Lexical Representation in Wubuy

Peter Christopher Nyhuis
ORCID: 0000-0002-4163-729X

Submitted in total fulfillment
of the requirements for the degree of
Doctor of Philosophy
in Linguistics

December 24, 2023

School of Languages and Linguistics
The University of Melbourne,
Wurundjeri country

Lexical representation in Wubuy

Peter Nyhuis

Abstract

Wubuy is a Gunwinyguan (non-Pama-Nyungan) language traditionally spoken in the lands around the southern part of Blue Mud Bay in Southeast Arnhem Land, Australia. Like other languages of the Gunwinyguan family, Wubuy is highly polysynthetic, with noun classification, productive noun incorporation, complex valency-changing operations, and obligatory realisation of up to two arguments on the verb. What distinguishes Wubuy from its neighbours is its intricate array of phonological alternations and the often extremely complex and opaque interactions between them. The result of this is a language in which words can be very large, consisting of many smaller elements, and each morphological element can have many different phonological realisations depending on its environment.

This raises the important question of how such complexity is represented in the minds of speakers. Do speakers store a minimal set of invariant atomic elements — i.e. morphemes, in the form of abstract underlying representations — and apply phonological operations to them during production? Or do they store words as wholes, in their surface phonological forms, and make generalisations about the relations between them? Mainstream generative phonology has almost always assumed the former (e.g. Chomsky & Halle 1968; Kenstowicz 1994), while the latter is espoused by much recent work in morphology (e.g. Blevins 2006; Booij 2010).

This thesis uses data drawn from text corpora and my own field elicitation to analyse several areas of Wubuy phonology in detail, including: 'non-optimising' empty morph insertion, sonority-conditioned reduplicant shape allomorphy, and several opaque alternations involving putative 'archiphonemes' and other abstract segments. In each case, I show that the analytical tools supplied by mainstream generative phonology fail to provide an adequate account of the facts, and that arriving at a solution requires a serious reckoning with the limitations of many of its underlying assumptions. Building on Jackendoff & Audring's (2020) theory of Relational Morphology, as well as on other currents in the tradition of Construction Grammar, I defend a model of phonology without morphemes, and without underlying representations. In this model, there is no ontological separation between representations and rules (or constraints). The lexicon stores whole words (and larger units) in their surface forms, as well as general schemas that abstract over words that have parts in common. These schemas take over the functions normally attributed to rules or constraints, and phonological alternations are encoded as relations of similarity and difference between full surface word forms.

In addition to these theoretical concerns, the thesis significantly extends our understanding of some key empirical phenomena in Wubuy. For one, I present evidence that there is widespread variation in the operation of an opaque deletion process, and that the variation is structured in a way suggestive of a sound change in progress, spreading by lexical diffusion. Further, I present the first detailed account of Wubuy prosody, by means of an Autosegmental-Metrical analysis of a corpus of unscripted narrative texts. In line with work on other languages of the region with complex morphologies, I find that intonational pitch events tend to align with the left edges of major morphological constituents in the verb. I also find evidence of a pitch accent that aligns to the right edge of the verb, and does so in a way indicative of quantity-sensitivity: it anchors to the final syllable of the word if it is heavy, and to the penultimate otherwise. This thesis thus substantially enriches our understanding of phonology-morphology interactions, in the context of a language where both the phonology and the morphology occupy a space close to the upper bounds of complexity.

Declaration

This is to certify that

(i) this thesis comprises only my original work towards the degree of Doctor of Philosophy;

(ii) due acknowledgement has been made in the text to all other material used;

(iii) the thesis is less than 100,000 words in length, exclusive of tables, maps, bibliographies and appendices.

Peter Nyhuis
ORCID: 0000-0002-4163-729X
December 24, 2023

Acknowledgements

Learning is nothing if not a collective practice, and one of the most rewarding things about this PhD project has been the many generous and knowledgeable people I got to know as a result.

Without question, my first thanks must go to the many Wubuy teachers who shared their language with me and brought me up in their world, as well as all the people in Numbulwar who supported this project in ways big and small. In particular, I must give special mention to two of my consultants, Hilda Ngalmi and Jangu Nundhirribala, who were also my colleagues at Numbulwar School. Hilda is an intelligent and enthusiastic teacher, and I feel honoured to have worked beside her in the often exhausting work of language revitalisation — and, even better, to have her as my *gaagu*, a role she takes on with more adroitness than anyone: her genius for crafting bespoke insults for every occasion, across five languages, has kept me chuckling for years. Jangu is a walking encyclopaedia of Nunggubuyu history and culture, a great storyteller, and is also one of the kindest and most unflappably cheerful people I've ever met. Her warmth made Numbulwar feel like a second home. She also recruited many new consultants for my work, making fieldwork run much more smoothly as a result. My other Wubuy consultants have been no less generous in sharing their time, their knowledge, their hospitality, and their enthusiasm for their language: thank you to †Leonie Murrungun, Jijila Turner, Mara Murrungun, Nanette Murrungun, Yidangka Ngalmi, David Murrungun, Reg Numamurdirdi, Susan Murrungun, Florence Murrungun. Thanks also to all the younger mob that brought me out to *wungubal* and encouraged my learning.

At the University of Melbourne, I have benefited from a supervision committee of the highest calibre. Brett Baker nudged me to take my first steps along the path that led to this project, and took me on my first field trip to Numbulwar. I would not be the linguist that I am without the benefit of his sharp attention to detail, his extensive knowledge of virtually every corner of linguistic theory, his ability to recall specific bits of data from languages all over Arnhem Land, and his encouragement to run with ideas he was initially skeptical of, to see where they land. John Mansfield came onto my supervision panel midway through the project, but by that point he had already made a big impact on my work through the many thoughtful conversations we'd had together in the Babel tearoom over lunch. As a supervisor, he offered much-needed encouragement, and gave detailed and perceptive feedback that greatly improved the final result.

Janet Fletcher is a brilliant teacher and her enthusiasm for phonetics is infectious. She offered crucial guidance on analysing the prosodic data, and helped me keep my eyes on the prize when I was lost in the weeds. At several junctures, Rachel Nordlinger injected pointed and timely interventions that helped with the nuts and bolts of writing a thesis, as well as plenty of fun banter.

I am very fortunate to have been surrounded by many talented and inspiring fellow linguistics PhD students over the last several years. Sasha Wilmoth and Katie Jepson deserve special mention for being both invaluable friends and comrades through some difficult times, keeping me sane when things could have gone otherwise, and for discussing with me practically every aspect of the present thesis at one time or another, as well as practically every other domain of linguistics not covered here. The Language Documentation Lab has been a convivial, intellectually stimilating and fun environment owing to the presence of Vanya Kapitonov, Lucy Davidson, Jonathan Moodie, Marco Espinoza, Katie Bicevskis, Tula Wynyard, Kate Charlwood, David Felipe Guerrero Beltrán, Eleanor Jorgensen, Giovanni Ma, and Phil Dwerk. The Phonetics Lab provided good humour, thoughtful discussion and chunks of R code, for whom I thank Rosie Billington, Josh Clothier, Alex Kilpatrick, Catalina Torres Orjuela, and Eleanor Lewis. Hywel Stoakes saved my skin with technical hurdles on untold occasions, and imparted his wisdom on the acoustic and articulatory phonetics of Arnhem languages. Tearoom chats with Reuben Brown, Bill Forshaw, Tom Watson, Debbie Loakes, and Ruth Singer offered important points of contact with other applications of linguistics. Beyond the Babel building, I am grateful for the feedback and encouragement of Rikke Bundgaard-Nielsen, Andrew Tanner, Tom Ennever, David Osgarby, Jacqui Cook, Mitch Browne, Mike Proctor, Jayden Macklin-Cordes, Alice Gaby, Mark Harvey, Mae Carroll and Nick Evans.

My first solo field trip would not have been possible without the material support of Helen Flanders at Numbulwar School, who found me accommodation in exchange for helping out with the Wubuy classes. That hands-on experience ignited my interest in bilingual education, and landed me an invitation to come back to fill in for her when she went on leave. This was a life-changing experience in several ways, but most importantly in giving me a much deeper knowledge of Wubuy and the Numburindi community than I ever could have hoped to get from research alone. For being an incomparable mentor in that role, I thank Chris O'Neill, as well as Ailsa Purdon back at HQ.

Also in the NT, I have had the good fortune to be able to draw on the expertise of a number of talented and committed linguists. On my first solo field trip, Jackie van den Bos at Ngukurr Language Centre helped me get my bearings in the Roper universe, and got me up to speed on Kriol jokes. Dean Austin Bara, *na-murruyung*, enthusiastically brought me into the Wubuy family and made sure I understood and participated in the culture; I owe much to his dedication

to keeping language and song alive. Celeste Humphris, a fine linguist and all-round good egg, reminded me why I loved the work we do, at a time when I nearly forgot. Melanie Wilkinson has been invaluable both as a mentor in navigating the ups and downs of working in bilingual education, and as a sounding board on all manner of fine-grained points of language documentation and analysis. Brian Gray taught me a great deal about psycholinguistic dimensions of phonology, and was a valued interlocutor on the theory and practice of Aboriginal languages in education. I also thank David Moore for some enlightening discussion of the big questions of linguistics, and Marg Carew and Rebecca Green for being supportive colleagues and sharing their knowledge of Maningrida languages.

In other parts of the world, thanks to audiences at the Annual Meeting on Phonology in 2019 and 2020, and to Raphael Winkelmann for assisting with his indispensable `emuR` software, and for quickly squashing the bugs I found in it.

Closer to home, I wish to thank all the friends and family who have learned not to ask how the thesis is going, and who have kept my soul well-fed and my spirit engaged in the world beyond linguistics. While I cannot acknowledge everyone who contributed in this way, thank you to Adam, Ainsley, Carolyn, Dom, Finn, Javed, Jerome, Lauren, Luqmaan, Mac, Margaux, Nadia, Rohan, Thomas, and my beloved dog Hector. Thanks to Sarah for her generosity, patience and unparalleled good humour. Thanks also to my parents, my sister, my step-siblings, and especially Carolyn and Gareth for all those science and stats chats. If you're reading this and think any of the above lists should include you, you're probably right.

I acknowledge also the Australian Government Research Training Program Scholarship that supported me throughout my candidature, the Faculty of Arts at the University of Melbourne for the fieldwork grant that took me to Numbulwar and paid my consultants, and a grant from the School of Languages and Linguistics that allowed me to attend the 2019 AMP conference in New York.

This thesis was mostly written on country belonging to the Wurundjeri people of the Kulin nation, and I pay my respects to their elders, past and present. I acknowledge their continuing care for country, and their ongoing efforts to restore their language to it. Sovereignty has never been ceded.

Contents

Abbreviations **xii**

List of Figures **xiv**

List of Tables **xvi**

1 Introduction **1**

1.1 The phonology of complex words in Wubuy 1

1.2 Overview of the thesis 4

1.3 Theoretical background 7

1.3.1 A lexicon of minimal, invariant morphemes 7

1.3.2 The lexicon as a network of constructions 11

1.3.3 Relational Morphology 17

1.4 Data collection 22

2 An overview of the Wubuy language **25**

2.1 Wubuy, its speakers, and its country 25

2.2 Morphology 28

2.2.1 Nominal morphology 29

2.2.2 Verbal morphology 33

2.2.3 Adjectival morphology 40

2.2.4 Closed classes: pronouns, demonstratives, kin terms, particles 42

2.2.5 Clitics, 'postpositions' 42

2.3 Phonology 43

2.3.1 Segmental inventory 44

2.3.2 Productive alternations 46

2.3.3 Lexically-restricted alternations 60

2.3.4 Rules which assume a maximally abstract analysis 62

2.3.5 Prosody 63

3 Alternations in the lexicon: ŋu-Epenthesis and Hardening 65
3.1 Introduction . . . 65
3.2 Descriptive Overview . . . 67
3.2.1 ŋu-Epenthesis . . . 67
3.2.2 Hardening . . . 70
3.2.3 Continuants undergo Hardening, stops trigger ŋu-Epenthesis . . . 74
3.2.4 Linking ŋu- . . . 81
3.2.5 Linking *ŋu-* as compounding morphology . . . 92
3.2.6 Minimal pairs . . . 95
3.2.7 Summary . . . 101
3.3 Theoretical issues . . . 102
3.3.1 Where lies *ŋu-*? . . . 102
3.3.2 Phonological epenthesis, natural and unnatural . . . 103
3.3.3 Empty morphs and allomorphy in an optimising grammar . . . 106
3.3.4 Opacity . . . 109
3.3.5 Outwards-sensitivity . . . 116
3.3.6 Constructive and abstractive morphotactics . . . 117
3.4 An abstractive analysis of *ŋu-* . . . 120
3.4.1 Empty morphs and the Paradigm Cell Filling Problem . . . 120
3.4.2 Phonological classes . . . 123
3.4.3 Uncertainty reduction and linking *ŋu-* . . . 128
3.5 Relational Morphology . . . 131
3.5.1 Introduction . . . 131
3.5.2 Words, schemas, and subcategorisation . . . 132
3.5.3 Phonologically-conditioned allomorphy, transparent and opaque . . . 140
3.5.4 Summary . . . 145
3.6 Analysing ŋu- in Relational Morphology . . . 146
3.6.1 Prefixal ŋu- . . . 146
3.6.2 Linking ŋu- . . . 153
3.6.3 Nasals and other non-alternating sonorants . . . 155
3.7 Conclusion . . . 157

4 Sonority-conditioned reduplication 160
4.1 Introduction . . . 160
4.2 Descriptive overview . . . 162
4.2.1 Functions of reduplication . . . 163

4.2.2 Sonorant-initial disyllabic reduplication and stop-initial monosyllabic reduplication . . . 169
4.2.3 Reduplicant position . . . 184
4.2.4 Reduplication with Hardening and ŋu-Epenthesis . . . 190
4.2.5 Reduplication of vowel-initial stems . . . 192
4.2.6 Summary: reduplicant shape alternations in Wubuy . . . 202
4.3 The emergence of sonority-conditioned reduplication . . . 203
4.3.1 Reduplication in Anindilyakwa . . . 204
4.3.2 Reduplication in other Gunwinyguan languages . . . 206
4.3.3 Reduplication, lenition, and informativity . . . 211
4.4 Sonority-conditioned reduplicant shape in Relational Morphology . . . 214
4.4.1 Orientations . . . 214
4.4.2 Reduplication in Relational Morphology . . . 216
4.4.3 Existing analyses of reduplication in RM . . . 217
4.4.4 Reduplication of consonant-initial verb stems in Wubuy . . . 222
4.4.5 Reduplication of vowel-initial verb stems . . . 227
4.4.6 Summary of RM analysis . . . 231
4.5 Comparison with other models . . . 232
4.5.1 Base-Reduplicant Correspondence Theory . . . 232
4.5.2 Unifying the rankings . . . 237
4.5.3 Morphologically indexed constraints . . . 239
4.5.4 Making the alternation optimising . . . 242
4.5.5 Resorting to parochial constraints . . . 244
4.5.6 Summary of BRCT . . . 248
4.5.7 Morphological Doubling Theory . . . 250
4.5.8 Summary of MDT . . . 257
4.6 Conclusion . . . 258

5 Variability and sound change in the lexicon: velar stop deletion 260
5.1 Introduction . . . 260
5.2 Background . . . 261
5.2.1 Stop Deletion, Stop insertion, and Hardening . . . 261
5.2.2 The archiphoneme /K/ . . . 266
5.2.3 Formalising deletion in RM . . . 269
5.2.4 Variability . . . 271
5.2.5 Parallels with velar nasal deletion and assimilation . . . 277

5.3 Dorsal (co-)articulation in Australia and elsewhere . . . 279
5.3.1 The weakness of coda consonants . . . 279
5.3.2 VC transitions in Australian languages . . . 282
5.3.3 The special weakness of dorsals . . . 284
5.3.4 Informativity effects . . . 286
5.3.5 Summary . . . 287
5.4 Phonetic properties of dorsal stops in deletion contexts . . . 288
5.4.1 Introduction . . . 288
5.4.2 Method . . . 289
5.4.3 Results and discussion . . . 297
5.5 Discussion . . . 313
5.5.1 Explaining the results . . . 313
5.5.2 Frequency, lenition, and the automatisation of articulatory routines . . . 314
5.5.3 Relative frequency of word-final vs. word-nonfinal velar stops . . . 316
5.5.4 Deletion in native noun roots in nouns and verbs . . . 319
5.5.5 Limitations . . . 320
5.6 Conclusion . . . 320

6 Prosody and morphological complexity 326
6.1 Introduction . . . 326
6.2 Background . . . 328
6.2.1 Prosody, prominence, and rhythm . . . 328
6.2.2 Phonological accounts of prosody in Australian languages . . . 330
6.2.3 Phonetic correlates of stress in Australian languages . . . 331
6.2.4 Intonational phonology of Australian languages . . . 334
6.2.5 Intonation and morphological complexity in Arnhem Land . . . 336
6.2.6 Previous work on Wubuy prosody . . . 341
6.3 Aims, methods and materials . . . 347
6.3.1 Hypotheses concerning verbal prosody in Wubuy . . . 347
6.3.2 Corpus materials and transcription . . . 348
6.3.3 AM annotation procedure . . . 348
6.3.4 Querying and analysis . . . 351
6.4 Results: Intonational tune types . . . 351
6.5 Results: Alignment of tune to text . . . 360
6.5.1 Pitch accents in the pronominal prefix . . . 365
6.5.2 Stem-initial accent . . . 368

6.5.3 The separability of *ma* within the pronominal prefix 371
6.5.4 Accents within the stem 372
6.5.5 Right edge accents 372
6.6 Discussion 382
6.7 Conclusion 384

7 Conclusion **386**
7.1 What does the Wubuy lexicon sound like? 386
7.2 Directions for future research 390
7.3 Concluding remarks 391

Bibliography **392**

Abbreviations

Abl	ablative
All	allative/dative
Anaph	anaphoric (demonstrative)
Appl	applicative
Aug	augment
Aux	auxiliary
Caus	causative
Centrif	centrifugal
Centrip	centripetal
Coll	collective
Comit	comitative
cons	consonantal
cont	continuant
Contrs	contrastive
Dist	distal
Distr	distributive
du	dual
Evit	evitative
Fact	factitive
Fem	feminine
Foc	focus
Hum	human
Ident	identity
Imm	immediate
Inch	inchoative
Instr	instrumental
Irr	irrealis
Loc	locative
Masc	masculine
Max	maximality (constraint)
Mult	multiple
NC	noun class
Neg	negative
Neut	neuter
Npst	non-past
Num	number
Obl	oblique
Orig	originative
Past	past
Perg	pergressive
Pers	person
pl	plural
Prox	proximate
Purp	purposive
Q	interrogative
R	realis
Recip	reciprocal
Red	reduplicant
Refl	reflexive
Rel	relative
Resid	residual
RPerg	retrospective pergressive
sg	singular
Simil	similative
son	sonorant
syll	syllabic

Top	topic
Veg	vegetable

List of Figures

2.1 Gunwinyguan language family map 26

5.1 Labelled waveform of *a-jinak-ŋu-cuɻaŋ* ‘push my head’ or ‘push me with your head’, from Bundgaard-Nielsen & Baker (2020) 275
5.2 Spectrogram of *ama-maṯaɭak-tuc* ‘at the beach’ 276
5.3 Spectrogram and waveform for *a-ɻaŋa(k)-tuc*, no [k] 295
5.4 Spectrogram and waveform for *a-ɻaŋa(k)-tuc*, with [k] 295
5.5 Spectrogram and waveform for *jiː-wiːk-tuc* ‘Masc.Obl-grey.hair-Loc’ 296
5.6 Rates of deletion for /k/ in native roots and loanwords 299
5.7 Rates of deletion for /k/ and /K/ preceding stops (other than /k/) in verbs . . . 299
5.8 Duration of (underlyingly) /k-t/ clusters in native roots and loanwords, by whether /k/ is deleted or retained. 302
5.9 Durations of intercontinuant stops and stop sequences. 304
5.10 Durations for stops and stop sequences between continuants in verbs 305
5.11 Durations for k~K-ṯ, following just the low vowels /a/ or /aː/ 307
5.12 Averaged, time-normalised trajectories of F1, F2 and F3 in vowels preceding /tV/ and /kV/ in nouns 308
5.13 F2 frequency in Hertz at 80% of vowel duration, preceding intervocalic /t/ and /k/. 309
5.14 Locus equations (Sussman et al. 1991) for /t/ and /k/ in VC transitions 310
5.15 Trajectories of F1, F2 and F3 in vowels preceding underlying /k-t/ in nouns, with and without deletion 311
5.16 Locus equations for VC transition into /k-t/ cluster 312
5.17 The /a/ subset of the data in Fig. 5.16. 312
5.18 Proportion of deleting and non-deleting tokens of the same noun root 317

6.1 Example of unaccented material before the verb. 352
6.2 Declarative ‘hat’ tune in two APs 354
6.3 Simple rise-fall (LHL) pattern on *wu-mawuɻaːti* ‘it’s cold’. 355
6.4 IP-unique token of *nuɾu* ‘1pl’, with final late rise at LH% boundary. 356

6.5 IP ending in late rise at LH% boundary: *anubanila-wala nuɾa:kina* 'from along there, we came back' . . . 357
6.6 IP token showing upstepped high boundary (^H%) . . . 358
6.7 High plateau tune, followed by final AP with rapid fall to L% boundary. . . . 359
6.8 Example of IP consisting of 3 APs, each with one verb . . . 359
6.9 Alignment of grammatical verbs to either AP or IP. . . . 361
6.10 Heatmap visualisation of contingency table, showing alignment of grammatical verbs to prosodic constituents . . . 362
6.11 Number of pitch accents per grammatical word. . . . 363
6.12 Number of pitch accents by word length, in syllables. . . . 363
6.13 Distribution of pitch accents at different positions of the verb. . . . 364
6.14 IP-unique token of *niɾima-ɻiɾi:* 'We dried it out [pandanus leaves]', with delayed word-initial H*, notated <. . . . 366
6.15 Frequency of pitch accents on either the first or second syllable of the pronominal prefix . . . 367
6.16 Verb tokens with stem-initial and other accents . . . 369
6.17 Locations of pitch accents among verb tokens with a stem-initial and/or other accent . . . 369
6.18 Token of *wa:-jila-ŋi* 'they sliced it (Coll)', with penultimate pitch accent. . . . 373
6.19 Token of *naɾa-ŋuni* 'we ate it (Coll)', with penultimate pitch accent. . . . 374
6.20 Token of *waɾa-l̪alwu̪la-ŋi* 'they soaked it (Coll)', with penultimate pitch accent. 375
6.21 Token of *waɾa-jami-cka:* 'they told them', with final pitch accent. . . . 376
6.22 Token of *niɾima-wul̪-i:* 'we cut it (Veg)', with final pitch accent, and accent on *-ma-* in prefix. . . . 377
6.23 Token of *wuɾ-a:kamp-a:* 'they roasted it', with final pitch accent. . . . 378
6.24 Token of *wu:-ɻama-ɻam-a:* 'they were abundant, everywhere', with final pitch accent. . . . 379
6.25 Distribution of stem-final pitch accent, by final syllable weight. . . . 380
6.26 Distribution of stem-penultimate pitch accent, by final syllable weight. . . . 380
6.27 Frequency of pitch accents at the right edge of the stem (penultimate or final) by prosodic juncture . . . 381
6.28 Number of pitch accents by number of prosodic words within verb token. . . . 383

List of Tables

2.1 Inflectional noun class prefixes 30
2.2 Derivational noun class prefixes 32
2.3 Case suffixes 33
2.4 Interactions between prefixal mood inflection and suffixal tense-aspect-mood inflection 36
2.5 The consonant inventory of Wubuy, adapted from Heath (1984, p. 12). Segments in slashes are transcribed in IPA, their counterparts in italics are the symbols used in the community orthography. 45
2.6 Vowel inventory 45
2.7 Contrasts in manner features 45
2.8 Productive phonological rules in Heath (1984) 47
2.9 Continuant-stop pairs 48
2.10 Summary of productive phonological rules 59
2.11 Comparison of surface forms of /api/ 'jump', /ɻa/ 'spear (V)' and /puɾa/ 'sit' in different preceding segmental contexts. 60
2.12 Lexically restricted phonological rules in Heath (1984). 61
2.13 'Maximally abstract' phonological rules in Heath (1984) 62

3.1 Minimal pairs distinguished by initial stop vs. continuant 97

4.1 Reduplicant shape, sonority of the base-initial segment, and skeletal composition of the base. 180
4.2 Necessary conditions for prefixal *ŋu-* and reduplicated linking *ŋu-* 192
4.3 Derivations illustrating interactions between reduplication and other processes 193
4.4 Main reduplication patterns for Anindilyakwa 206

5.1 Derivations illustrating Hardening and Velar Stop Deletion 263
5.2 Derivations illustrating Insertion and Deletion of /k/ 266
5.3 Surface products of underlying sequences involving /k/ 268
5.4 Stimulus items used in noun elicitation task 290

5.5 Derivational elements (prefixes or incorporable noun roots) and verb roots used to construct stimuli in verb elicitation task. 292
5.6 Summary of results of fitting a generalised linear mixed effects model 301
5.7 Token frequencies in NMET (Heath 1980), separated by whether the root appears word-finally or not. 317
5.8 Morphs analysed as /K/-final in Heath (1980) 325

6.1 AM labelling scheme 349
6.2 Frequency table for Intonational Phrase tune types in the corpus 353
6.3 Pitch accents, grammatical words and Accentual Phrases per Intonational Phrase 353
6.4 Result of fitting generalised linear model (binomial) on the data presented in Figure 6.15. 367
6.5 Results of logistic regression fitted to the data in Figure 6.27 381

1 Introduction

1.1 The phonology of complex words in Wubuy

Wubuy ['wubuj] is a Gunwinyguan (non-Pama-Nyungan) language traditionally spoken in the lands around the southern part of Blue Mud Bay in Southeast Arnhem Land, Australia. Like other languages of the Gunwinyguan family, Wubuy is highly polysynthetic, with noun classification, productive noun incorporation, complex valency-changing operations, and obligatory realisation of up to two arguments on the verb. What distinguishes Wubuy from its neighbours is its intricate array of phonological alternations and the often extremely complex and opaque interactions between them. The result of this is a language in which words can be very large, consisting of many smaller elements, and each morphological element can have many different phonological realisations depending on its environment.

The morphological and phonological exuberance of Wubuy raises two questions that have long been at the very heart of linguistic inquiry: how are linguistic constructs stored in memory, and how do speakers know how to form new linguistic constructs and be understood? The verbal morphological template contains several slots for roots, prefixes and suffixes, some of which are obligatorily filled, some of which are not. Meanwhile, each slot can be filled by a vast array of possible elements; the pronominal inflectional prefixes alone number in the hundreds (of course, a more precise figure depends on how one chooses to count them). Further, each of these elements takes multiple forms, depending on what appears alongside it, due to the various phonological alternations that apply at each morphological juncture within the word. Taken together, the combinatorial possibilities are staggering. This means that encountering a word one has never heard or spoken before is highly likely, just as it is highly likely that in a language like English, one is likely to encounter completely novel sentences with great frequency. Take the Wubuy words in (1), each of which appears in Heath's (1980) text corpus only once.[1]

(1) a. wuɾu-munṯu-mun-ṯaɾm-aɲi
3pl>3Resid.R-Red-foot-follow-Past2
'They followed its foot' (NMET 44.5.2)

[1](1c) appears once and is then immediately repeated; it does not appear elsewhere in the texts. The total size of this corpus is about 20,835 words.

b. niwu-jika-l̪aɾm-iɲ
3sgM>3Neut.R-fire-follow-Past1
'He followed the fire' (NMET 16.11.1)

c. wuɾaː-'-cika-l̪aɾm-aɲc-iːni
3pl.R-Appl-fire-follow-Recip-Past2
'They followed each other by fire (seeing each others' fires in the distance)' (NMET 46.1.1)

d. wiŋi-ʈ̪iɾit-aŋi
3duF-descend-Npst2
'The two women went down' (NMET 25.4.4)

e. wuɾaː-ɲci-ŋu-ʈ̪iɾit-aŋi
3pl.R-Comit-Ø-descend-Npst2
'They went down with it [a load of carcasses]' (NMET 59.10.2)

f. waː-ŋkaː-ŋkaɾa-ŋu-puɾ-i
3Resid.R-Red-Mult-Ø-sit-Past2
'They [nonhuman] sat together' (NMET 36.3.1)

g. wuɾu-maka-ɲc-iːni-juŋkujuŋ
3pl.R-tell-Recip-Past2-Purp
'In order to tell each other' (NMET 46.2.1)

h. wuː-l̪aka-l̪akaɾa-ŋaw-iː-ni-jiɲuŋ
3pl.R-Red-Mult-die-Refl-Past2-Rel
'They who all died' (NMET 47.9.2)

Hapaxes though they are, the surface word forms in (1) clearly aren't *completely* novel. The Past2 suffix *-aŋi* appears in three of these examples. There is a root glossed 'follow' in (1a,b,c), although it takes different forms: *ʈ̪aɾm* and *l̪aɾm*. Indeed, all of the morphs in these examples are well attested outside of these particular tokens, and there is nothing to suggest that the forms they take in these word contexts are in any way unexpected.

Any plausible theory of grammar must be able to express the fact that *l̪aɾm* and *ʈ̪aɾm* are 'the same' at some level of analysis, given their total identity in meaning and their near-total identity in form. But it must also account for what is different between the two — i.e. the manner of articulation in the first segment — and where we expect to find one alternant and not the other. Further, it must also be able to account for the difference between the behaviour

of these allomorphs and that of another verb root, such as *t̪iɾit* 'descend' in (1a) and (1b), which also begins in the dental stop /t̪/ but has no corresponding form **l̪iɾit*. We also need some characterisation of where in the word we expect to find these elements.

The answer provided by generative phonology is that the mental lexicon contains individual pieces in invariant form (morphemes), which are assembled together to form a word, whose form is then readjusted according to general well-formedness criteria. As I demonstrate in this thesis, these claims cannot grasp the complexities of Wubuy phonology and morphology. Rather, they lead inevitably to representations which are too abstract, rules which are too unnatural, and rule interactions which are too complex.

These issues are exhibited in Heath's (1984) descriptive grammar of Wubuy, the most comprehensive source we currently have on Wubuy phonology. For instance, there are claimed to be two labiovelar glide phonemes, /w$_1$/ and /w$_2$/, which are phonetically identical but participate in different alternations. There are stops that alternate with continuants at the same place of articulation, and stops that are always stops, as well as some continuants that are always continuants. There is an abstract 'archiphoneme' /K/, which never appears on the surface, and can only be inferred by the alternations it triggers. There is an empty morph, with the form /ŋu/, which is phonologically conditioned, but in a way that does not serve any obvious phonological aim. In the productive reduplication pattern, the reduplicant copies either a single light monosyllable or a disyllabic sequence, depending on whether the base begins in a stop or a sonorant — again, the alternation here is phonologically conditioned, but there is no sense in which it is phonologically optimising. Moreover, all of the processes just mentioned, which account for only a handful of the 50 transformational rules posited by Heath (1984), frequently interact with one another in the same word. Those interactions, more often than not, are opaque.

Each one of these phenomena, on its own, poses significant obstacles to a satisfactory analysis within mainstream approaches to phonology. But the real genius of Wubuy lies in the fact that all of these problematic phenomena interact together, in a single language, and in the context of a morphological system at the upper bounds of complexity. Wubuy thus offers an extraordinarily stringent test for the predictions of any theory of phonology, morphology, and their interface, since whatever approach is taken for one part of the system must be able to account for the whole.

This study aims at a comprehensive description and analysis of the realisation of complex words in Wubuy, encompassing their morphotactics, their segmental phonology, and their prosodic organisation, and paying particular attention to the ways that all of these interact. To that end, it adds new data based on my original fieldwork, as well as new analysis of existing data. In each of the main analysis chapters, the focus is on adequately describing the phenomena, understanding how they fit within the broader linguistic system, situating the problems they

pose within major debates in linguistic theory, and proposing analyses of my own to overcome the problems identified. While much of the argumentation is formal, I am interested throughout in engaging with strands of linguistics outside of formal theory, particularly the experimental psycholinguistic literature and the usage-based tradition.

I therefore expect that the work will interest not only Australianists, or phonologists, or those who share my particular theoretical inclinations; rather, I take care to show how each of the problems addressed speaks to much broader concerns within linguistics as a whole — and indeed, how linguistic theory cannot afford to ignore the kinds of structures we find in Wubuy. In doing so, I go some way to addressing the fact that — despite the ongoing centrality of older data from Australian languages in theoretical argumentation, in phonology especially — detailed studies of Australian phonological systems have become increasingly rare over the last couple of decades, as Baker (2014b) points out (some exceptions are Baker 2008b; Mansfield 2019; Panther 2021; Round 2009). Given the critical endangerment of many of these languages, including Wubuy, crucial questions, both for analysts and for younger heritage speakers who wish to reclaim their language, may never be answered if descriptive work is not undertaken now.

The remainder of this chapter is organised as follows. The next section, §1.2, provides a condensed summary of the entire thesis. In §1.3, I review the debates motivating this thesis' theoretical outlook, attempting to characterise the state of contemporary phonology, particularly in its relationship to newer developments in morphological theory. There, I also introduce Relational Morphology, which plays a central role in the analyses of the main chapters.

1.2 Overview of the thesis

Following this introductory chapter, I give a general overview of what Wubuy is like in Chapter 2. I provide some brief commentary on the social and linguistic fabric in southeast Arnhem Land, of which Wubuy is a part, before reviewing the current state of our knowledge on the morphology and phonology. Given the complexity of these domains of Wubuy grammar, this general overview provides invaluable context for the much more detailed treatments of particular phenomena that make up the rest of the thesis.

This is followed by the main analysis chapters. While each of them aims at both enriching the description of the phenomena at issue and working through their theoretical implications, the balance between these two demands varies: half concentrate mainly on formal issues (Chapters 3 and 4), and the other half (Chapters 5 and 6) make a primarily empirical contribution. In all of them, the analysis I offer differs almost completely from what has been presented in previous literature, particularly that of Heath (1984); whereas his grammar offers a chapter-length general

overview of phonological alternations and some brief, tentative discussion of prosody, this thesis goes into far greater detail on a narrower set of topics, and makes use of the fruits of the four decades of linguistic theory and methodology that have followed that grammar's publication.

At the heart of the thesis, especially Chapters 3 and 4, is a formal analysis couched in Relational Morphology (Jackendoff & Audring 2020). A relatively new theory of language, it is rooted in Construction Grammar (Bybee 2006; Croft 2001; Fillmore et al. 1988; Goldberg 1995; Jackendoff 2002), particularly Construction Morphology (Booij 2010; Gurevich 2006), and therefore takes not the morpheme but the *construction* as the basic element of linguistic knowledge. At its simplest, a construction is a word: a pairing of meaning and form. A construction can be relatively concrete, corresponding to the memory trace of an actual token of experience, or relatively abstract, generalising over a set of such concrete tokens. Stemming from the logical primacy of concrete tokens, the form part of the Saussurean equation, the phonology, has only one level, namely the surface form; there are no underlying or intermediate representations as in generative grammar. A crucial innovation of Relational Morphology compared to its theoretical predecessors lies in its claim that the lexicon contains not only constructions but relations between them, expressed as links between the parts they hold in common. I exploit this feature to formalise phonological alternations as expressions of similarity and difference between related surface word forms. A significant result of this approach is that it allows for a fully surface-based account of phonological opacity, something which has failed to receive a satisfying treatment in Optimality Theory (Prince & Smolensky 2004), despite decades of attempts to do so (Baković 2007; Kager 1999; McCarthy 2007; Mielke et al. 2003; Pruitt 2023). With few exceptions, the place of phonology in construction-based grammar has been seriously under-theorised, and more complex polysynthetic morphology has largely been ignored. Thus, in addition to deepening our understanding of the Wubuy language, the present thesis makes a substantial contribution to the further development of the theory, extending its empirical coverage and clarifying its predictions. Later in this chapter, I review the fundamental theoretical concerns animating Relational Morphology, introduce the basic elements of the formalism, and preview the way I use it to address phonological alternations in subsequent chapters.

In Chapter 3 I apply the Relational Morphology formalism to the problem of ŋu-Epenthesis, a productive process that inserts the meaningless string /ŋu/ preceding a stop-initial morph when it follows either a nasal-final inflectional prefix or a derivational element of any form. There is another productive process, Hardening, which makes continuants harden to a stop at the same place of articulation, also following nasals (and stops). Stops derived by Hardening, however, do not trigger ŋu-Epenthesis. Generative theory generally has very little to say about empty morphs like these, and even less so when they do not optimise phonological well-formedness. Worse still, ŋu-Epenthesis is conditioned by a two-sided context, which is ruled out in most

generative theory, and each of those contexts involves opacity, of the 'overapplication' type on one side and 'underapplication' on the other. Drawing on recent Word and Paradigm morphology (Blevins 2006), I argue that *ŋu-* nonetheless serves an important function, so long as we analyse it as part of the whole words it occurs in: it makes the rest of a word's paradigm predictable based on a single cell of the paradigm. I then show that Relational Morphology allows us to express those paradigmatic relationships as allomorphy that subcategorises for the phonology of a related word form.

In Chapter 4, I examine reduplication, an extremely pervasive feature of Wubuy discourse. I show that the reduplicant alternates in size according to the sonority of the base-initial segment: CV- if it is a stop, CV(C)(C)CV- otherwise. Other than Wubuy's closest relative Anindilyakwa, I know of no other language where reduplicant size depends on segmental features of the base. Although the pattern is straightforward to describe, it turns out to present a number of significant theoretical challenges. Some of those challenges stem from opacity, much as we saw for ŋu-Epenthesis, but others are unique to reduplication. The pattern is impossible to derive in the dominant OT approach to reduplication, Base-Reduplicant Correspondence Theory (McCarthy & Prince 1995), as well as its descendants, insofar as they are committed to the view that all alternations are optimising, and alternations targeting prosodic shape must optimise prosody. Although reduplication is only given perfunctory attention by Jackendoff & Audring (2020) in their formulation of Relational Morphology, I show that its basic mechanisms are once again sufficient to lead us beyond the impasse: each reduplicant shape allomorph subcategorises for the phonology of the base it copies. The same links that are used to express shared parts of structure across related word forms are used in reduplication constructions to relate part of a word (the reduplicant) to another part of the same word (the part of the base it copies).

