text{ː}_{1,2}$> as opposed to <V_1V_2>.[9]

(182) Accidental postlexical coalescence

|ka$_1$-ji$_2$t-a|	*G_iV_i	MAX-C	ONSET	IDENT(high)
a. ka$_1$.jiː$_2$.ta	*!			
b. ☞ keː$_{1,2}$Ø.ta		*		*
c. ka$_1$.Øiː$_2$ta		*	*!	

It is the case that some proposed versions of Stratal OT make an explicit assumption of the Strict Cycle Condition (Mascaró 1976; Kiparsky 1985; Kaisse & Shaw 1985), which specifies that structures derived by lexical rules (or, in this case, outputs selected by lexical rankings) cannot be operated on at later lexical levels. However, this still won't allow us to reintroduce absolute neutralization for Mushunguli. First, the Strict Cycle Condition only applies to lexical operations; postlexical ones are generally assumed to be able to apply across the board; as such, there is again no reason why *postlexical* coalescence should not apply once the glide is neutralized. Second, the Stratal OT analysis of Mushunguli crucially relies on the ability for operations to apply to derived structures–otherwise, we have no avenue by which we can account for the regular patterns of glide formation in postconsonantal contexts. Finally, this would seem to eliminate the possibility of neutralizing an abstract segment at all. This form of analysis assumes that these segments belong to the non-coalescing stems; however, remember that glide formation and identity coalescence apply to the non-coalescing stems, but the former at least does not apply post-lexically. Thus, there needs to be a lexical level where the segment is neutralized, and so neutralizing the segment itself means acting upon the derived stem.

[9]Note also here that this is illustrative, and should not be read as a formal proposal for handling variable postlexical coalescence in Mushunguli; an account of the postlexical phonology is an important next step for the description and analysis of Mushunguli, but is outside the scope of this dissertation.

Assuming strict cyclicity will also be problematic for the walljumping exceptions, which rely on derivation to account for their irregular behavior (and whose behavior supports a derivational analysis). It is slightly less problematic for the TNP exceptions, but not wholly; we will still need to eliminate the abstract consonant at some level above the stem, or else we again expect coalescence to variably apply to them. For absolute neutralization of whole segments to apply, then, we must assume that strict cyclicity cannot hold. If this is the case, then absolute neutralization can apply, but only in a very restricted set of cases. It is necessary that if an operation is blocked by an abstract segment, it cannot apply at the level that segment is neutralized or any level thereafter; elsewise, the protection of the segment is neutralized. This is similarly true for an exception that only blocks one of multiple dependent operations (e.g. a case like Mushunguli); if this occurs, then again, the blocked operation has to occur at a level earlier than the segment's neutralization, and crucially cannot apply at any level thereafter.[10]

6.2 Underspecification

We have seen that absolute neutralization of whole segments is not a good avenue for capturing synchronic exceptions in Mushunguli. A more comparable competitor to indexed constraints is some form of underspecification. For this purpose, we will adopt prespecification (Inkelas et al. 1997).

Prespecification is a form of contrastive underspecification that relies on a principled formal mechanism of assigning feature values proposed by Inkelas (1995) (which she refers to as Archiphonemic Underspecification). Prespecification assumes that regularly-behaving forms that predictably alternate are unspecified for predictable structure, while

[10]This is not unlike the behavior of indexed *COMPLEX^{E} from the end of §5.1.3.2; essentially, the introduction of any kind of exception feature into lexical levels raises the possibility of the feature's protection being "undone."

exceptions (or at least, exceptions that do not alternate), as well as predictable non-alternating structure, are fully specified. This means that constraints that target the relevant feature can only 'see' (and be violated by) specified segments.[11]

Prespecification relies on the principles of Lexicon Optimization (Prince & Smolensky 1993/2004) and Grammar Optimization (Kiparsky 1993) to assign feature values. Lexicon Optimization requires that, given a surface phonetic instantiation of some morpheme, and the set of all hypothetical underlying representations that could in principle lead to that surface form, the true underlying representation should be the one that incurs the fewest violations of highly-ranked constraints. Inkelas (1995) adjusts this to rely on alternations to determine feature specifications, rather than universal principles of markedness. Because of this, prespecification allows ternarity (i.e. individual segments can be marked for any of $[\pm f]$ and $[0f]$).

Grammar Optimization requires that optimal grammars are "transparent," i.e. maximally structure-filling, which means that feature deletion is penalized over feature insertion. To capture this, Inkelas (1995); Inkelas et al. (1997) adopt the theory of markedness proposed by Kiparsky (1994), which assumes that every constraint that refers to a feature value is split into a general version and a specific version, with the latter crucially only ever being able to refer to the marked value of that feature. FILL (or DEP) constraints are technically subject to the same restriction: DEP$[f]$ » DEP$[\pm f]$ will presumably fill in the unmarked value, while the reverse ranking will fill in only the marked value. However, neither Inkelas (1995) nor Inkelas et al. (1997) parameterize feature insertion, and the default assumption seems to be that unmarked values are filled in. Neither Inkelas (1995) nor Inkelas et al. (1997) substantively sketch out what an analysis of a language would

[11] It is worth pointing out that while there are important differences in the representational assumptions used by the proposed lexical indexation analysis and prespecification, their implementations in OT (i.e. the way that constraints are evaluated) in this case are equivalent: constraints looking for the specified feature cannot 'see' segments unspecified for that feature, just as lexically-indexed constraints cannot 'see' violations incurred entirely by morphemes lacking an index.

look like (i.e. determination of underlying representations *and* exploring the behaviors of those underlying forms in other contexts), instead primarily focusing on demonstrating the operation of Lexicon Optimization on selecting underlying forms.

To truly evaluate the predictions of a prespecification analysis against the current analysis of Mushunguli would require it to be rebuilt from the ground up, as they rely on fundamentally different representational assumptions. However, we can adopt some of these assumptions to guide us to a skeletal analysis of the underlying representations of the non-coalescing stems. Our previous representational assumptions were that all vowels in Mushunguli are fully specified for height and place features; the main exception being the low vowel, which is placeless. Now, we will revise this assumption to so that featural representations are dependent on alternation. As applied to coalescence, this cashes out to the following representations: both low prefix vowels and high root-initial vowels alternate with mid vowels; specifically, the height feature ([-high]) is donated by the initial low vowel, while the place feature is donated by a following high vowel. An alternation from a low to a mid vowel changes the [±low] feature, therefore prefix low vowels that participate in coalescence (/A/) must be [0low]; similarly, high-mid alternations change [±high], therefore root-initial high vowels participating in coalescence (/I,U/) must be [0high]. There are no alternations from a low vowel to a high vowel or vice versa; to capture this, we assume that they are specified as [-high] and [-low], respectively. Finally, underlying root-initial mid vowels do not alternate for any feature; thus, they are fully specified as [-high, -low]. These specifications are summarized in Table 6.1.

The illustrations in Inkelas 1995; Inkelas et al. 1997 generally rely on PARSE(*f*) and FILL constraints to follow Kiparsky's model. I will replace these with equivalent MAX(*f*) and DEP[*f*] constraints. The relevant constraints are given as follows:

(183) MAX[high]: any specified value for [αhigh] in the input should have a correspondent in the output

Table 6.1: Prespecification of Mushunguli vowels in coalescence contexts

Vowel	Position(s)	Specification
/A/	Any Prefix	[-high, 0low]
/I, U/	/V-/, $Root_1$ (Coalescing)	[0high, -low]
/i, u/	$Root_1$ (Non-coalescing)	[+high, -low]
/e, o/	Root	[-high, -low]

(184) MAX[-high]: any [-high] specification in the input should have a correspondent in the output

(185) DEP[*f*]: any value for [α*f*] in the output should have a correspondent in the input

Coalescence of all sequences is guaranteed provided that any MAX constraints covering specified features are dominated by ONSET, and in turn dominate DEP[*f*].

(186) Prespecification of regular high coalescence

	$A_P + I_S$	ONS	MAX[-hi]	UNIF	DEP[*f*]
a.	$a_P.i_S$	*!			***
b. ☞	$e_{P,S}$			*	
c.	$Ø_P i_S$		*!		*

In (186), the coalescence candidate beats faithful hiatus due to the fact that ONSET dominates DEP[*f*]. Because deletion of the initial low vowel, which is specified as [-high], would involve deleting that specification, deletion also loses due to the fact that MAX[-hi] dominates UNIFORMITY (and DEP[*f*]).

We can block high coalescence in the case of exceptional stems by ranking MAX[high] above ONSET.

(187) Prespecification blocks coalescence

$A_P + i_S$	MAX[hi]	ONSET	MAX[-hi]	UNIFORMITY	DEP[f]
a. ☞ $a_P.i_S$		*			**
b. $e_{P,S}$	*!			*	
c. $Ø_P i_S$	*!		*		

Here, coalescence would involve deleting the [+high] specification (in favor of the [-high] specification of the low vowel); thus it is blocked. Similarly, deletion of the initial low vowel would violate MAX[high] and MAX[-high]; thus, it too is blocked.

This analysis will still require us to assume that coalescence is favored over deletion. This is because /A+e/ → [eː] involves a redundant [-high] feature. If we assume that redundant features are deleted, then we incorrectly predict that /A+e/ will also result in unresolved hiatus, as in (188).

(188) Feature deletion blocks mid coalescence

$A_P + e_S$	MAX[hi]	ONSET	MAX[-hi]	UNIFORMITY	DEP[f]
a. ☞ $a_P.e_S$		*			**
b. $e_{P,S}$	*!		*	*	
c. $Ø_P e_S$	*!		*		

The same is true for identity coalescence, which for /A/ will involve the merger of a segment that is [-high, 0low, 0back] and [-high, -low, 0back].

(189) Feature deletion blocks identity coalescence

$A_P + a_S$	MAX[hi]	ONSET	MAX[-hi]	UNIFORMITY	DEP[f]
a. ☞ $a_P.a_S$		*			**
b. $aː_{P,S}$	*!		*	*	
c. $Ø_P a_S$	*!		*		

Assuming an undominated MAX-V will not help in this case; if redundant features are deleted, then any form of coalescence besides high coalescence will violate top-ranked MAX[high]. We can solve this by assuming that features are part of the set of elements covered by the definition of UNIFORMITY (46); if this is the case, then redundant features merging is simply a second UNIFORMITY violation.[12] We will make this assumption going forward.

(190) Feature merger guarantees mid coalescence

	$A_P + e_S$	MAX[hi]	ONSET	MAX[-hi]	UNIFORMITY	DEP[f]
a.	$a_P.e_S$		*!			**
b. ☞	$e_{P,S}$				**	
c.	$Ø_P e_S$	*!		*!		

(191) Feature merger guarantees identity coalescence

	$A_P + a_S$	MAX[hi]	ONSET	MAX[-hi]	UNIFORMITY	DEP[f]
a. ☞	$a_P.a_S$		*!			**
b.	$a_{P,S}$				**	
c.	$Ø_P a_S$	*!		*		

Prespecification with the constraint set proposed by (Inkelas 1995) and assuming (identical) feature merger can capture the two critical generalizations. Here, the specified feature is [+high] (which satisfies the "only high vowels" generalization), and high coalescence is still the only hiatus resolution process in the language in which a MAX[high] violation incurred by the stem vowel could come into play (as all MAX[high] violations for deletion candidates in other coalescence contexts are violated by deleting the prefix vowel's [-high] specification).

However, a fully fleshed-out analysis is going to need to account for a couple of issues. One is that our current analysis of simplification relies on the assumption that a long vowel can only obtain on the surface if its moras came from featurally non-identical

[12] For this to work, we will also need to assume that non-identical features cannot merge; otherwise there is an available interpretation of high coalescence where the merger of [+high] to [-high] is not penalized, and thus the non-coalescing stems cannot actually be blocked.

sources (in satisfaction of *MATCH-μ). This analysis cannot be easily translated over to the prespecification account, as illustrated by the data in (192).

(192) Simplification examples

a.	ihe:ma	/I+hem+a/	'it (cl 9) breathed'	
b.	ne:ni:ɟe	/nA+I+ni+di+e/	'it (cl 9) will eat me'	<neeníje>[D]
c.	ita:ŋga	/I+Itaŋga/	'it (cl 9) called'	
d.	siti:sa	/si+it+is+a/	'I went a lot'	
e.	sitaŋga	/si+Itaŋg+a/	'I called'	

Examples (192a-c) feature the class 9 subject prefix, which under Lexicon Optimization would have to be /I-/, or [0high, -low, -back]. The first example is before a C-initial stem, which illustrates that the prefix is in fact a high front vowel. The second example is a complex construction that demonstrates that this prefix can participate in coalescence; thus, it needs to have the same feature specification as initial root vowels. The specification of the /A/ in the non-past prefix in this case must similarly be [-high, 0low, 0back], because it exhibits the same alternation patterns as other /A-/ prefixes (e.g. ***nani:ɟe*** 'I will eat').

Notice here that compensatory lengthening applies; this means that antagonism between [0] and [-] has to be sufficient to constitute "non-identical" for the purposes of satisfying *MATCH-μ. We were already making that assumption implicitly to account for penult lengthening, which inserts a mora that is associated with no features, so this concords with the proposed lexical indexation analysis of Mushunguli.

(193)0 /[-] antagonism sufficient for compensatory lengthening

nA_P+I_S...	ONS	*MATCH-μ	MAX-μ	*V:
a. $na_P.i_S$	*!			
b. ☞ $ne\text{:}_{P,S}$				*
c. $ne_{P,S}$			*!	

The third example from (192) is the class 9 prefix concatenated with a regular coalescing stem. This too comes out as predicted by the analysis; the initial vowel of a regular coalescing stem is also [0high, -low, -back]. All features match, so a long vowel here would violate *MATCH-μ; thus, simplification results in a short vowel. So far, so good.

(194) Total matching drives mora pruning

$I_P + I_S$...	ONS	*MATCH-μ	MAX-μ	*V:
a. $i_P.i_S$	*!			
b. $iː_{P,S}$		*!		*
c. ☞ $i_{P,S}$			*	

Examples (192d-e) are problematic, however. The vowel of /si-/ does not alternate for any features, so we will assume that it is fully specified.[13] Example (d) is this prefix concatenated with a non-coalescing stem, which as we know, participates in simplification. Again, provided that we assume the vowel of /si-/ is fully specified, a short vowel is exactly what we predict here. The problem is example (e); here, the two vowels have disagreeing values for [high], with the first vowel being [+high] and the second being [0high]. If (192b) satisfies *MATCH-μ, then presumably (e) must as well, yet, it does not.

(195)0 /[-] mismatch incorrectly predicts long vowels

$si_P + I_S$...	ONS	*MATCH-μ	MAX-μ	*V:
a. $si_P.i_S$	*!			
b. ☞ $siː_{P,S}$				*
c. ☹ $si_{P,S}$			*!	

[13]Technically, /Ci-/ prefixes alternate with [Ø]. However, we cannot assume this is the underlying specification of these *i* vowels, as it would require paramaterizing DEP to ensure that the ranking fills in unmarked values for [±high] and [±back], but marked feature values for [±low]. This will minimally conflict with filling in *unmarked* values for unspecified low vowels, if not also partially unspecified back round vowels. Moreover, assuming the vowel in *si-* is wholly unspecified, or specified in such a way that makes it distinct from both *ita* and *Iva* will only make the problem we are about to encounter worse.

Any equivalent analysis to the one proposed in §3.4 will need to account for this. Note that under Lexicon Optimization, it will not be possible to do this by making a principled assumption that any of these vowels have an underlying difference in their moraic specifications. It is the case that high vowels alternate with moraic and non-moraic forms ([i]~[j] and [u]~[w]), so, in principle, one could assume that they are underlying glides that are vocalized when they occur between consonants. However, an analysis of prefix vowels as anything besides moraic is precluded by the fact that there *are* underlying glides that do not trigger compensatory lengthening when they are syllabified as onsets (consistent with the fact that they have no moras to spread, i.e. they are *glides*). Both /CV-/ and /V-/ prefixes trigger compensatory lengthening in prevocalic contexts (e.g. *saːsaːma* 'I gaped'; *jaːsaːma* 'it (cl 9) gaped'). Even if we assume a low-ranked DEP-μ, it is difficult to imagine what markedness constraint would be satisfied by inserting a mora into an adjacent syllable when an underlying representation is unchanged. Similarly, we cannot distinguish vowels in roots this way; the initial vowels of both the exceptional and regular roots in this language predictably undergo (or fail to undergo) compensatory lengthening in the exact same contexts (e.g. *kʷiːviːsa* 'to hear a lot'; *kʷiːtiːsa* 'to go a lot', and examples (d-e) above), and similarly undergo penult lengthening in the same contexts (e.g. *kʷiːva* 'to hear'; *kʷiːta* 'to go'). That these vowels never alternate with glides means that under prespecification, they must be underlyingly moraic.[14]

A second issue is that it will be necessary to come up with a compelling account of the behavior of the TNP exceptions, and it is unclear exactly what that analysis could be. The paradigm from (168) is repeated in (196).[15]

[14]Inkelas (1995) notes that assuming full specification of moras even in a language with predictable vowel length does in fact conform to Lexicon Optimization.

[15]The underlying representation of example (k) under the assumptions of the prior analysis is /i + a + iɟariʃ^{L3} + w^{L2} + a/ (or /i + a + iɟariʃ^{L3} + u + a/; the lexical indexation analysis allows for the passive suffix to either be a consonant protected by MAX-C or a vowel that undergoes hiatus resolution.) The BLP dictionary does not mark compensatory or penultimate lengthening in entries, so example (k) here is assumed to exhibit compensatory lengthening.

(196) Total non-participants

a.	oːna	'see!'	
b.	ku.oːna	'to see'	
c.	si.oːna	'I saw'	
d.	i.oːna	'it (cl 9) saw'	
e.	u.oːna	'it (cl 3) saw'	
f.	ka.oːna	'(s)he saw'	
g.	naː.oːna	'I am seeing'	
h.	k^wiːoːna	'to see it/them (cl 9/cl 4)'	
i.	ku.oneːsa	'to see a lot/too much'	
j.	ku.iɟariːʃa	'to rent'	<kuidyarisha>[B]
k.	jaː.iɟariːʃʷa	'CL9 is being rented'	<yaidyarishwa>[B]

The Mushunguli TNP exceptions are problematic for a view of exceptionality that relies on representational prespecification. Recall from §6.1.2 that absolute neutralization is not a viable option in Mushunguli due to the fact that postlexical coalescence applies; this is true here, as well. It is actually doubly true: whole segment absolute neutralization seems to be impossible under prespecification, due to the fact that this theory prefers insertion grammars to deletion ones. Even if we assume that there is a root-initial consonant specified as [0features], the same grammar that inserts unmarked features on segments that do surface will presumably need to insert those features on a wholly underspecified one. The alternative is also non-viable; if we assume that there is a fully specified but illicit segment present in the underlying representations, neutralizing that segment to a pronounceable one will violate fewer MAX constraints than deleting it altogether. This means that we have two possible options for this pattern: either there is a featural specification that all of these vowels share, or there is a prosodic (i.e. moraic) specification that blocks the application or triggering of any hiatus resolution process.

The former option is out; as discussed in §5.2, most TNPs begin with mid or low vowels, and all mid- and low-vowel-initial exceptions to hiatus resolution are TNPs. However, mid vowels in this language are already fully specified; they never alternate for place or height features, and even in coalescence, mid vowels are vacuous participants. Furthermore, example (196i), *ku.oneːsa*, indicates that *ona* has an underlying [-high] feature to spread; again, suggesting that it is featurally identical to regular roots. Finally, as evidenced by the examples in (196), which begin with *high* vowels, vowel height (and quality) is apparently immaterial for these exceptions, anyway.