Chapter 5 also involves a phonological interaction which is classically opaque: morph-final velar stops trigger Hardening in a following continuant, but are then deleted, making Hardening appear unmotivated on the surface. According to existing documentation, Wubuy also possesses a 'stop archiphoneme' /K/, which shows the same behaviour. While this kind of problem can also be addressed by Relational Morphology, that is not the principal focus of the chapter. I present evidence that velar deletion is actually subject to variation, contra previous description, and that the variation is structured primarily according to the kind of morph it appears in, and secondly by what type of word that morph appears in. I argue that this can be explained as a sound change in progress, where it is the relative frequency with which a given morph appears word-finally versus word-nonfinally that determines its deletion profile.

Having given extensive (though certainly less than exhaustive) treatment of Wubuy's segmental alternations in the preceding three chapters, in Chapter 6 I turn to prosody. I review previous attempts to describe lexical stress in Wubuy, arguing that the nature of the prosodic system

remains elusive in part because these attempts fail to address the intonational phonology, which may ultimately be the source of observed prominence, as has been argued in a number of recent accounts of Australian prosodic systems. To remedy this, I conduct an Autosegmental-Metrical analysis of a corpus of texts. I find that pitch accent is attracted to the left edge of morphological constituents in Wubuy verbs, but is in each case optional. This much is in keeping with what has been described in related languages. Additionally, I find that there is also an optional pitch accent attracted to the right edge of major phrases, and it appears to be quantity sensitive: it docks to the final syllable if it is metrically heavy, and to the penult otherwise. While evidence like this is often invoked to motivate the existence of metrical feet, I argue that this is not a necessary conclusion, and the data are also compatible with a 'footless' analysis. In this, I follow other recent literature on prosodic constituency in morphologically complex languages (Gordon 2014, 2016; Mansfield 2023).

Chapter 7, the final one, presents my conclusions. It looks back over the findings of each chapter and how they contribute to the overall picture of the phonology and morphology of complex words in Wubuy. This, in turn, leads to a consideration of how that picture contributes to our understanding of language in general. Since a project of this kind is always but the beginning of a much larger research program, in this final chapter I suggest some directions that future work might take.

1.3 Theoretical background

This thesis is concerned with the phonology and morphology of complex words in Wubuy. It is therefore important to clarify what I mean by 'phonology' and 'morphology', since each of these are highly contested terms. To do that, I begin by distilling the essential features of what might be called the hegemonic position on these questions in §1.3.1. I then trace a series of interrelated challenges to that position in 1.3.2, leading up to the specific proposal of Relational Morphology in 1.3.3.

1.3.1 A lexicon of minimal, invariant morphemes

Returning to the main question of this thesis, what do speakers need to know in order to generate novel words like those in (1)? Vaux (2003) terms the two dominant answers to this question, which lie at opposite ends of the possibility space, the 'Morpheme-Invariant Theory' and what he calls, somewhat pejoratively, the 'Redundant Lexicon Theory'. The answers differ in their account of the trade-off between lexicon and grammar, or storage and computation, or words and rules; they also differ in what they see as the units of storage in the lexicon, and what role redundancy plays.

The Morpheme-Invariant Theory has enjoyed near hegemonic status in generative grammar. It assumes a categorical split between the duties of the lexicon and of the grammar, whereby anything that *can* be computed by the grammar *must not* be stored in the lexicon: "lexical storage is maximally efficient and therefore does not encode entire words; rather, it encodes their constituent morphemes, in an invariant and minimally redundant form" (Vaux 2003, p. 92).

This view packages together several interrelated assumptions. First, that the task of morphological analysis is one of identifying morphemes — that is, word-parts extracted from the word contexts they appear in — and the rules that combine them to make complete words. Such morphemes may take on different forms depending on the context they appear in; where such form variation is predictable, adherents of the Morpheme-Invariant Theory attempt to posit abstract 'underlying representations' from which the surface variants can be derived. This brings us to the second assumption, which is that the task of phonological analysis is to discover what these underlying representations are, and by what procedures the underlying representation is transformed into the surface realisation, from which it may differ substantially. Further, there is assumed to be a strictly serial relationship between the operations of morphology and those of phonology, where the former crucially precede the latter. Finally, while surface word forms may be compared with one another in order to deduce underlying representations and the rules that operate on them, once this deduction has been accomplished, relationships between words play no role in the analysis: surface words are derived in total isolation from other words.

The Morpheme-Invariant position is stated quite explicitly in Chomsky and Halle's (1968) *The Sound Pattern of English* (SPE), setting the tone for generative phonology in the decades to follow. Chomsky and Halle assert that regular variations in the phonological form of a morpheme "are not matters for the lexicon, which should contain only idiosyncratic properties of items, properties not predictable by general rule", and that a lexical entry "must contain no indication of the effect of context on the phonetic form" (p. 12). Their lexical entries, *qua* maximally nonredundant underlying representations of morphemes, are 'impoverished' in a crucial sense (Jackendoff 1975). For example, observing that all three-consonant syllable onsets in English begin with /s/, Chomsky and Halle go so far as to posit the underlying representation of *spray* as /Cpray/, so committed are they to excising all predictable information from the lexicon. Certainly, not every generativist is committed quite so wholeheartedly to removing all redundancy from the lexicon, but the assumption remains a crucial one in most generative theory.

Di Sciullo and Williams characterise this nonredundant lexicon, rather famously, as "like a prison — it contains only the lawless, and the only thing that its inmates have in common is lawlessness" (Di Sciullo & Williams 1987, p. 3). The 'laws' of the grammar were assumed in early generative phonology to consist of rewrite rules, ordered in series, which take an input string

(starting with the underlying representation) and produce an output. This output may then serve as the input to further rules, there being no limit on the number of rules intervening between an underlying representation and surface realisation — indeed, rules can apply repeatedly if part of a 'cycle'. Grammars are evaluated according to the generality and simplicity of the rules; in other words, the best grammar is the one that requires the fewest symbols (Chomsky 2015, p. 43).

A given linguistic unit may be 'listed' in the lexicon, or derived by the grammar, but not both. Therefore, units larger than the morpheme may be listed if they cannot be derived by the grammar — that is, because they are sufficiently idiosyncratic as to escape analysis in terms of generally-applicable rules. Thus 'listed' tends to be virtually synonymous with 'exceptional' in such work. However, opinions differ widely on what the appropriate threshold of generality is for a pattern to be formalised as a rule: various 'rules' in *The Sound Pattern of English* apply to only one or two words (a practice also followed in Heath's (1984) grammar of Wubuy).

The ascendance of Optimality Theory (OT; Prince & Smolensky 2004) within generative phonology constituted a seismic shift in the mechanisms employed to describe phonological alternations, eschewing transformational rules applied in series in favour of constraints evaluated in parallel. However, many of the more foundational assumptions of SPE remained in place. Certainly, in principle, the input to an OT grammar need not be an underlying representation (Burzio 2005), and some extensions of the theory have introduced other types of input (Alderete 2001; Benua 1997; Łubowicz 2003; McCarthy 2004; Wolf 2008) or have sought to problematise the notion of an abstract UR (Golston 1996; Kirchner 1997). But, in practice, the goal of the vast majority of OT phonological analysis has been to characterise the set of constraints and their ranking that will produce, from a given abstract underlying form consisting of one or more minimal invariant morphemes, the attested surface form (Kager 1999, p. 413; Pierrehumbert 1999; Van der Hulst & Ritter 2000). Indeed, Prince & Smolensky (2004, p. 5) explicitly equate 'input' with 'underlying representation' in their original formulation of the theory. That is because the engine at the heart of OT analysis is the antagonistic relationship between faithfulness constraints, which enforce identity between underlying representation and surface form (or base and reduplicant), and markedness constraints, which enforce conformity with putatively 'unmarked' structures. To the extent that input-output faithfulness constraints depend on underlying representations as their point of comparison against which to evaluate surface forms, it is difficult to conceive of classical OT without assuming a lexicon of invariant morphemes. Even in OT models that allow multiple stages of evaluation, such as Stratal OT (Bermúdez-Otero 1999; Kiparsky 2000), each new layer of morphology is introduced by adding another morpheme, in the guise of its invariant underlying representation, and then submitting this as the input to the EVAL component. In other words, an OT tableau nearly always functions to replace the

SPE rules mediating underlying and surface forms, while leaving the basic underlying/surface distinction intact, and with it the respective roles of lexicon and grammar. And, just as in rule-based serialism, the UR in OT can be remote from its surface realisation: McCarthy's (2003a) account of Barrow Inupiaq, for instance, assumes historically antecedent forms as the input to the OT grammar, to which he adds powerful and complex new machinery in order to recapitulate the historical changes that led to contemporary surface forms (see Mielke et al. 2003).

One major difference between OT and earlier SPE rule-based serialism, and indeed one of the chief motivations for developing OT in the first place, is that SPE phonology advances a purely formal characterisation of the language faculty, completely insulated from any substantive grounding in physical, psychological or biological constraints; according to some, this weakens its plausibility as a model of what is really going on in speakers' minds, or in their speech (Hayes 1999; Hooper 1976; Kager 1999; Lakoff 1993; Prince & Smolensky 2004). A 'natural' rule such as one that devoices obstruents word-finally, modelling a process known to be cross-linguistically widespread and straightforwardly relatable to phonetic pressures, is no better in SPE than an 'unnatural' one which voices all obstruents word-finally, which is typologically rare and, on its own, phonetically inexplicable.

In contrast, OT postulates the existence of markedness constraints, which drive surface word forms to deviate from their underlying representation in accordance with typological tendencies, on the assumption that these tendencies reflect known phonetic or learnability pressures — though there is certainly wide-ranging debate on just how 'grounded' or 'substance-based' OT constraints should be (Archangeli & Pulleyblank 1994; Beguš 2018; Hayes 1999; Mortensen 2006). This theoretical decision not only asserts the importance of such pressures to modelling phonological behaviour, but encodes those pressures in the working of the formalism itself. But, certainly, other ways of grounding phonological processes in the pressures of the real world have been proposed — by locating them in sound change, for instance (Blevins 2004).

To summarise, the Morpheme-Invariant position assumes a strict ontological separation between lexicon and grammar. The lexicon contains minimal, invariant morphemes. These morphemes are stored in the form of Underlying Representations — strings of phonemes — which are combined to make words. The phonological grammar — whether consisting of rules or constraints — takes an Underlying Representation as input, and computes a Surface Realisation as output. Phonological analysis involves reducing the observed variation in the surface forms of morphemes across different word contexts down to the smallest possible set of invariant URs (the lexicon), and relating underlying and surface forms by means of the most parsimonious computation procedure (the grammar). Since redundancy is held to be antithetical to formal simplicity, it is stringently avoided.

1.3.2 The lexicon as a network of constructions

Criticisms of the Morpheme-Invariant position come from several sources. Firstly, from theoretical arguments aiming to eliminate the distinction between lexicon and grammar, or between representations and the rules that operate on those representations. Secondly, from growing calls in the morphological literature to dispense with the morpheme as a unit of analysis. And thirdly, from phonetic and psycholinguistic evidence for the storage of various kinds of linguistic phenomena that had previously been dismissed as 'redundant', including predictably-derived words. All of these have significant implications for phonology.

In this section, I summarise these converging lines of argumentation; in the next, I introduce one theory which has emerged out of these considerations — Relational Morphology (Jackendoff & Audring 2020) — and show how it can be used to capture phonological alternations. This forms the necessary background for the analyses in the main chapters of this thesis, where I argue that the core assumptions of the Morpheme-Invariant Theory pose insurmountable obstacles to analysing several aspects of Wubuy phonology and morphology in satisfactory fashion, to which I offer a redress in terms of Relational Morphology.

As we have seen, the Morpheme-Invariant Theory assumes a strict separation between lexicon and grammar. Theories in the tradition of Construction Grammar (e.g. Booij 2010; Bybee 2006; Croft 2001; Culicover & Jackendoff 2005; Fillmore et al. 1988; Goldberg 1995; Hoffmann & Trousdale 2013; Jackendoff 2002; Jackendoff & Audring 2020; Langacker 1987) typically seek to reject this dualistic split. The lexicon, they propose, contains 'constructions': conventionalised pairings of meaning and form. Constructions range from the relatively concrete — words like (2a,b) *home* and *homeless*, or multi-word sequences like (2e,f) *oh my god* or *kick the bucket* — to the relatively abstract — like the verb phrase schema (2h) $[\text{V NP}]_{\text{VP}}$. Or, they may sit somewhere between these two poles — as in the deadjectival noun schema $[[\ldots]_{\text{Adj}}\text{-}ness]_{\text{N}}$ or the partially filled phrasal schema in (2g). If the lexicon contains things that are like words and things that are like rules, as well as a spectrum of intermediate things, then there is no fundamental separation between lexicon and grammar; the functions of both are absorbed into a single entity.

(2) Examples of constructions

a. *home*

b. *homeless*

c. *homelessness*

d. $[[\ldots]_{\text{Adj}}\text{-}ness]_{\text{N}}$

e. *oh my god!*

f. *kick the bucket*

g. [SUBJ$_i$ [V [POSS$_i$ *way*] OBL]]

h. [V NP]$_{\text{VP}}$

One typical line of reasoning for this proposal can be reconstructed as follows.[2] Consider that in English there exist a vast number of multi-word expressions with conventionalised meanings — 'constructions' in the traditional sense — including idioms like *kick the bucket*. As an idiomatic expression, its meaning is not compositional, as it has nothing to do with either kicking or buckets: it means simply 'die'. Since this meaning is not predictable from the elements of which it is composed, the whole construction must be 'listed' in the lexicon, if we assume that the lexicon is the repository of irregularities. But it still behaves like an ordinary, regular VP: the V within it can be tensed (*kicked the bucket*, not ∗*kick-the-bucketed*) and the NP can be modified within certain limits (*kick the old/proverbial bucket*). As Jackendoff (2002, p. 169) emphasises, this presents a paradox if we assume a strict separation between lexical storage and computation: "although idioms must be stored, the way they must be stored cannot be reconciled with a theory in which only individual words are inserted into sentences".

Other idioms have more flexible structure, with apparent 'slots' that can be filled by any structure meeting their subcategorisation requirements. For example, *take* [NP] *for granted* or *put* [NP] *on ice* can take an NP direct object just like other transitive verbs, but they have special, non-compositional meaning. Other idioms are more flexible still. Take the '*way*-Construction' (Goldberg 1995), as exemplified in (3a) and (3b). The general form of this construction can be expressed as a *schema* — that is, a lexical entry with one or more variable slots. In the *way*-Construction schema in (3c), only one terminal element is phonologically specified, and the rest are variables. Again, the phrase structure is entirely unremarkable, but the meaning is non-compositional; it means to move along the path specified by the oblique, while doing the action of the verb.

(3) The '*way*-Construction' in English (Goldberg 1995)

a. Frank dug his way out of prison.

b. Their customers snorted and injected their way to oblivion and sometimes died on the stairs.

c. [SUBJ$_i$ [V [POSS$_i$ *way*] OBL]]

Note that (3c) essentially describes the form of an ordinary English sentence with a subject NP and a VP consisting of an NP and PP. The only crucial difference is in prespecifying the

[2] My reconstruction here largely follows Jackendoff (2017).

content of the object NP — as well as, of course, the non-compositional interpretation of the whole construction. If constructions such as these must be listed in the lexicon, then it seems unavoidable that ordinary syntactic phrase structure must be listed also. When stated schematically, a phrase structure rule only differs from these constructional idioms in having no elements already specified: it is a structure consisting entirely of open variables. A verb phrase can be represented therefore as $[\text{V NP}]_{\text{VP}}$. As Jackendoff (2008, p. 15) argues, this leads to replacing the dualistic lexicon-grammar split with a continuum:

> "construction grammar makes no principled distinction between words and rules: a lexical entry [i.e. 'construction'] is more wordlike to the extent that it is fully specified, and more rulelike to the extent that it contains variables that have to be filled by other items in the sentence. The extreme case is a phrase structure rule like VP → V-NP (or, in constraint-based format, $[_{\text{VP}}$ V NP]), which is made up entirely of variables and carries little or no independent meaning."

Construction Grammar responds to the problem of non-compositional syntax not by drawing an arbitrary line between 'listed in the lexicon' and 'computed by the grammar', but by embracing the gradience (see Fedorenko et al. 2020 for evidence from neuroimaging in support of this view). A corollary of this lexicon-grammar continuum is that the relation between morphology and syntax is also revealed to be gradient; this approach is especially compelling in the case of polysynthetic languages like Wubuy, where verbs can express a complete proposition, like a sentence in more analytic languages (Baker 2018).

Non-compositionality is, if anything, all the more apparent in morphology. While some words, like *homeless*, are analysable as a function of their constituent parts (*home+less*), others present problems. Some may have the same suffix, but the semantic relation to the base is not transparent (*shiftless*), or the base may not exist as an independent word (*hapless*, *gormless*). In morphemic theory, the semantic idiosyncrasy of *shiftless* requires it to be stored in the lexicon, whereas *homeless* must be computed on the fly: the language user searches the lexicon for *home* and *-less* and assembles them together during production (or disassembles them during processing). But how did *shiftless* appear in the language in the first place? Presumably it first arose by the same kind of derivation: *shift* + *less*. If a derived form is able to develop a meaning independent of its base, it must also have had its own entry in the lexicon to store that semantic information in the first place (Booij 1999; Bybee 1999). Likewise, *hapless* descends from a productively derived word, but its base, *hap* 'good fortune' (cf. *happy*), has completely fallen out of use. Morphemic theory would seem to predict that speakers keep a root *hap* in their mental lexicons just for those occasions when they also retrieve *-less* from the lexicon.

Instead, Construction Grammar claims that language users store words and multi-word

expressions in long-term memory not on the basis of their internal complexity or degree of idiosyncrasy, but rather as a function of frequency, attention, interference, surprisal, among other factors (Ambridge 2020; Baayen, Tomaschek et al. 2017; Harmon & Kapatsinski 2017; Olejarczuk et al. 2018). These concrete instances of language are the basis for generating schemas:

> "Language users first acquire words, and only once they have acquired a sufficiently large set of words of a certain type can they conclude by abstracting morphological patterns. This pattern will be memorized beside the set of memorized words on which it is based, and the abstract pattern serves as a recipe for coining new complex words" (Booij 2016, p. 425).

So, one may acquire *homeless*, *hairless*, *careless*, and *propertyless*, as well as *home*, *hair*, *care*, and *property*, and notice that they stand in a paradigmatic relationship to one another. On this basis, the learner can postulate a general pattern for denominal adjectives, stored in the lexicon in the same format as the words that instantiate the pattern: $[[\ldots]_{N}$-*less*$]_{Adj}$. This more abstract knowledge is useful when encountering an adjective ending in *-less* that one has not heard before, or when deriving a new adjective from a noun; but for already known words, one simply retrieves it from memory. The existence of idiosyncratic words like *shiftless* or *gormless* may reduce the confidence that a new word matching the schema also implies a related free noun root, but it otherwise has the same ontological status as its transparent cousins.

Importantly, affixes here do not have an existence independent of the constructions that instantiate them: the lexicon contains no lexical entry *-less* (since it cannot be its own utterance and will therefore never be encountered in this form), only words and schemas that contain *-less* following some other string. Because constructionist approaches to morphology take the word as the smallest unit of storage, the meaning of the whole word is primary, and whatever form-meaning relationships might be discovered among the parts of the word derive from their ability to distinguish the word from other words. That is, to the extent that the *-less* part of *hairless* has any meaning, it emerges from the comparison with *hair*.

In short, words are not strung together out of morphemes. It is here that Construction Grammar most directly interfaces with recent work in 'whole-word' approaches to morphology, particularly Word and Paradigm morphology (e.g. Ackerman, Blevins et al. 2009; Baayen, Chuang et al. 2019; Blevins 2006), which have drawn attention to the numerous problems inherent in the structuralist notion of the morpheme, conceived as a biunique one-to-one mapping of form and function. As Jackendoff & Audring (2020, p. 18) state, "practically all the difficulties with morphology — and all the fun and danger — lie in noncanonical correspondences, in which the phonological form of complex words does not split cleanly into identifiable morphosyntactic

constituents". In inflectional systems especially, it is common to have identifiable pieces of form with no meaning (empty morphs), meanings with no form (zero morphs), meanings distributed across several pieces of form (distributed exponence), several meanings tied to one form (cumulative exponence), or some simultaneous combination of the above, not to mention all the various kinds of non-additive morphology (Anderson 2015). Moreover, a language's inflectional morphology may be divided into a number of arbitrary lexical classes, such that a bare root or stem is not sufficient to determine the affixes that must be added to it when actually appearing in a real word (Blevins 2006). Instead of stripping stems and affixes out of their context and positing rules to put them back together, Word and Paradigm morphology argues that language users acquire inflectional systems by learning whole inflected words and paradigmatic relations between them. The question for morphological analysis concerns how elements of form reduce uncertainty about where a word sits in a paradigm, without requiring that element to associate directly to some meaning.

The foregoing arguments have approached the question of the respective roles of storage and computation by means of linguistic analysis, but the question is also, unavoidably, a psycholinguistic one, and has for that reason been a topic of intense focus in the psycholinguistic literature. Much of this work has attempted to study the mental lexicon using behavioural experimental methods, of which the lexical decision task has been perhaps the most popular and fruitful (Amenta & Crepaldi 2012; Cutler 2012; Gagné 2017). In this experimental paradigm, participants are presented with a word (the 'prime'), and then presented with another word (the 'target'), and must quickly decide whether the target is a real word of their language or not, registering their decision by pressing one of two buttons attached to a computer. The latency and accuracy of the responses are then used as an indication of processing difficulty. When participants are primed with a word that is morphologically related to the target (e.g. *quickly* and *slowly*), their reaction times are faster on average than when the prime is merely semantically similar (Boudelaa & Marslen-Wilson 2015; Creemers et al. 2020; Marslen-Wilson et al. 1994), even when the morphological relationship is irregular (e.g. *fall-fell*; Baayen & Moscoso del Prado Martín 2005; Crepaldi et al. 2010). Many of these studies, moreover, have shown that processing speed, in general, correlates with lexical frequency, including for predictably-derived complex forms (Baayen, Dijkstra et al. 1997; Baayen, McQueen et al. 2003; Baayen, Wurm et al. 2007; Schreuder & Baayen 1997; Sereno & Jongman 1997). The frequency results suggest that at least some regular surface words must be stored, directly contradicting the standard generative stance that the language faculty is minimally redundant, only storing what cannot be computed by rule.

In addition to these results from processing, evidence for storage of whole words can be found in production, where a range of word-specific effects on phonetic realisation have been

documented. For instance, Gahl (2008) shows that putatively homophonous word pairs in English like *time* and *thyme* show persistent differences in duration correlated with frequency, even when controlling for contextual factors such as speech rate, utterance position and local informativity: the more frequent member of the pair is consistently shorter in duration. Subphonemic acoustic differences have been also found in the production of English *like* depending on its function, whether used as an adverbial, discourse particle, quotative, and so on (Drager 2011; see also Jurafsky et al. 2002). Patterns of use can then become entrenched in particular words. Working from a large diachronic corpus spanning a 130-year period of New Zealand English, Sóskuthy & Hay (2017) find that specific words that increasingly appear utterance-finally over time, where they are subject to phrase-final lengthening, end up becoming longer in duration in all utterance positions. These findings show not only that words are stored in memory, but that they are not stored simply in a maximally nonredundant string of invariant abstract phonemes which exhaustively determines the word's phonetic realisation. Rather, they must be stored in some format that is very close to the surface form, and which bears the marks of their use over generations (Bybee 2001; Pierrehumbert 2008). This is not to deny the many obvious cases where production and perception behave categorically, but to acknowledge that categories are instantiated on the basis of gradiently varying exemplars (Pierrehumbert 2016).

While I am in no position to directly test most questions pertaining to cognitive representation in Wubuy, I see it as an important goal to ground linguistic theory in what we know about cognition, rather than bracketing off or even denigrating all such matters as mere 'performance' (Bybee 1999). The psycholinguistic and phonetic findings mentioned here, in tandem with the linguistic arguments from Construction Grammar and from Word and Paradigm morphology, make the notion of minimally redundant abstract morphemes much harder to sustain. Moreover, if whole words can be stored, there would seem to be no role for rules or constraints mediating underlying representations and surface forms. But the fact remains that for Wubuy or for any language with phonological alternations, the lexicon will contain some patterns that are irreducibly phonological in nature. The question for phonology through a constructionist lens then becomes, how do speakers manage variation in the form of morphs across different word contexts?

However, as has sometimes been noted (e.g. Archangeli & Pulleyblank 2022, p. 1), the place of phonology within Construction Grammar remains seriously underexplored, despite a handful of preliminary attempts (Czaplicki 2013; Höder 2014; Välimaa-Blum 2005; van der Spuy 2017; Vihman & Croft 2007). In this thesis, I present a number of novel analyses of Wubuy phonology within Relational Morphology (henceforth RM; Jackendoff & Audring 2020). Although Jackendoff & Audring give some consideration of how their theory might tackle some key issues in phonology and the morphology-phonology interface, many important problems

are left unexplored — their chapter on phonology is, as they acknowledge, mostly a proof of concept (p. 197). Moreover, with their focus on achieving broad coverage of cross-linguistic morphological issues, Jackendoff & Audring do not provide us with a detailed account of how the theory can address several systematic, interacting alternations within a single language.

One of the aims of the present thesis is to remedy this. As I will demonstrate, the complexity of Wubuy phonology and morphology provides ample grist for working through the implications of RM for phonology. But more importantly, RM is up to the task of carrying us beyond many of the impasses in the analysis of Wubuy phonology, and the manner in which it does so sheds a light on issues that have been at the forefront of phonological theorising for several decades. As this thesis demonstrates in detail, many of those impasses — non-optimising phonologically-conditioned allomorphy, empty morphs, irresolvably abstract underlying representations, sonority-conditioned reduplicant shape alternations, and pervasive opacity of several different kinds — cannot be resolved within the confines of the theoretical approaches currently dominant within mainstream phonology. I therefore hope that, beyond the narrow issue of formalising phonological alternations in RM, or indeed within Construction Grammar, the present work provokes creative responses and new insights even in those who are not of the constructionist persuasion, or at the very least brings the difficult problems posed by Wubuy within the domain of issues that any theoretical account of phonology and morphology must address.

1.3.3 Relational Morphology

Relational Morphology (Jackendoff & Audring 2020) is an attempt to account for morphology within the larger theory of the language faculty known as the Parallel Architecture (Jackendoff 2002). In this respect, it builds upon foundations laid down by Construction Morphology (Booij 2010), extending its empirical coverage and providing a more precise notation. In this section, I provide a very brief overview of RM and its notation. Since this is just an overview, my intention here is to present the material as simply as possible; matters of detail will be worked out much more comprehensively in the analyses in the main chapters.

According to the Parallel Architecture (Jackendoff 2002), linguistic structure is determined by (at least) three independent systems — phonology, (morpho)syntax, and semantics — each of which contains its own primitives and principles governing their combination. Structures on each of these levels are associated to one another via *interface links*. Prototypically, a word is a construction consisting of a bundle of linked structures on each of these levels, as in (4). The interface links are notated in RM by numerical coindices, but they can be thought of as equivalent to bidirectional, nonderivational association lines.

(4) Semantics: CAT_1
Morphosyntax: N_1
Phonology: $[kæt]_1$

Words that can be analysed as morphologically complex are also stored in the lexicon, as in (5). Once again, interface links join the pieces of structure on different levels. In addition to the interface link joining the noun root to its syntactic label and its semantic content (coindex 1), we now also have a suffix *-s* linked to a morphosyntactic feature value PL (coindex 2). Together, these elements form a larger constituent, namely the word *cats*, which receives its own interface link, known as an 'outer interface link' (coindex 3), which can unify the entire word with a noun slot in a syntactic construction.

(5) Semantics: $[\text{PLURAL}(CAT_{12})]_3$
Morphosyntax: $\{ N_1, PL_2 \}_3$
Phonology: $[[kæt]_1\ s_2]_3$

The root in (5) is coindexed with the same number as the root in (4). This is because, in addition to linking structures at different levels, numerical coindices serve an additional function of linking structures between one lexical entry and another. In this role, they are called *relational links*. Coindex 1, for instance, links the string [kæt] in the Phonology of (5) to the equivalent string [kæt] in the Phonology of the singular entry (4); it also serves as a relational link between the relevant structures in the Morphosyntax and Semantics.

Because RM holds that morphological patterns are abstracted from actual whole words stored in the lexicon, affixes have no existence outside the lexical entries that instantiate them. An English-speaker's lexicon might contain *cats*, *parrots*, and *parsnips*, but not the suffix *-s* alone, shorn of its context. Moreover, the motivation for having internal structure in plural words at all, as depicted in (5), only comes about by comparison with another set of stored words: the corresponding singulars *cat*, *parrot*, *parsnip*. That is, we establish internal structure on the basis of analogical comparison: *cat* is to *cats* as *parrot* is to *parrots*, and so on. Identity between the roots in the singular and plural word forms, or between the suffixes across all the plurals, is expressed not by reference to some underlying morpheme, but by relational links between instances. These can also be given lexical entries, as below.

(6)

		a.		b.	
Semantics:		$PARROT_4$		$[\text{PLURAL}(PARROT_4)_2]_5$	
Morphosyntax:		N_4		$\{ N_4, PL_2 \}_5$	
Phonology:		$[pæɹət]_4$		$[[pæɹət]_4\ s_2]_5$	

(7)		a.		b.	
	Semantics:		PARSNIP_6		$[\text{PLURAL}(\text{PARSNIP}_6)_2]_7$
	Morphosyntax:		N_6		$\{\text{N}_6, \text{PL}_2\}_7$
	Phonology:		$[\text{pɐːsnɪp}]_6$		$[[\text{pɐːsnɪp}]_6\ \text{s}_2]_7$

In order to generalise across all of these instances, we can posit a pair of *sister schemas* for singular-plural noun pairs. We do this by retaining the parts that stay constant across (4–7), and replacing any parts that change with variables, as in (8). If a variable replaces an element bearing an interface link, it is replaced with a *variable interface link*, notated with an alphabetic coindex $\{x, y, z, \dots\}$; if it bears a relational link, this is replaced with a *variable relational link*, notated with a Greek coindex $\{\alpha, \beta, \gamma, \dots\}$.[3] In the present case, the outer interface links on the plural forms do not link relationally to any other schemas; they are replaced with the variable interface link *x*. The coindexation on the roots relationally links the base of the singular to the related free form, and is therefore replaced with a variable relational link, α. Meanwhile, the plural suffix remains constant, and retains its numerical coindex, 2, thereby linking the plural schema to the instances that motivate it in (5), (6b) and (7b).

(8)		a.		b.	
	Semantics:		INDIVIDUAL_α		$[\text{PLURAL}(\text{INDIVIDUAL}_\alpha)_2]_x$
	Morphosyntax:		N_α		$\{\text{N}_\alpha, \text{PL}_2\}_x$
	Phonology:		$[\dots \begin{bmatrix} -\text{VOICED} \\ -\text{STRIDENT} \end{bmatrix}]_\alpha$		$[[\ \dots \begin{bmatrix} -\text{VOICED} \\ -\text{STRIDENT} \end{bmatrix}]_\alpha\ \text{s}_2]_x$

In simple terms, these sister schemas can be interpreted as follows: for a given noun, expect there to be a corresponding plural form ending in *-s* (and vice versa). But in fact, there is more to it than this, as the root slot in (8) is not simply a completely open variable. Rather, it specifies subcategorisation features on each of the linked levels. The Morphosyntax specifies that the base it unifies with must be a noun. The Semantics specifies that that noun must denote an individual, since all the plural nouns in the lexicon have this feature — collective nouns like *cattle* do not take plural inflection, and nor do mass nouns unless coerced into a count noun interpretation. The Phonology also contains a subcategorisation feature, namely that the segment preceding the plural suffix must be voiceless and non-strident — note that this is the only aspect of the phonology held in common between the likes of *cat, parrot, parsnip*. English has two other plural suffix allomorphs — *-z* following voiced non-stridents and *-əz* following stridents — each of which would require its own schema in RM.

[3] I make a slight notational departure from Jackendoff & Audring (2020): whereas they give distinct alphabetic coindices to variable interface links in each new schema they propose (j in one schema, k in the next), I always restart coindexation from x, in every schema. Either way, interface variables, by definition, can never link or refer to other schemas, so their numbering is only relevant to the individual schema they appear in. I also use square brackets rather than slashes to indicate constituency on the Phonology level, so as to be able to distinguish left and right edges, which becomes a necessity as constructions get longer and more complex.

Note that there is no requirement here to have a one-to-one association between structures on different levels. For instance, (9) shows a lexical entry for the German inflected adjective *großes* 'big.SG.NEUT.NOM'. The suffix *-es* realises three features simultaneously, without any possibility of associating the different features with distinct elements of form. In such a case, we simply notate multiple interface links on the one piece of phonological structure. Moreover, these morphosyntactic features do not correspond to any structure in the semantics, since they are purely grammatical categories controlled by the noun they agree with.

(9)

Semantics:	BIG_{8}
Morphosyntax:	$\{ \text{A}_{8}, \text{SG}_{9}, \text{NEUT}_{10}, \text{NOM}_{11} \}_{12}$
Phonology:	$[[\text{gʁo:s}]_{8}\ \text{əs}_{9,10,11}]_{12}$

We are not limited to affixing morphology, however. To address non-affixing processes like vowel ablaut in English 'irregular' verbs, Jackendoff & Audring (2020) introduce another device in RM — the 'same-except' relation (see also Culicover & Jackendoff 2012). This is notated by placing asterisks around any part of a constituent that differs in its realisation in different contexts. This is the general strategy for addressing any discrepancy between relationally linked word forms — including, as we shall see, those that arise from phonological alternations. In (10), the phonological forms are indicated to be the same in the present and past forms (via coindex 13), *except for* the vowel, as denoted by the asterisks. This sub-pattern can also be generalised with a schema (Jackendoff & Audring 2020, p. 147), where the phonology subcategorises for a final nasal, final velar consonant, or both (as in Albright & Hayes 2003; Bybee & Moder 1983).

(10) English verbal ablaut in *sing-sang*

Semantics:	a.	$[\text{SING(X)}]_{13}$	b.	$[\text{PAST}([\text{SING(X)}]_{13})]_{14}$
Morphosyntax:		V_{13}		$\{\text{V}_{13}, \text{PAST}\}_{14}$
Phonology:		$[\text{s}\ {}^{*}\text{i}{}^{*}\ \text{ŋ}]_{13}$		$[\text{s}\ {}^{*}\text{æ}{}^{*}\ \text{ŋ}]_{13,14}$

We have seen that the basic unit in RM is the construction, which consists of a set of linked structures on parallel levels — typically phonology, (morpho)syntax and semantics — each with their own conditions on well-formedness. Pieces of structure are linked between levels by *interface links*, and across entries by *relational links*, both of which are notated with numerical coindices. In schemas (as opposed to fully-specified word entries), at least some structure is made into a variable, as are any links to that variable. Interface links thereby become *variable interface links*, notated with alphabetic coindexation, while relational links become *variable relational links*, coindexed with Greek letters. If there is any part that changes between two or more structures that are linked relationally, the change is surrounded by asterisks in each entry.

As I show in much more detail in Chapter 3, these basic properties make an entirely surface-based analysis of phonological alternations possible, including those with opaque interactions. To offer a brief preview of that argument, take the case of 'self-destructive feeding' (Baković 2007). Suppose a language in which, under a morphemic analysis, underlyingly vowel-final noun roots take a suffix allomorph /-fu/ in the plural, while underlyingly consonant-final roots take /-ba/. However, a later rule deletes coda consonants when they are not word-final. The immediate surface environment is identical between each allomorph: the form of each root in the suffixing context is vowel-final. How do we express the fact that the allomorphy is sensitive to underlying and not surface phonology, when RM has no underlying representations?

In RM, we express these alternations as relations of correspondence between surface forms, without intermediate levels or derivational stages. Opacity emerges when the phonological form of a constituent in one word context depends on its form in some other word context that it is relationally linked to. The self-destructive feeding scenario is depicted in (11); though note that I have omitted the Semantics level since it is not relevant to the present discussion. In (11a), there is some constituent bearing a variable relational coindex γ, which links it to the equivalent constituent in (11b). This part is not problematic, since it involves no alternation. Between (11c) and (11d), however, there is a difference in the relationally-linked roots: as a free form, the root ends in a consonant, but preceding a suffix, this consonant is absent. The difference is indicated by same-except asterisks around the part that changes between linked constituents (the final consonant). The schemas therefore tell us that for any noun root which ends in a consonant in the singular, expect the root in the corresponding plural form to lack this consonant, and to be followed by /-ba/; for any noun which ends in a vowel in its singular form, expect it to be followed by /-fu/ in the plural (with no change in the form of the root).

(11) Allomorphy conditioned by **underlying** phonology, rendered opaque by a phonological alternation:

Morphosyntax:	a.	$\{N_{\gamma}, PL_{15}\}_x$	b.	N_{γ}
Phonology:		$[[...V]_{\delta}\ fu_{15}]_x$		$[...V]_{\gamma}$
Morphosyntax:	c.	$\{N_{\varepsilon}, PL_{16}\}_x$	d.	N_{δ}
Phonology:		$[[...V]_{\delta}\ ba_{16}]_x$		$[...V\ {*C*}]_{\delta}$

So, instead of relations between underlying and surface form, we have here only relations between surface forms (or between schemas motivated by those surface forms, as here). That is, RM allows us to address opacity without requiring any serial mechanism, and indeed the tools it requires are just those that are argued for on grounds other than opacity: relational links, phonological subcategorisation, and the same-except relation. While the presentation

here is left fairly simplistic, Chapters 3, 4, and 5 all address phonological interactions in Wubuy that involve more complex cases of opacity (the first of these involves two kinds of opacity simultaneously), and I refer the reader there for a much more comprehensive treatment of the way RM deals with these cases, and its advantages over alternatives from the generative literature.