The one option we have for representational specifications is to stipulate that these vowels have an underlying vowel feature that is otherwise never attested on the surface, such as [±ATR]. However, this goes against the principles that prespecification was built upon, which is to avoid opportunistic uses of underspecification (Steriade 1995). There is nothing measurably different about the vowel quality of *ona* than the quality of e.g. *omala*, and Mushunguli does not have a surface ATR contrast. Moreover, while I am unaware of any languages where ATR specifications block the application of glide formation, it is imaginable (e.g. a language which penalizes [-ATR] or [+ATR] glides). However, I am also unaware of any languages where the *featural* specifications of glide formation *triggers* is sufficient to block hiatus resolution *entirely*. There are, of course, languages that block the application of glide formation in favor of deletion–this is (descriptively) true for Mushunguli, and a more robust example comes from Gichode, where round vowels (bare or postconsonantal) cannot glide before other round vowels (Casali 1996). However, Gichode instead exhibits vowel deletion in this context; it does not block hiatus resolution entirely. This is why the analysis of the TNP exceptions relies on ALIGN–the generalization to be made here has nothing to do with featural representations, but rather strong edges.

To illustrate this further, in the lexical indexation analysis of Mushunguli, glide formation is guaranteed except for when it could potentially create a marked segment (such as a non-high glide) or sequence (such as a homorganic glide-vowel sequence). In the prespecification analysis of Mushunguli, glide formation of e.g. /I+a/ is similarly guaranteed, only violating DEP[*f*] (to insert unmarked [+high]). This is illustrated in (197); note that I add MAX[-low] above DEP[*f*] to block deletion of the prefix.[16]

[16]Under the prespecification analysis, at least as it is currently formulated, the existence of the non-coalescing stems does not guarantee the impossibility of vowel deletion, only the impos-

(197) Glide formation under prespecification (bare vowel contexts)

	$I_P + a_S$	ONSET	MAX[low]	UNIFORMITY	DEP[f]
a.	$i_P.a_S$	*!			*
b.	$e_{P,S}$		*!	*	
c.	$Ø_P a_S$		*!		
d. ☞	$j_P a_S$				*

With this being the case, it truly does seem to be impossible to determine what *featural* specification could possibly be shared by vowels of any quality that could uniformly block hiatus resolution.

The second option is moraic prespecification; we can, for example, assume that these vowels are underlyingly long, and then stipulate that long vowels neither trigger nor undergo hiatus resolution processes. However, this suffers from a similar problem: while vowel length alternations exist in Mushunguli, they are all predictable. The behavior of the initial vowels of the TNPs conforms entirely with the behavior of other vowels in e.g. C-initial roots in the language; they are short outside of the penultimate syllable, and long within it. This behavior also conforms with that of the initial vowels of the non-coalescing exceptions–they are short outside of the penultimate syllable if they are blocking hiatus resolution, and long otherwise.

However, given that surface long vowels *do* obtain in Mushunguli, this is at less stipulative than the previous option. Note, however, that the formalization of this alternation we will have to be more substantive than to simply stipulate that having more than two moras involved in a hiatus context will block application of hiatus resolution; triplicate sequences of vowels (including those with no identical sequences) do resolve hiatus and do result in long vowels. For example, the form /ku + a + u + mu + ti/ 'to the tree' surfaces as *koːm̩ti*. I have not substantively discussed three-vowel sequences in this

sibility of deleting vowels with an underlying specification for [±high]. Whether this is good, bad, or neutral depends on the prespecification analysis of onset structure, post-consonantal glide formation, and walljumping palatalization; all three of which I leave to future research.

dissertation, as I have relatively few examples that do not contain an underlying sequence of identical vowels, and/or which occur outside of phrase-final contexts. However, for triplicates involving identical vowel sequences (e.g. /a + a + itaŋg + a/ → [eːtaːŋga] '(s)he is calling'), mora pruning to a single long vowel is guaranteed due to the activity of *MATCH-μ. For non-identical sequences, my assumption is simply that there is a markedness constraint penalizing super-long vowels (i.e. trimoraic or greater) that dominates MAX-μ; this would allow for moras to be pruned in this context. However, we do not want this constraint to actually block hiatus resolution from applying when a super-long vowel could in principle be made.

Regardless of these issues, prespecification certainly is at the very least a superior representational option to absolute neutralization in a rule-based framework or in Optimality Theory. Provided that these issues (and any others not yet identified) can be accounted for satisfactorily and in a non-stipulative way, prespecification may be a valid alternative to lexical indexation for Mushunguli exceptionality. If this proves to be the case, then the choice between a (wholly) representational and a diacritic solution will almost certainly have to rely on evidence from outside of formalism, such as learnability.

6.3 Summary

This chapter has looked at three alternative representational approaches to the patterns of exceptionality in Mushunguli, with special emphasis on the non-coalescing stems. Two of these involved absolute neutralization of underlying abstract segments, one of which was based in rules and relied on rule ordering (§6.1.1), and the other of which was based in lexical strata in OT (6.1.2). The rule-based analysis was able to capture the critical generalizations from (90), but failed to unify them. The Stratal OT analysis was demonstrated to simply be unworkable, due to the presence of both lexical and post-lexical coalescence.

In §6.2, we then looked at an analysis based in prespecification, a principled form of underspecification. This analysis fared better than either absolute neutralization analysis–it was capable of capturing the unified generalization about the non-coalescing stems, albeit with some stipulation regarding the handling of redundant feature values. However, we noted two avenues of concern for an analysis in this vein that would need to be addressed in a fuller discussion. The first was that the proposed analysis of simplification sans compensatory lengthening, which assumes that mora pruning is a response to having moras that are associated with identical feature content, is not workable for an analysis that crucially relies on phonetically identical segments to have distinct underlying representations. The second was that the behavior of the total non-participants, exemplified by *ona*, are not analyzable via appeal to featural representation. An alternative account relying on moraic prespecification has the potential to fare better, but will need to be carefully crafted so as to conform to the principled restrictions imposed by prespecification, and so as to not undermine the resolution of triplicate vowel sequences in the language.

This chapter, in part, contains material that has previously appeared in San Diego Linguistics Papers 6. Hout, Katherine, 2016. It also contains material that has previously appeared in *Africa's Endangered Languages: Documentary and Theoretical Approaches*. Hout, Katherine, Oxford University Press, 2017. This chapter, in part also contains coauthored material currently being prepared for submission for publication of the material. Hout, Katherine; Baković, Eric. The dissertation author was the primary investigator and author of all of this material.

7 Conclusion

7.1 Summary of the Analysis of Mushunguli

This dissertation presented an in-depth description and analysis of the hiatus resolution and onset structure conspiracies in Mushunguli, an endangered, under-described Bantu language. As discussed in §2.2.4, word-internal hiatus is generally not tolerated in this language, and Mushunguli exhibits unusually diverse patterns of hiatus resolution in the same morphosyntactic contexts. Descriptively, there are four repairs with eight surface instantiations: glide formation, coalescence, elision, and secondary articulation (summarized in Table 7.1).

Table 7.1: Summary of (regular) hiatus resolution results (descriptive)

Sequence		Descriptive Repair
Low + High	a + i → eː a + u → oː	Coalescence
Low + Mid	a + e → eː a + o → oː	Elision
$V_i + V_i$	a + a → a i + i → i u + u → u	Elision
High + V	i + V → jVː u + V → wVː	Glide Formation
Cu + $V_{[\text{-round}]}$	Cu + V → C^wVː C_{lab}u + V → $C^ɣV$ː	Labialization Velarization
Cu + $V_{[+\text{round}]}$	Cu + o → Coː	Elision
Ci + V	Ci + V → CVː	Elision

However, because these repairs appeared in the same morphosyntactic contexts, certain common strategies for handling such a diverse set of surface results, such as appeal to morphosyntactic domains, was not possible for Mushunguli. For some patterns, this proved unproblematic. In §3.3, we saw that any analysis of coalescence of low + high sequences (high coalescence) ruled out the possibility of deletion for other low + nonlow sequences in the same morphosyntactic context; thus, low + mid sequences were more parsimoniously analyzed as coalescence.

More troublesome was the analysis of identical vowel sequences; these, too, were more parsimoniously analyzed as coalescence than deletion, but this left the failure of them to undergo compensatory lengthening unexplained. Indeed, that Mushunguli has compensatory lengthening at all is somewhat unexpected given its lack of contrastive vowel length. In §3.4, I introduced a novel analysis for handling both the lack of contrastive vowel length and the failure of identical sequences to exhibit compensatory lengthening as a result of coalescence: Mushunguli exhibits a dispreference for moras associated with identical featural content in the input to be associated to the same vowel in the output. This resolved the asymmetry, and once again, a case of elision was ruled out in favor of coalescence. However, at no point could we definitively rule out deletion as a strategy.

This proved problematic for dealing with the results of postconsonantal prevocalic high vowels (§3.5). Indeed, we were left with a paradox: the ranking necessary to guarantee that labialization (and velarization) could apply in appropriate contexts was incompatible with the ranking necessary to guarantee that elision could apply in the complementary set of contexts. With no evidence from the regular patterns of the language to help us, we appeared to be stuck.

However, in §4, I introduced three exceptions to hiatus resolution in Mushunguli, each representing a separate typological instantiation of exceptional blocking. The first type was simple blocking, in which an exception fails to participate in in one applicable process, but regularly undergoes other applicable ones within the conspiracy. This was represented by the set of high vowel-initial stems in §4.2, which uniformly participated in glide formation and identity coalescence, but which exceptionally and exclusively blocked high coalescence, leaving hiatus unresolved. I presented an analysis of these stems using lexically-indexed constraints, which crucially capitalized on shared structural characteristics to both explain and unify two critical generalizations: these stems were the only exceptions that *exclusively* blocked coalescence, and they all began with high vowels. It was seen that lexically indexed constraints actually predict this to be the case: the only stems that *could* ever block exclusive coalescence are those which begin with high vowels.