Finally, I wish to note here that by not separating lexicon from grammar, RM makes no attempt to directly incorporate substantive constraints into the formalism. That is, it assumes that, from the perspective of the language user, it does not matter whether a pattern of their language is typologically expected or 'marked' or judged impossible by grammarians — once they discover a pattern, they should be able to use it, and the pattern should be expressible by a schema. Schemas abstract over the lexicon in its current state, which is always in flux, and they are therefore constrained only indirectly, by whatever factors act to shape and reshape the lexicon: phonetic variation and change, language contact, changing patterns of usage, and so on (Yu 2007, p. 61). In that respect RM is broadly in accordance with the evolutionary program of Blevins (2004), who states:

> "Synchronic constraints incorporating naturalness or markedness are misguided, since whether a sound pattern is natural, crazy, or somewhere in between, is wholly a function of its history. Sound change is phonetically natural, but phonological constraints need not be." (Blevins 2004, p. 71).

We will see in Ch. 3 and 4 that incorporating markedness constraints into the grammar consistently fails to explain many of the alternations we find in Wubuy.

1.4 Data collection

Much of the data for this thesis was collected in Numbulwar, where I made field trips in 2016 (2 weeks), 2017 (3 months over two trips), and 2020 (3 weeks). The bulk of this field research consisted of elicitation in the traditional mode. This sometimes took the form of requesting translations from the English (in which all the participants are reasonably fluent), or suggesting Wubuy forms to obtain grammaticality judgements, but wherever practicable, I sought to describe a concrete scenario which I expected to elicit a particular expression. These recordings run to about 63 hours, though only about a third of that is transcribed, using Elan (Brugman & Russell 2004). On my last trip, I recorded people narrating a couple of silent short films, but as these proved difficult to transcribe without native speaker assistance and I did not have time to do so in the field, I have only used a couple of examples from this material. Examples from my field recordings are cited in the format (HN; PN2017-06-27.1 01:09:04), where the first part gives

the initials of the speaker, the second part indicates that it comes from my recordings, followed by the date, the number of the recording from that date if there was more than one, and the timestamp of the utterance. In addition to the elicitation corpus, I assembled a corpus specifically focussed on velar stop deletion, as described in Chapter 5. This corpus is instantiated as an `emuR` database (Winkelmann, Harrington et al. 2017), and contains 1,077 utterances transcribed at the segment level.

I also make extensive use of Heath's (1980) *Nunggubuyu Myths and Ethnographic Texts*, using a version digitised by Nick Thieberger and Simon Musgrave, and cleaned and marked up in XML by myself. The texts run to about 20,835 orthographic words. Examples taken from there are cited as 'NMET' followed by the text number, section, and line, as in (1) above. I also include the recordings of two of Heath's texts in the corpus used to analyse prosody in Chapter 6, alongside a text that I recorded in 2017. These were first transcribed at the utterance level by Kate Horrack; I re-transcribed them at the segment level and prosodically annotated them with an AM-style ToBI scheme, following the procedure described in that chapter.

In addition to these field trips, I spent most of 2018 working at Numbulwar School in the Wubuy revitalisation program, as the teacher-linguist. In 2020 I again spent some time working for the Northern Territory Department of Education, this time as the consulting linguist for the East Arnhem region, and part of that involved work at Numbulwar School. Although I did not conduct research during these periods, my time working there, being surrounded by Wubuy and forced to use it all day, substantially deepened my knowledge of the language and the social and cultural contexts of its use, and strengthened my connections with the community.

My main consultants were all either people who worked at the school in its Wubuy revitalisation program or had worked there in the past. Some of them have also been involved in Bible translation and Wubuy-language church activities. All have strong command of the local variety of Aboriginal English, including literacy, as well as Kriol, which is now the dominant language in the community; many also regularly speak Anindilyakwa with family, and some also speak a variety of Yolŋu Matha. The casual conversations between them that I have observed tend to show frequent code-switching between Wubuy, Kriol and English. Their level of Wubuy literacy varies considerably, from those who can only read it, to those who write it inconsistently, to those who are very confident and consistent in their Wubuy writing. The speakers I worked with most often — Jangu Nundhirribala, Hilda Ngalmi, Yidangga Ngalmi, Leonie Murrungun, Mara Murrungun, Nanette Murrungun, Jijila Turner — have all had fairly extensive prior experience working with linguists. Some of these consultants have a great deal of metalinguistic awareness, being able to not only give grammaticality judgements, but also to pinpoint the source of ungrammaticality, i.e.: "You've got the wrong prefix", or, "That one you have to curl your tongue" (for retroflex sounds). Nearly all of my consultants were women, and the two men

I worked with (David Murrungun and Reg Numamurdirdi) only got involved towards the end of my last field trip; they had previously been very involved in documenting traditional songs and cultural knowledge, but had never been involved in linguistic research. While this is unlike the situation in most of Aboriginal Australia, where research or other working relationships tend to be segregated by gender, it seems to be quite normal in the Roper region.

Of course, since I am a non-Indigenous scholar working within the academic discipline of linguistics, the questions and problems I entertain are those of importance to the discipline, and are not directly motivated by the concerns of the Nunggubuyu community. That said, the work in this thesis does not represent the entirety of my commitments to the Nunggubuyu people or their language, and I have been overwhelmed by the often very vocal support and encouragement people in Numbulwar have given and continue to give me. Moreover, the potential insights gained from studying the language linguistically are not lost on Wubuy language owners, who, as mentioned, show a high degree of metalinguistic awareness — on several occasions, I have run training sessions on linguistic topics at their request. I also note that understanding Wubuy phonology and morphology bears on questions of how one might go about teaching the language, especially to children, which is undoubtedly relevant to the aspirations of Nunggubuyu people.

2 An overview of the Wubuy language

2.1 Wubuy, its speakers, and its country

Wubuy ['wubuj], also known in the literature as Nunggubuyu,[1] is a highly polysynthetic 'prefixing' language of East Arnhem Land in Northern Australia. It is a member of the large Gunwinyguan family (Alpher et al. 2003; Figure 2.1), and is probably most closely related to Anindhilyakwa, spoken on Groote Eylandt, which many Wubuy speakers are also fluent in (van Egmond & Baker 2021). It is traditionally associated with the lands around Harris Creek and the southern half of Blue Mud Bay, though most speakers today reside in the small community of Numbulwar (formerly Rose River Mission), a little further south on the coast of the Gulf of Carpentaria, in traditionally Warndarrang country.

It is difficult to get an accurate assessment of the present number of speakers, but it is without doubt critically endangered. I would estimate around 30 fluent L1 speakers, all of whom are over fifty years of age; some but not all still use it as their primary language of daily communication. Sadly, several Wubuy speakers passed away during the course of preparing this PhD. The community has been undergoing a shift to Kriol — an English-lexifier creole (Dickson 2015; Meakins 2014) — over at least the last three generations, just as many other language communities across the region (particularly heading inland in the direction of Katherine) have already done so. The older Wubuy speakers I have worked with are quite comfortable in English, having learnt it on the mission as children, and I have overheard code-switching between Wubuy, English and Kriol in casual conversations between them. Among Kriol-dominant heritage speakers of Wubuy, English is very much a foreign language, and is used only for communicating with outsiders in the town, like police, shop staff, teachers, and doctors. Some of these heritage speakers appear to have a surprising amount of passive Wubuy knowledge, and many of the young men are heavily involved in local ceremonial practice, where they command a repertoire of traditional Wubuy songs. After I had spent several months in Numbulwar, various people in their late 20s and early 30s approached me to start up basic conversations in Wubuy, though I had only seen them speak Kriol to the older Wubuy speakers. The older, fluent

[1] 'Nunggubuyu' ['nuŋgubuju] is in fact the speakers' ethnonym; it means, "those who speak Wubuy". Wubuy is the name people use when speaking the language, and is also the name adopted by the traditional language program at Numbulwar School.

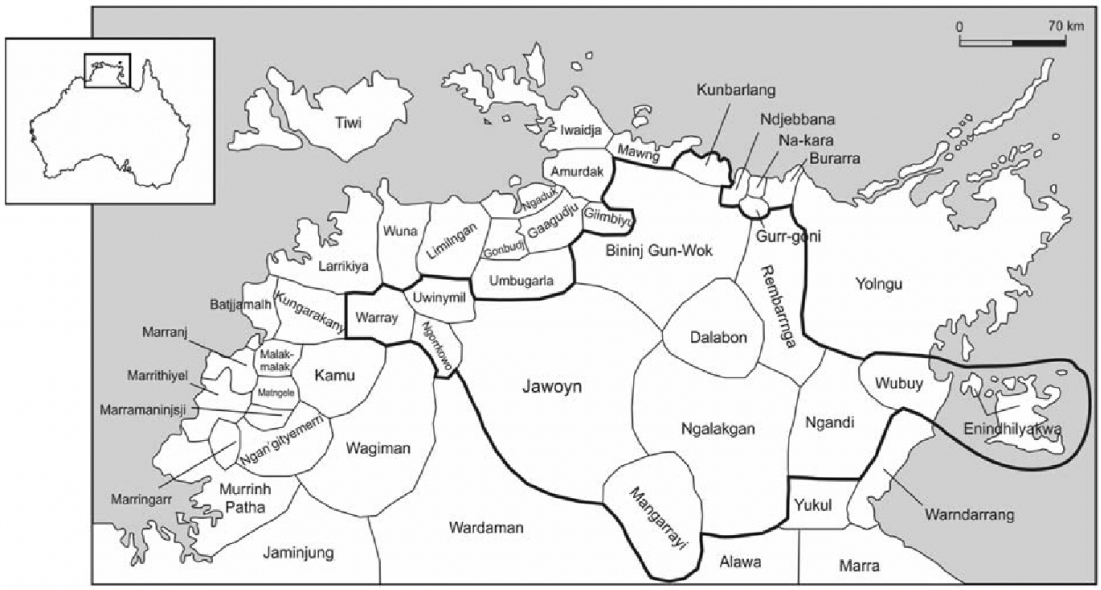

Figure 2.1: Map showing Gunwinyguan family, taken from van Egmond & Baker (2021). Wubuy is at the eastern end of the family's mainland extent. The 'Yolŋu' bloc contains numerous named varieties, some of which are substantially different from one another (see Jepson 2019; Wilkinson 2012a), though not as different as perhaps any two non-Pama-Nyungan languages on this map are from one another.

speakers and the younger heritage speakers alike may also have proficiency in Anindhilyakwa and 'Yolŋu Matha' (generally either Djambarrpuyŋu or Ritharrŋu) to varying degrees, due to extensive contact with these neighbouring groups. Wubuy is not known to be divisible into dialects, and, given the very small number of speakers, it is difficult to assess whether the (relatively small) differences between the language of particular individuals can be attributed to any sociolinguistic variable. There is, however, a 'song language' register featuring special vocabulary.

The precarious situation of Wubuy is felt keenly by Nunggubuyu people, who see their language as an indelible aspect of belonging to their country, and part of one's inheritance as a Nunggubuyu person — as is common across Aboriginal Australia. There has been a language revitalisation program running at Numbulwar School since the late 1970s, when a group of elders, increasingly concerned about the decline of Wubuy use among young people, demanded some school time be devoted to teaching local knowledge, stories and songs in Wubuy (see Carr et al. 2017). The Wubuy program was once central to the school curriculum, but now makes up at most two hours of the students' week. Under present conditions, it is highly unlikely that the school program will result in young children acquiring fluent command of Wubuy (though fluency is not necessarily the goal for those children).

Among linguists, Wubuy is best known through the formidable body of work undertaken by Jeffrey Heath, including, most notably, a collection of 171 glossed and translated texts (*Nunggubuyu Myths and Ethnographic Texts*, Heath 1980), a dictionary (*Nunggubuyu Dictionary*, Heath 1982), and an extremely extensive and detailed grammar (*Functional Grammar of Nunggubuyu*: Heath 1984). Early work was also undertaken by the missionaries Hughes (1969, 1970, 1971) and Hore (1978, 1981). More recent documentation and analysis has been undertaken by Brett Baker and colleagues (Baker 2008a, 2009, 2014a, 2018; Baker, Bundgaard-Nielsen, Nyhuis et al. 2020; Baker, Horrack et al. 2010; Baker & Nordlinger 2008; Bundgaard-Nielsen & Baker 2020), and Kate Horrack (Horrack 2010, 2018).

The aim of this chapter is to provide an overview of Wubuy grammar, mostly based on these prior sources but interspersed with my own commentary, providing crucial background to the analyses I develop over the subsequent chapters. Some of the typological features of Wubuy that have attracted attention, and which will feature in the discussion here, include: 'polysynthetic' verbal morphology, of sufficient richness that any well-formed verb can stand as a complete utterance; the existence of nine noun classes or genders, with extensive agreement morphology across most word classes (including verbs); productive noun incorporation in verbs and adjectives; and realisation of two core arguments on the verb, or a non-core argument by means of valence-changing derivational morphology. Indeed, the morphology of Wubuy is sufficiently rich, and word order is sufficiently free, that Heath (1986) questions whether one

can find evidence of syntactic constituency at all. The morphology of the different word classes and the patterns of agreement among them are discussed in the next section, §2.2.

While most of these features are fairly typical of Gunwinyguan languages, Wubuy stands out for the far greater proliferation of phonological alternations one finds at morphological boundaries. Many of these alternations are typologically unusual, even within Australia, and are not predicted by the dominant strands of phonological theory. We find, for instance, a number of alternations that do not seem to improve well-formedness, several highly opaque process interactions, a reduplicant shape alternation unlike anything found in other languages (except neighbouring Anindhilyakwa), a phonologically-conditioned but non-optimising syllable epenthesis process, and the apparent existence of abstract phonemes, according to previous description. In §2.3, particularly §2.3.2, I summarise what has so far been documented about these alternations, pointing out some of the theoretical questions they raise. On the other hand, the present state of our knowledge about Wubuy prosody is quite poor, as I discuss in §2.3.5.

As this thesis argues, the co-occurrence of all of these unusual phonological alternations interacting together within an already complex morphological system represents a unique challenge, with the potential to drive linguistic theory forward.

2.2 Morphology

Australian languages are famed for their 'non-configurational' properties: free word order, null anaphora, and discontinuous NP constituents (see Nordlinger 2023). Wubuy is particularly extreme in this regard: it places practically no demands on word order, possibly to an even greater extent than most Australian languages, and determining basic syntactic constituents like S, NP and VP is no straightforward matter (Heath 1986). The morphology, by contrast, is extremely rich, and provides the means by which elements in the utterance are grammatically related to one another.

Since this thesis is concerned with how the knowledge of complex words can be used to generalise to new contexts, I concentrate exclusively on the open classes. The open word classes are nouns, verbs and adjectives — the last being a formally distinct category in Wubuy, unlike in many Australian languages (see Kim 2023). Among the closed classes, there is a very elaborate set of pronouns, demonstratives and kin terms, as well as a number of particles and interjections.

Note that, following Haspelmath (2020; see also Blevins 2006), I generally use the term 'morph' to denote the fairly informal notion of some recurrent element of the form of words which is associated with a particular meaning. The more familiar term 'morpheme' has come to acquire a great deal of theoretical baggage, and I only use it when discussing generative work that explicitly depends on the concept of morphemes as the atoms from which words are constructed,

and that analyse phonology as acting upon the underlying representations of said morphemes.

2.2.1 Nominal morphology

Nominal words in Wubuy minimally consist of a noun root. Depending on their grammatical role and discursive function, they may also appear with a noun class prefix, as well as a case suffix. A noun may thus be as simple as that in (12a), or be slightly more complex (12b).

(12) a. *root*
wuɻuku
billabong
'(a) billabong'

b.

noun class	*root*	*case*
a-	wuɻuku	-wac
NEUT.OBL-	billabong	-PERG

'around the billabong'

Pronouns, demonstratives, kin terms, verbs and adjectives all show agreement for noun class. Although Heath (1984) identifies nine distinct noun classes that all pertain to one overarching nominal classification system, Horrack (2018) argues for separating these into two coexisting but distinct systems on the basis of animacy: human 'gender' and nonhuman 'noun class'. Both authors agree on all the basic facts that differentiate human and nonhuman nouns, and so the difference in analysis is not particularly consequential for my purposes. In this thesis, partly for the sake of convenience, I follow Heath in treating nominal classification in Wubuy as a single system, but I adopt Horrack's labels for the classes so distinguished: human Male (M), Female (F), and Plural (PL), plus nonhuman Masculine (MASC), Feminine (FEM), Neuter (NEUT), Residual (RESID), Vegetable (VEG), and Collective (COLL). The forms of these prefixes are set out in Table 2.1 below, and their distribution is elaborated in the discussion to follow.

The fact that each of the human noun class prefixes (M, F and PL) is homophonous with one of the nonhuman class prefixes (MASC, FEM, COLL) should be obvious from even casual observation of Table 2.1, and one may rightly ask what evidence there is that these are distinct classes at all. For this, we can observe differences in the marking of number and discourse status in human and nonhuman nominal expressions, as well as in the way that these arguments are cross-referenced by other constituents in the clause (which will be treated in the subsections to follow).

Firstly, human nouns obligatorily carry NUMBER features, while nonhuman nouns are numberless. That is, a human noun with a singular or dual referent is either MALE or FEMALE, and

Topic form	Oblique form	Class	Gloss	Animacy
na-	—	Male	M	human
ŋaɾa-	—	Female	F	human
waɾa-	—	Plural	Pl	human
na-	ji:-	Masculine	Masc	nonhuman
ŋaɾa- (∼ŋa:)	ji:-	Feminine	Fem	nonhuman
waɾa-	wa:- (∼waɾa-)	Collective	Coll	nonhuman
ana-	a-	Residual	Resid	nonhuman
ana-	a-	Neuter	Neut	nonhuman
mana-	ama-	Vegetable	Veg	nonhuman

Table 2.1: Inflectional noun class prefixes, based on Heath (1984). Only the nonhuman classes have Oblique forms. Where a noun is under focus or contrast and is not suffixed for case, no prefix is used. Note that the Resid and Neut classes are formally distinguished in the verbal agreement prefixes.

any human noun — male or female — with a plural referent (i.e. more than two) takes a Plural noun class prefix. The dual is expressed by means of a suffix (*-wa:* ∼ *-pa:*), and there is also an optional but generally preferred suffix for human singulars (*-juŋ* ∼ *-cuŋ* ∼ *-ɲuŋ*). Human plurals often feature reduplication, but it is also not obligatory. Thus we find sets like *ŋaɾa-ŋaciwaŋku-ɲuŋ* 'the woman', *ŋaɾa-ŋaciwaŋku-wa:* 'the two women', *waɾa-ŋaca-ŋaciwaŋku* 'the women'. Nonhuman referents typically do not carry number features, though Heath (1984, p. 195) notes that the dual suffix appears on humans and nonhumans alike. Even so, semantically plural nonhuman referents are not treated as grammatically plural, for instance in verbal inflection.[2]

Besides the obvious function of identifying which class a noun belongs to, the noun class prefixes also have functions linked to their discourse status interpretations (Baker 2008a; Heath 1984; Hore 1978). For both human and nonhuman nouns, the Topic prefixes are typically used when the speaker infers that the hearer will be able to identify the referent they have in mind (Baker 2008a). They can also be used to refer back to previously introduced referents, and in this respect they bear similarity to definite articles in European languages. Unlike definite articles, however, they can also be used to introduce new discourse topics, and are obligatory when they come under the scope of negative and interrogative operators; they are also used to refer to

[2]Human noun morphology, including number-marking and the use of prefixes matching biological sex can be applied to nonhuman animal referents insofar as one can coerce an interpretation of them as human-like. This strategy commonly applies in relation to named dogs, which are generically of the Coll class but can feature all the gender and number-marking morphology (singular and dual suffixes, pluralising reduplication) appropriate to humans when relating to a specific canine referent. I do not have data to indicate whether the noun class prefixes in such cases are really human M/F or rather nonhuman Masc/Fem. It is also found applying to totemic ancestors in mythological narratives — see, for instance, the first story presented in Heath (1980), where the narrator frequently slips between a F/Fem and Neut classification of the snake mother character.

kinds (non-specific referents). Where the noun is in contrastive focus, no topic-marking noun class prefix is used (or, in Baker's terms, a 'non-topic' zero prefix is used instead). I gloss the noun class prefixes in this thesis following Baker's analysis, but note that Heath (1984, p. 169) calls the Topic and Oblique forms 'continuous' and 'punctual' respectively.

Further, the nonhuman noun class prefixes have an additional series of Oblique (Obl) forms, which are used with nouns in an oblique case role. Thus a nonhuman Fem noun like *ŋaɾukaɭic* 'dugong' takes a topic-marking prefix under the conditions mentioned in the previous paragraph, yielding *ŋaɾa-ŋaɾukaɭic* 'Fem.Top-dugong / the dugong', but must instead take an oblique prefix when marked for, say, Allative case: *jiː-ŋaɾukaɭic-kuj* 'to the dugong'. However, in certain contexts where the case role can be easily inferred, the suffix can be left out: *jiː-ŋaɾukaɭic* '[to] the dugong'. For humans, there are no special Obl forms: when suffixed for case, human nouns take the topic-marking prefixes if they are topics, and no prefix otherwise.

These differences in the marking of case and number on humans and nonhumans go some way towards motivating the distinctions between the M/Masc, F/Fem and Pl/Coll classes, despite the clear syncretism in their Topic forms. But the syncretism between the Resid and Neut noun class prefixes cannot be motivated on this basis, since they are both nonhuman. For this, we can observe the way these classes are cross-referenced by other elements in the clause. Specifically, among the verbal pronominal prefixes — which are portmanteau forms encoding the subject, the mood category, and up to one other argument of the verb — there is one series for Resid class objects and another for Neut class objects: e.g., *nuɾu-* '1pl>3Resid.R' contrasts with *niɾi-* '1pl>3Neut.R'. Thus, we are justified in dividing two classes with identical forms so long as they are formally distinguished in the inflectional paradigms of at least one part of speech (see also Harvey 1997a,b).

The noun class prefixes we have considered so far are those Heath (1984) designates as 'inflectional', but there remains a set of 'derivational' noun class prefixes also, shown in Table 2.2, which take a noun of one class and derive a new noun with another class (or even with the same class as its base). This process most commonly serves the purpose of indicating part-whole possession, whereby the part takes on the derivational noun class prefix matching the noun class of the whole. It is often used in relation to plant and animal parts, as well as the parts of inanimate objects — including items of non-Indigenous origin that have entered Nunggubuyu material culture such as cars and telephones — though the noun controlling the agreement (the whole) may not actually appear overtly in the surrounding discourse (Heath 1984, p. 173). Derivational noun class is exemplified in (13), where both *pakaɭaŋ* 'eye (Veg)' and *ɻaŋak* 'wood (Neut)' appear with the Neut derivational noun class prefix *u-*, since they here denote parts of the strychnine tree, which is of the Neut class. A noun derived in this way can then also take whichever inflectional noun class prefix of Table 2.1 matches the derived class: see for example

Form	Gloss
ni-	M
ŋaɾi-	F
u- (with redup.)	Pl
ni-	Masc
ŋaɾi-	Fem
u-	Coll
u-	Resid
u-	Neut
ma-	Veg

Table 2.2: Derivational noun class prefixes, based on Heath (1984) but with the prefix-final archiphoneme /K/ omitted. This modification is motivated in Chapter 3. All prefixes presented in this table subcategorise for a following constituent beginning with a [–cont] segment (i.e. a stop or a nasal).

an-u-ŋu-pakaɭaŋ in (13), where the first prefix is the Neuter Topic prefix.

Note that all derivational noun class prefixes condition Hardening in the following segment if it is underlyingly continuant, and ŋu-Epenthesis if it is underlyingly stop-initial. For that reason, Heath analyses them as ending in the 'stop archiphoneme' /K/, which triggers the relevant phonological changes but is deleted from the surface. I introduce the archiphoneme and the alternations it is implicated in in 2.3.2 below.

(13) winiwu-ma-ŋi an-**u**-ŋu-pakaɭaŋ, ana-wumpaḻpuḻ,
3duM>3Neut-get-Past2 Neut.Top-**Neut**-Ø-eye strychnine tree (*Strychnos lucida*)

wa:ɻi an-**u**-ʈaŋak
not Neut.Top-**Neut**-wood

'They got the nuts (the 'eyes') of the strychnine tree, not the wood' (NMET 138.2.1)

There is also a number of bound roots which can only appear in a nominal word if bearing one of these derivational noun class prefixes (1984, p. 175). For example, the bound root *-ḻamik* 'taste' cannot stand on its own, but must take a derivational noun class prefix in order to appear in a nominal expression, like *mana-ma-ʈamik* and *an-u-ʈamik* in (14). Note that the initial lateral surfaces as a stop in this context. It can also be incorporated into a verb — like *ma-ḻami-cama-na* in (14). More detailed accounts of part-whole possession in Wubuy can be found in Heath (1984, pp. 173–177), Baker et al (2010) and Horrack (2010), the latter two couched in Lexical-Functional Grammar.

(14) ana-wuʈuɲmanci, niɾi-ŋu-ni wuʈuɲmanci, ŋuɲcu amaː-ʾŋ,
NEUT.TOP-wild.potato we ate it wild potato similar VEG.OBL-whatchamacallit

ama-*potato*, ma-jama-na **ma-l̪ami-cama-na** ŋuɲcu ama-*potato*-jiː
potato it (VEG) does that its taste does that similar VEG.OBL-potato-SIMIL

mana-ma-ʈamik wuʈuɲmanci, **an-u-ʈamik**
its taste (VEG) wild potato its taste (NEUT)

'We ate wild potatoes. They are sort of like potatoes, they have a taste like potatoes.'
(NMET 113.6)

As mentioned, Wubuy nouns use suffixes to express case relations; see Table 2.3. These suffixes typically mark non-core arguments and adjuncts, while the subject and object of the clause is typically identified by prefixation on the verb. Most of the case suffixes usually co-occur with the Oblique series of noun class prefixes when the referent is nonhuman, but the RELative case suffix tends to appear with the Topic series, and the INSTRumental tends to appear without any noun class prefix at all. All nominal case suffixes can also appear on pronouns and demonstratives, and most can appear with verbs, marking subordinate relations of various kinds, though their interpretation is quite different in that context.

Suffix	Case	Gloss	Expected nonhuman class prefix
-ɾuc ∼ -tuc	Locative	LOC	Oblique
-wuj ∼ -kuj	Allative/Dative	ALL	Oblique
-wala ∼ -kala ∼ -jala	Ablative	ABL	Oblique
-wac ∼ -pac	Pergressive	PERG	Oblique
-wac ∼ -kac	Retrospective Pergressive	RPERG	Oblique
-juŋkujuŋ ∼ -cuŋkujuŋ	Purposive	PURP	Oblique
-jiː ∼ -ciː	Similative	SIMIL	Oblique
-miɾi	Instrumental	INSTR	none
-miɾaːʈu	Originative	ORIG	(rare)
-jiɲuŋ ∼ -ciɲuŋ ∼ -iɲuŋ	Relative	REL	Topic

Table 2.3: Wubuy case suffixes, following Heath (1984, pp. 164, 199). The rightmost column lists which of the nonhuman prefix series each case suffix is most strongly correlated with.

2.2.2 Verbal morphology

A key characteristic of Gunwinyguan languages is the complexity of their verbal morphology, and therein lies the main motivation for their designation as 'polysynthetic' languages. Indeed,

in Wubuy, as in other languages of the family, an inflected verb is a complete proposition, and can stand on its own as a well-formed clause. Since verbs are also the most common site for the many phonological alternations examined in this thesis, many of which are sensitive to morphological environment, this section provides crucial groundwork for the analyses in the following chapters.

At its most minimal, a verb consists of a pronominal prefix and a stem consisting of a verb root and its suffixal tense-aspect-mood (TAM) inflection. An example is given in (15).

(15) wiɾima- l̪alwul̪ -aŋi
3pl>3VEG.R- soak -PAST2
[prefix [root suffix]$_{\text{stem}}$]$_{\text{word}}$
'They soaked it (VEG)' (NMET 122.4.5)

The verb root and its TAM inflection can be approached in different ways. Verb roots are listed in the dictionary (Heath 1982) in a citation form, which includes a final vowel: for example, *cuɻa*. These citation forms are abstract: verb roots are always inflected in actual words, so the root *cuɻa* in fact only appears as *cuɻiɲ, cuɻaŋi, cuɻaŋ, cuɻiː, cuɻi*, etc., but never as *cuɻa*. In Heath's analysis, these abstract roots are the underlying representations, and the deletion or change of quality in this abstract final vowel is the result of applying phonological rules after suffixing the TAM inflection (see Heath 1984, pp. 101–105, for example); these rules' behaviour is nonetheless quite idiosyncratic. For convenience, I often refer to roots in this citation form (transliterated into IPA) in the running text. In my glosses, however, I treat the root as the part preceding the final vowel of the citation form — segmented *cuɻ-iɲ, cuɻ-iː*, and so on — essentially assuming that these abstract citation forms are not part of speakers' mental lexicon (as Heath seems to suggest, p. 145).

The pronominal prefix encodes one of two MOOD categories (Realis or IRREALIS), and depending on the transitivity of the verb, up to two arguments. Arguments are categorised on the basis of their person, number (for humans) and noun class features. The person marking distinguishes first exclusive (1), first inclusive (12), second (2), and third person (3). Human participants can be singular (SG), dual (DU) or plural (PL) in number, and the singular and dual categories further distinguish Male and Female — the human plural being, as always, effectively an additional noun class. Compare, for instance, *ni-* '3sgM' vs. *ŋi-* '3sgF', or *wini-* '3duM' vs. *wiŋi-* '3duF'. Dual referents of mixed gender (i.e. one male and one female) default to male, so far as I have observed. There are also special 'trial' forms (very rarely used, in my experience) for first inclusive subjects that include the speaker, addressee, and one additional referent. On that basis, the number marking is better described as a minimal/augmented system, but I use the

singular/dual labels for their intuitive appeal.[3]

Unlike other Gunwinyguan languages, Wubuy verbal pronominal agreement distinguishes noun class for third person participants. However, there are a number of syncretisms. In the subject marking, we find the same syncretisms as found in the Topic series of noun class prefixes: M/Masc, F/Fem, Pl/Coll, and Neut/Resid. In the object marking however, human M and nonhuman Masc are kept distinct, as are Neut and Resid (both nonhuman) — the Resid object prefixes are identical with the intransitive prefix series. For example, to command someone with the phrase, *banu-ɻajaŋ* 'spear him!' (2sg>3sgM.Irr-spear.Npst1), means the intended victim is a human male, whereas with *bani-ɻajaŋ* 'spear it [Masc]!' (2sg>3Masc.Irr–spear.Npst1) it is a nonhuman masculine object (a fish of that class, for instance).

Throughout this thesis, I treat all pronominal prefixes as portmanteau forms, as Heath (1980, 1984) does when analysing Wubuy synchronically. I therefore provide one gloss for the entire string preceding the verbal stem. Even though it is occasionally possible to parse out some component of the prefix and associate it with some element of meaning, such associations are only ever tendencies. That is, we cannot reliably predict components of the form of the prefix from the components of its meaning, nor vice versa. For example, the substring *ŋa* is often associated with a 1sg subject, as in the 1sg instransitive realis prefix *ŋa-* or 1sg>3sgM realis prefix *ŋanu-*, but it is not present at all in the 2sg>1sg realis *numba-* or the 1sg>2sg irrealis *wa-*. The one exception to this general lack of form-meaning isomorphism is in the association between a Veg participant (either subject or object) and the substring *ma* (cf. *ŋama-* '1sg>3Veg.R'), where the element of form and the element of meaning uniquely predict each other. This characteristic turns out to be relevant to intonational phrasing, and is discussed in Chapter 5.

Further, there are two series of pronominal prefix for each possible combination of subject and object features, one Realis (R) and one Irrealis (Irr). This choice of mood category, together with the choice of inflectional suffix on the verb root, as well as the presence or absence of a polarity operator in the same clause, supplies the full set of tense-aspect-mood-polarity functions. This phenomenon — typical of non-Pama-Nyungan languages — is what Verstraete (2005) dubs 'composite mood marking', and he includes Wubuy in his sample. The combinations are given in Table 2.4. The suffixal inflection makes a basic distinction between past and nonpast tenses, but provides more than one form for each: two past tense suffixes and three nonpast tense suffixes. There is also an evitative category, separate from these. As can be seen in the table, TAMP expression is not straightforwardly compositional, as each suffix has restrictions on which prefix series and what kind of negative operator (if any) it can co-occur with.

[3] An additional justification for this comes from the similarities in form between the plural 1 and 12 forms: e.g. *nuɾu-* '1pl.R' and *ŋuɾu-* '12pl.R'

Suffixal inflection	Prefix series	Negative element	Meaning	Example of rough English equivalent
Past 1	Realis	-	Past Punctual	'I put it down', or 'I have put it down'
Past 2	Realis	-	Past Continuous	'I was putting it down'
	Irrealis	-	Past Potential	'I could have put it down' 'I was going to put it down' 'I should have put it down'
	Irrealis	waːɻi	Past Negative Actual	'I did not put it down'
	Irrealis	jaki	Past Negative Potential	'I would not have put it down'
Nonpast 1	Irrealis	-	Future Potential	'I will put it down'
	Irrealis	waːɻi	Present Negative	'I do not put it down'
Nonpast 2	Realis	-	Present	'I am putting it down'
	Irrealis	-	Future Continuous	'I will be putting it down'
Nonpast 3	Realis	jaki	Future Negative	'I will not put it down'
	Realis	-maki	Evitative (rarer form)	'If I had put it down (I would have broken it)'
Evitative	Realis	-maki	Evitative	'Lest it be put down'
	Realis	jaki	Negative Evitative	'Lest it not be put down'

Table 2.4: Interactions between prefixal mood inflection and suffixal tense-aspect-mood inflection. Table adapted from Horrack (2018, p. 35), who draws on Heath (1984) and Wilkinson (2012b).

Note that my labelling of the two prefix series as realis and irrealis departs from Heath (1984) and Horrack (2018), who label them more abstractly as 'A' and 'B' respectively. They claim that the B series is not canonically irrealis, since it is not used with negative future readings. However, this is not especially unusual in non-Pama-Nyungan languages. As Verstraete's (2005) typology shows, the Irrealis prefix encodes potentiality when used with nonpast tense categories, encompassing epistemic, deontic, and desiderative-intentional modalities. Therefore, the negation of a future event — the statement that an event will not happen — is a negation of its potentiality. Following his analysis, I do not regard it as problematic to label Heath's 'B' prefix series as irrealis.

There are also some different approaches in the literature to glossing the TAM suffixes. Heath (1980) glosses each suffix according to the label in the first column of Table 2.4, giving each allomorph just one label: PAST1, PAST2, NPST1 etc. Horrack (2018) glosses the suffixes according to their composite tense-aspect-mood category, as in the fourth column of Table 2.4, and so has several labels for each allomorph depending on which prefix series and negative element

it appears with. The focus of this thesis is on phonological alternations and allomorphy; I therefore follow Heath's approach, since having a consistent label for each suffix brings a degree of clarity to that end.

The forms of these TAM suffixes are not uniform throughout the lexicon — rather, each verb belongs to one of 26 lexically specified inflection classes (though 8 of these have just one member). Moreover, verbs can undergo a number of derivational processes, which generally either change the argument structure of the root verb or derive verbs from nouns and adjectives. Inflectional TAM suffixes always occur outside any derivational ones. Where these derivational processes involve an overt suffix — as with the CAUsative *-(i)cka*, RECIProcal *-ɲci*, INCHoative *-ma* or FACTitive *-wa* — the suffix then determines which inflection class the derived verb belongs to. But there are also some valency-changing derivational processes that involve no overt suffix, only a change from one inflection class to another. A change to the I1 inflection class, for instance, is associated with reflexive readings (or passive, or antipassive, depending on the verb). See Horrack (2018, pp. 39–41) for further discussion on some of the difficulties of distinguishing the two kinds of derivation. Note that although the inflectional pronominal prefixes can have multiple allomorphs according to their phonological environment, they show no variation according to the lexical inflection class of the verb they attach to.

(16) a. wu-ɻap-**iːni**, an-upa-ni ana-l̪uːɲ
3NEUT.R-fall-**PAST2** that NEUT.TOP-stringybark
(= *I1 class*)
'That stringybark fell down' (NMET 17.1.5)

b. niː-ɻap-ick-**a**ː
1pl>3NEUT.R-fall-CAUS-**PAST2**
(= *A1 class*)
'We made it [i.e. the tree (NEUT)] fall' (NMET 143.1.4)

An even more pervasive kind of derivation in Wubuy is that of noun incorporation. Baker (2014a, p. 231) defines noun incorporation as "the process whereby a morpheme which functions as a noun root under some circumstances can be found as part of a verb word under others, where the result is more-or-less compositionally interpreted". In many ways this can be considered a kind of compounding, where the noun attaches to the left of the inflected verb root, forming a complex predicate headed by the verb. Like more familiar compounding processes in European languages, the range of possible interpretations of a N-V compound in Wubuy is vast, as Baker (2014a) sets out in detail. In (17), the incorporated noun *mun-* 'foot' is the object of the first (transitive) verb, but the subject of the second (intransitive) verb. Notice also that the pronominal

agreement in the prefixes on both verbs reflect the possessor of the foot (a singular Male), not the foot itself (which is Neuter). This kind of possessor raising is generally restricted to body parts, as well as more abstract but similarly inalienable parts of a person, such as *jaŋ* 'voice', *malŋuc* 'shadow', *mun* 'foot; footprint' and *muwac~mic* 'name' (Horrack 2010, p. 29).

(17) taki ŋunu-**mun**-na-ɲ, ni-**mun**-ŋalŋala-ŋi
he is there 3sgF>3sgM.R-foot-see-Past1 3sgM.R-foot-shine-Past2

'There! She saw his foot [sticking out of the water, or visible just under the surface]. His foot was clearly visible.' (NMET 17.8.2)

But agreement with the noun class of the incorporated noun is also possible in Wubuy, and is available for alienable and inalienable possessions alike. See (18b) for an example, where the object marking on the verb agrees with the noun class of the incorporated noun (again, an inalienable body part), not its possessor. Aside from transitive objects and intransitive subjects, a notably broad range of oblique arguments can also be incorporated: themes, locations, and instruments are all widely attested. Baker (2014a) claims that the last of these appears to be unique to Wubuy.