The non-coalescing stems also introduced evidence that proved crucial for untangling our ranking paradox: the fact that these stems left hiatus unresolved in the same contexts in which other applicable repairs meant that all applicable repairs needed to be blocked. This clarified the status of vowel deletion in Mushunguli: namely, because MAX-V was undominated, an alternative path to accounting for the behaviors of postconsonantal prevocalic high vowels was necessary. The strategy adopted was to treat them as uniformly undergoing glide formation. The fact that these glides did not surface unchanged (or at all) indicated that there was a derivational characteristic to glide formation in Mushunguli, which would need to be accounted for in some fashion. It also meant that in order to account for these patterns, it would be necessary to substantially add to the analysis of syllable structure.

To deal with this, in §5.1, I adopted Stratal OT to account for the derivational nature of glide formation in the language. To begin motivating the analysis of glide formation, I first developed an analysis of word-level onset structure repairs, specifically focusing on consonant clusters and secondary articulations, both underlying and derived (5.1.1). In §5.1.2, I moved on to the postlexical level, to account for the surface results of glide formation in postconsonantal contexts. The ranking established there allowed for labialization (and velarization) to apply provided the following vowel was not round; in all other contexts, glides were deleted.

Thereafter, I introduced a second set of exceptions, a typologically unusual pattern which I referred to as *walljumping*. Walljumping exceptions are formally the opposite of simple blocking exceptions: they block applicable processes, but instead undergo an alternative repair that has not been ruled out. In the case of Mushunguli, the walljumping exceptions exhibited an otherwise unattested repair: palatalization, which crucially can only apply in exceptional contexts and only at morpheme edges. The introduction and analysis of these exceptions had several consequences. First, it revealed that although palatalization is not a regular pattern in the grammar, it is nevertheless permissible over leaving a consonant cluster unrepaired. Second, it underscored and reinforced the derivational character of glide formation in Mushunguli: as I demonstrated in §5.1.3.2, there is no available solution that would allow secondary articulations to be formed in one step, even in exceptional contexts.

Finally, in §5.2, I introduced a third set of exceptions, which I referred to as total non-participants, or TNPs. TNP exceptions uniformly block *all* applicable repairs. I demonstrated that these exceptions behave the way they do due to having especially strong left edges; this is in fact the only way to analyze such exceptions. That these exceptions, and *only* these exceptions, appear to avoid hiatus resolution entirely still does not treat them as extragrammatical: they conform to generalizations imposed upon them by the grammar and by the form and behavior of other exceptions in the language.

To summarize further: what appear to be eight surface results of four hiatus repairs were instead demonstrated to be eight surface results of *two* hiatus repairs, four of which were actually demonstrated to instead be onset structure repairs. This analysis crucially relied upon the predictable forms and patterns of exceptionality in the language, indicating that while exceptions are lexical, this is not equivalent to extragrammatical: rather, exceptions are reifying and reinforcing agents within a unified grammar.

In §6, I examined three alternative representational approaches to exceptionality in Mushunguli. Representational solutions to exceptionality take "unity" as a given, as most of these theories do not actually distinguish exceptions from regular forms. Rather, exceptions are just deviously regular forms with more abstract underlying representations than non-exceptional lexical items. I demonstrated that of the three approaches examined (absolute neutralization of segments with rules, absolute neutralization of segments in Stratal OT, and representational prespecification), only prespecification showed any promise for accounting for the exceptional and regular patterns in Mushunguli, and even this analysis struggled to capture some generalizations that were accounted for by the lexical indexation analysis, which generally assumed that underlying representations were fully specified. Whether these weaknesses can be accounted for is a topic for future research.

7.2 Future Research

There are several interesting points raised by the case of Mushunguli that require future research, both in terms of descriptive gaps and exceptionality in general.

7.2.1 Description of Mushunguli

As I discussed in §2, the grammar fragment of Mushunguli is just that: a fragment. While I strove to provide as complete a description and analysis as possible (without excessive speculation), there a number of aspects of the language that I did not address. Of these, tone is probably the most important; however, as I discussed in §2.2.5, it likely will be necessary to find a speaker who has better metalinguistic awareness of tone than the consultant that I worked with. The behavior of verbal extensions (especially with respect to reduplication), the behaviors of mid vowels in hiatus resolution, and other verbal paradigms also needs to be accounted for in future description. Similarly, some aspects of the grammar of Mushunguli are simply non-existent in either formal or descriptive literature, in particular description of syntax and morphosyntax, and again whether (and/or how) tone or intonation may play a role.

One final gap that I would like to note, because it is a potential problem for the proposed analysis of Mushunguli, are the behaviors of object prefixes. As I mentioned in §2.3.2.3, there are examples of object prefixes that are single high vowels resisting some aspect of simplification, e.g. /si + i + hisab + a/ → [si:hisa:ba] 'I counted them (cl 4)' (cf. *sihisa:ba* 'I counted'). They do not resist glide formation, e.g. /si + i + az + a/ → [sija:za] 'I lost them (cl 4)'. I did not provide an analysis of these prefixes because available data on anything but the human object prefixes (none of which are single vowels) is extremely limited. This was due to a combination of simply accidentally not eliciting certain hiatus contexts (e.g. present tense with non-human pronominal object prefixes, single object prefixes before disyllablic V-initial verb roots), and the fact that my primary consultant strongly preferred overt object marking for non-humans, possibly due to the discourse context of elicitation.

If this was demonstrated to be a consistent trend (and especially, if it could be demonstrated that these prefixes do *not* resist high coalescence), this would provide a challenge to the analysis presented in this dissertation.[1] However, there are options available for analysis. As I noted in §2.3.2.3, one is to capitalize on the fact that without an overt object, an object marker is obligatory. If this is the case, then it would be possible to index single vowel object prefixes to a constraint that penalizes eliminating all surface phonetic reflexes of the underlying morpheme. Another option would be to eschew indexation and instead rely on a transderivational constraint such as BASE-ID (Kenstowicz

[1]If the prefixes uniformly resist high coalescence, then they would be indexable to UNIFORMITY.

1996), which evaluates outputs based on their similarity to other possible words. A third option would be to assume that object prefixes form a new domain (i.e. the macrostem) or are enclitic; if this is the case, their apparently "exceptional" behavior may actually be representative of their morphosyntactic category or domain, and thus subject to different restrictions than e.g. subject prefixes. The correct answer to this problem will likely not come exclusively from phonology, but will require considerably more investigation into the morphosyntax and pragmatics of object marking and agreement in Mushunguli.

7.2.2 Exceptions as part of the grammar: Yine

In the case of Mushunguli, the dependencies between exceptional and regular forms were only possible to be seen because the analysis did not treat the exceptions to hiatus resolution in a vacuum, but rather contextualized them in the discussion of the larger patterns of the grammar. Analyses like this are rare, but as I discuss below, there are similarly morphologically-complex languages with similarly complicated patterns of exceptionality (or morpheme-specific phonology), which have not received this sort of formal treatment, and as I discuss briefly here, present challenges to common assumptions made by exceptionality theories. One notable case is Yine (Matteson 1965), which currently has multiple formal proposals, none of which actually account for all patterns of morpheme-specific phonology *and* regular cases of opacity born of the morphological complexity of the language.

Yine (formerly Piro) (Matteson 1965) is one of the seminal cases in morpheme-specific phonology, as it was the first language used to demonstrate a typological distinction between exceptional (non-)triggering and (non-)undergoing patterns (Kisseberth 1970b) (as opposed to only the latter, discussed in SPE (Chomsky & Halle 1968)). In relatively theory-neutral terms, the basic generalizations are as follows: Yine is a primarily-suffixing, agglutinating language with a pattern of vowel deletion that is sensitive to both structural and morphological factors. In general, the concatenation of a suffix with some preceding structure causes the final vowel of that structure to delete; for example, /heta + lu/ → [het_lu] 'see it'.[2] This can happen multiple times, as exhibited by the form /nika + ya + waka + lu/ → [nik_yawak_lu] 'to eat it there.' However, deletion is blocked if it would create a triconsonantal cluster, as was exhibited by the failure of the vowel in the indirective suffix /-y**a**/ to delete in the preceding example.

In addition to this predictable blocking, there are some suffixes which consistently fail to trigger deletion, which I will indicate via underline. An example of this is the anticipatory suffix /-<u>nu</u>/, which fails to trigger deletion when concatenated to a preceding verb stem, even if it would not create an illicit cluster (e.g. /heta + <u>nu</u>/ → [het**a**nu] 'going to see'). Most non-triggering suffixes still undergo deletion if they are followed by a trigger, e.g. /heta + <u>nu</u> + lu/ → [hetan_ru] 'going to see him'. Importantly, there are no obvious structural characteristics that distinguish triggering from non-triggering suffixes, as illustrated by the behavior of the homophonous abstract noun suffix /-nu/, e.g. /hata + nu/ → hat_nu 'light, shining'. As such, this is properly viewed as a morpheme-specific phenomenon.

There have been multiple attempts to formalize the Yine facts since Matteson's (1965) grammar of the language was published. The earliest analysis of the asymmetrically behaving suffixes is rule-based: Kisseberth (1970b) notes that there are no striking phonological generalizations about the forms in question, and so elects to treat the non-triggering suffixes as exceptional failures to condition deletion (which he refers to as "vowel drop"). His solution is to mark them diacritically as being [-context], which blocks them from conditioning deletion, but does not prevent them from undergoing it.