(18) a. ŋawu-nakiː-na ana-**jaɾka** jiː-ŋaɾukaɭ̪ic-iɲuŋ
1sg>3Neut.R-cook-Npst2 Neut.Top-flipper Fem.Obl-dugong-Rel

'I'm cooking the dugong's flipper' (Horrack 2010, p. 32)

b. ŋawu-**jaɾka**-nakiː-na jiː-ŋaɾukaɭ̪ic
1sg>3Neut.R-flipper-cook-Npst2 Fem.Obl-dugong

'I'm cooking the dugong's flipper' (Baker, Horrack et al. 2010, p. 16)

A few more points about noun incorporation are worth bearing in mind:

- incorporation is in most cases optional, so a syntactic paraphrase is almost always possible
- incorporation does not change the valency of the root verb
- the incorporated element can be 'doubled' by an independent noun in the same clause; this may be the same root as the incorporated noun, or one with substantial semantic overlap
- incorporated nouns can be externally modified, as in (18b), where the IN is modified by the possessor noun (the dugong)
- several nouns have a suppletive incorporated form

- some verb roots appear to require an incorporated noun; where no specific noun is available, the MULT derivational prefix *(w)aɾa-~kaɾa-* is used (see below)

As should be apparent from this very brief sketch, the syntax and semantics of noun incorporation in Wubuy is an extremely rich topic. Nonetheless, such issues are not addressed in any detail in the present work. Rather, incorporation constructions will be investigated primarily for their value in probing the complexities of Wubuy's many phonological alternations, and what they suggest about the structure of the lexicon.

Wubuy also features a number of derivational prefixes, some of which are highly frequent in discourse. These attach at the left edge of the verb stem — that is, after the pronominal prefix and before the incorporated noun, if there is one, or preceding the verb root otherwise. The most common of these are the APPLicative *a(K)-* and the MULTiple prefixes *(w)aɾa-~kaɾa-* and *ŋaɾa(K)-*. The applicative increases the verb's valency by one: intransitive verbs become transitive, while transitive verbs raise an 'affectee' argument (benefactive or malefactive) to the status of primary object, as reflected in the pronominal agreement marking on the verb. The multiple prefixes quantify the participants of the verb…

Finally, this entire verb stem, encompassing everything to the right of the pronominal prefix, may be reduplicated, such that a partial copy of the stem is prefixed to (in the most productive pattern) its left edge. Reduplication lends the verb a pluractional (plurality of events) or distributed (greater distribution in time or space) sense. Of particular interest is the fact that the form of this partial copy (the reduplicant) varies according to the sonority of the initial segment of the stem of which it is a copy: initial stops condition a monosyllabic reduplicant, while initial sonorants condition a disyllabic reduplicant. Some examples are given in §2.3.2. As we saw in the previous subsection, reduplication can also target noun stems, in which case it gives a plural reading, and follows the same pattern of sonority-conditioned shape allomorphy.

Loanwords can be assimilated into Wubuy morphology by suffixing a kind of dummy auxiliary *-ṯ(a)* to the loan, followed by TAM inflectional suffixes from the A2 class, as in (19). Most loan verbs come from Kriol, and they are borrowed together with the Kriol transitive suffix *-im*, if present.[4] Such constructions are encountered with great frequency in contemporary Wubuy conversation, though they are avoided in more careful speech, and speakers sometimes would not admit to using them.

[4] I have often heard these speakers produce what appear to be direct calques from English when speaking Kriol, by attaching *-im* to an English verb root. If my suspicions are correct, then we might speculate that some verbs borrowed in to Wubuy with *im-ṯa* come directly from English, rather than Kriol, and that *-im* has simply become part of the default verb borrowing construction. Also, note that in contemporary Kriol spoken in the region, including by contemporary Wubuy speakers, the transitive suffix is usually reduced to *-i*, as in *gaji* 'get', from earlier *gajim*.

(19) nani- *look-after-im-* ṯa -ŋi -jac
3sgM>1pl.R- look.after- Aux -Past2 -Perg
'Back when he looked after us (Excl.)' (NMET 138.1.5)

Having progressed from simple verbs to ones of substantial complexity, we have established the basis for the verbal morphological template in (20). The template is also valid for adjectives, except that for those, no tense-aspect inflection is possible (see next section, §2.2.3). For exemplification and discussion of Wubuy verbs of even greater morphological complexity, including stacking of derivational affixes and interactions between multiple valency-changing operations, see Horrack (2018).

(20) Wubuy morphological template (adapted from Heath 1984, p. 145)

-4	-3	-2	-1	0	+1	+2	+3
[[Pron	[Redup	[Der_1	[Cpd	[[Root	Der_2]	Infl]]]]]	X]
Subj	Red	Appl	N	V	Caus	Tense	Postpositions
Obj		Mult		Adj	Refl	Aspect	Case
Mood		Comit			Recip		

This thesis is concerned with the phonological organisation of morphologically complex words. Verbs are the most complex structures in Wubuy, and phonological alternations apply at the boundaries between each constituent in the template, and sometimes the applicable alternations are different for different domains. For these reasons, most of my analyses focus to a large extent on verbs, and I frequently refer to elements of the template in (20).

2.2.3 Adjectival morphology

Adjectives may be regarded as a distinct word class in Wubuy, sharing formal properties with both nouns and verbs (Heath 1984, p. 152). Adjectival roots can either be used in a referential function, where they take the same noun class inflectional prefix (with or without case suffix) as the noun they modify; or in a predicative function, where they take verbal inflectional prefixes, but without the TAM suffixes that verbs take.

In (21), the adjective root *ɻuŋkal* 'big' appears twice, modifying a different noun in each instance: *wuluk* 'object for slurping honey' and *pulaji* 'billycan'. The first noun is of the Veg class, while the second is of the Resid class; the prefixes on each token of *ɻuŋkal* reflect this.

(21) ŋicaŋ wiɾima-ma-ŋi, **mana-ɻuŋkal**, wuɾu-wulu-wulu-pi-ni::: ŋa,
more 3pl>3Veg.R-get-Past2 Top.Veg-big 3pl>3Resid.R-Red-*wulug*-hit-Past2 and

ana-ɻuŋku-ɻuŋkal-iɲuŋ ana-pulaji wiɾima-jap-a:
Top.Resid-Red-big-Rel Top.Resid-billycan 3pl>3Veg.R-put.in-Past2

'They got some big ones [*wulug* (Veg): material, usually bark, used to slurp honey]. They used the *wulug*. They put them into some big billycans (Resid).' (NMET 142.4.2)

In their predicative use, adjectives take the same inflectional pronominal prefixes as verbs and, like verbs, can take incorporated nouns. These prefixes typically agree with the noun class of the noun they modify, as in (22a), where *wu-~wa:-* '3Neut.R' agrees with the Neut noun *ana:-kuku* 'water'. If there is an incoporated body part, as in (22b), it typically agrees with the possessor. If the incorporated noun is an intransitive subject, the prefix agrees with that noun (22c). This behaviour illustrates the similarity between N+Adj incorporation and N+V incorporation (see Baker & Nordlinger 2008). Note that in these and other examples I follow the practice of Heath (1984) in using an apostrophe to indicate a missing vowel in the surface form. This is used where the underlying representation has a vowel that has coalesced with a preceding vowel, due to a regular process targeting vowel hiatus over morpheme boundaries; thus V-V becomes V:-' on the surface. This apostrophe does not correspond to any unit of sound, but makes the identification of morphological boundaries in the gloss somewhat easier to follow.

(22) a. "kiɾcak!, wa:-'la:ti ana:-kuku" ŋa-jamp-i:ni
no! 3Neut.R-bad Top.Neut-water 1sg.R-speak-Past2

"'No! The water is bad", I said' (NMET 161.13.3)

b. mana-pakaɭaŋ ma:-'la:ti, ŋa-pa-kala:ti
Top.Veg-eye 3Veg.R-bad 1sg.R-eye-bad

'My eyes are bad, I'm blind.' (NMET 47.13.3)

c. ana-l̪a:wu wu-l̪awu-wala:ti anatanu ala:ti
Top.Neut-words 3Neut.R-words-bad Neut.that bad

'words, bad words' (NMET 69.3.4)

Verbs can also be derived from adjectives by means of Factitive and Inchoative suffixes, as can be seen in the derivations headed by *waʈawaʈaʈ* 'strong' in (23b,c) below.

(23) a. wa:ɻi ana-waʈawaʈaʈ ana-l̪a:wu ampiɾi-l̪aɾma-ŋi
NEG TOP.NEUT-strong TOP.NEUT-words 3pl>3NEUT.IRR-Chase-PAST2
'They did not go after the strong (i.e. secret) words' (NMET 47.14.1)

b. niɾima:-'ɾka:-'ɾk-i:ni::: ŋa kiɾcak! ma-waʈawaʈaɳ-ma-ɲ[6]
1pl>3VEG.R-RED-pull-PAST2 and no! 3VEG.R-strong-INCH-PAST1
'We tried to pull it (vehicle = VEG) out but it was too strong' (NMET 163.1.5)

c. waŋku-waɾa-waʈawaʈaʈ-ka:
3NEUT>3pl.R-MULT-FACT.PAST2
'It made them strong all over' (NMET 139.8.9)

For the purposes of the present thesis, it is important to recognise that all of the phonological behaviour associated with verbal morphological domains apply in exactly the same way to the equivalent structures in predicative adjectives. For instance, if a process applies at the boundary between the pronominal prefix and the stem in a verb, it will also apply at the the pronominal prefix-stem boundary in a predicative adjective.

2.2.4 Closed classes: pronouns, demonstratives, kin terms, particles

There are a number of other word classes in Wubuy that the Australianist descriptive tradition usually groups under a broad 'nominal' category, since they tend to function similarly to nouns and take similar morphology: pronouns, demonstratives and kin terms. Each of these has elaborate morphological paradigms in Wubuy, but since they are closed classes, they are seldom discussed in this thesis; my focus is instead on productive alternations. For extensive discussion, the interested reader is invited to consult the relevant chapters in Heath (1984).

For much the same reason, I will have little to say on particles, another large but basically closed word class. However, the NEGative particles *wa:ɻi* and *jaki* do come up fairly frequently in the examples throughout this thesis, given their crucial involvement in morphosyntax; their distribution was set out in Table 2.4 above.

2.2.5 Clitics, 'postpositions'

There are also a number of elements that can follow a word of any class: clitics, and what Heath (1984) calls 'postpositions'. The former are prosodically bound to the preceding word — and so are not separated by pause (cf. Bundgaard-Nielsen & Baker 2020) — but do not undergo any of the segmental alternations proper to the grammatical word domain; the main example is *jampa* 'because'. The postpositions are fully bound, and therefore do undergo alternations — on that

basis, they may just as well be called 'suffixes'. The most frequent of these are *-wukic∼-pukic* 'only; still', *-maci:* 'if', *-maki* 'Evit' and *-maka:* (used in certain negative contexts). These appear in many of the examples discussed here.

2.3 Phonology

The picture I have presented for Wubuy morphology and morphosyntax is not dissimilar from its neighbours, such as Anindhilyakwa (van Egmond 2012), Bininj Kunwok varieties (Evans 2003), Ngalakgan (Baker 2008b), or Ngandi (Heath 1978b). However, these other languages are mostly agglutinative: morphs typically have the same realisation across different phonological environments.[7] This makes the morphological segmentation of complex words relatively straightforward, and makes it easy to predict the phonological outcome of a morphological operation — it generally involves simple concatenation.

Wubuy is strikingly different. It is overwhelmingly the norm for affixation to condition changes in the phonological realisation of the affix, the base, or both. The most comprehensive account we have of these alternations is to be found in Chapter 3 of Heath (1984). His account presents no less than 50 transformational rules in the mode of *SPE* (Chomsky & Halle 1968), ordered in a series, in some cases cyclically.

This analysis, however, has achieved a certain reputation for its difficulty. In part, this is a genuine reflection of Wubuy's complexity. But it is surely also due to certain of Heath's analytical decisions. In particular, his grammar places a high value on using phonological rules to eliminate any and all redundancy, and on keeping the lexicon as small as possible, much like *SPE*. About a quarter of the rules describe alternations that are general and productive: they apply practically wherever their conditions are met, and their conditions are met fairly frequently. The other 'rules' are either restricted to a small number of lexical items (in some cases, just one), or apply at a very abstract or historically remote level of analysis, which even Heath does not claim to be synchronically active. That all these 'rules' are interspersed among one another tends to frustrate efforts to glean a picture of what Wubuy phonology is like in the general case.

This section aims at a more gentle introduction to Wubuy phonology, separating out the productive alternations (§2.3.2) from the rest: lexically restricted phonology (§2.3.3), and rules that work at a more abstract level of analysis (§2.3.4). Then, in §2.3.5, I summarise what is known about Wubuy prosody, an area that is only given brief commentary in Heath (1984). These overviews thus provide the background for the more detailed analyses of particular

[7] Anindhilyakwa phonology is certainly strikingly different from these other GN languages, but it has few alternations beyond vowel assimilation and epenthesis of various kinds.

phonological processes and their interactions in Chs. 3, 4 and 5, and the analysis of prosodic phonology in Ch. 6. However, before we turn to these alternations, we first need to present the segments they implicate.

2.3.1 Segmental inventory

Among vowels, Wubuy exhibits a 'symmetrical' triangular vowel system (Fletcher & Butcher 2014, p. 92) contrasting /i, a, u/ and their long counterparts /iː, aː, uː/, as in Table 2.6. One also finds /oː/ in free variation with /aw/ or /awu/, but only in certain frequent lexical items (Heath 1984, p. 33). Only the /a-aː/ pair is contrastive morph-internally, though there are only a few minimal pairs, and the difference in duration is fairly inconsistently realised. Long vowels are more frequently encountered at the boundaries between morphological constituents — the product of vowel hiatus resolution — and there they tend to be easier to recognise as long.

The consonant inventory (see Table 2.5) is in many ways typical of Australian languages, though somewhat unusual for the Top End region. There are six constrastive places of articulation for stops and nasals, no contrastive voicing, and no fricatives. I use voiceless IPA symbols for the single stop series, though there is wide variability of voicing in production. Unusually amongst Australian languages, the four coronal place contrasts are maintained even in word-initial position; the acoustic, articulatory and perceptual characteristics of the stop place contrasts has been the subject of some detailed phonetic investigation (Bundgaard-Nielsen, Baker, Harvey et al. 2012; Bundgaard-Nielsen, Baker, Kroos et al. 2012; Proctor et al. 2010). The stop and nasal place contrasts are not quite parallel, since the interdental nasal /n̪/ is very infrequent in the lexicon, as Heath notes (Heath 1984, p. 12), though they can arise by assimilation to a following dental stop; his dictionary (Heath 1982) only contains four entries with this segment, none of them high-frequency items. The lateral series contrasts lamino-interdental /l̪/, apico-alveolar /l/ and apico-postalveolar (retroflex) /ɭ/ places of articulation, and there are two rhotics: an alveolar tap or short trill /ɾ/ and a retroflex glide /ɻ/. There is also a palatal glide /j/ and, depending on the analysis, two labiovelar glides /w_1/ and /w_2/, which have exactly the same phonetic qualities but are associated with distinct morphophonological behaviour: under the conditions where Hardening applies, for instance, /w_1/ becomes [k], while /w_2/ becomes [p]. The problem of the two /w/ phonemes is an important part of the analysis in Ch. 3, where I argue for doing away with underlying representations and instead mapping surface forms directly to other surface forms.

The diacritised labiovelar glides are not the only abstract element in the consonant inventory. Heath's documentation also makes use of the 'stop archiphoneme' /K/ and the 'nasal archiphoneme' /N/. These have the equivalent distinctive features and participate in the same alternations as /k/ and /ŋ/, respectively, both of which can trigger various phonological alter-

	labial		lamino-dental		apico-alveolar		apico-retroflex		lamino-alveolar		velar		
stop	/p/	*b*	/t̪/	*dh*	/t/	*d*	/ʈ/	*ḍ*	/c/	*j*	/k/	*g*	(/K/)
nasal	/m/	*m*	(/n̪/	*nh)*	/n/	*n*	/ɳ/	*ṇ*	/ɲ/	*ny-yn*	/ŋ/	*ng*	(/N/)
lateral			/l̪/	*lh*	/l/	*l*	/ɭ/	*ḷ*					
tap/trill					/ɾ/	*rr*							
approximant	/w$_2$/	*w*					/ɻ/	*r*	/j/	*y*	/w$_1$/	*w*	

Table 2.5: The consonant inventory of Wubuy, adapted from Heath (1984, p. 12). Segments in slashes are transcribed in IPA, their counterparts in italics are the symbols used in the community orthography.

	front		central		back	
high	/i/	*i*			/u/	*u*
	/iː/	*ii*			/uː/	*uu*
low			/a/	*a*		
			/aː/	*aa*		

Table 2.6: Wubuy vowel inventory, adapted from Heath (1984, p. 9), given in IPA and in the community orthography.

class	SON	CONT	SYLL	segments
'stops'	–	–	–	{p, t̪, t, ʈ, c, k, K}
'sonorants'	+	–	–	{m, n̪, n, ɳ, ɲ, ŋ, N}
'continuants'	+	+	–	{l̪, l, ɭ, ɾ, ɻ, j, w$_1$, w$_2$}
'vowels'	+	+	+	{a, aː, i, iː, u, uː}

Table 2.7: Terms used to refer to the major manner distinctions in this thesis, the segments they encompass, their features. The stops and sonorants can also be grouped together as 'non-continuants'.

nations word-internally, but must then always succumb to deletion or assimilation. The key difference between the archiphonemes and their true phoneme counterparts is that the latter can appear at the ends of words; the former appear at the ends of bound prefixes and suppletive incorporated nouns, and therefore cannot appear at the right edge of the grammatical word. The archiphonemes are discussed to some extent in the next subsection, in the context of the alternations they are implicated in. The /K/ is also a major focus of Ch. 5, where I show that, whereas deletion of /k/ is variable, deletion of /K/ is absolute; there, I present an analysis that does not make use of archiphonemes, deriving the observed behaviour by subcategorisation constraints instead. As noted earlier, another somewhat abstract symbol I make use of, also borrowed from Heath (1984), is the apostrophe, used where the underlying representation has a morph-initial vowel, but it has coalesced with a preceding morph-final vowel due to a regular process described in the next section.

This thesis assumes that segments bear features (in the sense of Mielke 2008), and that these are targeted by phonological operations. The most consequential features for the analysis are CONTinuant and SONorant, as is indicated by the vertical grouping of the consonants in Table 2.5. Stops are [−SON, −CONT], nasals are [+SON, −CONT], while the laterals, the tap, and the glides are all [+SON, +CONT] — I often refer to this latter category of consonants as 'continuants'. I assume vowels are also [+SON, +CONT], but unlike the continuants, are [+SYLL] rather than [−SYLL]. Stops and nasals can be grouped together as 'non-continuants', since they share the [−CONT] feature. Note also that, since there are no native fricatives, there are no segments that are [−SON, +CONT], and the term 'continuants' is never intended to refer to these. These groupings are summarised in Table 2.7.

For the most part, these divisions are unproblematic. The stops and nasals pattern together as [−CONT] in triggering Hardening, while the nasals and continuants pattern together in determining the appropriate pattern of reduplication. However, there are some wrinkles, particularly in trying to capture the set of consonants that undergo Hardening, which includes all the continuants except /l/ and /l̪/; this issue is considered in more detail in Ch. 3.

2.3.2 Productive alternations

I turn now to the most pervasive of the segmental alternations described by Heath (1984). Note that, when discussing the work of Heath and others, I first present the relevant facts in terms of underlying representations and surface realisations, with rules or constraints mediating between them. I do this out of convenience only, and none of the claims about the content of underlying representations or of rules should be taken for granted. Indeed, one of the major arguments in this thesis is that positing underlying representations at all leads to unwelcome conclusions in Wubuy, and that surface forms should instead be related to one another directly.

By my reckoning, based on close reading of Heath's documentation as well as my own direct investigation, only 15 of his 50 rules can be considered productive, in the informal sense of being generally applicable in their conditioning environment, and being encountered with reasonable frequency; these are listed in Table 2.8.[8] They include all the rules in Heath's grammar which are conditioned only by their phonological environment: P2, P18, P25, P26, P27a, P29, P30ab, P31, P33 and P49. The remaining productive rules are conditioned by both their phonological and morphological environments. There are no fully productive rules that are morphologically conditioned only. Note also that three of these rules could be considered 'process morphology' rather than phonology as such: P39 really just describes the phonology of vocative or interjection forms, P2 and P40 describe different aspects of the reduplicative pattern.

no.	name	relevant chapters in this thesis
P1	ŋu-Epenthesis	3
P2	Initial Reduplication	4
P8	Velar-Insertion	5
P9	w_1-Deletion	(4)
P18	Hardening	3,4,5
P20	Pronominal t-Insertion	
P26	Palatal-Deletion	
P27(a)	Nasal-Assimilation	3,5
P29	Stop-Deletion	5
P30(a,b)	Nasal-Deletion	3,5
P31	Geminate Contraction	5
P39	Interjection-Formation	
P40	Rdp-Shortening	4
P49	VV-Contraction	4

Table 2.8: Productive phonological rules in Heath (1984)

One of the most frequent and general rules is P18 **Hardening**, whereby certain underlying continuant consonants in morph-initial position become hardened to homorganic stops on the surface when they follow either a stop or a nasal — that is, when they follow a [−CONT] segment. In (24b), the underlying palatal approximant /j/ hardens to a palatal stop /c/, conditioned by the final nasal of the pronominal prefix preceding it. In (24a), the pronominal prefix ends in a vowel, which does not condition Hardening, so the underlying continuant surfaces unchanged. Sequences of [−CONT][+CONT] consonants never arise morph-internally, so the Hardening rule can be seen as making derived environments obey morph-internal phonotactics.

[8] Note that I have re-transcribed segments from Heath's orthography to IPA in the names of his rules. Rule P1, for example, is titled 'n^{g}u-Epenthesis' in his work.

w_2	⟷	p
l̪	⟷	t̪
(ɾ	⟷	t)
ɻ	⟷	ʈ
j	⟷	c
w_1	⟷	k

Table 2.9: Continuant-stop pairs

(24) Hardening

a. ŋa-jamp-iːna
/ŋa-jamp-iːna/
1sg.R-speak-NPST2
'I speak' (NMET 7.21.2)

b. ŋaɲ-camp-iːna
/ŋaN-jamp-iːna/
1sg.IRR-speak-NPST2
'I will speak' (NMET 13.1.1)

The continuants which are subject to this rule are shown in Table 2.9, paired with the stops they harden to. There are a couple of peculiarities to be noted about these pairings. Firstly, one notes that the segments which undergo the rule do not form a natural class; the set cannot be captured by [+CONT] as ordinarily understood, since the continuants /l/ and /ɭ/ are not subject to the rule, but /l̪/ is. The alveolar tap/trill /ɾ/ also behaves like these two laterals in root-initial position, but alternates with /t/ in the very frequent Locative suffix: *-ɾuc~-tuc*. The other peculiarity is that underlying /w/ in morph-initial position can harden to either [p] or [k], depending on the lexeme. As mentioned earlier, this leads Heath to posit two underlying labiovelar glide phonemes, /w_1/ and /w_2/, where the former hardens to [k] and the latter hardens to [p] (Heath 1984, p. 12). This distinction is a purely abstract one, since, when they surface as [w], they are phonetically identical: on hearing a morph-initial [w], one cannot predict whether it will harden to [k] or [p]. A reasonable question to ask is why the stop allophones are not taken as underlying, leniting to continuants when they follow [+CONT] segments (see Round 2023). That way, underlying /k/ and /p/ would neutralise to [w].

In reply, Heath points out that there are also roots which *always* have an initial stop (that is, in all environments). To analyse the continuant-stop alternations as lenition, rather than fortition, we would need to explain why some stops lenite and others do not, potentially requiring us to posit two kinds of underlying stop at each place of articulation: /k_1/ and /k_2/, /p_1/ and /p_2/, /$t̪_1$/

and /t̪₂/, etc. Thus, this approach would only amplify the very problem it was meant to solve.

Moreover, underlyingly (i.e. non-alternating) stop-initial stems feed another phonological rule, P1 **ŋu-Epenthesis**. This rule inserts an empty morph /ŋu/ before an underlyingly stop-initial morpheme when it follows either a consonant-final inflectional prefix (of which there exist only three), or a derivational prefix or incorporated noun of any phonological shape (Heath 1984, pp. 35–37). In Chapter 3, I term the morph inserted in the former environment *prefixal ŋu-*, and in the latter environment *linking ŋu-*. Examples (25a) and (25b) show a stop-initial stem following a vowel-final and a consonant-final inflectional prefix, respectively. P1 only applies in the latter case, where the stop follows a prefix-final consonant, though the consonant is then deleted by P30 Nasal-Deletion. In (25c), the same root appears following a derivational prefix, exemplifying the other condition for ŋu-Epenthesis to apply. The root /puɾa/ 'sit' always surfaces with an initial [p], never a [w].

(25) ŋu-Epenthesis

a. ŋa-puɾ-i a:-'ŋa-ɾuc
/ŋa-puɾ-i/ /a-aŋa-ɾuc/
1sg.R-sit-PAST2 OBL.NEUT-camp-LOC
'I was sitting at home' (JT; PN2017-06-29, 01:42:36)

b. nakaŋ nu-**ŋu**-puɾ-i ŋajawi-ɾuc a:-'ŋa-ɾuc
/nakaŋ/ /nuN-(Ø)-puɾ-i/ /ŋajawi-ɾuc/ /a-aŋa-ɾuc/
2sg 2sg.R-(Ø)-sit-PAST2 1sg-LOC OBL.NEUT-camp-LOC
'You were sitting down at my place' (HN; PN2017-06-27, 00:33:52)

c. numpu:-waɾa-**ŋu**-puɾ-aŋaŋ wa:kaka:la
/numpuɾu-waɾa-(Ø)-puɾ-aŋaŋ/ /wa:kaka:la/
2pl.IRR-MULT-(Ø)-sit-NPST1 that side
'All of you sit down on that side!' (MM; PN2017-07-06, 00:04:05)

P1 ŋu-Epenthesis only applies to underlying stops, not ones derived by Hardening. We can see this in (24b) *ŋaɲ-camp-i:na* above: the surface form has a prefix-final nasal followed by a stem-initial stop, but ŋu-Epenthesis has not applied. This illustrates the crucial difference in the behaviour of underlying stops and continuants. As a result of it, minimal pairs which differ only in whether their initial underlying segment is a stop or continuant, such as /cuɻa/ 'to push' and /juɻa/ 'to transport', can still be contrasted on the surface (Heath 1984, p. 37). Following a stop or nasal, /juɻa/ undergoes Hardening, while /cuɻa/ undergoes ŋu-Epenthesis; following a vowel, /juɻa/ remains /juɻa/, while /cuɻa/ either remains unaffected or, after derivational elements,

has /ŋu/ inserted. These two rules conspire to maintain an underlying lexical distinction on the surface. The rule interaction can also be taken as a classic case of opacity (Baković 2007; Kiparsky 1973), an ongoing topic in phonological theory. In rule-ordering terms, the Hardening rule could create new conditions for ŋu-Epenthesis to apply, but is ordered too late to do so. As a result, ŋu-Epenthesis appears to have 'underapplied' in the Hardened forms, as the surface form shows the condition targeted by the rule, and yet the rule has failed to apply.

(26) Realisation of the /cuɻa/, /juɻa/ minimal pair in hardening and non-hardening contexts at the prefix-stem boundary

	/juɻ-iː/ 'transport-NPST2'	/cuɻ-iː/ 'push-NPST2'
Following continuant	[juɻ-iː]	[cuɻ-iː]
Following (underlying) nasal	[cuɻ-iː]	[ŋu-cuɻ-iː]

In Chapter 3, I analyse the interactions between Hardening and ŋu-Epenthesis in detail. There, I show that Heath's functional insight into the way that alternations may support lexical or paradigmatic distinctions is poorly served by most versions of generative phonology and morphology, but can be profitably analysed using concepts drawn from 'abstractive' Word and Paradigm morphology (Blevins 2006). Among the problems for generative phonology are that ŋu-Epenthesis is unnatural (not driven by well-formedness), opaque and outwards-sensitive; alternatively, treating it as insertion of a morpheme instead of phonological epenthesis has its own problems, since it contributes no lexical or grammatical meaning. Addressing all of these problems simultaneously turns out to lead towards a solution, couched in Relational Morphology (Jackendoff & Audring 2020), that offers a novel surface-based account of opaque phonological interactions.

Several of the productive rules are triggered by underlying vowels at morphological boundaries. In general, when two vowels in the input straddle a morpheme boundary, the output is a long monophthong, as per P49 **VV-Contraction**. The output of two identical adjacent vowels is, unsurprisingly, a single long vowel: an /a-a/ sequence becomes [aː], /i-i/ becomes [iː], /u-u/ becomes [uː]. When the input vowels differ, the quality of the output vowel is more or less arbitrarily stipulated by the grammar for each possible combination. For any input involving /a/ and /i/ — /a+i/ or /i+a/ — the output is /iː/. When the input involves an /a/ and the other high vowel /u/, the quality of the output is that of the second vowel in the sequence: /a+u/ becomes /uː/ and /u+a/ becomes /aː/. In practice, the second vowel in a VV sequence is nearly always /a/, as there are so few items that have an initial underlying /i/ or /u/: for /i/, there are seven low-frequency verb roots attested, some of them probably related, and two incorporated nouns; for /u/, there are just three inflections of the verb root *wu* 'hit', and a special compounding form

for ‘give’, *u*, which I have not been able to elicit. A subset of the possible vowel combinations is illustrated in (27); I discuss issues around vowel hiatus resolution in greater detail in Chapter 4.

(27) VV-Contraction

a. nimaː-’l̪ŋa-wanŋ-aː
/nima-al̪ŋa-wanŋ-aː/
3sgM>3Veg.R-mud-step-Past2
‘He stepped in the mud (Veg)’ (NMET 10.15.1)

b. ŋunaː-’nt̪ar̪ma-na
/ŋunu-ant̪ar̪ma-na/
1sg>2sg.R-imitate-Npst2
‘I’m copying you’ (HN; PN2020-10-30; 00:03:43)

c. niː-’ŋa-ŋu-puɾ-i
/ni-aŋa-Ø-puɾ-i/
3sgM.R-camp-Ø-sit-Npst2
‘He was camping’ (NMET 3.1.1)

d. wiɾimuːmana
/wiɾima-u-ma-na/
3pl>3Veg.R-hit-Ø-Npst2
‘They were hitting it (Veg)’ (NMET 155.1.5)

When an underlyingly vowel-initial element follows a noncontinuant (stop or nasal), one of two things must happen. If the preceding morph is derivational, a velar stop /k/ is inserted, according to P8 **Velar-Insertion**.[9] If it is an inflectional pronominal prefix, an alveolar stop /t/ is inserted, and the following vowel is lengthened, according to P20 **Pronominal t-Insertion**. Consider the inflected verb root *ap-iɲ* ‘jump-Past1’. Following a vowel, it undergoes VV-Contraction (28a). Following a consonant-final inflectional prefix, as in (28b), the root is realised on the surface as [taːpiɲ], whereas after a consonant-final derivational element like in (29a,b), it will be [kabiɲ] on the surface. (As an alternative to inserting /t/ and lengthening the vowel, as Heath (1984) suggests, forms like (28b) could be derived by inserting /ta/; due to the VV-Contraction rules, the surface result would be a long vowel with the same quality as the morph-initial underlying vowel.) These two stop insertion processes may be seen as participants in the same conspiracy as Hardening, as they also repair a heteromorphemic sequence of rising

[9] Heath (1984, p. 49) notes that sometimes /ŋ/ is inserted instead, but I have found this behaviour to be much rarer than he suggests.

sonority: although the details of implementation are different in each case, they all target a [−CONT][+CONT] sequence and yield a [−CONT][−CONT] in its place.

(28) Pronominal t-Insertion

a. ŋaː-ʾp-iɲ
/ŋa-ap-iɲ/
1sg.R-jump-Past2
'I jumped' (HN; PN2017-06-25 00:24:40)

b. nun-taːp-iɲ
/nuN-ap-iɲ/
2sg.R-jump-PAST1
'You jumped' (HN; PN2017-06-25 00:24:02)

(29) Velar-Insertion

a. ma-ŋaɾaɳ-**k**api-ɲ
/ma-ŋaɾaɳ-api-ɲ/
3VEG.R-MULT-jump-PAST1
'It all jumped up' (NMET 11.9.1)

b. ana-ʾɾcampal wu-wulŋ-**k**api-ɲ a-ɻaːɾ-wuj
/ana-aɾcampal/ /wu-wulŋ-abi-ɲ/ /a-ɻaːɾ-wuj/
TOP.NEUT-kangaroo 3sg.Neut.-area-jump-Past1 OBL.NEUT-burnt_grassland-ALL
'The kangaroo moved out into the burnt-out grassland' (Heath 1982)

Complicating the picture for velar stops, rule P29 **Stop-Deletion** deletes /k/ when it occurs before any other consonant. Consider a noun root like (30a) *jinak* 'head'. In (30b), this noun is incorporated into a verb, where its final /k/ occurs before the consonant /l/, triggering deletion. (The change in the prefix is handled by another rule which is not immediately relevant here.)

(30) Stop Deletion

a. ana-jinak
/ana-jinak/
TOP.NEUT-head
'(the) head' (NMET 47.11.4)

b. wuː-jina-ɭaɭaːt-aɲi
/wuɾu-jina**k-**ɭaɭaːt-aɲi/
3pl.R-head-peel-PAST2

'The skin on their heads was peeling' (NMET 132.10.2)

When the consonant following the coda /k/ is one of the set of alternating continuants, this /k/ will trigger Hardening, but will still undergo Stop Deletion. Since Heath analyses Wubuy phonology within an *SPE*-style framework, where rules are ordered in series and the output of one rule serves as input to another, this interaction is explained by rule ordering: the Stop-Deletion rule must follow the Hardening rule so that the /k/ is still present at the stage at which Hardening applies. To derive a form like that in (31a), Hardening applies first to the initial /ɾ/ of the locative suffix (since it follows a noncontinuant): the /k-ɾ/ sequence becomes /k-t/. Then Stop Deletion applies to the root-final /k/ which triggered Hardening (since it precedes a consonant): /k-t/ becomes just /-t/. This interaction is classically opaque: the trigger for Hardening, the /k/, is absent from the surface (due to a later rule), so its application appears unmotivated at the surface.

(31) Hardening + Stop Deletion

a. ama-maṯaɭa-tuc
/ama-maṯaɭa**k-ɾ**uc/
OBL.VEG-beach-LOC

'at the beach' (JT; PN2017-06-29 01:38:52)

b. ma-cinak, ma-jina-ṯa-j
/ma**K-j**inak/ /ma-jina**k-ɭ**a-j/
VEG-head 3VEG-head-stand-PAST2

'its head (of paperbark = VEG), its head stood' (NMET 158.1.4)

Recall that Velar Insertion inserts [k] between a [−CONT] segment and a vowel. By definition, then, a morph-final /k/, being itself a non-continuant, should be able to trigger Velar Insertion, just as it triggers Hardening, which would then condition Stop Deletion. As illustrated in (32), the result of Velar Insertion and Stop Deletion applying in series is a single [k] on the surface: /…k-V…/ becomes /…k-kV…/, which finally becomes /…-kV…/. Although the deletion and insertion processes are independently motivated, the result of applying both rules is essentially identical to applying neither — a type of derivation known as a 'Duke-of-York gambit' (McCarthy 2003b; Pullum 1976). In the incorporation construction headed by the adjective *aṇṯurk* 'straight', in (32b), the surface form of that root is transcribed with an initial [k] not present in its underlying

representation (cf. 32a), while the surface form of *jinak* is shown without the final /k/ of its underlying representation. Both the *jina* and *kanṯuɾ(k)* variants are independently motivated, even if their appearance together is brazenly vacuous, as the surface segmental string would also seem to be consistent with a segmentation like *jinak-anṯuɾk*.[10]

(32) Velar Insertion + Stop Deletion

a. aʈapa maː-'nṯuɾk ma-wajama-ŋi
/aʈapa/ /ma-anṯuɾk/ /ma-w$_2$ajama-ŋi/
now 3VEG.R-straight 3VEG.R-proceed-PAST2
'Now it's going straight' (YN; PN2017-07-07 00:20:46)

b. ama-jin**a-k**anṯuɾ-kiː-ni waːɻi
/ama-jina**k-a**nṯuɾk-w$_1$iː-ni/ /waːɻi/
3VEG.IRR-head-straight-FACT.REFL-PAST2 NEG
'its head (of canoe = VEG) wasn't going straight' (YN; PN2017-07-07 00:15:09)

Because a morph-final /k/ must be deleted before any consonant, and Velar Insertion ensures that even underlyingly vowel-initial morphs will be consonant-initial at the time that Stop Deletion applies, /k/ is always deleted when some other element follows in the word, whatever its phonological form. The only place it can surface intact is at the right word edge. There are certain prefixes and bound roots that always appear with a final continuant, and yet also condition Hardening and Velar Insertion in the following constituent. In this respect, they resemble /k/-final morphs, but we can never observe a /k/ directly because, by definition, these bound elements do not appear at the right word edge. For these morphs, Heath's analysis posits a 'stop archiphoneme' /K/, a stop that cannot be observed but must be inferred on the basis of its effects. We have already seen an example of /K/ in the part-whole possession construction in (31b), /maK-jinak/ 'its head'.

As I discuss further in Chapter 5, Stop Deletion turns out to be variably applied amongst contemporary speakers, contra Heath's description, and this variability is structured. It is not applied at all to recent borrowings (as was also noticed by Baker 2009), but is uniformly applied to the stop archiphoneme /K/. Among native roots, there is greater variation, partly conditioned by word class: at the right edge of nouns incorporated into verbal words, deletion of /k/ is applied almost always, while in suffixed nouns deletion fails to apply in the majority of tokens. The difference between the categorical behaviour of /K/ and the variable behaviour of /k/ means that the difference between the two segments is not merely abstract, but is rather

[10]Note also that the final /k/ of *anṯuɾk* is itself deleted in (32b), but only after conditioning Hardening of the initial /w$_1$/ in the FACTitive derivational suffix to [k].

of empirical consequence. I suggest that, instead of positing an abstract underlying segment that is necessarily unobservable, speakers' lexical representations of the morphs previously analysed as having a final /K/ include phonological subcategorisation frames that require the following segment to be [–CONT].

Two of the productive rules in Table 2.8 target the velar nasal /ŋ/: P27 **Nasal-Assimilation** and P30 **Nasal-Deletion**. The first of these causes /ŋ/ to assimilate in place to any following consonant; the second causes it to delete before another sonorant. The assimilation rule is seen in (33); these also show Hardening of the verb root-initial continuant to a stop. The deletion rule is seen in (34). Nasals at other places of articulation do not undergo assimilation or deletion.

(33) Nasal assimilation

a. nu-ŋu-pakaɭaɳ-ʈaji: [ŋ → ɳ]
/nuN-pakaɭaŋ-ɻaji:/
2sg>3RESID.R-eye-spear.NPST2
'You're spearing it in the eye' (LM, PN2017-10-03 00:05:04)

b. jaki nuɾu-pakaɭam-paɭp-i [ŋ → m]
/jaki/ /nuɾu-pakaɭaŋ-w$_2$aɭp-i/
NEG 2sg.R-eye-paint-REFL.NPST3
'Don't paint your (pl.) eyes' (HN, PN2020-11-10 00:03:08)

c. ŋama-maɻaŋ-caɾt-iɲ [ŋ → ɲ]
/ŋama-maɻaŋ-jaɾt-iɲ/
3VEG>1sg.R-hand-poke-PAST1
'Something (VEG) poked me on the hand' (JN, PN2017-10-01.2 00:20:13)

(34) Nasal deletion

a. ŋa-pakaɭa-ŋaɾkiw-aɲ [ŋ → Ø]
/ŋa-pakaɭaŋ-ŋaɾkiw-aɲ/
1sg>3RESID.R-eye-pull_out-PAST1
'I pulled its eye out' (HN, PN2020-10-30 00:31:25)

b. nara-maɻa-na-ni [ŋ → Ø]
/nara-maɻaŋ-na-ni/
1pl>3pl.R-hand-see-PAST2
'We saw their hands' (NMET 157.4.4)

Both of these noun roots, *pakaɭaŋ* 'eye' and *maɻaŋ* 'hand', can appear at the right edge of the word, such as when used as an ordinary (not incorporated) noun — there, their final segment is [ŋ]. This is our firmest piece of evidence for determining that the underlying segment is /ŋ/, since it is the only place that assimilation and deletion are not implicated. There also exist prefixes and bound roots with a final nasal which undergoes deletion and assimilation exactly like /ŋ/. For such cases, following a line of reasoning parallel to the one motivating /K/, Heath posits an underlying 'nasal archiphoneme', /N/. We have already seen examples of /N/ in the pronominal prefixes of (24b), (25b) and (33a), where the surface forms show evidence of assimilation before obstruents and deletion before sonorants.

Rather peculiarly, Heath (1984) acknowledges the assimilation and deletion of these prefix-final nasals, but is inconsistent in his analysis of their underlying representations — sometimes he regards the prefix-final nasals as /ŋ/ or /n/ depending on their historical source (pp. 65, 70), while on other pages they are all /n/ (p. 348), even though their behaviour exactly meets his criteria for /N/, and the alveolar nasal /n/ is not elsewhere subject to assimilation or deletion. This forces him to append stipulations for all of his phonological rules targeting /ŋ/ and /N/, saying that they must also apply to /n/ just where it occurs at the end of an inflectional pronominal prefix. His reasoning for this is that it would make the Pronominal t-Insertion rule appear more phonetically natural, suggesting that the choice of alveolar /t/ as the inserted stop can be traced to a (presumably historical) process of progressive place assimilation, even though place assimilation is otherwise regressive in Wubuy; the prefix-final nasals have to assimilate in place to stem-initial consonants anyway, as we saw in (24), where the prefix-final nasal is /ɲ/ before /c/. This part of Heath's analysis is therefore needlessly complex, and when presenting his work, I treat the pronominal prefixes as underlyingly /N/-final, and ignore the special stipulations in the nasal assimilation and deletion rules. While I ultimately seek to move beyond underlying representations altogether in the analysis pursued throughout this thesis, recasting these prefixes as /N/-final allows us to examine and critique a version of Heath's analysis that is more consistent with itself than his original presentation.

Another frequently encountered process is P2 **Initial Reduplication**, though it is perhaps better seen as a morphological process. It describes the productive reduplication pattern, which is conditioned by the sonority of the initial segment of the base: stop-initial bases reduplicate the first CV, and sonorant-initial bases reduplicate a minimal disyllable (everything up to the second vowel). For disyllabic reduplication, both vowels of the reduplicant are short, and the second vowel is the same as the first. Moreover, if the first vowel of the base is long, it is shortened when preceded by a disyllabic reduplicant, as per P40 **Rdp-Shortening** — the status of this rule is somewhat dubious though, given the extreme paucity of bases that can feed it.[11] Some

[11]There may also be other explanations for perceived vowel shortening, such as the general tendency for words

examples are given in (35) and (36); these also show that reduplication can feed ŋu-Epenthesis and Hardening.[12] Reduplication on verbs generally indicates greater distribution in time and/or space, or pluractionality, and is often used to background an event in discourse; on nouns, it indicates plurality; on adjectives, it is less common but can indicate plurality or intensification. That said, the semantics of reduplicative structures is not particularly well understood, and would benefit from further research.

(35) Reduplication (stop-initial base)

a. na-ŋu-puɾ-aː
/naN-puɾ-aː/
12du.Irr-sit-Npst2
'We will sit' (NMET 20.18.4)

b. na-ŋu-pu-puɾ-aː
/naN-pu-puɾ-aː/
12du.Irr-Red-sit-Npst2
'We will sit (for some time)' (NMET 20.18.5)

(36) Reduplication (sonorant-initial base)

a. ŋa-jampiː-na-jiɲuŋ
/ŋa-jampiː-na-jiɲuŋ/
1sg.R-speak-Npst2-Rel
'What I'm talking about' (NMET 47.23.7)

b. ŋa-jampa-jampiː-ni-jiɲuŋ
/ŋa-jampa-jampiː-ni-jiɲuŋ/
1sg.R-Red-speak-Past2-Rel
'What I was talking about' (NMET 47.23.2)

c. waːɻi ŋaɲ-campa-jampiː-ni
/waːɻi/ /ŋaN-jampa-jampiː-ni/
not 1sg.Irr-Red-speak-Past2
'I didn't speak [about such things]' (NMET 47.22.4)

with more segments to be shorter in duration. Vowel length within morphs is, in any event, quite elastic.

[12] Where linking *ŋu-* appears inside a stem following a monosyllabic, sonorant-initial element, and this stem is reduplicated, *ŋu-* will appear in the reduplicant. Thus, linking *ŋu-* feeds reduplication.

To the best of my knowledge, the only other language reported to have a reduplicant shape alternation conditioned by the sonority of the base-initial segment is neighbouring Anindhilyakwa, Wubuy's closest genetic relative. In addition to being typologically unusual, the pattern is not predicted by the most influential theories of reduplication in the phonological literature, such as Base-Reduplicant Correspondence Theory. The challenges that this reduplication pattern poses for phonological theory are examined in detail in Chapter 4.

We have seen from the preceding discussion that producing a phonologically correct Wubuy word is never a simple matter of filling out a morphological template with invariant roots or affixes, taken directly from a minimal lexicon. Almost always, some phonological modification will have to apply at each morphological boundary. For ease of reference, these modifications are summarised in Table 2.10, and some example derivations, illustrating the most frequent alternations, are given in Table 2.11.

If left side is:	ending in:	and right side is:	beginning with:	Output will be:	Heath's rule
any	non-continuant	any	continuant*	stop	P18 Hardening
pron. prefix	C	any	stop	C-[ŋu]-stop	P1 ŋu-Epenthesis
deriv. prefix	any	any	stop	X-[ŋu]-stop	P1 ŋu-Epenthesis
pron. prefix	C	any	V	C-tV:	P20 Pronominal t-Insertion
deriv. prefix	C	any	V	kV	P8 Velar-Insertion
any	ŋ	any	sonorant	sonorant	P30 Nasal Deletion
any	ŋ	any	stop	homorganic nasal-stop	P27a Nasal Assimilation
any	C_i	any	C_i	C_i	P31 Geminate Contraction
any	V_i	any	V_i	V_iː	P49 VV-Contraction
any	a	any	u	uː	P49 VV-Contraction
any	u	any	a	aː	P49 VV-Contraction
any	a	any	i	iː	P49 VV-Contraction
any	i	any	a	aː	P49 VV-Contraction
any	k	any	C	C	P29 Stop-Deletion

Table 2.10: Summary of the most productive phonological alternations, showing the conditions they target and their surface effects. ∗The 'continuants' that participate in Hardening are only those in Table 2.9 — it excludes /l,ɭ/ and, depending on the item, /ɾ/.

Citation form	V_	N]$_{\text{Infl}_}$	−CONT]$_{\text{Der}_}$	+CONT]$_{\text{Der}_}$
/api/	ni:pi-ni	nun-ta:pi-ni	wu-wulŋ-kapi-ɲ	ni-caŋkal-api-ɲ
‘jump’	3sgM.R-jump-PAST2	2sg.R-jump-PAST2	3NEUT.R-area-jump-PAST1	3sgM.R-ankle-jump-PAST1
/ɻa/	nu-ɻa-ni	nuɳ-ʈa-ŋ	ŋaŋku-pakaɭaɳ-ʈa-ji:	nu-ɻaŋka-ɻa-ni
‘spear’	3sgM>3sgM.R-spear-PAST2	2sg.R-spear-PPunc	1sg>3NEUT.IRR-eye-spear-NPST2	3sgM.R-asleep-spear-PAST2
/puɾa/	ŋa-puɾ-i	nu-ŋu-puɾ-i	ni-wuɾic-ŋu-puɾ-i	ma-ɳuka-ŋu-puɾ-i
‘sit’	1sg.R-sit-PAST2	2sg.R-sit-PAST2	3sgM.R-chest-Ø-sit-PAST2	3VEG.R-stone-Ø-sit-PAST2

Table 2.11: Comparison of surface forms of /api/ ‘jump’, /ɻa/ ‘spear (V)’ and /puɾa/ ‘sit’ in different preceding segmental contexts.

2.3.3 Lexically-restricted alternations

A large number of rules appear to be motivated only to handle some idiosyncratic behaviour relating to particular lexical items. These are listed in Table 2.12.

In many cases these rules seek to account for some unusual behaviour observed in a single root. For example, P43 **Length-Shift** handles the exceptional behaviour of just one root, /ku:ku/ ‘water’. When preceded by the NEUT noun class prefixes /a-/ or /ana-/, the result is [a:-kuku] or [ana:-kuku], respectively, rather than the expected forms [a-ku:ku] and [ana-ku:ku]:[13] in the attested forms, the initial stem vowel is shortened and the preceding vowel in the prefix undergoes compensatory lengthening. No other nouns with initial long vowels exhibit this behaviour. Likewise, P6 **w$_1$-Insertion** only applies to the verb root *a:ju* in derived environments; P17 **ʈiku-Lenition** has the purpose of deriving the incorporated noun *ɻiku* or *ɻuku* ‘raw; dead; unconscious’ from the adjectival root *ʈiku* ‘unripe, raw’; P47 **ji-Elision** applies to the NPST2 inflection of the verb *ji* ‘give’, when not preceded by an overt pronominal prefix; P32 **l̪-Delateralisation**’s sole motivation is to derive an optional PAST2 inflection of the verb *pil̪la* ‘to be tilted’.

Some apply to more than one item, but are not generalisable beyond the few items they apply to. P37 **V-Assimilation**, for instance, bundles together various unsystematic and isolated examples of vowel harmony in the inflectional morphology, “but it is not suggested that they cohere into a single rule in any psychologically or even historically real sense” (Heath 1984, p. 79). Two of the lexically restricted rules are generally unproductive, but do shed some light on the nature of the productive rules, in that they point out a small class of items to which the productive rules fail to apply. P16 **Lenition** and P19 **Leftward Hardening** both seek to account for lexical items to which the more productive rules P1 ŋu-Epenthesis and P18 Hardening either over- or under-apply (see Chapter 3), as well as to cases where P2 Initial Reduplication applies

[13] In fact, in my experience it is not entirely unusual to hear [a-ku:ku] and [ana-ku:ku], and I have recorded one token of the latter, uttered spontaneously (MM; PN2017-07-06 00:57:14).

no.	name
P3	Final Reduplication
P4	Homorganic Semivowel-Deletion
P5	Initial Semivowel-Insertion
P6	w_1-Insertion
P12	Stem w_2-Deletion
P13	w_1-Palatalisation
P14	jɾ-Contraction
P16	Lenition
P17	ʈiku-Lenition
P19	Leftward Hardening
P21	Derivational t-Insertion
P25	ɾ-Deletion
P28	Semivowel-Deletion
P32	l̪-Delateralisation
P34	Feminine Prefix-Truncation
P35	Prefix-Truncation
P36	ɾV-Truncation
P37	V-Assimilation
P42	Adverb-Shortening
P43	Length-Shift
P45	Verb-Lengthening
P47	ji-Elision

Table 2.12: Lexically restricted phonological rules in Heath (1984).

in an idiosyncratic fashion (see Chapter 4).

Of course, 'morphological irregularity [is] a well-known fact of life' (Burzio 2002, p. 162), and we should not be in the business of ignoring facts. However, to treat these sundry lexical exceptions — genuine as they may be — as reflecting the application of symbolic transformational rules, on the same footing as the productive rules, reflects a fairly extreme theoretical commitment to eliminating redundancy — in this case by capturing all observed contextual variation in form by means of the grammar, not the lexicon. This demonstrates the extent of Heath's (1984) commitment to Morpheme-Invariance, in the terms of Chapter 1. I will consider various idiosyncratic lexical exceptions to the main productive patterns when I analyse them in Chapters 3, 4 and 5, but I do not treat them as separate rules or processes as Heath does.

no.	name
P3	Final reduplication
P11	Prefixal w_2-Deletion
P12	Stem w_2-Deletion
P23	Denasalisation
P24	rn-Contraction
P27b	Nasal-Assimilation
P30c	Nasal-Deletion
P32	lh-Delateralisation
P37	V-Assimilation
P38	V-Ablaut
P41	Nonsg-Shortening
P42	Adverb-Shortening
P45	Verb-Lengthening
P47	yi-Elision
P48	w2an-Contraction

Table 2.13: 'Maximally abstract' phonological rules in Heath (1984)

2.3.4 Rules which assume a maximally abstract analysis

For some inflectional paradigms, Heath provides what he calls a 'maximally abstract analysis', deriving almost all inflections from a small set of invariant atoms, though it comes with the (double-underlined) disclaimer that he makes 'no strong claims about the psychological reality of this analysis'. So, for the pronominal prefix forms, elsewhere treated in his work as unanalysable portmanteaux, his maximally abstract analysis breaks them down into 'component morphemes': for example, *niɾi-* '2pl>3NEUT.R', breaks down into /na-w_2V-ɾi-Ø-/. The derivation of the portmanteau forms from its abstract components requires a special set of rules, both morphological and phonological. Phonological rules that implement this maximally abstract analysis are listed in Table 2.13.

It is difficult to know how to take rules such as these: even within the realm of his maximally abstract analysis, the rules are far from completely general, usually applying just to certain component morphs or combinations of them. Heath often qualifies these rules with a comment like 'synchronically this relationship is at best opaque' (Heath 1984, p. 68), or that they should be 'regarded as at most of historical (not synchronic) interest' (p. 53). I do not discuss any of these rules in this thesis, since the component morphemes they operate on are extremely abstract, and I have not seen evidence that they have any bearing on speaker behaviour. Indeed, Bundgaard-Nielsen & Baker (2020) present experimental evidence that these abstract morphological boundaries have no salience for Wubuy speakers, unlike the boundaries between the

pronominal prefix and the stem, or between an incorporated noun and verb root.

2.3.5 Prosody

We have seen from the preceding discussion that the rich segmental alternations in Wubuy complex words have been given significant attention in the existing literature. On the suprasegmental phonology, however, the existing literature leaves us largely in the dark.

We have three main sources, for which I provide only brief summaries here; see Chapter 6 for more extensive commentary. Heath (1984) provides an impressionistic description of the main intonational tune types and some analysis of their role in organising discourse, based on one of the texts in Heath (1980). He offers a few general comments on word-level prosody, but is reticent to completely detach the prominence he identifies from intonation — perhaps wisely.

Hore's (1981) article is a serial rule-based account of Wubuy stress — though again, it is unclear whether the prominences identified relate to word-level stress or higher-level intonation, given that, as far as we can tell, the analysis appears to proceed from individual tokens of words uttered in isolation (see Gordon 2014). There are some further difficulties. Like Heath's (1984) account of phonological alternations, Hore (1981) attempts to capture everything in rules, but it is difficult to glean any broad generalisations from these rules: many of them come with a long list of exceptions (and exceptions to the exceptions), and their interactions are often complex. Also, the work pre-dates many of the major developments in theories of metrical structure (e.g., Hayes 1995) and the rules therefore depend on some unwieldy theoretical constructs — for instance, each syllable is organised into a 'stress group', an elastic unit which is somewhere between a foot and a prosodic word, and has some unclear relation to morphological structure.

Approaching prosodic organisation from a different angle, Bundgaard-Nielsen & Baker (2020) show that Wubuy speakers sometimes pause inside the grammatical word, but only at major morphological boundaries, such as between a pronominal prefix and verb stem, or between incorporated noun and verb root. This is reflected in speaker judgements of verb tokens with artifically-inserted pauses, where speakers tended to accept tokens with pauses at these 'licit' morphological boundaries (with some speakers preferring them to tokens without pause), while rejecting tokens with pauses at other locations in the word. On the basis of this, the authors argue that the grammatical word in Wubuy potentially consists of multiple planning units in production and perception. Their findings also align with an observation that has a time-honoured pedigree in Australian linguistic description, namely that prosody is constrained by morphological structure (see Baker 2014b).

What remains lacking is a comprehensive account of the intonational phonology. This is the aim of Chapter 6 of the present thesis. In that chapter, I provide an Autosegmental-Metrical analysis of Wubuy intonation based on a corpus of three texts, establishing an inventory of

pitch accents and boundary tones, a prosodic hierarchy to which these tones may be anchored, and an analysis of how these tunes are mapped to morphological constituents.

3 Alternations in the lexicon: ŋu-Epenthesis and Hardening

3.1 Introduction

In Chapter 2, I briefly introduced what Heath calls ŋu-Epenthesis, a process which inserts the meaningless string /ŋu/ before underlyingly stop-initial roots under certain conditions. In this chapter, I examine the phenomenon of ŋu-Epenthesis and its interaction with other phonological processes — especially continuant-stop alternations — more closely.

Given that this *ŋu-* element is not associated with any lexical or grammatical meaning, it raises an important question: if it doesn't mean anything, why should it exist at all? As is well known ever since Hockett (1954), 'empty morphs' such as these are troubling for post-Bloomfieldian morphemic theories that regard words as built up out of morphemes, where morphemes are considered the 'smallest units of meaning'. Empty morphs, since they have no meaning, fail this criterion.

This process is interesting for several reasons. Chief among them is the fact that the appearance of *ŋu-* is phonologically conditioned but does not improve phonological well-formedness — it is clearly not the result of any known articulatory or perceptual phonetic pressures, and it does not repair any phonotactically dispreferred or 'marked' underlying sequence. Moreover, it interacts with other phonological processes in a manner which is opaque. Opacity is a well-known problem for classical Optimality Theory (Prince & Smolensky 2004), but, as I show in this chapter, the particular opaque interactions analysed here cannot be overcome by any of the the various extensions of OT that have been proposed to address opacity, because *ŋu-* is opaque in two different ways at the same time. And, because ŋu-Epenthesis is sensitive to a two-sided context, it violates the 'no lookahead' principle, a possibility ruled out by theories that derive words from the root outwards.

To be sure, none of these properties — emptiness, unnaturalness, opacity, and outwards-sensitivity — is unique to the alternation under consideration here, and each has been discussed separately in the literature. But the fact that ŋu-Epenthesis encompasses all these properties simultaneously forces us to reckon with problems at a much deeper level than the details of

this particular case, and precludes us from proposing solutions that address only a subset of the difficult properties in isolation from the others.

As I show in this chapter, ŋu-Epenthesis puts the contours of these deeper problems into relief, and provides crucial evidence for the formulation of an alternative approach capable of overcoming them. I diagnose the obstacles to a satisfactory analysis of ŋu-Epenthesis as stemming from two prevailing commitments in mainstream phonological theory: namely, that the job of morphology is to string invariant underlying morphemes together — from the root outwards, one layer at a time — and that the job of phonology is to modify the resulting string so as to 'repair' any violations of well-formedness therein.

In order to forge an alternative, I bring together the insights of some recent 'abstractive' approaches to morphology (Ackerman, Blevins et al. 2009; Blevins 2006; Booij 2010) and subcategorisation approaches to phonological alternations (Orgun 1996; Paster 2006; Yu 2007). I argue that *ŋu-*, despite being largely 'empty' of lexical or morphosyntactic content, is nonetheless informative, so long as we consider it in the context of the constructions it appears in, and the paradigmatic relations holding between those constructions. Like other empty morphs that have been analysed in the literature — theme vowels in Latin (Aronoff 1994), stem augments in Pitjantjatjara/Yankuntjatjara (Wilmoth & Mansfield 2021) — *ŋu-* in Wubuy provides a disambiguating function by reducing paradigmatic entropy in a context threatened by widespread phonological neutralisation. It thus supports the phonological organisation of the lexicon, in a way which cannot be captured by comparing the underlying and surface forms of a single word in isolation.

I then integrate these insights into an analysis couched in Relational Morphology (Jackendoff & Audring 2020), a declarative, constituent-structure based theory that extends Construction Morphology (Booij 2010; Gurevich 2006) in several important respects. In this model, morphologically complex words are conceived of as constructions that are stored in the lexicon, and that motivate abstract schemas that generalise morphological patterns to new contexts. These constructions are related to one another, forming paradigms.

However, the place of phonology, especially productive phonological alternations, has long remained seriously under-theorised in construction-based models of grammar. One of the aims of this chapter is to fill that gap. I show that, although the original formulation of Relational Morphology (Jackendoff & Audring 2020) pays little attention to phonological alternations, the formal tools it develops for other purposes are sufficient to reach a satisfactory analysis of ŋu-Epenthesis and its interaction with other processes in Wubuy. As it turns out, the solution to the ŋu-Epenthesis problem leads towards a general solution to analysing opacity without underlying representations or serial transformations, a result that has so far been impossible to achieve in standard generative phonology. Indeed, the solution points to a way of doing

phonology that does not fundamentally rest on morphemes — a formal object that has long been largely abandoned by morphological theory.

The chapter is organised as follows. In the next section, I set out the data that my analysis seeks to explain, beginning with the ŋu-Epenthesis pattern on its own (§3.2.1), then in relation to other relevant phonological alternations, especially Hardening (§3.2.2). The empirical presentation also encompasses some exceptions to the general process (§3.2.4), and the extent of minimal pairs of roots or stems in the lexicon that are distinguished by whether they can be the input to ŋu-Epenthesis (§3.2.6).

The rest of the chapter is more theoretical in orientation. In §3.3, I examine what I regard to be the obstacles to a satisfying account within the confines of mainstream generative phonology, and the models of the phonology-morphology interface it assumes: unnaturalness, emptiness, opacity and outwards-sensitivity. I then argue, in §3.4, that these obstacles are in large part a consequence of assuming that morphology strings together morphemes, and phonology repairs the result. There, I show that *ŋu-* is far more profitably analysed if morphology instead concerns relations between whole words. Following that, in §3.5, I develop an account of opaque alternations as surface-to-surface correspondence, using Relational Morphology; this enables an explicit formal account of *ŋu-* in §3.6. I offer final thoughts and suggestions for further research in §3.7.

3.2 Descriptive Overview

3.2.1 ŋu-Epenthesis

Heath's (1984) grammar describes a process, ŋu-Epenthesis, whereby the string /ŋu/ is inserted preceding stops under certain conditions, as in (37). This epenthetic string does not appear to make any semantic or morphosyntactic contribution to the resulting word.

(37) The string /ŋu/ is inserted preceding a morpheme-initial stop, when it follows either:

a. An inflectional prefix ending in a consonant, or

b. A derivational prefix of any kind

In order to distinguish between these two conditions in the discussion to follow, I will refer to the *ŋu-* of (37a) as *prefixal ŋu-*, since it appears at the boundary between the inflectional prefix and the stem. The other, that of (37b), I will call *linking ŋu-*, since it bears some similarities to compound linking elements in other languages.

We can observe this process at work in the examples below. The examples in (38) show the stop-initial root *puɾa* 'sit' following vowel-final inflectional prefixes: the first singular realis

/ŋa-/ in (38a) and the third plural realis /wuɾu-/ in (38b).

(38) No *ŋu-* following vowel-final prefix

a. ŋa-puɾ-i a:-'ŋa-ɾuc
/ŋa-puɾ-i/ /a-aŋa-ɾuc/
1sg.R-sit-PAST2 OBL.NEUT-camp-LOC
'I was sitting at home' (JT; PN2017-06-29, 01:42:36)

b. wuɾu-puɾ-a:
/wuɾu-puɾ-a:/
3pl.R-sit-NPST2
'They're sitting [staying there]' (MM; PN2017-07-06, 00:43:43)

Contrast this with (39a), where this same root follows the first person singular irrealis prefix, underlyingly /ŋaN-/.[1] Since the /ŋaN-/ prefix is consonant-final and the following root is stop-initial, the string /ŋu/ must be inserted between them as per (37a), giving the intermediate derivation /ŋaN-ŋu-puɾ-i/. Because — according to a general phonological process — /N/ and /ŋ/ are deleted before other nasals, the prefix-final /N/ is deleted, giving the final output [ŋa-ŋu-puɾ-i]. Similar examples are given in (39b,c).

(39) Prefixal *ŋu-* following N-final prefix

a. wa:ɻi ŋa-**ŋu**-puɾ-i mana-l̪akaja-tuc
/wa:ɻi/ /ŋaN-(Ø)-puɾ-i/ /mana-l̪akajak-ɾuc/
NEG 1sg.IRR-(Ø)-sit-PAST2 TOP.VEG-saltwater-LOC
'I wasn't sitting by the water' (JT; PN2017-06-29, 01:42:08)

b. nakaŋ nu-**ŋu**-puɾ-i ŋajawi-ɾuc a:-'ŋa-ɾuc
/nakaŋ/ /nuN-(Ø)-puɾ-i/ /ŋajawi-ɾuc/ /a-aŋa-ɾuc/
2sg 2sg.R-(Ø)-sit-PAST2 1sg-LOC OBL.NEUT-camp-LOC
'You were sitting down at my place' (HN; PN2017-06-27, 00:33:52)

c. na-**ŋu**-puɾ-a: aɲca:pu-ɾuc
/naN-(Ø)-puɾ-a:/ /aɲca:pu-ɾuc/
12du.IRR-(Ø)-sit-NPST2 together-LOC
'Let's (you and I) sit together' (YN; PN2017-07-06, 00:16:47)

[1] The /N/ is a 'nasal archiphoneme', and functions like the velar nasal /ŋ/. Likewise, the /K/ here and throughout denotes a 'stop archiphoneme', which in Heath's analysis functions identically to /k/. For discussion of these see chapter 5.

This same inflected verb root, /puɾ-i/, appears in (40a), this time preceded by a derivational prefix, /ŋaraK-/ MULT.[2] This meets the condition for linking *ŋu-* as stated above; the result on the surface is thus [wuru-ŋara-ŋu-puɾ-i]. In (40b,c), we have further examples of linking *ŋu-*, where the derivational element — a different MULT prefix in (40b), an incorporated noun in (40c) — ends in a vowel. (An additional, prefixal *ŋu-* does not apply in any of the examples in (40), since the inflectional pronominal prefix in each is vowel-final.)

(40) a. wuru-ŋara-**ŋu**-puɾ-i
/wuru-ŋaraK-(Ø)-puɾ-i/
3pl.R-MULT-(Ø)-sit-PAST2
'They all sat' (HN; PN2017-06-27, 00:37:23)

b. numpuː-wara-**ŋu**-puɾ-aŋaŋ wa:kaka:la
/numpuru-wara-(Ø)-puɾ-aŋaŋ/ /wa:kaka:la/
2pl.IRR-MULT-(Ø)-sit-NPST1 that.side[4]
'All of you sit down on that side!' (MM; PN2017-07-06, 00:04:05)

c. ma-ɳuka-**ŋu**-puɾ-i
/ma-ɳuka-(Ø)-puɾ-i/
3VEG.R-stone-(Ø)-sit-PAST2
'The stone was sitting' (HN; PN2017-06-27, 00:31:57)

Heath (1984) refers to this inserted *ŋu-* as a 'meaningless morpheme'. Unlike other morphemes, it cannot be inserted into words by the morphological component of the grammar, since it has no grammatical or semantic meaning to realise. Instead, its appearance is conditioned by phonological factors (though the phonological conditioning differs slightly depending on its morphological environment). Heath therefore analyses *ŋu-* as the outcome of a phonological rule, which can be formulated as follows: (41a) concerns what I am calling prefixal *ŋu-*, while (41b) conerns linking *ŋu-*. Both rules target a stop in the environment on the right-hand side, but only (41a) also requires that the left-hand environment end in a consonant.

(41) ŋu-Epenthesis (following Heath 1984, p. 36)

a. $\emptyset \rightarrow$ ŋu / C]$_{\text{Infl}}$ _ $\begin{bmatrix} -\text{CONT} \\ -\text{SON} \end{bmatrix}$

b. $\emptyset \rightarrow$ ŋu /]$_{\text{Deriv}}$_ $\begin{bmatrix} -\text{CONT} \\ -\text{SON} \end{bmatrix}$

[2] Recall that anything between the inflectional pronominal agreement prefix and the root counts as 'derivational' for present purposes.

The two kinds of /ŋu-/ differ not only in which part of the word they apply to and under what conditions, but also, as we'll see, in their generality of distribution: while there are various lexical exceptions to linking *ŋu-*, prefixal *ŋu-* appears to be inescapable for stop-initial stems. Note also that *ŋu-* occurs in nouns, verbs and adjectives, but because the structural condition is met far more frequently in verbs, I will concentrate primarily on verbs in this chapter. I repeat here the verb template from Chapter 2, with slots for *ŋu-* added. Prefixal *ŋu-* appears after slot -4, while linking *ŋu-* can appear between slots -2 and -1, or between -1 and 0. A noun template is given in (43).

(42) Wubuy verb/adjective template (adapted from Heath 1984, p. 145)

-4	-3	-2	-1	0	+1	+2	+3
[[Pron **ŋu**$_1$	[Redup	[Der$_1$ **ŋu**$_2$	[Cpd **ŋu**$_2$	[[Root	Der$_2$]	Infl]]]]]	X]
Subj	Red	Appl	N	V	Caus	Tense	Postp.
Obj		Mult		Adj	Refl	Aspect	Case
Mood		Comit			Recip		

(43) Wubuy noun template

-3	-2	-1	0	+1
[[N class	[Redup	[Deriv class **ŋu**$_2$	[Root]]]	Case]

The next section is concerned with alternations applying between the verbal/adjectival pronominal prefix and the stem: in addition to prefixal *ŋu-*, there are continuant-stop alternations affecting the stem-initial consonant, and assimilation and deletion processes applying to the prefixes. Linking *ŋu-* is treated in the section to follow, where some lexical exceptions to the predominant pattern are attested. Then, in §3.2.6, I consider some of Heath's functional explanations for ŋu-Epenthesis, and set out some of the general theoretical issues that this functional account raises, issues that will be addressed in detail throughout the rest of the chapter.

3.2.2 Hardening

Although this chapter is primarily concerned with ŋu-Epenthesis, a key plank of the analysis is the proposition that ŋu-Epenthesis is only intelligible in relation to other productive alternations. And so, in this section I provide the necessary background on the most important of those other processes: Hardening, Nasal Assimilation, and Nasal Deletion.

Hardening realises underlying continuant consonants as stops on the surface when they follow an underlying stop or nasal. Hardening can be seen in word pairs like those in (44): the stem which has an initial palatal glide /j/ in (44a) appears instead with an initial palatal stop /c/ in (44b). (The behaviour of the nasal will be discussed later on.)

(44) a. ŋa-jamp-iːna
/ŋa-jamp-iːna/
1sg.R-speak-NPST2
'I speak' (NMET 7.21.2)

b. ŋaɲ-camp-iːna
/ŋaN-jamp-iːna/
1sg.IRR-speak-NPST2
'I will speak' (NMET 13.1.1)

The continuant-stop alternation can be formulated as a rule like that in (45). It takes any continuant consonant following a stop or nasal and turns it into a stop. It applies in (44b) because the stem-initial underlying segment is the continuant /j/, and it follows a prefix-final nasal. Crucially, note that a nasal-stop sequence over the prefix-stem boundary is exactly the condition targeted by prefixal *ŋu-*, as we saw above. However, as examples like (44b) demonstrate, ŋu-Epenthesis does not apply to stops derived by Hardening.

(45) Hardening (after Heath 1984)

$$\begin{bmatrix} -\text{SYLL} \\ +\text{SON} \\ +\text{CONT} \end{bmatrix} \rightarrow \begin{bmatrix} -\text{SYLL} \\ -\text{SON} \\ -\text{CONT} \end{bmatrix} / \; [-\text{CONT}] - _$$

The rule as formulated above, however, is an oversimplification. In fact, the rule does not apply to all continuant consonants; the continuants that participate in the rule are only those listed in (46) along with the 'paired stop' they alternate with. We see that while all of the stop phonemes in Wubuy are paired with some continuant, not all continuant phonemes are paired with some stop — /l/ and /l̪/ are exempt from Hardening, and pattern instead with the sonorant non-continuants (i.e., nasals), which do not alternate. The rule in (45) therefore does not apply to a natural class.

(46) Alternating continuant-stop pairs

w_2	⟷	p
l̪	⟷	t̪
(ɾ	⟷	t)
ɻ	⟷	ʈ
j	⟷	c
w_1	⟷	k

The second peculiarity to note about the pairings in (46) is that underlying /w/ can harden to either [p] or [k], depending on the lexeme. This leads Heath to posit two underlying labiovelar glide phonemes, /w_1/ and /w_2/, where the former hardens to [k] and the latter hardens to [p] (Heath 1984, p. 12). In most other contexts, including word-initially, both /w_1/ and /w_2/ surface as [w]. So far as I can tell, there is no acoustic or articulatory difference between the [w] from /w_1/ and the [w] from /w_2/.[5] This is a potentially difficult task for the learner of Wubuy: presumably, on hearing a morpheme-initial /w/, one cannot know in advance whether it will become a [p] or a [k] in the conditioning environment. And although this issue only concerns two phonemes, roots with initial /w_1/ and /w_2/ make up as much as a third of the Wubuy verb lexicon (Baker 2009).

As Baker (2009) points out, such an analysis is clearly in violation of Kiparsky's Revised Alternation Condition, which asserts that obligatory neutralisation rules apply only in derived environments (Kiparsky 1973).[6] That is, the putative /w_1/–/w_2/ contrast is *only* realised in derived environments, wherever Hardening applies, while being neutralised everywhere else.[7]

Indeed, these alternations probably reflect a historical process of lenition, not fortition. The available evidence from comparative reconstruction suggests that, roughly speaking, fortis

[5] Although there does appear to be a morphophonological difference, in that /w_1/ can undergo an optional deletion rule where intervocalic, while /w_2/ tends to resist this (see Heath 1984, p. 50ff). Wubuy is not entirely unusual here. In Round's (2023) typology of lenition/fortition alternations in Australian languages, alternations of /w/ with both /k/ and /p/ are the most common type of one-to-many alternation across the continent, while alternations of /k,w,Ø/ are more common than /p,w,Ø/.

[6] Though not a violation of his earlier Alternation Condition: "Obligatory neutralization rules cannot apply to all occurrences of a morpheme" (Kiparsky 1982b).

[7] Baker & Harvey in press, p. 159 mention several other Australian languages showing similar continuant-stop alternations in which the paired stop allophone is wholly or partly unpredictable, and yet the existing analyses posit the continuant as underlying. In Mangarrayi, for instance, there is claimed to be not only a /w_1/ and /w_2/, as in Wubuy, but also a /j_1/ and /j_2/, which alternate with /k/ and /c/, respectively; note that the continuant forms are the ones that appear in the 'elsewhere' environment. Similar claims have been made for Worrorra, Marra, and some Yolŋu varieties — though accounts of the Yolŋu alternations usually analyse them as lenition, not fortition (Chong 2011). Relatedly, Mansfield & Green 2021 discuss the stop/fricative contrast in the Western Daly language Marri Tjevin. This manner contrast is also neutralised word-initially but maintained following prefixation. The methodological and conceptual issues of analysing a given alternation as lenition or fortition are discussed in some detail by Round (2023), in the context of a detailed typology of such alternations.

stops in the proto-language became lenis stops, while historical lenis stops were reduced to continuants (Heath 1978a, p. 38; see also van Egmond & Baker 2021). These correspondences are set out in (47). This lenition process was inhibited following stops and nasals, thus giving rise to the contemporary continuant-stop alternation, where the historically antecedent realisation (the stop) occurs in the more restricted environment.

(47) Historical phoneme correspondences (Heath 1978a, p. 38)

Language	*Fortis → lenis stop						*Lenis stop → continuant					
proto-Ngandi-Wubuy	pː	t̪ː	tː	ʈː	cː	kː	p	t̪	t	ʈ	c	k
Wubuy	p	t̪	t	ʈ	c	k	w	l̪	ɾ	ɻ	j	w

The temptation here, then, would be to regard the stops (e.g. /k/ and /p/) as underlying, and the [w] as a neutralising allophonic surface realisation of both whenever they do not follow a non-continuant — following the historical lenition process that actually occurred (at least for these alternating forms). This would cast a stem like [jamp-iːna] in (44a) as derived from underlying */camp-iːna/.