The remaining three analyses of interest are all based in Optimality Theory. First,

[2] I follow the practice adopted by Lin (1997) of marking deleted segments with a <_>, as the morphological complexity of this language can otherwise make it difficult to identify which segments have been deleted.

Lin (1997) proposes a cyclic analysis of *regular* vowel deletion (implying again that the non-triggering suffixes are the ones which are exceptional), as well as concomitant patterns of opacity and repair of consonant clusters. She acknowledges the existence of the non-triggering suffixes, but does not offer an analysis of them specifically. However, this analysis is still worth considering for two reasons: first, cyclicity in Optimality Theory is a successor to rule-based Lexical Phonology (Kiparsky 1982b,a, 1985; Mohanan 1986), which is itself a theory of how to handle morpheme-specific phonological patterns; as such, the fact that this is not adopted by subsequent analyses is a specific analytical choice. Second, Lin's contribution accounts for some related patterns that later contributions do not.

Pater (2010) offers a spiritual successor to Kisseberth's analysis, and refers to the non-triggering suffixes as "exceptions". However, unlike Kisseberth, Pater assigns a diacritic to the *triggering* suffixes, linking them to a copy of the markedness constraint ALIGN-SUF^{L}-C, which penalizes suffixes whose left edge does not align wit ha consonant.

Finally, Zimmermann (2013) proposes a (mostly) representational analysis: she assumes that non-triggering suffixes have a mora, while triggering suffixes deficiently lack one. The "defective" underlying representation suggests that the triggers are considered "exceptional" (cf. Inkelas (1995); Inkelas et al. (1997), who associate exceptional behavior with specification). In this case, exceptions are in some sense marked by their lack of a mark, though it is worth noting again that representational analyses of this sort are more in line with the approach that Moravcsik (2011) describes as "regularizing" exceptions; by reducing exceptional behavior to an underlying structural distinction, it becomes roughly as exceptional as any other alternation in a language–that is to say, not exceptional at all.[3]

[3]Matteson's (1965) description is agnostic as to which set is exceptional, instead simply noting the disparity in behaviors. Hanson's (2010) more recent description is similarly neutral, only noting that outside of the predictable avoidance of three-consonant sequences, a suffix's ability to trigger deletion is lexically-determined.

So far, we have seen that all four authors agree that deletion in Yine involves some form of morphological conditioning, but disagree both on its nature and its implication for the analysis. This is partially because unlike e.g. Mushunguli, it is difficult to actually identify which set of suffixes is exceptional. That said, there is a morpheme-specific pattern in Yine that all authors seem to agree is exceptional (insofar that it is specifically marked).

There is a suffix /-*wa*/ 'yet, still', which is homophonous with the non-triggering intransitive verb theme suffix /-wa/. This suffix neither triggers *nor* undergoes deletion (e.g. /heta+*wa*+lu/ → [het**a**w**a**lu] 'going to see him yet'). This morpheme, and *only* this morpheme, is identified as an exception by all four authors, and of the three who specifically address exceptionality within their paper, all choose to mark it in some way.

Kisseberth again relies on a rule diacritic, marking the suffix as [-rule], which accounts for both non-undergoing and non-triggering behavior. Pater links it to an indexed faithfulness constraint; since non-triggering is treated formally as the unmarked pattern, this is the only diacritic necessary to account for this suffix's behavior. Finally, Zimmermann assigns the suffix two moras, one attached and one floating. This time, it is the addition of structure that marks the suffix as exceptional, rather than that lack of it.

All of these authors make strong arguments for their analyses; in particular, Zimmermann's analysis can account for positional deficiencies in Pater's (namely that Pater's analysis incorrectly predicts excessive vowel deletion when multiple triggering suffixes are concatenated, see Zimmermann (2013) for details). However, it is also the case that none of these analyses actually account for all patterns discussed by each author, and it is also the case that Yine has patterns that *no* analysis considers or accounts for.

In the former case, Lin (1997) points out (and provides an analysis for) an interesting fact about the apparent ban on triplicate consonant clusters; namely, it is not as general as other analyses tend to treat it. Sometimes, when two triggering suffixes are concatenated, an opaque pattern of *consonant* deletion and compensatory lengthening can emerge. An example of this is /kna + mta$_1$s$_2$a$_3$ + x$_4$e$_5$ + kaka/ 'lanky' (lit. 'pole-hollow-always-COLLECTIVE'), which surfaces as [knamta: $_1$_$_2$_$_3$x$_4$_$_5$kaka], with segments /a$_3$/ and /e$_5$/ lost due to the application of vowel deletion and segment /s$_2$/ lost due to the repair of an illicit sequence of fricatives.

Under Lin's cyclic analysis, whether this pattern applies is predictable based on the type of consonant cluster that is created by the initial application of deletion, with fricative clusters being one such predictable context. So, for 'lanky,' the initial round of deletion produces an illicit intermediate fricative cluster (/knamta$_1$s$_2$a$_3$ + x$_4$e$_5$/ → |knamta$_1$s$_2$_$_3$x$_4$e$_5$|). When the final suffix is added on the next cycle, deletion is *not* blocked; instead, a triplicate sequence is formed (|knamta$_1$s$_2$_x$_4$e$_5$| + /kaka/ → |knamta$_1$s$_2$_x$_4$_$_5$kaka|). The unsyllabifiable consonant cluster is then later repaired by deleting the first of the offending consonants; because consonants are assumed to be moraic in Yine, the lingering mora is reassociated with the preceding vowel, resulting in compensatory lengthening and a surface output [knamta:xkaka].

Again, only Lin discusses this pattern in significant detail or tries to account for it formally. Pater (2010) acknowledges that there are complications he chooses to abstract away from (fn 5); this is understandable given that his primary goal is not to challenge Lin's analysis, but rather to demonstrate a possible application of indexed markedness constraints. However, for Zimmermann (2013), the presence of these examples demonstrates a weakness in an otherwise elegant analysis. Zimmermann's proposal relies on the assumption that deletion-triggering suffixes, which are underlyingly moraless, can "steal" the mora from the vowel of the preceding syllable, causing the segmental content to be deleted. More concisely, deletion is a repair operation intended to fix deficient underlying representations. In order for this analysis to function, it is critical that triggers be *moraless*. As alluded to above, in the rare cases where two triggers are concatenated together, as was the case for /nika + ya + waka + lu/ 'to eat it there' from above, the analysis is saved by by a high-ranked constraint penalizing CCC sequences, which prevents deletion from applying.

However, the opaque example of 'lanky' does not conform to this generalization. Here, the vowel of the triggering suffix /-xe/ 'always' is deleted, despite not having a mora to steal. And while Zimmermann's goal *does* appear to be to challenge both Lin's and Pater's analyses (probably because she is proposing a representational analysis as opposed to the latter authors' diacritic approaches), she chooses to dismiss examples like 'lanky' as "colloquial" exceptions on the grounds that they are rare and not discussed in detail by Matteson (1965), but instead appear in the texts exemplifying the actual use of the language (Zimmermann 2013:138, fn 15).

This is doubly problematic because it relies on exceptionality both as a formal device and a rhetorical shortcut: morpheme-specific behavior amenable to a representational analysis is accounted for, but anything that is not amenable to this analysis is an exception that does not need to be accounted for (discarded to the lexicon, memorized, etc.). That this opaque behavior conforms to other generalizations about the repair of geminates (Lin 1997) is immaterial.[4]

Like Mushunguli, Yine is a morphologically complex language, and this morphological complexity seems to be a breeding ground for complex interrelationships between exceptional and regular forms. As such, I would suggest that the failures of these analyses to either capture (or address) every pattern stems from the lack of any unified analysis of the language more generally. In light of this, I would like to return to the point taht there are potentially important patterns in Yine noted in descriptions by Matteson (1965) and Hanson (2010), which may prove challenging for the extant analyses of the language, as well as some assumptions adopted by of exceptionality theories.

[4]Indeed, this is a case where an important phonological generalization is nevertheless missed by a representational analysis, despite their purported superiority in this regard.

First, it is the case that in Yine, it is not only morpheme boundaries that can trigger deletion; deletion applies optionally at *word* boundaries, as well. For example, /wanepnute koca/ 'next also,' variably surfaces as [wanepnute koca]~[wanepnut_ koca] (Hanson 2010:(33), p31). This pattern is potentially problematic for a representational analysis like Zimmerman's, which assumes moras can be stolen from preceding syllables due to a lack of a mora in following ones. While it is possible that one could stipulate that every root-initial vowel is potentially moraless, such an argument would be more compelling if there was evidence for it from multiple sources. This pattern is also problematic for Pater's (2010) skeletal analysis, which relies on indexed ALIGN-SUF^{L}-L to trigger deletion. Minimally, this analysis will need to be modified to account for the fact that it is not only suffixes that can trigger deletion.[5]

Problematic for all extant analyses is the existence of the elative suffix /-pa/, which triggers the deletion of what Matteson describes as the "noncontiguous preceding vowel" (Matteson 1965:80), and what Hanson refers to as the "penultimate stem vowel" (Hanson 2010:30). An example from each source is given below; note that I am leaving each glossing and transcription intact from each description, but I have added an underscore to illustrate the deleted vowel.

(198) rap_kapyalo
r-hapoka-pa-ya-lo
3-arrive-ELV-APPL-she
'It came and arrived there where she was.' (Matteson 1965:80)

(199) rethim_tapanrɨna...
-heta-hima-ta-pa-nɨ-lɨ-na
3-see-QUOT-VCL-ELV-ANTIC-3SGM-3PL
'(they went) to see (...the jaguars)' (Hanson 2010:386)

[5]It is probably unproblematic for Lin's proposal, given that she is already assuming the presence of cyclicity; thus it would not be difficult to build postlexical deletion into it.

Note that the anticipatory suffix /-ta/ here is a non-trigger, as evidenced by the following example.

(200) rɨylahimatanro...
r-hiyla-hima-ta-na-lo
3-kill-QUOT-VCL-CMPV-3SGF
'he killed (the black fly)' (Hanson 2010:69)

Both Matteson and Hanson note that deletion caused by this suffix is blocked when it would create an impermissible triconsonantal cluster; that is, this suffix conforms to the same phonotactic generalizations used to motivate all analyses of the typically-discussed (non-)triggering suffixes. Note also from the first example that this really does seem to be a purposeful deletion of the "non-contiguous" preceding vowel, as in that case it is the root-medial /o/ that is being deleted.

That this affix can trigger deletion in a non-adjacent vowel (in some cases skipping a morpheme entirely) makes it problematic for the version of lexical indexation proposed by Pater (2010). It is also problematic for the mora usurpation analysis proposed by Zimmermann (2013), as a mora usurpation analysis would require the crossing of an association line–remember, non-triggering /-ta/ must have a mora in order to account for why it does not trigger deletion in the third example.

That patterns of this complexity exist in Yine underscores a final point: again, much work on exceptionality is restricted to major world languages, frequently European ones, which may exhibit significant overlaps in both their morphological systems and their lexica. The case study of Mushunguli, as well as the unresolved case of Yine, attest to the fact that typologically unusual and complicated patterns of exceptionality exist in understudied language families, especially those with complex agglutinating morphology. If we would truly like our formal theories—whatever they may be—to capture attested cross-linguistic patterns of exceptionality, it behooves us as linguists to broaden our search space, both by engaging in more in-depth work on a much more typologically rich sample of languages, and by investigating the typology of exceptional patterns on its own.

Bibliography

ALBRIGHT, ADAM; ARGELIA ANDRADE; and BRUCE HAYES. 2001. Segmental environments of Spanish diphthongization. UCLA working papers in linguistics 7, ed. by Adam Albright and Taehong Cho, 117–151. Los Angeles, CA: UCLA.

ALBRIGHT, ADAM, and BRUCE HAYES. 2002. Modeling English past tense intuitions with minimal generalization. Proceedings of the sixth meeting of the ACL special interest group in computational phonology, ed. by Michael Maxwell, Philadelphia, PA: ACL, 58–69.

ALBRIGHT, ADAM, and BRUCE HAYES. 2003. Rules vs. analogy in English past tenses: A computational/experimental study. Cognition 90.119–161.

ALBRIGHT, ADAM, and BRUCE HAYES. 2006. Modeling productivity with the Gradual Learning Algorithm: The problem of accidentally exceptionless generalizations. Gradience in grammar: Generative perspectives, ed. by Gisbert Fanselow, Caroline Fery, Mathias Schlesewsky, and Ralf Vogel, 185–204. Oxford: Oxford University Press.

ALDERETE, JOHN D. 2001. Dominance effects as trans-derivational anti-faithfulness. Phonology 18.201–253.

ANTTILA, ARTO. 2002. Morphologically conditioned phonological alternations. Natural Language and Linguistic Theory 20.1–42.

AOKI, PAUL K. 1974. An observation of vowel contraction in Xhosa. Studies in African Linguistics 5.223–241.

ARCHANGELI, DIANA, and DOUG PULLEYBLANK. 2015. Tonal allomorphy in Kinande. Capturing phonological shades within and across languages, ed. by Yuchau E. Hsiao and Lian-Hee Wee. Cambridge Scholars Publishing.

BAKOVIĆ, ERIC. 2000. Harmony, dominance, and control. New Brunswick, NJ: Rutgers dissertation.

BAKOVIĆ, ERIC. 2006. Hiatus resolution and incomplete identity. Optimality-Theoretic studies in Spanish phonology, ed. by Sonia Colina and Fernando Martínez-Gil, 62–73. Amsterdam/Philadelphia: John Benjamins. ROA-813.

BAKOVIĆ, ERIC. 2011. Opacity and ordering. The handbook of phonological theory, ed. by John Goldsmith, Jason Riggle, and Alan Yu. Blackwell.

BARLEW, JEFFERSON. 2013. Point of view in Mushunguli locatives. Selected proceedings of the 43rd Annual Conference on African Linguistics, 115–129. Somerville, MA: Cascadilla Proceedings Project.

BARLEW, JEFFERSON. 2016. Point of view and the semantics of spatial expressions in mushunguli and english. Journal of Semantics 33.215–267.

BASTIN, YVONNE; ANDRÉ COUPEZ; EVARISTE MUMBA; and THILO C. SCHADEBERG. 2002. Bantu lexical reconstructions 3 / reconstructions lexicales bantoues 3. http://linguistics.africamuseum.be/BLR3.html. Online database.

BECKER, MICHAEL. 2003. Lexical stratification of Hebrew: The disyllabic maximum. Proceedings of the Israel Association for Theoretical Linguistics, ed. by Yehuda Falk, vol. 19.

BECKER, MICHAEL. 2006. Learning the lexicon through the grammar - evidence from Turkish. Ms., University of Massachussetts, Amherst.

BECKER, MICHAEL. 2009. Phonological trends in the lexicon: The role of constraints. Amherst, Massachussetts: University of Massachussetts, Amherst dissertation.

BECKER, MICHAEL, and MARIA GOUSKOVA. 2016. Source-oriented generalizations as grammar inference in Russian vowel deletion. Linguistic Inquiry 47.391–425.

BECKMAN, JILL. 1997. Positional faithfulness, positional neutralization, and Shona vowel harmony. Phonology 14.

BEEDHAM, CHRISTOPHER. 2005. Language and meaning: The structural creation of reality, Studies in Functional and Structural Linguistics, vol. 55. Amsterdam/Philadelphia: Benjamins.

BERMÚDEZ-OTERO, RICARDO. 2012. The architecture of grammar and the division of labour in exponence. The morphology and phonology of exponence, ed. by Jochen Trommer, Oxford Studies in Theoretical Linguistics, vol. 41, 8–83. Oxford: Oxford University Press.

BERMUDEZ-OTERO, RICARDO. 2011. Cyclicity. The blackwell companion to phonology, ed. by Marc van Oostendorp, Colin J. Ewen, Elizabeth Hume, and Keren Rice, vol. 4, 2019–48. Malden, MA: Wiley-Blackwell.

BOERSMA, PAUL, and BRUCE HAYES. 2001. Empirical tests of the gradual learning algorithm. Linguistic Inquiry 32.45–86.

BOSTOEN, KOEN, and YVONNE BASTIN. 2016. Bantu lexical reconstruction. Oxford handbooks online, 1–31. Oxford University Press. Online: http://dx.doi.org/10.1093/oxfordhb/9780199935345.013.36.

BRAME, MICHAEL K. 1972. On the abstractness of phonology: Maltese ʕ. Contributions to generative phonology, ed. by Michael K. Brame, 22–61. Austin: University of Texas Press.

CASALI, RODERIC F. 1996. Resolving hiatus. Los Angeles, CA: UCLA dissertation.

CASALI, RODERIC F. 1997. Vowel elision in hiatus contexts: Which vowel goes? Language 73.493–533.

CASALI, RODERIC F. 2011. Hiatus resolution. The blackwell companion to phonology, ed. by Marc van Oostendorp, Colin J. Ewen, Elizabeth Hume, and Keren Rice, vol. 3, 1434–1460. Malden, MA: Wiley-Blackwell.

CHANOFF, SASHA. 2002. After three years: Somali Bantus prepare to come to America. Refugee reports, ed. by Shannon Dennett and Alyson Springer, vol. 23, 1–11. Periodical distributed by Immigration and Refugee Services of America.

CHOMSKY, NOAM, and MORRIS HALLE. 1968. The sound pattern of english. New York: Harper & Row.

CLEMENTS, G. N. 1981. Akan vowel harmony: a nonlinear analysis. Harvard journal of phonology, vol. 2, 108–177.

CLEMENTS, G.N. 1991. Vowel height assimilation in Bantu languages. Working Papers of the Cornell Phonetics Laboratory 5.37–76.

COMPTON, RICHARD, and B. ELAN DRESHER. 2011. Palatalization and "strong i" across Inuit dialects. Canadian Journal of Linguistics/Revue canadienne de linguistique 56.203–228.

DAVIDSON, LISA, and ROLF NOYER. 1997. Loan phonology in Huave: Nativization and the ranking of faithfulness constraints. Proceedings of the 15th west coast conference on formal linguistics, Department of Linguistics, University of Toronto, 65–79.

DAYLEY, JON P.; MWALIKO MBERWA; and MICHAL TEMKIN MARTINEZ. 2020. Chizigula of Somalia - English dictionary. Webonary.org. SIL International.

DOLPHYNE, F. 1988. The Akan (Twi-Fante) language: its sound systems and tonal structure. Ghana Universities Press.

DOWNING, LAURA. 1997. Correspondence effects in Siswati reduplication. Studies in the Linguistic Sciences 25.81–95.

DOWNING, LAURA. 2005. The emergence of the marked: Tone in some african reduplicative systems. Studies on reduplication, ed. by Bernard Huch, 87–108. Berlin: Mouton de Gruyter.

DOWNING, LAURA. 2006. The prosody and syntax of focus in Chitumbuka. ZAS Papers in Linguistics 43.55–79.

EBERHARD, DAVID M.; GARY F. SIMONS; and CHARLES D. FENNIG. 2019. Ethnologue: Languages of the world. Published online at http://www.ethnologue.com. Version 22, SIL International.

ENO, OMAR, and MOHAMED ENO. 2007. The journey back to the ancestral homeland: The return of the Somali Bantu (wazigwa) to modern Tanzania. From Mogadishu to Dixon: The Somali diaspora in a global context, 13–43. Red Sea Press.

ERNESTUS, MIRJAM, and R. HARALD BAAYEN. 2003. Predicting the unpredictable: Interpreting neutralized segments in Dutch. Language 79.5–38.

FINLEY, SARA. 2010. Exceptions in vowel harmony are local. Lingua 120.1549–1566.