Synchronically, however, Heath rejects a lenition analysis on the grounds that there are also morphemes which *always* have an initial stop (that is, in all environments) — these are the ones subject to ŋu-Epenthesis, such as we saw above. A lenition analysis would require us to explain why some stops lenite and others do not. In a sense then, the small degree of arbitrariness in positing two phonemes with identical phonetic properties, /w$_1$/ and /w$_2$/, avoids the greater arbitrariness that would be invoked by positing a whole class of leniting stops and a whole class of non-leniting stops: /p$_1$, p$_2$/ and /k$_1$, k$_2$/, but also /t$_1$, t$_2$/, /t̪$_1$, t̪$_2$/, /ʈ$_1$, ʈ$_2$/ and /c$_1$, c$_2$/. But clearly neither the hardening nor the leniting analysis is entirely satisfying, since one is forced to accept diacritised segments either way.[8] Later on in this chapter, I will develop an analysis that does not force us to choose one stem allomorph as basic to the other, treating them instead as surface variants of each other, with equivalent ontological status. As a matter of convenience, the discussion in the remainder of this section will assume the Hardening analysis, as per Heath (1984), though in following sections I will occasionally consider the implications of assuming a Lenition analysis where relevant.

I note also that the historical chain shift is only part of the story here anyway. It is unlikely that morph-initial underlying (non-alternating) stops in Wubuy descend from historical fortis stops, simply because the Top End languages that have a fortis/lenis contrast only allow the fortis stops morph-medially, not morph-initially where *ŋu-* is implicated. Instead, contemporary

[8] See Round (2023) for a general statement of this problem in Australian phonological description.

underlying stops may simply be the reflex of lenis stops that were never lenited in the first place. According to Van Egmond & Baker's (2021) reconstruction of the phonology of the shared ancestor to Wubuy and Anindilyakwa, lenition applied on a lexeme-by-lexeme basis, never covering the entire lexicon in either language, and affecting different lexemes in each.

3.2.3 Continuants undergo Hardening, stops trigger ŋu-Epenthesis

Hardening, then, produces a stem-initial stop following a prefix-final nasal. Of course, this is exactly the structural condition targeted by ŋu-Epenthesis. Crucially, however, /ŋu-/ is never inserted before stops derived by such means; ŋu-Epenthesis only applies to stems that are stop-initial in their underlying representation. Thus, the stems that undergo Hardening and the stems that undergo ŋu-Epenthesis form disjunctive sets. In Heath's SPE-type analysis, this result is achieved by rule ordering. Because ŋu-Epenthesis applies before all other rules[9], the stops derived by Hardening miss their chance to undergo ŋu-Epenthesis. This is therefore a case of counterfeeding opacity — the Hardening rule *would have fed* ŋu-Epenthesis if the ordering were reversed. That is, /ŋaN-jamp-iːna/ from (44b) would become |ŋaɲ-camp-iːna| after Hardening, and then ŋu-Epenthesis would apply to this form to yield (after deleting the prefix-final nasal) a surface form *[ŋa-ŋu-camp-iːna]. Examples attesting to derivations like this are not found anywhere in Wubuy, and my consultants uniformly rejected my attempts to produce them.

We thus see two basic classes of stem or root in Wubuy, defined according to the behaviour of the initial segment: one class is stop-initial always, while the other alternates between a stop-initial and a continuant-initial realisation. (Though there is also a sonorant-initial class, which shows no alternation, and a vowel-initial class, which is subject to a different type of consonant epenthesis under the same conditions as Hardening and ŋu-Epenthesis.) In Heath's work, this amounts to a difference between stop-initial and continuant-initial underlying representations. The systematic behaviour of these two types of stems also correlates with a systematic difference in the behaviour of the inflectional pronominal prefixes that attach to them.

Consider the examples below. The surface word forms in the left column and the words in the right column differ along two different dimensions simultaneously: the form of the stem and the form of the pronominal prefix. The stems in the words of the left column alternate between a continuant-initial realisation — the (a) form in these examples — and a stop-initial realisation — the (b) form. These are the stems subject to Hardening. Thus in (48a), the stem meaning 'transported' (or 'carried', or 'took away') is *ju.ɻ-aŋi*, while in (48b) it is *cu.ɻ-aŋi*. This difference in the stem realisation is accompanied by a difference in the prefixes: in all the (a) forms the

[9]Except Reduplication, as linking *ŋu-* can appear inside the reduplicant under certain conditions; see ch. 4.

prefixes are vowel-final, in the (b) forms they are nasal-final, and the nasal always assimilates in place to the following consonant. Thus we find (48a) *ŋa-juɻ-aŋi* and (48b) *nuɲ-cuɻ-aŋi*. The subsequent examples show that the prefix-final nasal in the (b) forms always takes on the place features of the following stop.

By contrast, the examples in the right column show different behaviour in both the prefixes and the stems. Firstly, the stems are stop-initial across all their realisations. Thus *cuɻ-aŋi* 'push' has the same form in (49a), where it follows a vowel-final prefix, and in (49b), where it follows a prefix that is consonant-final at least in the underlying representation. Secondly, the prefixes are the same as in the left column for the (a) forms, but the (b) forms are /ŋu/-final rather than nasal-final. Thus we can contrast (48b) *nuɲ-cuɻ-aŋi* 'You transported it', where the 2sg.R prefix has the form [nuɲ], with (49b) *nu-ŋu-cuɻ-aŋi* 'You pushed it', where the prefix is instead [nuŋu]. Note that these two roots form a minimal pair, distinguished by the continuancy of the initial consonant: /j/ versus /c/. Hardening neutralises this contrast by transforming /j/ into /c/ in (48b); however, this difference in prefix form — [nuɲ] in (48b) and [nuŋu] in (49b) — ensures words with these roots are kept distinct in those contexts where the stem-initial contrast would otherwise be neutralised.

(48) a. ŋa-**j**uɻ-aŋi
/ŋa-juɻ-aŋi/
1sg.R-carry-Past2
'I transported it'
(Baker & Harvey in press, p. 339)

b. nu**ɲ-c**uɻ-aŋi
/nuN-juɻ-aŋi/
1sg.Irr-carry-Past2
'You transported it'
(Baker & Harvey in press, p. 339)

(49) a. ŋa-**c**uɻ-aŋi
/ŋa-cuɻ-aŋi/
1sg>3Resid.R-push-Past2
'I pushed it'
(Baker & Harvey in press, p. 339)

b. nu-**ŋu-c**uɻ-aŋi
/nuN-cuɻ-aŋi/
1sg>3Resid.Irr-push-Past2
'You pushed it'
(Baker & Harvey in press, p. 339)

(50) a. ŋa-**j**alṯi-ɲ
/ŋa-jalṯi-ɲ/
1sg.R-go_past-Past1
'I went past'
(JN; PN2020-10-28.2 00:00:34)

b. nu**ɲ-c**alṯi-ɲ
/nuN-jalṯi-ɲ/
2sg.R-go_past-Pst1
'You went past'
(JN; PN2020-10-28.2 00:00:38)

(51) a. pa-**c**ulup-aŋ
/pa-culup-aŋ/
2sg>3Resid.Irr-conceal-Npst1
'Hide it!'
(HN; PN2020-11-08.1 00:02:32)

b. nu-**ŋu-c**ulup-i
/nuN-culup-i/
2sg>3Resid.R-conceal-Npst3
'(Don't) hide it'
(HN; PN2020-11-08.1 00:02:44)

(52) a. ŋa:mpu-**j**amiŋkaɾa-ŋ
/ŋa:mpu-jamiŋkaɾa-ŋ/
12pl.Irr-do_what-Npst1
'What will we (pl.inc) do?'
(NMET 41.21.4)

b. na**ɲ-c**amiŋkaɾa-ŋ
/naN-jamiŋkaɾa-ŋ/
12du.Irr-do_what-Npst1
'What will we (du.inc) do?'
(NMET 43.9.1)

(53) a. ni-**c**aɾaɾa-ŋi
ni-Ø-caɾaɾa-ŋi
3sgM.R-walk-Past2
'I will walk'
(NMET 5.5.3)

b. ŋa-**ŋu-c**aɾaɾa-ŋ
ŋaN-Ø-caɾaɾa-ŋ
1sg.Irr-walk-Npst1
'I will walk'
(NMET 62.3.4)

(54) a. ŋa-l̪apaɻa-ŋu-puɾi-ɲ
/ŋa-l̪apaɻa-puɾi-ɲ/
1sg.R-coolamon-put-Past1
'I put it in the coolamon'
(JN; PN2017-10-22.2)

b. nu**n̪-ṯ**apaɻa-ŋu-puɾi-ɲ
/nuN-l̪apaɻa-puɾi-ɲ/
2sg.R-coolamon-put-Past1
'You put it in the coolamon'
(JN; PN2017-10-22.2)

(55) a. ŋa-ṯ̪apaɭajiti-ɲ
/ŋa-ṯ̪apaɭajiti-ɲ/
1sg.R-slip-Past1
'I slipped over'
(MM; PN2017-07-06.1 01:00:50)

b. nu-**ŋu-**ṯ̪apaɭajit-i:
/nuN-ṯ̪apaɭajit-i:/
2sg.R-slip-Past1
'You slipped over'
(MM; PN2017-07-06.1 01:00:38)

(56) a. nu-**ɻ**aŋ
/nu-ɻaŋ/
3sgM>3sgM.R-spear.Past1
'He speared him'
(JN; PN2017-09-30.1 00:05:08)

b. nuɳ-**ʈ**aŋ
/nuN-ɻaŋ/
2sg>3Resid.R-spear.Past1
'You speared it'
(JN; PN2017-09-30.1 00:02:08)

(57) a. ŋa-**t**an-kaːjuː
/ŋa-tan-aːjuː/
1sg>3Resid.R-guts-cut.Past2
'I cut up the guts'
(HN; PN2017-06-25, 00:49:50)

b. ŋa-**ŋu-t**an-kaːjuː
/ŋaN-tan-aːjuna/
1sg>3Resid.R-guts-cut.Npst2
'I'll cut up the guts'
(HN; PN2017-06-25, 00:49:58)

(58) a. pa-**w**ulʈ̪-iː
/pa-wulʈ̪-iː/
2sg>3Resid.Irr-cut-Npst2
'Cut it! (fish)'
(HN; PN2020-11-11 00:25:29)

b. ŋa**ŋ-k**ulʈ̪-iː
/ŋaN-w$_1$ulʈ̪-iː/
1sg>3Resid.Irr-cut-Npst2
'I'll cut it (fish)'
(HN; PN2020-11-11 00:25:19)

(59) a. ŋa-**k**ɑrac-iː
/ŋa-kɑrac-iː/
1sg.R-jump-Npst2
'I'm jumping'
(HN; PN2017-06-25 00:13:43)

b. ŋa-**ŋu-k**ɑrac-iː
/ŋaN-kɑrac-iː/
1sg.Irr-jump-Npst2
'I will jump'
(HN; PN2017-06-17 00:14:32)

(60) a. ŋi-**w**anka:kiʈ̪u
/ŋi-w$_2$anka:kiʈ̪u/
3sgF.R-crazy
'She's crazy'
(JN; PN2020-11-05.2 00:26:19)

b. nu**m-p**anka:kiʈ̪u
/nuN-w$_2$anka:kiʈ̪u/
2sg.R-crazy
'You're crazy'
(JN; PN2020-11-05.2 00:26:32)

(61) a. pa-**p**aɻawut-aŋ
/pa-paɻawut-aŋ/
2sg>3Resid.Irr-discard-Npst1
'Throw it away!'
(HN; PN2020-11-03 00:07:41)

b. nu-**ŋu-p**aɻawuti-ɲ
/nuN-paɻawuti-ɲ/
2sg>3Resid.R-discard-Past1
'(Did you) throw it away?'
(HN; PN2020-11-03 00:07:54)

(62) a. ŋa-**j**ina-ŋu-paḻ-u:
/ŋa-jina-(Ø)-paḻ-u:/
1sg>3Resid.R-head-Ø-cut-Past2
'I cut up the head'
(LM; PN2017-10-03.2 00:43:21)

b. ŋa**ɲ-c**ina-ŋu-paḻ-aŋ
/ŋaN-jina-(Ø)-paḻ-aŋ/
1sg>3Resid.Irr-head-Ø-cut-Past2
'I will cut up the head'
(LM; PN2017-10-03.2 00:41:45)

(63) a. ŋa-**p**akaɭaɳ-ʈaji:
/ŋa-pakaɭaɳ-ɻaji:/
1sg>3Resid.R-eye-spear.Npst2
'I'm spearing it in the eye'
(LM, PN2017-10-03.2 00:06:13)

b. nu-**ŋu-p**akaɭaɳ-ʈaji:
/nuN-pakaɭaɳ-ɻaji:/
2sg>3Resid.R-eye-spear.Npst2
'You're spearing it in the eye'
(LM, PN2017-10-03.2 00:05:04)

As the examples show, the prefixes that trigger Hardening and those that trigger ŋu-Epenthesis are the same, and there are only three of them — the 1sg.Irr, 2sg.R, and the 12du.Irr — while the hundreds of other inflectional pronominal prefixes are always vowel-final.[10] It is worth noting that, owing to the person categories involved, all three of these have extremely low token frequency in Heath's (1980) collection of texts, which are all monologic narratives. Given the underrepresentation of stop-initial roots more generally, only a small fraction of the tokens that do occur in the texts trigger prefixal ŋu-Epenthesis, the vast majority triggering Hardening instead. In conversational discourse, on the other hand, they are among the most frequent prefixal categories: the speaker, addressee and their combination (first person dual inclusive) being far more frequent subjects in dialogic, as opposed to monologic speech. Note also that the second singular realis /nuN-/ is used in negative imperatives.

These three prefixes show systematic form variation in tandem with the stem allomorphy. Preceding an underlyingly continuant-initial stem, they are nasal-final, with the nasal assimilating in place to the following consonant; preceding an underlying stop, they are instead /ŋu/-final. We can treat the prefixes in question as having underlying representations as in (64). I have treated the final nasal of these forms as /N/, the 'nasal archiphoneme', which in Heath's analysis functions identically to /ŋ/, in that it always assimilates in place to a following stop, and always deletes before another sonorant consonant. The only place in which a morpheme-final /ŋ/ can appear 'as itself' — that is, the only place where it is not subject to assimilation or deletion — is at the end of a word. But because inflectional prefixes necessarily cannot appear at the end of a word, these prefix-final nasals cannot ever be directly observed as /ŋ/.[11]

[10] A precise quantification is difficult, as not all logically possible combinations of grammatical features are distinguished formally, resulting in widespread syncretism. Even collapsing most of these syncretisms leaves us with 296 distinct prefixes.

[11] These are not the same underlying representations as adopted by Heath (1984). See §2.3.2 of the previous chapter,

(64) a. /ŋaN-/ b. /nuN-/ c. /naN-/
1sg.Irr 2sg.R 12du.Irr

Nasal Deletion, as in (65), deletes /ŋ/ or /N/ preceding other sonorant consonants — including the non-alternating laterals /l/ and /ɭ/. Nasal Assimilation targets /ŋ,N/ as well, making them assimilate to a following non-sonorant, as in (66).

(65) Nasal Deletion:
$$\begin{bmatrix} +\text{NASAL} \\ +\text{DORSAL} \end{bmatrix} \rightarrow \emptyset \; / \; _ \; \begin{bmatrix} -\text{SYLL} \\ +\text{SON} \end{bmatrix}$$

(66) Nasal Assimilation:
$$\begin{bmatrix} +\text{NASAL} \\ +\text{DORSAL} \end{bmatrix} \rightarrow \begin{bmatrix} +\text{NASAL} \\ \alpha \text{ PLACE} \end{bmatrix} \; / \; _ \; \begin{bmatrix} -\text{SON} \\ \alpha \text{ PLACE} \end{bmatrix}$$

These rules are general, though apparently somewhat optional for true /ŋ/ (see 5.2.5). In (67a), the noun *pakaɭaŋ* 'eye' is incorporated[12] before an underlying /ɻ/, which then hardens to its paired stop /ʈ/. The final /ŋ/ of *pakaɭaŋ* then assimilates to this stop, becoming /ɳ/. In (67c), the final /ŋ/ must instead delete, since the initial segment of the verb root that follows it is also a nasal. Similar examples are shown in (68) for *maɻaŋ* 'hand'.

(67) a. nu-ŋu-pakaɭaɳ-**ʈ**aji:
/nuN-pakaɭaŋ-ɻaji:/
2sg>3Resid.R-eye-spear.Npst2
'You're spearing it in the eye' (LM, PN2017-10-03 00:05:04)

b. jaki nuɾu-pakaɭa**m-p**aɭp-i
/jaki/ /nuɾu-pakaɭaŋ-w$_2$aɭp-i/
Neg 2sg.R-eye-paint-Refl.Npst3
'Don't paint your (pl.) eyes' (HN, PN2020-11-10 00:03:08)

c. ŋa-pakaɭ**a-ŋ**aɾkiw-aɲ
/ŋa-pakaɭaŋ-ŋaɾkiw-aɲ/
1sg>3Resid.R-eye-pull_out-Past1
'I pulled its eye out' (HN, PN2020-10-30 00:31:25)

where I argue that treating these prefixes as /N/-final makes Heath's analysis more consistent with itself, and eliminates some unnecessary complexity in the phonological rules.

[12]There exists a suppletive incorporated noun for 'eye', underlyingly /paK-/, but it is of very limited productivity; the free form *pakaɭaŋ* is highly productive.

(68) a. ŋama-maɻa**ɲ-c**aɾt-iɲ
/ŋama-maɻaŋ-jaɾt-iɲ/
3VEG>1sg.R-hand-poke-PAST1

‘Something (VEG) poked me on the hand’ (JN, PN2017-10-01.2 00:20:13)

b.

anika	ana-maːɾɲ	jaːni	ŋa-waɳaka-na
/anika/	/ana-maːɾɲ/	/jaːni/	/ŋa-w$_2$aɳaka-na/
where.RESID	TOP.RESID-snake	here.RESID	1sg>3RESID.R-hold-NPST2

ŋa-maɻa**m-p**aɳaki-na
/ŋa-maɻaŋ-w$_2$aɳaki-na/
1sg>3RESID.R-hand-hold.REFL-NPST2

‘Where’s the snake? It’s here, I’m holding it in my hand’

(DM, PN2020-11-12 00:11:28)

The following example shows deletion of the pronominal prefix-final nasal — what I have taken as /N/ — before another nasal. Although the surface form of the first verb, *ŋa-maɾaː-maciː*, could in other contexts be interpreted as having the 1sg realis prefix /ŋa-/ — which would have the same surface realisation here — the *-maciː* conditional suffix here requires an irrealis prefix. Additionally, the following future tense verb, *ŋaɲ-campiː-na-maciː*, clearly has the irrealis /ŋaN-/. This makes it even more obvious that the first verb should be interpreted as future tense also, which requires the irrealis.

(69)

ŋicaŋ	ŋa-maɾaː-maciː,	ŋaɲ-campiː-na-maciː	aː-’ɲpac-kuj
/ŋicaŋ/	/ŋaN-maɾawa-maciː/	/ŋaN-jampiː-na-maciː/	/a-aɲpac-w$_1$uj/
again	1sg.IRR-deviate.NPST2-if	1sg.IRR-speak-NPST2-if	OBL.RESID-other-ALL

‘If I go off course, if I speak of some other place…’ (NMET 47.22.1)

Having established that Nasal Assimilation and Nasal Deletion are general processes, we can consider how they are ordered with respect to Hardening and ŋu-Epenthesis. The crucial elements of the rule ordering in (70) are that ŋu-Epenthesis must precede both Hardening and Nasal Deletion; its relation to both of these rules is opaque, as will be discussed in §3.3.4. The relative ordering of the two nasal rules cannot readily be decided, since the correct output will emerge either way.

(70) Derivation of continuant- and stop-initial stems following Heath (1984).

Input:	2sg.R-carry.Npst2 /nuN-juɻi:/	2sg.R-push.Npst2 /nuN-cuɻi:/
ŋu-Epenthesis	—	nuN-ŋu-cuɻi:
Hardening	nuN-cuɻi:	—
Nasal deletion	—	nu-ŋu-cuɻi:
Nasal assim.	nuɲ-cuɻi:	—
Output:	nuɲ-cuɻi:	nu-ŋu-cuɻi:

Considering the example derivations above, we see how Hardening neutralises underlying distinctions on the surface — both /j/ and /c/ can surface as [c] — while the remaining rules result in different surface forms for the same underlying string — /nuN-/ here surfaces as both [nuɲ] and [ŋuŋu]. As a result, the surface words, as totalities, retain the distinctiveness that was present in their underlying forms, but the location of the contrast has shifted: instead of a lexical distinction centred on the continuancy of the stem-initial segment, we find that the two forms are distinguished by the material preceding the stem.

3.2.4 Linking ŋu-

In general, if a root is subject to prefixal *ŋu-* when it appears at the inflection-stem boundary, then it is also subject to linking *ŋu-* when it appears within a complex stem, as we already saw for the root *pura* 'sit' in (40). There are some exceptions to this statement, which will be examined later in this section.

The root *paḻu* is shown following a vowel-final prefix in (71a); in a context subject to prefixal *ŋu-* in (71b); and inside complex stems, requiring linking *ŋu-*, in (71c-e). The examples of derived stems show that a stop-initial verb is subject to linking *ŋu-* irrespective of the immediately preceding phonological context — a stop, a nasal, and a vowel are all represented here. Further examples are given for stops at other places of articulation in the examples following.

(71) a. pa-paḻ̈-una
/pa-paḻ̈-una/
2sg.Irr-cut-Npst2
'Cut it up [the fish]!' (LM; PN2017-10-03.2 00:38:02)

b. ŋa-ŋu-paḻ-aŋ
/ŋaN-paḻ-aŋ/
1sg.Irr-cut-Npst1
'I will cut it up [the fish]' (JN; PN2020-11-11.1 00:04:18)

c. pa-jina-**ŋu**-paḻ-aŋ
/pa-jinak-(Ø)-paḻaŋ/
2sg.Irr-head-(Ø)-cut.Npst1
'Cut the head!' (LM; PN2017-10-03.2 00:41:13)

d. ŋaŋu-tan-**ŋu**-paḻ-aŋ[14]
/ŋaŋu-tan-(Ø)-paḻa-ŋ/
1sg>3Fem.Irr-guts-(Ø)-cut-Npst1
'I will cut up its [the dugong's = Fem] guts' (NMET 111.5.3)

e. ŋaŋu-waɾa-jaɾka-**ŋu**-paḻ-uː
/ŋaŋu-w$_1$aɾa-jaɾka-(Ø)-paḻ-uː/
1sg>3Fem.R-Mult-flipper-(Ø)-cut-Past2
'I cut up all of the [dugongs'] flippers' (JN; PN2020-10-30.3 00:36:21)

(72) a. ŋa-kaɾac-iː
/ŋa-kaɾac-iː/
1sg.R-jump-Npst2
'I'm jumping' (HN PN2017-06-25 00:13:31)

b. ŋa-ŋu-kaɾac-iː
/ŋaN-kaɾac-iː/
1sg.Irr-jump-Npst2
'I will jump' (HN PN2017-06-25 00:12:47)

c. nu-ŋu-kaɾac-iː
/nuN-kaɾac-iː/
2sg.R-jump-Npst2
'You are jumping' (HN PN2017-06-25 00:14:04)

[14]The /ŋaŋu-/ prefix (1sg>3Fem.Irr) here is not an instance of the first singular plus *ŋu-* as in (71b), but is homophonous with it in this context.

d. wu-ŋaɾa-**ŋu**-kaɾac-iː
/wu-ŋaɾaK-kaɾac-iː/
3RESID.R-MULT-jump-NPST2
'They [kangaroos] are all jumping' (HN PN2017-06-25 00:34:06)

(73) a. ni-ʈuːlwiː-ni
/ni-ʈuːlwiː-ni/
3sgM.R-sing-PAST2
'He sang for it' (NMET 29.4.5)

b. ni-jaɾa-**ŋu**-ʈuːlwiː-ni
/ni-waɾaK-(Ø)-ʈuːlwiː-ni/
3sgM.R-MULT-(Ø)-sing-PAST2
'He sang for them all' (NMET 29.4.4)

(74) a. nimpaː-'-**ŋu**-paɻawut-aŋ
/nimpaː-aK-Ø-paɻawut-aŋ/
2sg>1sg.IRR-APPL-Ø-discard-NPST1
'Throw it away for me!' (JN; PN2020-11-11 00:16:41)

b. pa-waɾa-**ŋu**-paɻawut-aŋ
/pa-w_1aɾa-Ø-paɻawut-aŋ/
2sg>3RESID.IRR-MULT-Ø-discard-NPST1
'Throw it all away!' (JN; PN2020-11-11 00:16:53)

(75) a. nakaŋ ciːci a-maɳʈa-ɾuc
2sg sores OBL.NEUT-shin-LOC
'You've got sores on your shin' (HN, PN2017-06-27 01:06:05)

b. nu-mun-**ŋu**-ciːci
2sg.R-foot-Ø-sores
'You've got sores on your foot' (HN, PN2017-06-27 01:06:53)

(76) a. wunu-kaɾi-**ŋu**-cuɻ-iɲ
3duM>3sgM.R-behind-Ø-push-Past1
'Those two men pushed him from behind' (JT; PN2017-06-30 00:24:47)

b. ŋanu-canta-**ŋu**-cuɻ-iɲ
1sg>3sgM.R-lower_back-Ø-push-Past1
'I pushed him in the lower back' (HN; PN2017-10-30.2, 00:32:24)

In theory, prefixal *ŋu-* and linking *ŋu-* can appear in the same word. Heath's (1984, pp. 106–107) grammar asserts that stems with a stop-initial derivational element and stop-initial root condition both kinds of *ŋu-*, and he provides a derivation of /ŋaN-Red-paK-paḻu/ '1sg.Irr-Red-eye-cut' illustrating this. The status of this example is unclear, however, as the root is in its uninflected (citation) form, and no similar example appears in the texts. I was able to elicit an equivalent example, though only with the more productive incorporable noun for 'eye' *pakaɭaŋ*, rather than the suppletive *paK*, as (77a). I also recorded the phonologically similar (77a,b).

(77) Prefixal *ŋu-* + linking *ŋu-*

a. ŋa-**ŋu**-pakaɭa-**ŋu**-paḻ-una
/ŋaN-(Ø)-pakaɭaŋ-(Ø)-paḻ-una
1sg.Irr-(Ø)-eye-(Ø)-cut-Npst2
'I'm going to cut the fish's eye' (HN; PN2020-10-30 00:31:46)

b. ŋa-**ŋu**-kaɭa-**ŋu**-puɾaŋ ana-kakaɭaŋ ama-pacpaɾa-ɾuc
/ŋaN-(Ø)-kaɭaŋ-(Ø)-puɾaŋ/ /ana-kakaɭaŋ/ /ama-pacpaɾa-ɾuc/
1sg.Irr-(Ø)-egg-(Ø)-put-Npst1 Top.Neut-egg Obl.Veg-mat-Loc
'I'll put the egg down on the mat' (JN; PN2017-10-03)

c. na-**ŋu**-piɻa-**ŋu**-puɾta-ŋaŋ
/naN-(Ø)-piɻa-(Ø)-puɾta-ŋaŋ/
12du.Irr-(Ø)-anus-(Ø)-bend_down-Npst1
'We have to keep our bums down' (HN; PN2020-11-04 00:22:27)

d. ŋa-**ŋu**-canta-**ŋu**-puɾta-ŋaŋ
/ŋaN-(Ø)-canta-(Ø)-puɾta-ŋaŋ/
1sg.Irr-(Ø)-back-(Ø)-bend_down-Npst1
'I'll bend down (keeping my back out of sight)' (JN; PN2020-11-04 00:31:58)

It is also possible for linking *ŋu-* to occur twice in a single word. An example is given in (78), which has a complex stem featuring a derivational prefix, a stop-initial incorporated noun *canta* 'back', and a stop-initial root *purta* 'bend over (to hide)'.

(78) Linking *ŋu-* + linking *ŋu-*

a. numpuɾaː-ɾa-**ŋu**-canta-**ŋu**-puɾta-ŋaŋ
/numpuɾu-aɾa-(Ø)-canta-(Ø)-puɾta-ŋaŋ/
2pl.IRR-MULT-(Ø)-back-(Ø)-bend_down-NPST1

'All of you, bend over (keeping your backs out of sight)' (JN; PN2020-11-04 00:33:35)

In practice, however, such examples are extremely rare. This can probably be explained by the fact that stop-initial roots are generally rarer than sonorants in the lexicon (about 31% of verb roots according to Baker & Harvey (in press), Table 6.4), and stop-initial derivational elements only more so (about 20%, all of them incorporable nouns). This makes their combination in the same stem particularly uncommon. Further, any such combination of incorporated noun and stem must be semantically plausible, and, in order to trigger prefixal *ŋu-*, must be interpretable as either intransitive, or transitive with a RESID object, so as to require one of the three N-final prefixes that trigger *ŋu-*. I will assume for the remainder of this chapter that words with two instances of *ŋu-* are possible words of Wubuy, even if the empirical basis of this assumption is only suggestive at present. Note also that while theoretically possible, words with one prefixal *ŋu-* and two linking *ŋu-* cannot in fact occur, as it would require a complex stem consisting of a derivational prefix, incorporated noun, and verb root, all of which are stop-initial, but there happen to be no derivational prefixes that are stop-initial.

Turning to underlyingly continuant-initial roots, these show the same behaviour inside complex stems as they do at the left edge of the stem: they alternate between continuant-initial and stop-initial realisations, depending on the immediately preceding (underlying) phonological environment. For instance, we can see the inflected verb root for 'go.NPST2', underlyingly /jaːɾiː/, appear as its continuant-initial and stop-initial allomorphs, respectively, in (79a) and (79b). Both examples are simple stems, where the verb root directly adjoins the pronominal prefix. In (79c) and (79d), however, this inflected root appears with derivational prefixes in between; in the second of these, Hardening applies — opaquely, in this case, since the /K/ that conditions it must then be deleted by an independent process, making the application of Hardening appear unmotivated on the surface. Additional examples of complex stems are provided in (80) and (81) — compare the same roots in simple stems in (58) and (56) above.

(79) a. ŋa-**j**aːɾiː ama-maṯaɭa-kuj
/ŋa-jaːɾiː/ /ama-maṯaɭak-w$_1$uj/
1sg.R-go.Npst2 Obl.Veg-beach-All
'I'm going to the beach' (HN 2017-06-25 00:01:16)

b. ŋaɲ-**c**aːɾiː ŋampa-jaːɾik-ina aː-ɭanṯuɾk
/ŋaɲ-jaːɾiː/ /ŋampa-jaːɾik-ina/ /waː-ɭanṯuɾk/
1sg.Irr-go.Npst2 1sg>3Coll.Irr-bring-Npst2 Foc.Coll-dog
'I'm going out, taking the dogs' (JN PN2020-10-27 00:00:35)

c. wuɾaː-'ɲci-**j**aːɾiː
/wuɾu-aɲci-jaːɾiː/
3pl.R-Comit-go.Npst2
'They go with (it)' (NMET 155.1.5)

d. numpuː-ŋaɾa-**c**aːɾiː
/numpuɾu-ŋaɾaK-jaːɾiː/
2pl.Irr-Mult-go.Npst2
'All of you, go!' (YN PN2017-07-06 00:58:53)

(80) a. juka nuɳ-ʈa-**ɻ**aŋ
/juka/ /nuN-ɻa-ɻaŋ/
Q 2sg>Resid.R-tooth-spear.Past1
'Did you spear it in the teeth?' (LM PN2017-10-03.2 00:02:20)

b. nuŋ-kaɭa-**ʈ**aŋ
/nuN-w$_1$aɭak-ɻaŋ/
2sg>Neut.R-throat-spear.Past1
'You speared it in the throat' (LM PN2017-10-03.2 00:00:39)

(81) a. niː-ɻapaɾa-**w**ulṯ-iː
/niɾi-ɻapaɾa-w$_1$ulṯ-iː/
1pl>Resid.R-tail-cut-Npst2
'We cut up its tail' (NMET 153.1.6)

b. ŋama-tan-**k**ulṯi:
/ŋama-tan-w$_1$ulṯi:/
1sg>Veg.R-guts-cut-Npst2
'I was cutting up the guts' (LM PN2017-10-04.2 00:18:45)

While this is true of the majority of verb and adjective roots, there are some that are stop-initial in the typical fashion when they occur at the left edge of the stem, but instead show the characteristic behaviour of continuant-initial roots when appearing in (some) compounds.[15] That is, ŋu-Epenthesis does not apply, and the initial segment is either stop-initial or continuant-initial. Heath (1984, p. 60) reports 10 roots that show this behaviour across the board, and 4 that show it only in certain compounds and not others. One root of the former category is illustrated in (82)

(82) a. ŋa-ŋu-**c**aɻjaɻ-aŋ
ŋaN-**c**aɻjaɻ-aŋ
1sg.Irr-go.down-Npst1
'I will go down'

b. wini:-'lbuŋuŋ-**c**aɻjaɻ-iŋ
/wini-albuŋuN-**j**aɻjaɻ-iŋ/
3du.R-together-go.down-Past1
'Those two went down' (NM; PN2017-07-03)

The root *pilja* 'to be tilted' demonstrates some of this exceptional behaviour. In (83a,b), this stop-initial root is immediately adjacent to a pronominal prefix; in (83c) an incorporated noun appears between the prefix and the root, triggering linking *ŋu-*, as would be expected with a stop-initial root. In (83d), however, there is no *ŋu-*, despite the phonological conditions being largely identical to (83c). Heath (1982, 1984) does not record any other instances of this compounding element *.ɻac-*, so its unusual underapplication of ŋu-Epenthesis is perhaps indicative of its 'lexicalised' status.

(83) a. mana-paɾawu ma-pilja-.ɻa
Top.Veg-boat 3Veg.R-tilt-Npst2
'The boat is on its side' (MM, PN2017-07-07 00:04:26)

[15]Similar construction-specificity of alternations is reported for Kayardild (Round 2009, p. 147).

b. wiɾima-pilja-ŋi-cka-ɲ
3pl>3VEG.R-tilt-AUG-CAUS-PAST1
'They made it [the boat] capsize' (NMET 161.7.4)

c. ma-jina-**ŋu**-**p**ilja-ɻa-wukic
/ma-jinak-(Ø)-pilja-ɻa-wukic/
3VEG.R-head-(Ø)-tilt-NPST2-still
'The prow is still tilted' (NM, PN2017-07-07 00:14:52)

d. ma-ɻa**c**-**p**ilja-ɻa
/ma-ɻac-pilja-ɻa/
3VEG.R-?-tilt-NPST2
'The prow is tilted (from overcrowding)' (NM, PN2017-07-07.2 00:01:35)

For some of the other verb and adjective roots that Heath says do not take linking *ŋu-* when forming compounds, my data suggests that this exceptionality is less extensive than he describes: among the tokens I have recorded, some of these exceptional roots show variability in the application of *ŋu-*, some behave like ordinary stop-initial roots, and some behave like ordinary continuant-initial roots. For the root *pilwila* 'to be weak', for instance, the speakers I consulted varied back and forth between productions with and without linking *ŋu-*, as in (84a,b).

(84) a. ni-pakaɭa**m**-**p**ilwiliɲ
/ni-pakaɭaŋ-pilwiliɲ/
'His eyes were weak' (YN PN2017-07-06 00:26:03)

b. ni-pakaɭa-**ŋu**-**p**ilwiliɲ
/ni-pakaɭaŋ-pilwiliɲ/
'His eyes were weak' (YN PN2017-07-06 00:26:13)

Other roots simply do not appear to be exceptional amongst contemporary speakers. For instance, Heath records the compound *kari-l̪ita* 'shut from behind' in a narrative text, reproduced here in (85). This is a special derivation of the stop-initial root *t̪ita* 'to close', wherein the initial stop lenites to its paired continuant, /l̪/ — according to Heath (1984, p. 58), thanks to a rule of Lenition that only targets certain lexically-specified morphemes. However, none of the speakers I consulted recognised this 'lenited' form. One offered me a stop-initial alternative, with *ŋu-*, shown in (86). As can be seen from the other stems involving *t̪ita* in (87), my consultants treat it as a prototypical stop-initial root, subject to both prefixal *ŋu-* (87b) and linking *ŋu-* (87c-e).

(85) ni-ɻiwakiː-ni-jiɲuŋ wunu-kaɾi-l̪iti-ɲ
/ni-ɻiwakiː-ni-jiɲuŋ/ /wunu-kaɾi-ʈiti-ɲ/
3sgM.R-bushwalk-Past2-Rel 3pl>3sgM.R-behind-close-Past1

'The man who was walking through the bush, they closed him off from behind (blocked his escape)' (NMET 40.6.4)

(86) nuŋu-kaɾi-**ŋu-ʈ**iti-ɲ
/nuŋu-kaɾi-(Ø)-ʈiti-ɲ/
2sg>3sgF.R-behind-(Ø)-close-Past1

'Did you shut her in from behind?' (HN, PN2017-06-25 00:41:53)

(87) a. pawu-ʈita-ŋ
/pawu-ʈita-ŋ/
2sg>3Neut.Irr-close-Npst1

'Close it!' [the door = Neut] (LM, PN2017-10-03 00:49:05)

b. na-**ŋu-ʈ**iti-ɲ l̪iɾipala a-kaɻa-ɾuc
/naN-(Ø)-ʈiti-ɲ/ /l̪iɾipala/ /a-kaɻa-ɾuc/
12du.Irr-close.Refl-Npst1 inside Obl.Neut-hole-Loc

'Let's lock ourselves inside the room' (NM, PN2017-07-03 00:06:21)

c. pa-l̪akuɭa-**ŋu-ʈ**iti-ɲ
/pa-l̪akuɭa-ŋu-ʈiti-ɲ/
2sg.Irr-lip-(Ø)-close.Refl-Npst1

'Close your lips!' (HN, PN2017-06-27 00:24:25)

d. nuɳ-ʈakuɭa-**ŋu-ʈ**iti-ɲ
/nuN-l̪akuɭa-(Ø)-ʈiti-ɲ/
2sg.R-lip-(Ø)-close.Refl-Past1

'Did you close your lips?' (HN, PN2017-06-27 00:25:20)

e. nuŋku-waɾa-**ŋu-ʈ**ita-ŋi
/nuŋku-w$_1$aɾa-(Ø)-ʈita-ŋi/
2sg>3Neut.R-lip-(Ø)-close-Past2

'Did you close them all [the doors = Neut]?' (LM, PN2017-10-03 00:59:58)

There are also constructions in which ŋu-Epenthesis appears to overapply. Interestingly, nearly all these attestations are in nouns. There are no nominal inflectional prefixes that are consonant-final; therefore, there is no possibility for prefixal *ŋu-* to apply to nouns except in N-V or N-Adj incorporation constructions. However, certain nouns can take a series of 'derivational' noun class markers, as discussed in §2.2.1, which mark the noun as being the part in a part-whole possession relation to a whole of the specified noun class.