FORTUNE, GEORGE. 1955. An analytical grammar of Shona. Capetown and New York: Longmans, Green.

FUKAZAWA, H. 1999. Theoretical implications of ocp effects on features in optimality theory. University of Maryland, College Park dissertation. [ROA-307].

GOUSKOVA, MARIA. 2012. Unexceptional segments. Natural Language and Linguistic Theory 30.79–133.

GOUSKOVA, MARIA, and MICHAEL BECKER. 2013. Russian yer alternations are governed by the grammar. Natural Language and Linguistic Theory 31.735–765.

HANSON, REBECCA. 2010. A grammar of Yine (Piro). Melbourne, Australia: La Trobe University dissertation.

HARMS, ROBERT T. 1973. How abstract is Nupe? Language 49.439–446.

HAYES, BRUCE. 1989. Compensatory lengthening in moraic phonology. Linguistic Inquiry 20.235–305.

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

HAYES, BRUCE; KIE ZURAW; PETER SIPTAR; and ZSUZSA LONDE. 2009. Natural and unnatural constraints in Hungarian vowel harmony. Language 85.822–863.

HOUT, KATHERINE. 2011. The noun class system of Mushunguli. Unpublished Ms.

HOUT, KATHERINE. 2012. The vocalic phonology of Mushunguli. B.A. Thesis, The Ohio State University. Retrievable at http://kb.osu.edu/dspace/handle/1811/53373.

HOUT, KATHERINE. 2016. A lexical indexation account of exceptions to hiatus resolution in Mushunguli. San Diego Linguistics Papers 6.

HOUT, KATHERINE. 2017. Exceptions to hiatus resolution in Mushunguli (Somali Chizigula). Africa's endangered languages: Documentary and theoretical approaches, ed. by Jason Kandybowicz and Harold Torrence, chap. 17. New York: Oxford University Press.

HOUT, KATHERINE. 2019a. The disambiguating effects of phonological exceptions in grammar. Proceedings of the 2018 Annual Meeting on Phonology, ed. by Katherine Hout, Anna Mai, Adam McCollum, Sharon Rose, and Matt Zaslansky.

HOUT, KATHERINE. 2019b. Dominance-as-markedness: evidence from Bari. Studies in African Linguistics 48.206–224.

HSU, BRIAN, and KAREN JESNEY. 2018. Weighted scalar constraints capture the typology of loanword adaptation. Proceedings of the 2017 Annual Meeting on Phonology, Washington, DC: Linguistic Society of America.

HYMAN, LARRY. 1973. Nupe three years later. Language 49.447–452.

HYMAN, LARRY M. 1970. How concrete is phonology? Language 46.58–76.

HYMAN, LARRY M. 1985. A theory of phonological weight. Dordrecht: Foris. Reprinted 2003, Stanford: CSLI.

HYMAN, LARRY M. 1999. The historical interpretation of vowel harmony in Bantu. Bantu historical linguistics: Theoretical and empirical perspectives, ed. by Jean-Marie Hombert and Larry M. Hyman, 235–295. Stanford: CSLI Publications.

HYMAN, LARRY M. 2003. Segmental phonology. The Bantu languages, ed. by Derek Nurse and Gérard Philippson, 42–58. London: Routledge.

HYMAN, LARRY M. 2009. The natural history of verb-stem reduplication in Bantu. Morphology 19.177–206.

INKELAS, SHARON. 1995. The consequences of optimization for underspecification. Proceedings of the twenty-fifth Northeastern Linguistics Society, ed. by E. Buckley and S. Iatridou, 287–302. Amherst: GLSA.

INKELAS, SHARON; ORHAN ORGUN; and CHERYL ZOLL. 1997. The implications of lexical exceptions for the nature of grammar. Derivations and constraints in phonology, ed. by I. Roca, 393–418. New York: Oxford University Press.

INKELAS, SHARON, and CHERYL ZOLL. 2007. Is grammar dependence real? Linguistics 45.133–171.

ITO, JUNKO, and ARMIN MESTER. 1995a. The core-periphery structure in the lexicon and constraints on re-ranking. Papers in optimality theory, ed. by Jill N. Beckman, Laura Walsh Dickey, and Suzanne Urbanczyk, 181–210. University of Massachusetts, Amherst: GLSA.

ITO, JUNKO, and ARMIN MESTER. 1995b. Japanese phonology. The handbook of phonological theory, ed. by John A. Goldsmith, 817–838. Oxford: Blackwell.

ITO, JUNKO, and ARMIN MESTER. 1999a. The phonological lexicon. A handbook of japanese linguistics, ed. by Natsuko Tsujimura, 62–100. Oxford: Blackwell.

ITO, JUNKO, and ARMIN MESTER. 1999b. Realignment. The prosody-morphology interface, ed. by Rene Kager, Harry van der Hulst, and W. Zonneveld, 188–217. Cambridge University Press.

ITO, JUNKO, and ARMIN MESTER. 2001. Covert generalizations in optimality theory: the role of stratal faithfulness constraints. Studies in Phonetics, Phonology, and Morphology 7.273–299.

KABAK, BARIS, and IRENE VOGEL. 2011. Exceptions to stress and harmony in Turkish: co-phonologies or prespecification? Expecting the unexpected: Exceptions in grammar, ed. by Horst J. Simon and Heike Wiese. Berlin & New York: Walter de Gruyter.

KADENGE, MAXWELL. 2013. Hiatus resolution in Nambya: an Optimality Theory analysis. Language Matters 44.94–121.

KAISSE, ELLEN, and PATRICIA SHAW. 1985. On the theory of lexical phonology. Phonology yearbook 2, 1–30.

KEER, EDWARD. 1999. Geminates, the OCP, and the nature of Con. New Brunswick, NJ: Rutgers dissertation.

KENSTOWICZ, MICHAEL. 1989. Tone and accent in Kizigua – a Bantu language. Certamem phonologicum: papers from the 1987 cortona phonology meeting, ed. by Marco Pier Bertinetto and Michele Loporcaro, 177–188. Torino: Rosenberg and Sellier.

KENSTOWICZ, MICHAEL. 1996. Base-identity and uniform exponence: alternatives to cyclicity. Current trends in phonology: Models and methods, ed. by Jacques Durand and Bernard Laks, 365–395. CNRS, Paris-X and University of Salford Publications.

KENSTOWICZ, MICHAEL, and CHARLES KISSEBERTH. 1977. Generative phonology: Description and theory. New York: Academic Press.

KENSTOWICZ, MICHAEL, and CHARLES KISSEBERTH. 1990. Chizigula tonology: the word and beyond. The phonology-syntax connection, ed. by Sharon Inkelas and Draga Zec, 163–194. Chicago: University of Chicago Press.

KIPARSKY, PAUL. 1968/1982a. How abstract is phonology? Explanation in phonology, 119–164. Berlin & New York: Mouton de Gruyter.

KIPARSKY, PAUL. 1968/1982b. Linguistic universals and linguistic change. Universals in linguistic theory, ed. by E. Bach and R. T. Harms, 170–202. Holt, Reinhart, and Winston. Reprinted in Explanation in Phonology, pp. 13-55.

KIPARSKY, PAUL. 1982a. From cyclic phonology to lexical phonology. The structure of phonological representations (i), ed. by Harry van der Hulst and Norval Smith, 131–175. Foris.

KIPARSKY, PAUL. 1982b. Word formation and the lexicon. Proceedings of the Mid-America Linguistics Conference, ed. by F. Ingemann, Lawrence, KS.

KIPARSKY, PAUL. 1985. Some consequences of lexical phonology. Phonology Yearbook 2.85–138.

KIPARSKY, PAUL. 1993. Blocking in non-derived environments. Studies in lexical phonology, ed. by S. Hargus and Ellen Kaisse. San Diego: Academic Press.

KIPARSKY, PAUL. 1994. Remarks on markedness. Trilateral Phonology Weekend II. University of California, Santa Cruz.

KIPARSKY, PAUL. 2000. Opacity and cyclicity. The Linguistic Review 17.351–366.

KIPARSKY, PAUL. 2015. Stratal OT: A synopsis and FAQs. Capturing phonological shades, ed. by Yuchau E. Hsiao and Lian-Hee Wee. Cambridge Scholars Publishing.

KISSEBERTH, CHARLES. 1992. Metrical structure in Zigula tonology. African linguistic contributions, ed. by Derek Gowlett, 227–259. Via Afrika.

KISSEBERTH, CHARLES W. 1970a. On the functional unity of phonological rules. Linguistic Inquiry 1.291–306.

KISSEBERTH, CHARLES W. 1970b. The treatment of exceptions. Papers in Linguistics 2.44–58.

KISSEBERTH, CHARLES W. 2011. Conspiracies. The blackwell companion to phonology, ed. by Marc van Oostendorp, Colin Ewen, Elizabeth Hume, and Keren Rice, 1644–1665. Oxford: Blackwell.

KISSEBERTH, CHARLES W., and MOHAMMAD IMAM ABASHEIKH. 2011. Chimwiini phonological phrasing revisited. Lingua 121.1987–2013.

KISSEBERTH, CHARLES W., and FARIDA CASSIMJEE. n.d. Chizigula lexicon exemplified. Unpublished Ms.

KRASKA-SZLENK, IWONA. 1997. Exceptions in phonological theory. Proceedings of the 16th International Congress of Linguists. Oxford: Pergamon.

KRASKA-SZLENK, IWONA. 1999. Syllable structure constraints in exceptions. Phonologica 1996: Syllables?!, ed. by John Rennison and Klaus Kühnhammer, 113–131. The Hague: Thesus.

LIN, YEN-HWEI. 1997. Syllabic and moraic structures in Piro. Phonology 14.403–436.

MAHANTA, SHAKUNTALA. 2008. Directionality and locality in vowel harmony. Utrecht University dissertation.

MAHANTA, SHAKUNTALA. 2012. Locality in exceptions and derived environments in vowel harmony. Natural Language and Linguistic Theory 30.

MARLETT, STEPHEN A., and JOSEPH PAUL STEMBERGER. 1983. Empty consonants in Seri. Linguistic Inquiry 14.617–639.

MASCARÓ, JEAN. 1976. Catalan phonology and the phonological cycle. Cambridge, Massachussetts: MIT dissertation.

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

MCCARTHY, JOHN. 1997. Faithfulness and prosodic circumscription. ROA-201.

MCCARTHY, JOHN. 2003. OT constraints are categorical. Phonology 52.

MCCARTHY, JOHN. 2007. Hidden generalizations: Phonological opacity in Optimality Theory. Equinox.

MCCARTHY, JOHN, and ALAN PRINCE. 1993. Prosodic morphology i: Constraint interaction and satisfaction. Cambridge, MA: MIT Press.

MCCARTHY, JOHN H., and ALAN PRINCE. 1995. Faithfulness and reduplicative identity. Papers in optimality theory, ed. by Jill N. Beckman, Laura Walsh Dickey, and Suzanne Urbanczyk, 249–384. University of Massachusetts, Amherst: GSLA.

MCCARTHY, JOHN J. 1999. Sympathy and phonological opacity. Phonology 16.

MCLAREN, JAMES. 1955. A Xhosa grammar. 2 edn. Cape Town: Longmans, Green & Co.

MEEUSSEN, ACHIEL E. 1967. Bantu grammatical reconstructions. Africana Linguistica 3.79–121.

MOHANAN, K. P. 1986. The theory of lexical phonology. Dordrecht, Holland: Reidel.

MOORE-CANTWELL, CLAIRE, and JOE PATER. 2016. Gradient exceptionality in Maximum Entropy Grammar with lexically specific constraints. Catalan Journal of Linguistics 15.53–66, Proceedings of the OCP 2015 special workshop on exceptionality.

Moravcsik, Edith. 2011. Coming to grips with exceptions. Expecting the unexpected: Exceptions in grammar, ed. by Horst J. Simon and Heike Wiese. Berlin & New York: Walter de Gruyter.

Mudzingwa, Calisto. 2010. Shona morphophonemics: repair strategies in Karanga and Zezuru. Canada: University of British Columbia dissertation.

Mudzingwa, Calisto, and Maxwell Kadenge. 2011. Comparing hiatus resolution in Karanga and Nambya: An Optimality Theory account. Nordic Journal of African Studies 20.203–240.

Nurse, Derek. 2010. The decline of Bantu in Somalia. Essais de typologie et de linguistique générale. Mélanges offerts à Denis Creissels., ed. by F. Floricic, 187–200. Lyon: ENS Editions.

Odden, David. 1996. The phonology and morphology of Kimatuumbi. New York: Oxford University Press.

Odden, David. 2015. Bantu phonology. Oxford handbooks online (linguistics). Oxford.

Odden, David, and Mary Odden. 1999. Kihehe syllable structure. The syllable: Views and facts, ed. by Harry van der Hulst and Nancy A. Ritter, 417–445. Mouton de Gruyter.

Padgett, Jaye. 1994. Stricture and nasal place assimilation. Natural Language & Linguistic Theory 12.465–513.

Paradis, Carole, and Caroline Lebel. 1994. Contrasts from segmental parameter settings in loanwords: core and periphery in Quebec French. Proceedings of the mot conference on contrast in phonology, Department of Linguistics, University of Toronto, 75–95.

Paster, Mary. 2006. Aspects of Maay phonology and morphology. Studies in African Linguistics 35.73–120.

Pater, Joe. 1999. Austronesian nasal substitution and other NC effects. The prosody morphology interface, ed. by Rene Kager, Harry van der Hulst, and Wim Zonneveld, 310–343. Cambridge University Press.

Pater, Joe. 2000. Non-uniformity in english stress: the role of ranked and lexically specific constraints. Phonology 17.237–274.

Pater, Joe. 2004. Exceptions in Optimality Theory: Typology and learnability. http://people.umass.edu/pater/exceptions.pdf. Handout from the Conference on Redefining Elicitation: Novel Data in Phonological Theory.

Pater, Joe. 2010. Morpheme-specific phonology: Constraint indexation and inconsistency resolution. Phonological argumentation: Essays on evidence and motivation, ed. by Steve Parker, 123–154. London: Equinox.

PILLION, ELIZABETH. 2013. An autosegmental analysis of verbal tone in Mushunguli. B.A. Thesis, The Ohio State University.

PRINCE, ALAN, and PAUL SMOLENSKY. 1993/2004. Optimality Theory: Constraint interaction in generative grammar. Blackwell.

ROBERTS-KOHNO, RUTH. 1998. Empty root notes in Kikamba: Conflicting evidence and theoretical implications. Language description and linguistic history in Africa, ed. by Ian Maddieson and Thomas J. Hinnebusch. Trenton, NJ: Africa World Press.

ROSENTHALL, SAMUEL. 1997. The distribution of prevocalic vowels. Natural Language & Linguistic Theory 15.139–180.

SACIUK, B. 1969. The stratal division of the lexicon. Linguistics 1.464–532.

SANDE, HANNAH. 2019. A unified account of conditioned phonological alternations: Evidence from Guébie. Language 95.456–497.

SHIH, STEPHANIE. 2018. Learning lexical classes from variable phonology. Proceedings of AJL 2.

SHIH, STEPHANIE, and SHARON INKELAS. 2016. Morphologically-conditioned tonotactics in multilevel Maximum Entropy grammar. Proceedings of the 2015 Annual Meeting on Phonology. Washington, D.C.: Linguistic Society of America.

SIBANDA, GALEN. 2009. Vowel processes in Nguni: resolving the problem of VV sequences. Selected proceedings of the 38th Annual Conference on African Linguistics, 38–55. Somerville, MA: Cascadilla Proceedings Project.

SIMANGO, SILVESTER RON, and MAXWELL KADENGE. 2014. Vowel hiatus resolution in ciNsenga: An Optimality Theory analysis. Southern African Linguistics and Applied Language Studies 32.79–96.

SMITH, JENNIFER L. 2006. Loan phonology is not all perception: evidence from Japanese loan doublets. Japanese/korean linguistics, ed. by Timothy J. Vance and Kimberly Jones, vol. 14, 63–74. Stanford: CSLI Publications.

STERIADE, DONCA. 1995. Markedness and underspecification. The handbook of phonological theory, ed. by John Goldsmith, 114–174. Oxford: Blackwell.

TANAKA, YU. 2017. The sound pattern of Japanese surnames. Los Angeles, CA: UCLA dissertation.

TEMKIN MARTINEZ, MICHAL, and VANESSA ROSENBAUM. 2017. Acoustic and aerodynamic data for Somali Chizigula stops. Africa's endangered languages: Documentary and theoretical approaches, ed. by Jason Kandybowicz and Harold Torrence. New York: Oxford University Press.

TESAR, BRUCE, and PAUL SMOLENSKY. 1998. Learnability in Optimality Theory. Linguistic Inquiry 29.229–268, ROA-155.

TESAR, BRUCE, and PAUL SMOLENSKY. 2000. Learnability in Optimality Theory. Cambridge, MA: MIT Press.

TSE, HOLMAN. 2014. The emergence of retroflexion in Somali Bantu Kizigua: internal motivation or contact-induced change? Unpublished Ms. University of Pittsburgh. Available at http://d-scholarship.pitt.edu/25989/.

TSE, HOLMAN. 2015a. The diachronic emergence of retroflexion in Somali Bantu Kizigua: Internal motivation or contact-induced change? Selected proceedings of the 44th Annual Conference on African Linguistics, ed. by Ruth Kramer, Elizabeth C. Zsiga, and One Tlale Boyer, 277–289. Somerville, MA: Cascadilla Proceedings Project.

TSE, HOLMAN. 2015b. The role of shift-induced interference in the development of a typologically rare phonological contrast in Somali Bantu Kizigua. Travaux Interdisciplinaires sur la Parole et le Langage (TIPA) 31.

ŁUBOWICZ, ANNA. 2005. Locality of conjunction. Proceedings of the 24th west coast conference on formal linguistics, ed. by John Alderete, Chung-hye Han, and Alexei Kochetov, 254–262. Somerville, MA: Cascadilla Proceedings Project.

VAN LEHMAN, DAN, and OMAR ENO. 2003. The Somali Bantu: Their history and culture. Culture Profile No. 16. Center for Applied Linguistics, Washington, DC.

WILLIAMS, STEVEN TYLER. 2012. Gender conflict resolution in Mushunguli. B.A. Thesis, The Ohio State University.

YEARLEY, JENNIFER. 1995. Jer vowels in Russian. Papers in Optimality Theory ii, ed. by Jill Beckman, Laura Walsh Dickey, and Suzanne Urbanczyk, University of Massachusetts occasional papers in linguistics, 533–571. Amherst, MA: GLSA Publications.

ZIMMERMANN, EVA. 2013. Vowel deletion as mora usurpation: the case of Yine. Phonology 30.125–163.

ZURAW, KIE. 2000. Patterned exceptions in phonology. Los Angeles, CA: UCLA dissertation.

ZURAW, KIE. 2010. A model of lexical variation and the grammar with application to Tagalog nasal substitution. Natural Language and Linguistic Theory 28.417–472.

ZYMET, JESSE. 2018. Lexical propensities in phonology: corpus and experimental evidence, grammar, and learning. Los Angeles, CA: UCLA dissertation.

ZYMET, JESSE. 2019. Learning a frequency-matching grammar together with lexical idiosyncrasy: MaxEnt versus Hierarchical Regression. Proceedings of the 2018 annual meeting on phonology, ed. by Katherine Hout, Anna Mai, Adam McCollum, Sharon Rose, and Matt Zaslansky.