Let us first see the expected behaviour of stop-initial roots in constructions like these. In (88), the noun root *piɻa* 'anus; bottom part' is both incorporated into a verb, and used in a part-whole possession construction. In the latter case, linking *ŋu-* applies. Similar examples are given in (89) and (90); the latter also features an example of a continuant-initial noun in the same type of construction, which must undergo Hardening in this context.

(88) anupaku na-ni-**ŋu**-piɻa-jiɲuŋ wini-piɻa-ap-aː
/anupaku/ /na-niK-(Ø)-piɻa-jiɲuŋ/ /wini-piɻa-w$_1$ap-aː/
there Top.Masc-Masc-(Ø)-anus-Rel 3pl>3Masc.R-anus-wrap-Past2
'They wrapped up the base part [of the fish trap = Masc]' (NMET 108.3.1)

(89) wiɾiŋa-ci-ciɻ-waɲc-aː ŋaɾa-ŋaɾi-**ŋu**-ciɻ
/wiɾiŋa-ci-ciɻ-w$_2$aɲc-aː/ /ŋaɾa-ŋaɾiK-(Ø)-ciɻ/
3pl>3Fem.R-Red-fur-pull_out-Past2 Top.Fem-Fem-(Ø)-fur
'They plucked out all its [the emu's = Fem] feathers' (NMET 95.5.2)

(90) mana-ma-**ʈ**ulpu mana-ma-**ŋu**-canta-jiɲuŋ waːɻi ampiɾima-l̪alwul̪a-ŋ
/mana-maK-ɻulpu/ /mana-maK-(Ø)-canta-jiɲuŋ/ /waːɻi/ /ampiɾima-l̪alwul̪a-ŋ/
Top.Veg-Veg-lower_back Top.Veg-Veg-(Ø)-back-Rel Neg 3pl>3Veg.Irr-soak-Npst1
'The back part [of the pandanus = Veg], they don't soak it' (NMET 156.2.4)

Certain noun roots that are not underlyingly stop-initial appear with linking *ŋu-* when in a derivational (part-whole possession) construction. All of these examples appear to have initial nasals. The first of these here, (91), is reported as an exceptional root, along with *ŋuɻuʈ̪u* 'bird', by Heath (1984, p. 162), though he doesn't provide a full example for either root. While I do not have a recording of (91), it was used by several speakers on numerous occasions in the Wubuy classroom at Numbulwar school, while helping children colour in a picture of a crab. It is curious that all of the examples I know of with an unexpected linking *ŋu-* are nasal-initial nouns in a part-whole possession construction. The root *maɻaŋ* 'hand' appears in incorporation

constructions headed by a verb or adjective, and there behaves like any other nasal-initial root in not conditioning ŋu-Epenthesis.As far as I am aware, the roots in (92) and (93) cannot be incorporated, but they are plainly nasal-initial roots and the appearance of *ŋu-* here is therefore unexpected.

(91) wu-ŋu-maɻaŋ
NEUT-(Ø)-hand
'crab claw' (i.e. 'hand of NEUT') (unrecorded)

(92) ma-ŋu-ŋumic
VEG-(Ø)-ant
'scraps, bits and pieces' (Heath 1984, p. 37)

(93) wu-ŋu-malaɲ a-maɻpiti-jiɲuŋ
NEUT-(Ø)-spike FOC.NEUT-stingray-REL
'the stingray's spike' (NMET 166.11.2)

There is one attestation of linking *ŋu-* preceding an ordinarily continuant-initial verb root in the texts, namely that in (94c).Examples (94a,b) show its ordinary behaviour, which is that of a continuant-initial root. Interestingly, the example (94c) with unexpected *ŋu-* has the stop-initial allomorph of the verb root, thus conforming to the surface generalisation that *ŋu-* precedes stops.

(94) a. ma-kaɻa-wawaḷ̪-i: mana-waɻaŋ
/ma-kaɻa-w$_2$awaḷ̪-i:/ /mana-waɻaŋ/
3VEG.R-hole-open-NPST2 TOP.VEG-ear
'Its ears are open' (NMET 112.1.3)

b. wu-waɭa-pawaḷ̪-aɲi-jiɲuŋ
/wu-waɭak-w$_2$awaḷ̪-aɲi-jiɲuŋ/
3NEUT.R-hollow-open-PAST2-REL
'Hollow log which was open' (NMET 43.15.5)

c. amaː-ɻ-ŋu-pawaḷ̪-iː-maci
/ama-aɻ-(Ø)-w$_2$awaḷ̪-iː-maci/
3VEG.IRR-water-(Ø)-open-NPST2-COND
'If the sea is clear' (NMET 42.5.4)

Thus, we see that some constructions escape linking *ŋu-* where it would otherwise be expected (on the basis of its underived form), and a few constructions show linking *ŋu-* when it is not expected. In general though, if a form is stop-initial and subject to prefixal *ŋu-* at the prefix-stem boundary, the overwhelming expectation is that it will also be stop-initial and subject to linking *ŋu-* when following a derivational element.

The exceptions to linking *ŋu-* are certainly low in number, but even a small degree of lexical variation is enough to raise questions about its status as a phonological rule. Heath, in any case, is somewhat ambivalent about how to approach exceptional behaviour like this — his grammar (Heath 1984) provides special rules stipulated to apply only to a small set of lexical items, while his dictionary (Heath 1982) provides special alternative underlying forms for those items, indexed to their different morphological environments.

3.2.5 Linking *ŋu-* as compounding morphology

The occasional under- and overgeneration of linking *ŋu-* is less troubling if we regard it not as the outcome of a phonological rule but as morphology (though, as we will see in §3.3, much hinges on what we take morphology to be). Given that linking *ŋu-* appears inside the domain of derivation, it is perhaps unsurprising to find that it is less than completely predictable — unlike prefixal *ŋu-*, which appears only outside this domain, and is inescapable.

If we accept that noun incorporation constructions are essentially a form of compounding — as argued by Heath (1984) and Baker (2014a, 2018) — then the *ŋu-* in this domain appears not unlike the linking elements that occur in Germanic compounds, hence my decision to label this morph linking *ŋu-*. In English, these linking elements are restricted to a small handful of old compounds — *tradesman* (trade+s+man), *huntsman*, *lambswool*, etc. — and their usage is no longer productive: new compounds in contemporary English simply adjoin two noun roots without any intervening element. Compounds in German, however, show a range of linking elements. Among them are *-s-*, *-(e)n-*, *-es-*, *-e-*, *-er-*, though only the first two of these are productive — see (95,96) — and their productivity has been increasing over time (Nübling & Szczepaniak 2008). Each linking element is to a large extent predictable — though not completely determined — on the basis of a complex combination of morphosyntactic and phonological features of the left member of the compound, and speakers converge to a remarkable extent in which element they prefer when forming new compounds in production experiments (Krott et al. 2007).

(95) German compound linking element *-s-*

a. Abfahrt → Abfahrt-s-zeit
'departure' 'departure time'

b. Gesicht → Gesicht-s-punkt
'view' 'point of view'

c. Religion → Religion-s-unterricht
'religion' 'religious instruction'

d. Zukunft → Zukunft-s-angst
'future' 'fear of the future'

(96) German compound linking element *-(e)n-*

a. Blume → Blume-n-stängel
'flower' 'flower stem'

b. Suppe → Suppe-n-topf
'soup' 'stock pot'

c. Frau → Frau-en-haus
'woman' 'women's asylum'

d. Bär → Bär-en-hunger
'bear' 'ravenous appetite'

These elements originated as case and number suffixes, but have over time have become specialised as compound-forming elements. The *-s-* in (95), for instance, originated as the genitive singular suffix for masculine or neuter nouns, but now appears in compounds where the left member is a feminine noun. *Abfahrt*, for example, is a feminine noun, and cannot appear with an *-s* suffix on its own, but it now forms compounds with linking *-s-*, such as *Abfahrtszeit* in (95a). Over time, as Nübling & Szczepaniak (2008) show, linking *-s-* has lost its morphosyntactic function, and acquired a phonological function in its place: it is increasingly frequent in new compounds whose first constituent is a poor phonological word, for instance by not ending in a trochee. They argue that *-s-* increases the markedness of the right edge of the phonological word, and therefore its salience in auditory processing, alerting the listener to the location of the boundary between the two compounded constituents, and thereby to the fact that it is a compound and should be processed as such. Thus, as the title of their paper expresses it, German linking elements are 'on the way from morphology to phonology' (Nübling & Szczepaniak 2008).

While we do not have the documentary records to assess the history of *ŋu-*, as we do in German, it seems most probable that it also began as a meaningful inflectional affix that increasingly came to be associated with its phonological context and not whatever morphosyntactic

category it once expressed. Heath (1978a, 1984) suggests it may have originated as the oblique-case class-marking morpheme *-*ku*-, which appears in Marra, Alawa and Ngandi as -*ku*-, and Warndarrang as -*ŋu*-, and is also associated with particular phonological and morphological environments in those languages, alternating with zero. Baker & Harvey (2020) reconstruct a historical trajectory for these putatively 'epenthetic' syllables in Marra and Alawa. According to their reconstruction, once-meaningful morphemes — originally independent class-marked demonstratives or pronouns, then bound articles — came to be deleted, except in those contexts where deletion would have led to a disfavoured phonotactic sequence. The result in the modern languages is a synchronic alternation where these historical inflectional affixes are normally absent, but remain present in a restricted environment, giving the appearance of syllable epenthesis. A similar argument has been made for the 'irrealis augment' *ji*- in Jaminjung, which Schulze-Berndt (2000) regards as epenthetic preceding the irrealis prefix when it follows a consonant-final bound pronoun; Osgarby (2018, pp. 269–271) reconstructs this as the reflex of a Proto-Mirndi desiderative mood proclitic **jV*-, which lost its meaning and was deleted except where deletion would result in a dispreferred consonant sequence.

On the basis of the available evidence, it is difficult to establish a direct historical link between *ŋu*- and the epenthetic syllables in Marra, Alawa, Warndarrang and Ngandi — and, given the tragic fortunes of these four languages, new evidence is unlikely to be found. If Wubuy shares an ancestor with these languages, it is only in the very distant past (except with Ngandi), and while we cannot rule out direct borrowing, the distribution of *ŋu*- is substantially different from the other epenthetic morphs, as Heath acknowledges. However, it seems unavoidable that *ŋu*- also began life as a morphosyntactically meaningful morph of some kind, since it is entirely unlike anything that could emerge from the phonologisation of phonetic processes. One possibility is that it began as a marker of a Fem class object, given that most of the contemporary prefixes marking F/Fem objects contain the string /ŋu/ in some part of their form. Perhaps there were other changes in the structure of complex words that led to this /ŋu/ appearing in places where one could not interpret any F/Fem participant, and this precipitated a bleaching of its meaning. Perhaps this occurred around the same time as the lenition of obstruents, and this /ŋu/ appeared frequently enough before historical fortis stops that it came to be associated with a phonological environment and not only a morphosyntactic one — that is, it became phonologised.

Of course, we can only speculate about the origins that led to this juncture, but there is no doubt that in contemporary Wubuy, *ŋu*- serves a crucial disambiguating function in the context of the phonological system. As Heath (1984) points out, by only applying to invariant stops, and not to those resulting from Hardening, ŋu-Epenthesis keeps (underlying) lexical distinctions apparent on the surface in contexts where they would otherwise be threatened by neutralisation.

In the next section, I consider this claim in more detail.

3.2.6 Minimal pairs

In seeking an explanation for the appearance of /ŋu-/, Heath points to the fact that it allows lexical constrasts to be maintained in a context where neutralisation would otherwise occur. Hardening transforms underlying continuant consonants into a stop at the same place of articulation, which in all cases constitutes an existing stop phoneme in the inventory. For example, /juɻa/ 'carry, transport' hardens to [cuɻa] in the relevant environment, leading to potential confusion with /cuɻa/ 'push'. However, because /cuɻa/ becomes [ŋu-cuɻa] in those environments where neutralisation could occur (that is, where Hardening applies), the lexical contrast between the two roots is maintained (Heath 1984, p. 37). As we see in Table (97) below, the two verbs (shown with their NPST2 suffixal inflection) are never homophonous in the same environment thanks to the intervention of ŋu-Epenthesis:

(97) Realisation of the /cuɻa/, /juɻa/ minimal pair in hardening and non-hardening contexts at the prefix-stem boundary:

	/juɻ-iː/ 'transport-NPST2'	/cuɻ-iː/ 'push-NPST2'
Following continuant	[juɻ-iː]	[cuɻ-iː]
Following (underlying) nasal	[cuɻ-iː]	[ŋu-cuɻ-iː]

Inside complex stems, Hardening behaves just as it does at the prefix-stem boundary: underlyingly continuant-initial roots, if they follow a derivational element ending in an underlying stop or nasal, appear as stop-initial on the surface instead. And again, ŋu-Epenthesis avoids neutralisation between underlying and derived stops in this environment, by only appearing before the former. But, unlike at the prefix-stem boundary, ŋu-Epenthesis also applies to root-initial stops even when the preceding morpheme ends in a [+CONT] segment (i.e. not a Hardening environment). As such, the underlying continuant/stop contrast at this juncture is not only preserved but *overdetermined*. Notice that in (98), neither of the possible realisations of /cuɻa/ are homophonous with the possible realisations of /juɻa/.

(98) Realisation of the /cuɻa/, /juɻa/ minimal pair in derived environments:

	/juɻ-iː/ 'transport-NPST2'	/cuɻ-iː/ 'push-NPST2'
Following continuant	[juɻ-iː]	[ŋu-cuɻ-iː]
Following underlying stop or nasal	[cuɻ-iː]	[ŋu-cuɻ-iː]

Heath's argument that /ŋu/'s role is to avoid homophony resulting from the neutralisation of continuant-stop contrast in hardening environments makes intuitive sense. The notion that phonemic contrasts persist where they enable lexical contrasts has a long history in linguistics (e.g. Hockett 1967; Martinet 1952), and its credibility has been significantly strengthened by recent studies using broader typological samples and more mathematically refined methods (Blevins & Wedel 2009; Hall et al. 2016; Surendran & Niyogi 2006; Wedel, Jackson et al. 2013; Wedel, Kaplan et al. 2013; Winter & Wedel 2016). What all of these studies make clear is that it is not enough merely to point to the existence of a 'phonemic' opposition and assume speakers and listeners will have inherent motivation to preserve it; it will only be preserved to the extent that the contrast aids listeners in accessing their lexicons and successfully decoding speakers' messages (Hall et al. 2016). Therefore, the importance of maintaining the phonological contrast between continuant-stop pairs in morpheme-initial (or more particularly root-initial) position, and therefore the necessity of *ŋu-* for this purpose, depends on how much 'work' this contrast performs in aiding listeners. A significant part of this calculation concerns how many minimal pairs are distinguished on the basis of the opposition (i.e. its 'functional load'; Wedel, Kaplan et al. 2013), as well as other factors like how easily those pairs can be distinguished on the basis of context, and how frequent the minimal pairs are.

But in fact, other than /cura/∼/jura/, there are few other minimal pairs (or sets) of roots which differ only in the continuancy of their initial segment. I have only counted 13 such sets in the dictionary (Heath 1982), plus two near-minimal sets (there are 149 stop-initial verb roots listed in the dictionary); these sets are presented in Table 3.1.

Only six pairs consist of roots from the same inflection class (all such pairs happen to be members of the A2 class). One could imagine some of these pairs occurring in a context where — if not for the minimal difference in phonological form — ambiguity between the pairs might arise, especially the first three pairs. But for others, contextual cues would in many cases be sufficient to avoid confusion, as well as differences in inflection class and transitivity.

Other functional factors to mitigate the potential for homophony here include the existence of synonyms. The /paɭajita/ 'to slip, slide' ∼ /w$_2$aɭajita/ 'be handsome, smooth' pair was recognised by speakers I consulted, as in (100) and (99). But there is also another root with the meaning 'to slip, slide', /t̪apaɭajita/, which appears to be far more common than /paɭajita/, though the two forms are almost certainly related. Elicited examples are shown in (101). When viewing a scene in Charlie Chaplin's *The Immigrant* in which the characters slip over each other on the deck of the boat, all four Wubuy speakers I asked to narrate the film (separately) used the /t̪apaɭajita/ root. Two of these examples are shown in (102).

Stop-initial			*Continuant-initial*		
Root	**Infl.**	**Meaning**	**Root**	**Infl.**	**Meaning**
/cuɻa/	TrA2	‘to push’	/juɻa/	TrA2	‘to carry, transport’
/paɭajita/	IntrA2	‘to slip, slide’	/w$_2$aɭajita/	IntrA2	‘to be handsome, smooth’
/ṯiɾita/	A2	‘to go down’	/l̪iɾita/	A2	‘to scratch (something) deeply’
/ta/	A2	‘to defecate’	/ɾa/	A2	‘to swear at’
/ciɾa/	A2	‘to throw spear at someone’	/jiɾa/	A2	‘to decorate, adorn’
/puɭa/	IntrA2	‘to be happy, satisfied’	/w$_2$uɭa/**	A2	as in mun-buɭa/ ‘to have a flat tyre’
/paɾa/†	Nga1	‘to be crooked’	/w$_2$aɾa/	A2	‘to drop, throw down’
			/w$_1$aɾa/	Ya	‘to open up (native oven)’
/paɾka/†	A1	‘to set ablaze by blowing’	/w$_2$aɾka/	A1	‘to sew, mend’
			/w$_2$aɾka/	IntrN	‘to throw’
/puɾta/	A2	‘to dig’	/w$_2$urṯa/	A2	’to bury’
/puɾta/	Nga1	‘to bend over’			
/kaɾaɾawi/*	I1	‘(dugong or other animal) to be exhausted after struggling’	/w$_1$aɾaɾawi/**	I1	as in maɻaŋ-kaɾaɾawi ‘to gesture or signal with hand’
/kulṯa/*	A2	‘to have culprit face punishment’	/w$_1$ulṯa/	A2	‘to cut’
			/w$_2$ulṯa/	A2	occasional variant of /w1ulṯa/
/ʈamaṯan/	NAdj	‘full’	/ɻamaṯan/	N	‘mouth’
/-ṯaɻaŋ-/	IN	‘spear shaft’	/l̪aɻaŋ/	NAdj	‘ready to eat, ripe’
/ṯawal/	N	‘bone above anus’	/-l̪awal-/	IN	‘thorn splinter’
/-caŋ-/	IN	‘root (of plant)’	/-jaːŋ-/	IN	‘voice’

Table 3.1: Known minimal pairs (or sets) which are disambiguated by means of -ŋu- in hardening environments, according to Heath (1982). The /puɾta/~/w$_2$urṯa/ and /caŋ-/~/jaːŋ-/ sets are only near-minimal.
* denotes a form not recognised by the consultants I have worked with.
** denotes a form only attested as the right-hand member of a compound
†denotes a stop-initial root that takes a lenited form in compounds but retains its semantic contribution.

(99) a. ŋa-paɭajit-iɲ
/ŋa-paɭajit-iɲ/
1sg.R-slide-Past1
'I slipped over' (MM; 2017-07-06.1 00:50:56)

b. nu-ŋu-paɭajit-iɲ
/nuN-paɭajit-iɲ/
2sg.R-slide-Past1
'You slipped over' (YN; 2017-07-06.1 00:50:52)

(100) a. mana-muːŋ ma-waɭajit-iː
/mana-muːŋ/ /ma-w$_2$aɭajit-iː/
Top.Veg-hair 3Veg.R-be_smooth-Npst2
'Your hair is beautiful but not curly; falling down' [speaker translation]
(JN; 2017-10-22 00:39:04)

b. num-paɭajit-iː
/nuN-w$_2$aɭajit-iː/
2sg.R-be_smooth-Npst2
'You're smooth and shiny and beautiful' (JN; PN2017-10-22)

c. wuɾu-wanca-waɭajit-iː
/wuɾu-wanca-w$_2$aɭajit-iː/
3pl.R-arm-be_smooth-Npst2
'They have smooth arms' (JN; 2017-10-17.1)

d. ni-kulmum-paɭajit-iː
/ni-kulmuŋ-w$_2$aɭajitiː/
3sgM.R-belly-be_smooth-Npst2
'He has a smooth belly' (JN; 2017-10-17.1)

(101) a. wuɾu-t̪apaɭajit-iː
/wuɾu-t̪apaɭajit-iː/
3p.R-slide.Npst2
'They're sliding [on the mud]' (JN; PN2017-10-22 00:21:37)

b. nu-ŋu-t̪apaɭajit-iː
/nuN-t̪apaɭajit-iː/
2sg.R-slide-Npst2
'You're sliding' (JN; PN2017-10-22 00:27:49)

c. wuɾaː-'ɾa-ŋu-t̪apaɭajit-iː
/wuɾu-aɾa-ŋu-t̪apaɭajit-iː/
3pl.R-Mult-Ø-slide-Npst2
'They're sliding around everywhere' (JN; PN2017-10-22 00:25:09)

(102) a. aʈapa ŋaːt̪uɾapatiɲuŋ **ŋiʈ̪apaɭajitiɲ**
'and then the white lady slipped over' (FM; Immigrant, 00:02:26)

b. waːj! ŋiɻapini maniɲuŋ, ŋiːkaːkiɲ, **ŋiʈ̪apaɭajitiɲ**
'Oh no! A woman fell, went back, and slipped over' (JN; Immigrant, 00:03:05)

Interestingly, one of the roots in Table 3.1, /paɾka/ 'to set ablaze by blowing' also shows (somewhat variable) failure of linking *ŋu-* in compounds. There are 18 of these phonologically irregular roots attested, of which two are also included among the minimal pairs in Table 3.1. So, for example, /paɾka/ 'to set ablaze' differs minimally from /w$_2$aɾka/ 'to sew, mend', but /paɾka/ can fail to trigger ŋu-Epenthesis in compounds, as below.

(103) a. ŋaː-t̪umpala ŋaŋu-waɾk-aː
/ŋaɾa-t̪umpala/ /ŋaŋu-w$_2$aɾk-aː/
Top.Fem-sail 1sg>3Fem.R-mend-Past2
'I was mending the sail [= Fem]' (HN; PN2017-10-30)

b. ŋanaː-'-paɾkaː na-ni-l̪aɾi
/ŋanu-aK-w$_2$aɾkaː/ /na-ni-l̪aɾi/
1s>3sM.R-Appl-mend.Past2 your brother
'I was mending it for your brother' (HN; PN2017-10-30)

(104) a. ŋawu-pa-paɾk-aː ana-ŋuɻa
/ŋawu-pa-paɾk-aː/ /ana-ŋuɻa/
1sg>3Neut.R-Red-set_ablaze-Past2 Top.Neut-fire
'I kept blowing on the fire' (MM; PN2017-07-07.2 00.03.12)

b. ŋawaː-ʼni-paɾk-aː
/ŋawu-aniK-paɾk-aː/
1sg>3Neut.R-ashes-set_ablaze-Past2
‘I blew on the ashes’ (HN; PN2017-10-30)

The upshot of all this is that the number of genuine minimal pairs which differ in the continuancy of their initial segment is so small that full homophony between roots is not a particularly credible threat. But nor can we say that homophony is altogether avoided, since in some compounds an otherwise stop-initial root can fail to undergo ŋu-Epenthesis even when that makes it indistinguishable on the surface from another root. Thus [paɾkaː], without /ŋu-/, serving as the right-hand member of a compound, can be related to /paɾka/ ’set ablaze’ or /w$_2$aɾka/ ‘mend’, which in any case clearly have very different semantics.

Further, while the application of /ŋu-/ at the prefix-stem boundary does serve to avoid homophony among minimal pairs of roots, it actually *produces* greater homophony in the pronominal prefixes. The result of concatenating the three consonant-final intransitive prefixes with *ŋu-* is surface-identical with the prefixes expressing the same subject with a FFem object, as shown in (105). So, for example, before a stop-initial stem, [ŋaŋu-] could, absent other information, be interpreted as either the 1sg irrealis intransitive, or the 1sg irrealis with a F/Fem object.

(105) Homophony of pronominal prefix surface forms:

	intransitive	transitive (F/Fem obj.)
1sg.Irr	ŋaŋu-	ŋaŋu-
	/ŋaN-ŋu-/	/ŋaŋu-/
12du.Irr	naŋu-	naŋu-
	/naN-ŋu-/	/naŋu-/
2sg.R	nuŋu-	nuŋu-
	/nuN-ŋu-/	/nuŋu-/

One may object that, since the pronominal prefixes in the left column are all intransitive and those on the right are transitive, the forms can still be distinguished based on the transitivity of the verb they combine with. However, bear in mind that all intransitive prefixes can appear with transitive verbs, in which case they are interpreted as transitive prefixes with a Resid class object. Therefore, in a form like [na(-)ŋu-cuɻ-aŋ] ‘let’s (us two) push X’ there is no ambiguity that the root is /cuɻa/ ‘push’ and not /juɻa/ ‘carry, transport’, but there is ambiguity about

whether the object is of the Resid class or the F/Fem class, and this ambiguity is introduced precisely by the insertion of /ŋu-/. We can say then that if homophony is indeed to be avoided in roots, it is only achieved at the cost of engendering homophony in pronominal prefixes. Of course, homophony is commonplace among the pronominal prefixes even where /ŋu-/ is not involved: e.g. /ŋaŋi-/ could be either '1s>3Masc.R' or '3sM>1s.R' — hardly surprising given that only 3 vowel qualities and 8 consonantal phonemes can appear in pronominal prefixes, and yet these must sustain a huge array of grammatical feature distinctions.

However, there is another way of interpreting Heath's argument, such that /ŋu-/ does not avoid strict homophony between lexical roots so much as it instantiates contrast at a more abstract, system-wide level. The suggestion here is that phonological contrast may be an end in itself, even if it is only rarely necessary to support contrasts in meaning in practice. This position merits closer scrutiny. In §3.4, I show that if we reframe contrast as the reduction of uncertainty within paradigms of related word forms, we can capture the spirit of Heath's insight without running into the many problems that arise when we task a phonological grammar with generating /ŋu-/ from monolithic invariant underlying representations. These problems are discussed in detail in §3.3.

3.2.7 Summary

This section has set out the empirical foundations of the theoretical discussion to follow. We have seen that there is a crucial contrast in Wubuy morphs between those that are always stop-initial, and those that alternate between stop-initial and continuant-initial realisations. Heath (1984) places this contrast in their underlying representations: the latter set are underlyingly continuant-initial and become stop-initial due to a process he calls Hardening, which applies following a stop or a nasal. Morphs in the former set are instead underlyingly stop-initial, and undergo a different process, ŋu-Epenthesis, in the same environment where Hardening applies to underlying continuants. This process, the primary focus of this chapter, inserts the syllable /ŋu/, which does not appear to make any independent meaning contribution. Heath suggests a functional explanation: ŋu-Epenthesis ensures that lexical contrasts are not neutralised on the surface by Hardening.

However, this account appears to inevitably lead towards representations that are highly *abstract*, rules that are *unnatural*, and interactions between them that are very *complex*. Indeed, Heath suggests that, "[b]ecause of the opacity of the underlying representations and the complexity of our derivations, it could reasonably be objected that a more *surface-oriented analysis* might be called for … I am rather sympathetic to such a surface analysis" (Heath 1984, pp. 148–9, emphasis in original). I am also sympathetic to a more surface-oriented analysis, and the rest of this chapter is dedicated to discovering just what kind of theory is required to achieve that goal.

3.3 Theoretical issues

3.3.1 Where lies *ŋu-*?

We have seen that, barring some examples of lexicalised behaviour in certain compounds, the appearance of *ŋu-* is phonologically conditioned. However, it is not immediately obvious how we should analyse this [ŋu] string. If it is to be seen as the result of a phonological process, it is a peculiarly unnatural process, since it lacks any plausible phonetic grounding. If it is morphology, it is not a morph of the typical kind, since it lacks any meaning.

In this section, I discuss the obstacles that must be overcome in order to reach a compelling theoretical treatment of *ŋu-*. I identify four problems: unnaturalness, emptiness, opacity, and outward-sensitivity. The theoretical value of ŋu-Epenthesis, and what makes it so unique, lies in the fact that all of these problems occur simultaneously. As a consequence, patching a theory to handle just one or a subset of the problems will not lead to a solution, and we must instead reckon with questions at a much deeper level than would be necessary when dealing with comparatively simpler alternations.

The first of these issues is that ŋu-Epenthesis is phonologically conditioned but not phonologically optimising. It cannot be seen as a 'repair' strategy, modifying some dispreferred structure in the underlying input in order to yield a more well-formed (less marked) structure on the surface — it is, in other words, *unnatural*. This represents a particularly acute problem for classical OT, since the theory explicitly rules out non-optimising alternations (Moreton 2004). But it is also a problem for a number of OT sub-theories that attempt to model alternations as competition between lexically-supplied allomorphs (e.g. Bonet 2004; MacBride 2004; Mascaró 2007; Wolf 2008), though for somewhat different reasons in each case. I examine these issues in detail in §3.3.2.

Further, ŋu-Epenthesis cannot be seen as optimising even language-internally, since its interaction with other processes is *opaque*, as I demonstrate in §3.3.4. Of special interest is the fact that *ŋu-* demonstrates two kinds of opacity — 'overapplication' and 'underapplication' — at the same time. Again, this is especially problematic for OT, given its well-known inability to handle opaque alternations generally, due to its insistence on mapping underlying inputs directly to surface outputs, with no intermediate stages (Idsardi 1998; Kager 1999; McCarthy 2007; Pruitt 2023). Unfortunately, the various modifications and extensions that have been proposed to address this shortcoming lack a solution that is general enough to accommodate both these types of opacity at once (e.g. Benua 1997; Kirchner 1996; Łubowicz 2003; McCarthy 1999, 2003a, 2007; Moreton & Smolensky 2002).

Even if we jettison the assumption that *ŋu-* reflects a phonological repair and instead regard it as a morphological process, we must contend with the fact that it does not realise any

morphosyntactic feature: it is *empty*. This is clearly a problem for any theory in which morphs (or the rules that license them) necessarily bear features (e.g. Xu & Aronoff 2011). But the non-optimising phonological conditioning also falls outside the predictions of some theories that explicitly engage with the problem of empty morphs (Caballero & Inkelas 2013; Stump 2001). This issue is raised throughout §3.3.2 and §3.3.4.

Finally, the fact that ŋu-Epenthesis is sensitive to a two-sided context renders it *outward-sensitive*. That is, in order to know whether to insert *ŋu-*, the grammar would need to 'look ahead' from an inner morphological layer to an outer layer, but, paradoxically, that outer layer has not yet been applied at the time that *ŋu-* must be inserted. Outwards-sensitivity is ruled out by OT-based theories with morphologically-triggered constraints (MacBride 2004; Wolf 2008; Xu & Aronoff 2011), or with stratal or serialised architectures (Caballero & Inkelas 2013; Kiparsky 2000; Paster 2006; Wolf 2008), but also for certain 'amorphous' models of morphology (Stump 2001). I discuss the problem of outwards-sensitivity in §3.3.5.

As we shall see, however, all of these problems — emptiness, unnaturalness, opacity, and outwards-sensitivity — can be diagnosed as proceeding from a commitment to what Blevins (2006) calls 'constructive' models of morphology, whereby words are built up incrementally from the root or stem outwards, and where individual forms are derived in isolation from other forms in the system (Blevins 2006, p. 534). As an alternative, I propose in the following section, §3.4, that *ŋu-* should instead be seen in terms of paradigmatic alternations between full words, as in 'abstractive' models (Ackerman & Malouf 2016; Blevins 2006; Bonami & Beniamine 2016; Lévêque & Pellard 2023; Wilmoth & Mansfield 2021). Taking an abstractive perspective allows us to see that *ŋu-* makes the surface phonological form of one word predictable given the phonological form of a related word, which would otherwise be severely limited due to the surface neutralisation of continuants and stops. Then, in §3.5, I introduce the Relational Morphology formalism and show how it captures phonological alternations as relations between surface complex words stored in the lexicon, overcoming the obstacles identified here. I then apply it to the *ŋu-* problem in §3.6.

3.3.2 Phonological epenthesis, natural and unnatural

Given the fact that *ŋu-* does not contribute any meaning content and is conditioned by phonological factors, Heath (1984) posits that the appearance of *ŋu-* is governed by the phonological grammar, and expresses this in the form of a transformational rule, epenthesising the string [ŋu] when its conditions are met. This rule is serially ordered — in a rather complex, partly cyclical web of interactions — with respect to 49 other transformational rules. However, in light of the directions that phonological theory has taken over the past fifty years, such an explanation no longer seems satisfying.

Whether one appeals to innate universals of markedness (Prince & Smolensky 2004) or to universals of phonetically-based sound change (Blevins 2004; Bybee 2012), an alternation pertaining to phonology is typically expected to be grounded in some degree of 'naturalness' (Archangeli & Pulleyblank 1994; Beguš 2018; Hayes 1999; Mortensen 2006). That is, a plausible phonological rule should bear some relation to known phonetic processes.

Epenthetic processes can in some cases be analysed in such terms. Blevins (2008) provides a survey of consonant epenthesis, distinguishing between 'natural' patterns, which are phonetically conditioned, and 'unnatural' ones, which arise for other reasons. Cases of natural epenthesis include glide insertion between vowels, arising from the reanalysis of V-to-V formant transitions, or insertion of [h] and [ʔ] at prosodic boundaries, arising from the laryngeal activity associated with low pitch at the beginning and end of larger prosodic constituents, as well as with departures from a resting posture as one initiates or terminates speech. These are genuine cases of epenthesis in the sense that an additional segment appears in the pronunciation of a word that was not present at an earlier stage in the language, and represent the phonologisation of well-established phonetic processes. Such cases may lend themselves to an analysis in terms of phonological markedness.

But there are also 'unnatural' types of epenthesis. One major type arises when a consonant that existed in a historically prior form of the language is deleted in all but a restricted environment, where it comes to appear synchronically epenthetic. This is what was suggested in Harvey & Baker's (2020) reconstruction of Marra and Alawa 'epenthetic' prefixes, and indeed, what I suggested was a possible source of Wubuy *ŋu-* also. In such cases, the epenthetic material is not limited to vowels, glides and glottals, and may even be unpredictable: in many Oceanic languages, for instance, deleted root-final consonants reappear under suffixation, and one cannot predict the identity of this 'epenthetic' consonant from the form of the root alone (Blevins 2008). As Blevins makes clear, such unnatural (that is, non-historical) types of epenthesis cannot be generated by synchronic markedness constraints, as markedness considerations pertain only to the earlier lenition and deletion processes that led to the present situation.[16]

In the case of Wubuy *ŋu-*, it is quite clear that it is unnatural in Blevins' terms. In order to analyse *ŋu-* as a natural process, bringing surface forms into harmony with well-formedness constraints, we would have to show both that non-application would result in some marked or dispreferred structure, and that the inserted material is unmarked. Neither of these conditions holds. For one thing, most would regard the sequence prefixal *ŋu-* interrupts — a nasal followed

[16] Of course, one could always resort to ad-hoc constraints, or reanalyse the synchronically epenthetic element as emerging from universal markedness; this is the path taken in some OT accounts of English Intrusive [ɹ] (e.g., Baković 1999; McCarthy 1993; Uffmann 2007), a phenomenon that is well-known to be the historical result of conditioned deletion and is therefore unnatural in Blevins' sense. For further criticisms of the markedness account, see Vaux 2002; Vaux & Samuels 2017.

by a stop — as the least marked of all consonant sequences. Hamilton (1996) shows that these are the most frequent CC sequences in Australian languages, and notes the 'phonotactic symmetry' between homorganic nasal-stop clusters and singleton consonants. Indeed, as I discuss in the next section, the nasal-stop sequence that ŋu-Epenthesis avoids is exactly the sequence that Hardening demands, so we cannot regard ŋu-Epenthesis as a repair strategy even language-internally. We must also note that linking *ŋu-* appears in a broader range of contexts than prefixal *ŋu-*, so arguments for the markedness of the latter rule's structural description do not necessarily translate to the former.

In other contexts, insertion of an entire syllable has been argued to optimise prosody, such as in order to meet word minimality requirements, as in Martu Wangka (McManus 2008) or Lardil (Round 2011; Wilkinson 1988). The ŋu-Epenthesis pattern cannot be analysed in these terms, however, as *ŋu-* appears before a stop-initial root of any number of syllables, and those syllables may be open or closed, with a short or a long vowel, as is amply demonstrated by the examples throughout this chapter. We therefore cannot say that *ŋu-* repairs any prosodic defect.

At the same time, we do not have any independent evidence that the specific 'repair' that is initiated here is the optimal one. To the best of my knowledge, no one has ever suggested that a syllable of such a phonological shape could be predicted from, say, coarticulation through the overlap of the articulatory gestures involved in a cluster of non-continuants, or reanalysis of clusters due to perceptual ambiguity. For insertion of [ŋu] to constitute a good repair strategy, we would need to find some indication that [ŋu] is less marked than some other syllable — [ŋa] or [wa], for instance. If text frequency is any indication of relative markedness (Hume et al. 2011), the diphone [ŋu] is only the 46th most frequent within-word diphone in Heath (1980). Therefore, we do not have independent grounds for treating epenthesis as a repair for a marked structure, nor for treating the string /ŋu/ as the least marked material to insert in such an operation.

Neither can we avoid the naturalness problem by reversing the terms of the analysis, such that underlying forms with /ŋu-/ attached already are used to derive stop-initial morphemes without /ŋu-/. For instance, Baker (2009) suggests a phonological analysis according to which stop-initial roots all have /ŋu/ prefixed to them in their underlying lexical representation, while roots that would otherwise be subject to Hardening are underlyingly stop-initial, and lenite on the surface (as in the Lenition analysis that Heath rejects). A phonological rule then deletes this string under the inverse conditions to those of Heath's ŋu-Epenthesis rules. Under such an analysis, the root *puɾa* 'sit' would underlyingly be /ŋupuɾa/ (instead of /puɾa/ as in Heath's analysis), and would lose the initial syllable just when following an inflectional prefix ending in a vowel. However, there is no reason to regard a synchronic process deleting /ŋu/ under these conditions as any more natural than a process inserting it, and there is counter-evidence in the

form of /ŋu/-initial roots that do not undergo such deletion, like *ŋucica* 'fish', *ŋupunuŋ* 'cloud', *ŋuji:* 'eat.NPST2', *ŋucpi* 'to chase away', *ŋurta* 'to groan', *ŋuniŋ* 'sex' — though, admittedly, there are no incorporable nouns or verb roots of the form /ŋu/+stop+V.[17]

To the extent that the appearance of *ŋu-* is unnatural, it is resistant to any analysis in which alternations are assumed to be driven by phonological well-formedness. Yu (2007) calls this assumption Phonological Readjustment: roots and affixes are supplied (separately) by the lexicon in invariant form, and whatever phonological differences arise when they appear in surface words is the result of changes demanded by a phonological grammar (see also Paster 2006). This is precisely the assumption made by most varieties of Optimality Theory, and indeed by generative phonology more broadly. OT, in its original formulation (Prince & Smolensky 2004), predicts no unnatural alternations at all: any time a surface form differs from its underlying representation, it does so under the pressure of high-ranking Markedness constraints, which are supplied by Universal Grammar and reflect putatively universal preferences. As put succinctly by Moreton (2004, p. 148), "if the output is not identical to the input, it must be less marked than the input". Indeed, it is difficult to imagine what combination of markedness and faithfulness constraints could result in inserting [ŋu] just in the contexts required. And as we will see in §3.3.4, the opacity of *ŋu-* means it is difficult to analyse in terms of well-formedness even in a language-particular sense, especially in OT.

As we have seen, ŋu-Epenthesis would surely fail just about any criterion of naturalness or unmarkedness one could come up with: it does not resemble any of the various kinds of epenthesis (or deletion) processes argued to be universal in the literature, and relatable to concrete phonetic pressures. On that basis, we can conclude that *ŋu-* is not the reflection of phonological epenthesis.

3.3.3 Empty morphs and allomorphy in an optimising grammar

If, given the preceding discussion, ŋu-Epenthesis does not look like a strictly phonological process, it seems plausible that we may arrive at a better account by treating the form variation we see among surface words as more morphological in nature. That is, we might seek a solution by treating the [ŋu] we find at the surface not as the epiphenomenal result of an optimising phonological epenthesis process, but rather as realising a morphological constituent in its own

[17] There are other issues with Baker's (2009) Lenition proposal. One of them is reduplication: we would need to explain why reduplication targets the string *after* the prefixal /ŋu/ when there is one, and why its shape depends on the sonority of the initial segment of this string — in other words, why this /ŋupura/ stem reduplicates as *-ŋu-pu-pura* and not **-ŋupu-ŋupura*. See §4.2.4 of Ch. 4 for the interaction of ŋu-Epenthesis and Reduplication. Also, recall from (105) above that there are many pronominal prefixes ending in the string /ŋu/, not derived by ŋu-Epenthesis, which would presumably be subject to a deletion rule of this kind also, absent some morphological restriction.

right, albeit an empty morph. To that end, a number of proposals have attempted to integrate allomorph choice into an OT grammar. These proposals have diverse views on the nature of the phonology-morphology interface, however. One set of proposals locates allomorphs in the lexicon, and tasks an OT grammar with choosing the right one for that context; the other set uses language-particular constraints to realise morphs, and ranks the constraints in such a way as to select the correct morph for the context.

According to Bonet's (2004) Morphemes in the Input Hypothesis (MIH), the input to EVAL can include a root in its underlying representation together with a (disjunctive) set of allomorphs for any affixal material; EVAL then decides on the correct allomorph according to which incurs the lowest-ranked violations. The constraints in this case are the same markedness and faithfulness constraints as in standard OT. For example, an English plural output like *cats* would reflect an input /kæt, {-s, -z, -əz}/, with a root and a set of allomorphs. Any candidate with one of these allomorphs would satisfy identity constraints, and phonological markedness constraints penalising, say, voicing mismatches in clusters would take care of choosing [kʰats] as optimal. Mascaró (2007) pursues a similar argument, but his allomorph sets are intrinsically ordered: the first in the set is preferred unless it violates a sufficiently high-ranking constraint, favouring the next in the set, and so on.

The MIH is different to a strictly phonological, 'vanilla' OT account in the sense that it allows at least some form variation to be supplied by the lexicon, rather than emerging from the interaction of markedness and faithfulness (i.e. the grammar). However, since the choice of one allomorph in the set over another is still decided by universal markedness constraints, the MIH predicts, like standard OT, that allomorphy is necessarily phonologically optimising (Paster 2006) — it is ultimately Phonological Readjustment under a different guise.

Other models (Hammond 2000; MacBride 2004; Picanço 2002; Russell 1995; Xu & Aronoff 2011; Yip 1998) propose language-particular morpheme realisation constraints to drive allomorph selection. The input to EVAL is a root together with affixal morphemes of a more abstract nature, lacking form altogether. On this view, English plurals would be fed to the input as a root plus any morphosyntactic categories that need to be realised, e.g. /kʰæt/+Plural. The phonological form associated with the plural is governed by morphological constraints, such as Pl=z, Pl=s and Pl=əz (MacBride 2004, p. 26). Wolf's (2008) theory of Optimal Interleaving, while assuming a more unusual implementation of OT which evaluates 'chains' of serial derivations instead of surface forms (as in McCarthy 2007), also posits an OT grammar in which morphosyntactic categories are associated with phonological form through constraints.[18]

[18] There are some technical differences. Optimal Interleaving (Wolf 2008) maintains a separation between the lexicon and the grammar, and morphological constraints are a variety of identity constraint ensuring that the phonological form of the candidate accords with the underlying phonological representation of that morpheme as listed in the lexicon. MacBride (2004), Russell (1995), Hammond (2000) and Yip (1998) largely seek to obviate

Just as in the MIH, in these theories allomorphy and regular phonology are handled simultaneously, in the same part of the grammar. Because morphological constraints specifying the phonological content of affixes are essentially arbitrary, and can be ranked above phonological constraints, such theories predict that allomorphy will be phonologically non-optimising at least some of the time.

There are limits to how much arbitrariness these models will tolerate, however. In general, they seek to capture competition between phonological markedness and faithful feature expression in such a way that one or the other demand *must* provide the justification for any allomorph choice. How, then, can we generate the appearance of an empty morph like Wubuy *ŋu-*? Without any morphosyntactic feature to realise, we cannot specify any morphological constraint that *ŋu-* could satisfy, nor is there anything to include in the input that would make that morphological constraint become active. The only recourse is once again to generate it via the phonological constraints, which again is only possible if it is driven by markedness. Thus while regular contentful allomorphy can be non-optimising, empty morphs cannot be. This leaves us in a bind: a morphosyntactically featureless morph can surface only in order to satisfy some phonological markedness constraint, while a phonologically non-optimising morph must only surface in order to satisfy a morphological identity constraint. There is apparently no room for a morph which does neither. Interestingly, Wolf mentions Wubuy in passing as an example of a language with "non-meaningful affixes inserted to satisfy phonological requirements" (Wolf 2008, p. 139), apparently assuming that *ŋu-*, being generated by a phonological rule, must be a strategy for satisfying phonological markedness — which, as we have seen, it is not.

It is not only empty morphs that cause the breakage here. Any allomorphy whose conditioning is only statable in terms of its phonological environment, but which does not reduce phonological markedness, will create problems. This is because, both in theories utilising constraints to realise morphs as well as theories with input allomorph sets, affixes cannot specify their own phonological requirements directly (i.e., through subcategorisation). This characteristic can be observed in Wolf's discussion of Dyirbal ergative suffix allomorphy, which, similar to many other Pama-Nyungan languages, appears as /-ŋku/ if the stem is disyllabic and V-final, and as /-ku/ otherwise (Wolf 2008). The conditioning environment here is strictly phonological, even though it does not produce a less marked nor more faithful structure on the surface. Because of his theoretical insistence that a morph that does not serve phonology must serve morphology, Wolf is forced to suggest, without any independent evidence, that one of these morphs is a more faithful instantiation of the ergative category than the other (p. 62).

Bonet (2004) also discusses the Dyirbal problem, acknowledging that neither realising the

the rule/lexicon distinction; they posit language-particular constraints that specify phonological form and take over the functions of morphemes as previously understood.

suffixes through input allomorph sets nor through morphological constraints offers a particularly satisfactory account, as the phonological conditioning here is fundamentally arbitrary. The crucial similarity between the Dyirbal case and that of Wubuy is that in both, each allomorph — /-ŋku/∼/-ku/ in the former and /ŋu-/∼/-Ø-/ in the latter — appears in an arbitrary but predictable phonological environment, and each allomorph is morphosyntactically equivalent. I am hardly alone in calling attention to this problem. For several decades now, research on allomorphy has pointed to the fact that affixes can have phonological requirements that do not result in forms that meet the language's own phonological preferences (see esp. Booij 2012; Bye 2006; Carstairs 1990; Lapointe 2001; Paster 2006, 2009; Yu 2007). Carstairs (1990, p. 19) cautions that "the existence of a phonologically conditioned alternation does not by itself prove the existence of some synchronic phonological process giving rise to it", nor should we assume "a shared phonological underlier" holding between allomorphs, whether obviously suppletive or otherwise.[19]

We can conclude from all of this that we cannot reach a satisfactory OT analysis of *ŋu-* by either feeding allomorph sets to the input, nor by positing morphological realisation constraints. In the first case, we again run up against the problem of non-optimising phonological conditioning; in the second, we can have non-optimising allomorphy but only if there is some morphosyntactic feature the allomorph can realise. Moreover, we still have opacity to contend with, which is not addressed by any of the theories so far discussed. It is to the problem of opacity that we now turn.

3.3.4 Opacity

Even if we manage to get past the question of ŋu-Epenthesis' naturalness, its interaction with other phonological alternations in the language presents further difficulties for an analysis in terms of Phonological Readjustment, insofar as these interactions are *opaque*.

We can consider opacity in terms of Baković's (2011) definition in (106), following Kiparsky (1973) and McCarthy (1999). Note that, since opacity is often most clearly demonstrated by means of serial derivations of the *SPE* variety, I will discuss the Wubuy patterns in that idiom before engaging with any problems pertaining to parallel evaluation.

(106) A phonological rule *P* of the form A → B / C_D is opaque if there are surface structures with any of the following characteristics:

a. instances of A in the environment C_D (*underapplication*)

[19] See Czaplicki (2013) for discussion of a number of related issues in Polish, where three different palatalisation processes, while phonetically natural and therefore optimising when examined in isolation from one other, appear strikingly arbitrary when encountered in the same synchronic system.

b. instances of B derived by *P* that occur in environments other than C_D (*overapplication*)

Indeed, both *overapplication* and *underapplication* opacity are simultaneously in evidence in ŋu-Epenthesis: overapplication in the rule's left-hand environment and underapplication in its right-hand environment. I will discuss both of these in turn.

Recall that, in Heath's (1984) formulation, ŋu-Epenthesis at the pronominal prefix-stem boundary is conditioned by a two-sided phonological context: it inserts /ŋu/ preceding a stop and following a consonant. As we saw above, this consonant can only be the nasal archiphoneme /N/, which is always subsequently deleted due to an independent process, Nasal Deletion. As the derivation in (107) shows, from the input /ŋaN-puɾaŋaŋ/ 'I will sit', ŋu-Epenthesis produces the intermediate form |ŋaN-ŋu-puɾaŋaŋ|, as in (107). Nasal Deletion now has an opportunity to apply, since ŋu-Epenthesis has created a /N-ŋ/ sequence that was not present in the underlying form. The result of that transformation is the surface form [ŋa-ŋu-puɾaŋaŋ], where ŋu-Epenthesis has appeared to *overapply*, because the prefix-final consonant which triggered it is not visible on the surface. Indeed, the surface environment of *ŋu-* in the output of (107b) is no different to the output of (107a), without *ŋu-*.

(107) Overapplication opacity: ŋu-Epenthesis self-destructively feeds Nasal Deletion

Input:	a.	/ŋa-puɾaŋaŋ/	b.	/ŋaN-puɾaŋaŋ/	
ŋu-Epenthesis		—		ŋaN-ŋu-puɾaŋaŋ	Ø → ŋu
Nasal deletion		—		ŋa-ŋu-puɾaŋaŋ	N → Ø
Output:		[ŋa-puɾaŋaŋ]		[ŋa-ŋu-puɾaŋaŋ]	

This interaction between ŋu-Epenthesis and Nasal Deletion thus represents a particular kind of overapplication opacity that Baković (2007, 2011) terms *self-destructive feeding*. A rule is said to 'feed' another rule if it is ordered before it and changes the input string in such a way that the later rule is now able to apply. In self-destructive feeding, the application of an earlier rule feeds a later rule which changes the string in such a way that the earlier rule appears no longer justified. That is, if our generalisation says that ŋu-Epenthesis (at the prefix-stem boundary) is triggered by a preceding prefix-final consonant, and then Nasal Deletion deletes this consonant, the generalisation asserted by the structural description of ŋu-Epenthesis — that /ŋu/ should occur between a prefix-final consonant and a stem-initial stop — is apparently falsified by the surface realisation, since no prefix-final consonant is in fact present. Owing to the fact that the only possible pronominal prefix-final consonant is /N/, ŋu-Epenthesis at the prefix-stem

boundary cannot do otherwise than feed Nasal Deletion. In other words, the left-hand context of prefixal /ŋu-/ is necessarily opaque: it is not merely sometimes but *never* surface-apparent.

Linking *ŋu-*, on the other hand, will be inserted before stop-initial roots regardless of the phonological environment to their left (barring some lexical exceptions, as discussed). Over-application opacity is therefore not necessarily implicated for linking *ŋu-*. However, as I will discuss in §3.4.3, complex stems frequently involve an opaque interaction between Hardening and the deletion of velar stops, which will help us understand why linking *ŋu-*, unlike prefixal *ŋu-*, is not subject to phonological conditioning at its left edge.

The right-hand environment is no more transparent than the left. The opacity here concerns the interaction of ŋu-Epenthesis with Hardening, and is instead a case of underapplication (or counterfeeding) opacity. Recall that Hardening targets a continuant consonant following a stop or nasal and makes it a stop, resulting in a nasal-stop sequence. A nasal-stop sequence is precisely what ŋu-Epenthesis targets at the prefix-stem boundary. However, ŋu-Epenthesis cannot apply to the output of Hardening, as we have seen. Therefore, we have a situation meeting the definition of opacity in (106a): there exist word forms in which the target to ŋu-Epenthesis is present on the surface without apparently having undergone the transformation described by the rule. In serial terms, Hardening *counterfeeds* ŋu-Epenthesis, meaning that if the ordering of the two rules were reversed, Hardening would have fed ŋu-Epenthesis; the fact that there are morpheme-initial surface stops without evidence of ŋu-Epenthesis attests that Hardening must in fact come later in the derivation, where ŋu-Epenthesis can no longer apply. A derivation is illustrated in (108).

(108) Underapplication opacity: Hardening counterfeeds ŋu-Epenthesis

	2sg.R-push.NPST2	2sg.R-carry.NPST2	
Input:	/nuN-cuɻi:/	/nuN-juɻi:/	
ŋu-Epenthesis	nuN-ŋu-cuɻi:	—	
Hardening	—	nuN-cuɻi:	(*nuN-ŋu-cuɻi:)
(Nasal del./assim.)	nu-ŋu-cuɻi:	nuɲ-cuɻi:	
Output:	nu-ŋu-cuɻi:	nuɲ-cuɻi:	

In the phonology literature, much has been made of the theoretical centrality of opacity, in large part because the strictly parallel evaluation mechanism of OT, despite its many obvious successes, is in principle unable to capture many kinds of opaque interactions which are fairly straightforward to describe in terms of rule ordering in *SPE* serialism (Baković 2007; Bermúdez-Otero 2003; Idsardi 1998; Kager 1999; McCarthy 1996; Pruitt 2023). Under rule-based serialism, a

derivation may involve any number of intermediate steps between the underlying representation at the input and the surface realisation at the output. Opacity always involves positing some intermediate form which is different from both the input and the output, as seen in (70). Standard parallel OT, on the other hand, allows only an input and an output, with no intermediate stages. This quality is otherwise considered a point in OT's favour, as it allows for a compelling account of 'conspiracies', and seems to move phonological theory closer to psychological plausibility by removing unnecessary layers of abstraction and focusing on surface patterns. But the absence of intermediate stages in OT is also the property usually understood to make opacity impossible.

To see this, let us attempt to consider the counterfeeding dimension of the *ŋu-* problem in terms of surface markedness and faithfulness constraints. The continuant-stop alternations accounted for by Hardening in the rule-based analysis appear to demonstrate that [–CONT][–CONT] consonant sequences are less marked than [–CONT][+CONT] consonant sequences; indeed, the latter sequence never appears in any surface words in Wubuy. We can then infer some markedness constraint, perhaps appealing to the Syllable Contact Law (Murray & Vennemann 1983), which penalises heterosyllabic sequences of rising sonority; this is the operative constraint in Chong's (2011) analysis of Hardening in Wubuy, which unfortunately does not consider ŋu-Epenthesis. In order to favour candidates applying the alternation, this markedness constraint must be ranked above the faithfulness constraint penalising changes to the [CONT] feature of the input. This part of the analysis is possible in OT, as demonstrated by Chong (2011).

When we consider the roots subject to ŋu-Epenthesis, however, problems emerge. If ŋu-Epenthesis serves to avoid nasal-stop sequences in the input underlying form, this would entail that nasal-stop sequences are more marked than [ŋu]-stop sequences. That is, if we accept that nasal-stop is a less marked sequence than nasal-continuant (109a), and [ŋu]-stop is a less marked sequence than nasal-stop (109b), we have to accept that [ŋu]-stop is a less marked sequence than nasal-continuant (109c): therefore, whatever constraint ranking generates ŋu-Epenthesis for underlying stops will generate it for surface stops also.

(109) a. /nasal-cont./ → nasal-stop

b. /nasal-stop/ → [ŋu]-stop

c. **/nasal-cont./ → nasal-stop → [ŋu]-stop*

While we have already seen in the previous section that ŋu-Epenthesis does not resemble a markedness avoidance strategy, a high-ranking markedness constraint is the only way of favouring a candidate that departs from the input in OT. For the sake of the argument, we might imagine here that the epenthesising forms are driven by a ban on sequences of non-continuants. But the actual content of the constraint is irrelevant to the opacity issue. Whatever markedness

constraint(s) may be driving the insertion of [ŋu] before stops, the architecture of Optimality Theory furnishes us with no way of stopping this constraint from forcing the same repair onto derived stops as well as underlying ones.

Note that this opacity issue is not simply an artifact of Heath's (1984) analysis of the underlying representations and the readjustments imposed on them, as it emerges even if we assume the Lenition with ŋu-Deletion analysis suggested by Baker (2009). If this were the case, we would still have to contend with an undesirable 'fell swoop' change, this time going from an underlying /ŋu/-stop to a surface continuant, as in (110c), in the complementary environment to that of (109). In this case, the problem is that Lenition underapplies, since it does not apply to the output of (110b), as expected.

(110) a. /V-stop/ → V-cont.

b. /V-ŋu-stop/ → V-stop

c. *[V-ŋu-stop] → V-stop → V-cont.

We are retracing some well-worn terrain here. We are essentially concerned with how to accommodate a chain shift in OT — a question that has never been satisfactorily answered, despite decades of attempts to do so (see Idsardi 1997, 1998; Kiparsky 2000; Kirchner 1996; Łubowicz 2003; McCarthy 1996, 1999, 2007; Mielke et al. 2003; Moreton & Smolensky 2002). As diagrammed in (111), a chain shift is an alternation in which some underlying string A becomes surface string B, but underlying B (in the same environment) becomes surface C. In OT terms, $A \prec B$ (where $\prec$ = 'is less optimal than'), and $B \prec C$. A typical example can be found in Servigliano Italian, where there is a partial height harmony process such that /e,o/ become [i,u] before suffixal /i,u/, while /ɛ,ɔ/ become [e,o] in the same environment (Smith 2020). This parallels our present case: Hardening repairs an underlying nasal-continuant sequence (A) by realising a nasal-stop sequence on the surface instead (B), while ŋu-Epenthesis repairs an underlying nasal-stop sequence (B) by inserting an intervening /ŋu/ (C).[20] In rule ordering terms, as we saw, this is what emerges when the rule that derives $A \rightarrow B$ could feed the rule deriving $B \rightarrow C$, but it is ordered too late to do so, thus counterfeeding it. (Again, we come no closer to addressing the *ŋu-* problem by reversing the process, as suggested by the Lenition analysis: one simply reverses all the arrows in (111), such that we must instead explain why C does not become A.)

[20] Technically, the two patterns are different subtypes of underapplication in Baković's (Baković 2007) typology: the Servigliano Italian one is 'counterfeeding-on-focus', while the Wubuy one is 'counterfeeding-on-environment'. The distinction is not relevant to the present discussion.

(111) Underapplication (chain shift) opacity:

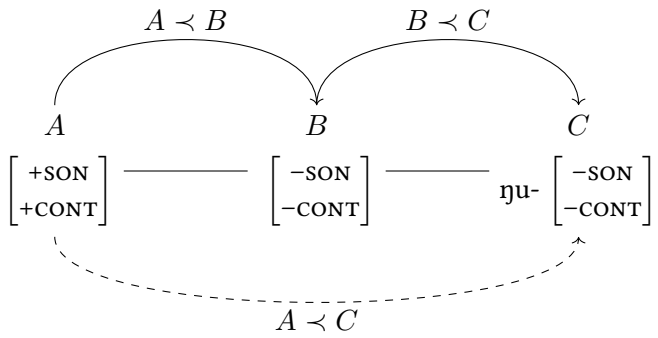

The difficulty for OT is that there appears to be nothing to stop A becoming C in one fell swoop — here, a nasal-continuant sequence becoming a *ŋu-* plus stop sequence — which is not the attested pattern. If B is more optimal than A, and C is more optimal than B, the inevitable conclusion is that C is more optimal than A. But in a chain shift scenario, C is not more optimal *as such*: it is only optimal when the input is B, and not when the input is anything else. Likewise B is only optimal when the input is A, but not otherwise. This is a serious problem for OT. A markedness constraint cannot — without substantially changing the fundamental workings of OT — restrict the scope of its operation by reference to phonological features of the input, e.g. 'A surface string C is preferred but only when the corresponding input string is B', or alternatively, 'A surface string B incurs a violation but only when its corresponding input string is B' (cf. Hauser & Hughto 2020).

The attempt to deal with opacity in parallel OT has led to a proliferation of modifications and extensions to the core theory, with very few of these being taken up with anything close to the enthusiam that greeted OT as originally conceived. The proposals can also be roughly divided into those that maintain parallel evaluation, but introduce new types of constraints, and those that introduce some serial mechanism to the evaluation procedure. In the rest of this section, I will first address some of the most relevant parallel theories of opacity before giving attention to serial ones.

A general observation to make of the parallel proposals is that none of them has an account of all types of opacity; most address either underapplication only or overapplication only, or even just certain subtypes of one or the other (as reviewed by Kager 1999, pp. 372–399). For instance, Transderivational Identity (Benua 1997) proposes an additional family of constraints, Output-Output (OO) Correspondence, enforcing identity between surface word forms that share the same underlying morpheme. It is capable of reanalysing some kinds of overapplication opacity accounted for under serialism by cyclic rule application, but it cannot account for types of opacity that depend on an abstract intermediate level of derivation that does not correspond

to some other existing surface word (Kager 1999, p. 386). Sympathy Theory (McCarthy 1999) broadens the empirical scope to certain other cases of overapplication opacity, while again not accounting for underapplication, and at the same time substantially overcomplicating the theoretical machinery (Bermúdez-Otero 2003; Idsardi 1997; Ito & Mester 2003; Kager 1999). Contrast Preservation Theory (Łubowicz 2003), on the other hand, offers a family of faithfulness constraints to account for cases of underapplication (chain shift) opacity, but they offer no purchase at all on overapplication opacity. A number of other proposals for addressing chain shifts/underapplication opacity draw on Local Conjunction (Kirchner 1996; Moreton & Smolensky 2002), though again it comes with a number of unwelcome side effects.

The inability to cover all types of opacity is enough to render all of the theories just mentioned incapable of deriving ŋu-Epenthesis, given that it shows simultaneous over- and underapplication. But all of these theories face a further hurdle, in that they still assume that alternations are driven by high-ranking markedness constraints — that is, they assume Phonological Readjustment. For the reasons given above in §3.3.2, *ŋu-* is not amenable to an analysis in those terms.

Arguably the greatest success in analysing opacity in OT has come from models that introduce a serial component to the grammar, such as those that developed out of Lexical Phonology and Morphology (LPM; Booij & Rubach 1984; Kiparsky 1982a,c, 1985; Mascaró 1976; Mohanan 1982). In LPM, morphology and phonology are handled by separate modules of the grammar. Words are assumed to be built from the root outwards, and phonology has a chance to apply each time a layer or 'stratum' of morphology has been added to the word-in-progress. Different morphological domains may be associated with different phonological behaviour, hence the distinction between 'lexical' and 'postlexical' phonology. Although LPM originally emerged from the *SPE*-derived rule-based formalisms of the 1970s, LPM's core concepts are assumed in a great number of OT sub-theories, including LPM-OT (Kiparsky 2000), Stratal OT (Bermúdez-Otero 1999; Itô & Mester 2001; Kiparsky 2009), Realization OT (Xu & Aronoff 2011), Cophonology Theory (Inkelas & Orgun 1995; Orgun 1996, 1998), Optimal Construction Morphology (Caballero & Inkelas 2013), and Morphological Doubling Theory (Inkelas & Zoll 2005). These theories generally retain the parallel evaluation component of OT, but allow evaluation to apply to each morphological domain in the derivation of a complex word. These domains can consist of just a Root or Stem level and a Word level (Baker 2008b; Borowsky 1993), or any number of strata (Inkelas & Orgun 1995; Mohanan 1982; Orgun 1996).

As is often acknowledged, any theory which allows at least two 'stages' at which a phonological change can be applied — whether by way of serially-ordered rules or serially-ordered phonological strata — can handle most kinds of phonological opacity without any great difficulty (Baković 2011; Bermúdez-Otero 2003; Kiparsky 2000). This is because a process that applies

to an outer morphological domain can erase the trigger for a process at an inner domain, as in overapplication, or can fail to apply to an inner domain where it seems to be otherwise expected, as in underapplication.[21] For that reason, any one of these LPM-derived theories should be better able to capture the opacity problems surrounding *ŋu-*, as discussed above, just as rule-based serialism can. The underapplication problem, for instance, where ŋu-Epenthesis fails to apply before surface stops derived from underlying continuants via Hardening, could be approached as a matter of assigning the former alternation to an inner Stem level and the latter to an outer Word level.

It is standardly assumed in these models that the phonology at each layer is optimising — that is, that alternations are the result of Phonological Readjustment. Thus, if *ŋu-* is generated by an OT tableau at the Stem level, there must be some way of construing it as optimal at that level. As we have already seen, this is not an attractive proposition. However, certain LPM-derived OT models adopt Phonological Subcategorisation in order to explain phonologically-conditioned allomorphy, either as a complement to Phonological Readjustment (Paster 2006; Round 2012; see also Xu 2016, pp. 567–568), or as a replacement for it (Yu 2007). Under Phonological Subcategorisation, the representation of an affix includes some specification of the kind of base it can attach to (as in Kiparsky 1982d; Orgun 1996; Paster 2006, 2009; Yu 2007). Since subcategorisation frames are provided by the lexicon, they need not be optimising. Under this view, *ŋu-* could be treated as an affixal allomorph (alternating with zero), whose subcategorisation frame requires that it be prefixed to a constituent beginning in a stop. This would make it the input to the phonological component rather than its output.

While this takes us close, there is fatal shortcoming to the serial approach to affixation: the two-sided conditioning of ŋu-Epenthesis appears to require the grammar at an inner layer to 'look ahead' to the contents of an outer layer, something strictly ruled out by stratal models (Xu 2016). This issue is addressed in the next section.

3.3.5 Outwards-sensitivity

Recall that the structural description for ŋu-Epenthesis refers to a two-sided context: it includes information about phonological properties on both the left-hand and right-hand sides of the point of its insertion. In the case of prefixal *ŋu-*, this information is phonological on both sides — it refers to a stop to its right and a consonant to its left. For linking *ŋu-*, the phonological characteristics of the segment on the left are irrelevant, but its morphological characteristics do matter; it will apply so long as the segment on the left of the morphological juncture comes from a derivational element.

[21] Though see Orgun (1996, esp. pp.15-28), who argues for the necessity of interpreting level-ordering in non-derivational terms; he nonetheless rejects any metaphor of temporal succession between levels.

If we assume, as LPM does, that words are built incrementally from the root outwards, this two-sided context is a problem. At the point at which a decision must be made about whether to insert *ŋu-*, the morphology must know about what material is already part of the word (on the right) as well as what material will be inserted later (on the left). In that sense, it violates the 'No Lookahead' constraint (Carstairs-McCarthy 2001; Simpson & Withgott 1986), since it requires that the morphological component be able to 'look ahead' to the result of a later stage of affixation in order to operate correctly. As Wolf (2008) summarises the issue:

> When determining whether or not to apply, a given rule R has access to (at least some) information about what the previous rules have done, but R cannot have access to information about the effect of rules ordered after R, for the very simple reason that those rules haven't applied yet. That is to say, in theories with ordered rules, a given rule can't 'look ahead' to the results of applying subsequent rules. (Wolf 2008, p. 148)

Regardless of the specific mechanism used to drive allomorphy, any theory which sees words as constructed layer by layer from the root outwards will not be capable of deriving any outwards-sensitivity. Paster (2009), for instance, argues for affixes with phonological subcategorisation frames, but these can refer only to the string they attach to at the relevant stage of affixation, and not to any later (outer) layers of affixation. An intolerance for lookahead is also a feature of some theories of morphology which reject the classical morpheme, such as that of Anderson (1992) or Stump (2001), since they share the same commitment to strictly root-outwards (or stem-outwards) morphotactics (Blevins 2006).

Almost all the theories mentioned so far see the inability to account for outward-sensitive allomorphy as a strength, not a weakness, since it is generally assumed not to exist (Xu 2016, p. 7). But as Xu (2016, pp. 9–10) points out, cases of outwards-sensitive phonologically-conditioned allomorphy are attested in Fulfulde, Italian, Sanskrit, Midob, and Kashaya. More detailed examinations of outwards-sensitive allomorphy can be found in Round (2009) for Kayardild, Caballero (2008) for Choguita Rarámuri, Deal & Wolf (2017) for Nez Perce, and Svevonius (2012) for Icelandic.

3.3.6 Constructive and abstractive morphotactics

We have seen that ŋu-Epenthesis is a case of phonologically-conditioned allomorphy which is phonologically non-optimising, is opaque in both directions, does not realise any morphosyntactic feature, and is outwards-sensitive. To the best of my knowledge, no case of allomorphy with this entire set of characteristics has been discussed in the theoretical literature on phonology-morphology interactions. While many of the theories of allomorphy discussed so far can handle

a subset of these characteristics, none can handle the entire set. The ŋu-Epenthesis problem thus provides us with an extraordinarily stringent test of any phonological and morphological theory, since it throws virtually the entire 'kitchen sink' of intractable problems at the analyst.

With that in mind, it is useful here to consider what assumptions are held in common among each of the models mentioned, and whether we can find a plausible way out of the dilemma by overturning them. To that end, Blevins (2006) draws a useful distinction between morphological models which are 'constructive' versus those which are 'abstractive' in their approach to morphotactics. Constructive models — what could also be called 'root-based' or 'morph-based' models — assume "an inventory of morphotactically minimal forms, from which surface forms are 'built' or 'derived'" (Blevins 2006, p. 532). A model is instead abstractive if it takes "surface word forms as the basic elements of the system, and regards roots, stems and exponents as abstractions over a set of full forms" (pp.531-2).

These different approaches also implicate the two different views of the lexicon discussed in Chapter 1. Constructive models try to achieve a 'redundancy-free' lexicon restricted to individual morphs and some exceptional word forms that cannot be derived by rule; predictable patterns are to be expressed elsewhere, by some sort of symbolic computation (rules, constraints, etc.). Abstractive models instead support the storage of fully-inflected word forms in the mental lexicon, including those exhibiting entirely predictable morphological patterns. The process of abstracting patterns from those word forms also takes place in the lexicon, by means of analogy, and analogical patterns can be extended to coin new words, or to interpret unfamiliar ones.

A further distinguishing feature is that constructive models typically derive an individual word form in isolation from other forms within the same system. According to such a perspective, "surface patterns of both words and networks of words (i.e. relations between surface alternants) are not regarded as proper objects of linguistic analysis, while the abstract elements and operations responsible for constructing these ephemera are" (Ackerman & Malouf 2013, p. 430). Abstractive models, on the other hand, analyse words only in relation to other words: whatever sub-word units can be identified are relevant to the extent that they illuminate contrasts between full word forms in a set of related words (i.e. a paradigm), and have no independent existence beyond the words that instantiate them.

The two basic assumptions of constructive models — that the task of morphology is to describe by what procedures a word is built out of minimal, invariant atomic elements until it achieves its full surface form, and that such procedures are best described for one word at a time in isolation from other words — are shared by all the theories mentioned in the preceding section. The Morphemes in the Input Hypothesis (Bonet 2004; Mascaró 2007) assumes that affixed words must necessarily be produced through the combination of a lexically-listed root (with an invariant underlying representation) and a set of possible allomorphs, which are also lexically

listed; the job of the grammar is to decide which of the atomic elements in the supplied set is the correct one for the given phonological environment. Such a question is resolved internally — no consultation with other word forms is required for EVAL to select the winning candidate, and indeed EVAL has no access to any kind of information about the distribution of a given affix. Realization OT (Xu & Aronoff 2011), which assigns morphological exponence to realisation constraints, is similarly commited to the view that morphological systems build complete word forms out of sub-word atomic units, but it limits the content of the lexicon even further so that it only contains roots. Even though Realization OT is hardly committed to the structuralist notion of the morpheme — it allows a given morphological exponent to realise multiple grammatical features, or several exponents to realise a single feature — the basic assumption that full words are built out of smaller parts remains intact. Likewise, Optimal Interleaving (Wolf 2008) tasks an OT grammar with evaluating word forms by evaluating the derivational chains which produced them, where only one change (such as morph insertion) may occur in each link of the chain. Finally, any model accepting the basic premises of Lexical Phonology and Morphology will view morphology as fundamentally constructive in nature, since this is what drives its key theoretical achievements: opacity and cyclicity.

Abstractive analyses instead analyse the form a given word takes by observing it in relation to other words with which it is in paradigmatic opposition. The crucial device here is the exemplary paradigm, which represents the actual forms that lexical items take while at the same time providing a model for the creation of new forms. Such analyses therefore "do not impose a radical separation between 'data' and 'patterns', but represent the morphological patterns of a language by actual forms that display those patterns" (Blevins 2006, p. 537). The fully inflected word here is primary, and whatever roots, stems or exponents might be recognised in the course of comparing full word forms to each other are still useful as objects of analysis, but they have no independent existence outside of the full words that instantiate them:

> 'An abstractive model is fully compatible with the practice of segmenting words into component morphs, as in [Item and Arrangement] models, provided that these morphs are regarded as abstractions over forms, not as the 'building blocks' from which the forms are constructed.' (Blevins 2006, p. 536)

Rather than expressing some meaning independently, parts of individual words can be more or less *diagnostic* of a word's position in a paradigm:

> 'While the meaning of a word form is not necessarily construed as a straightforward composition of individually meaningful parts, the meaning of the whole is associated with reliable arrangements of its constitutive elements.' (Ackerman & Malouf 2013, p. 430)

Where inflection classes exist, these word parts may also be more or less diagnostic of which class's paradigm a word belongs to. Morphology, on this view, thus emerges as the study of how full words are related to one another in paradigms, and how word parts contribute to identifying both where that word is in its paradigm, and which class of paradigm it belongs to.

Abstractive models have been profitably applied to analysing cases of empty morphs in inflectional paradigms, where they have been shown to support the discriminability of paradigm cells despite not being reliably associated to any particular 'meaning'. In the next section, I review some of these analyses, and show that the same thinking can be applied to Wubuy *ŋu-*, so long as we approach phonological alternations as also instantiating a kind of paradigm between full surface word forms, much like inflectional morphology.

3.4 An abstractive analysis of *ŋu-*

3.4.1 Empty morphs and the Paradigm Cell Filling Problem

We have seen that the common thread running through the many different theories discussed in the previous section is that they assume constructive morphotactics: words are assumed to be built up piece by piece, adding invariant abstract morphemes at each layer. Abstractive approaches, on the other hand, take full words as primary, and examine the relations holding between them.

This abstractive outlook leads one to ask a different set of questions than is typically pursued in constructive approaches. Instead of asking how speakers know which of the available morphs to select when assembling a word, we can ask instead how speakers use their existing knowledge of words to predict the form of a word they do not know. This leads to what Ackerman, Blevins & Malouf (2009) call the Paradigm Cell Filling Problem (PCFP), which they formulate thus:

(112) **The Paradigm Cell Filling Problem (PCFP)**
'What licenses reliable inferences about the inflected (and derived) surface forms of a lexical item?'

The idea here is that, for languages of even a moderate degree of morphological complexity, it is likely that a speaker can only be exposed to a subset of all the possible surface forms of the paradigm for a given lexeme; this matters not only during the early stages of acquisition but throughout a lifetime of language use.[22]

[22] Bonami & Beniamine (2016, p. 160) provide concrete evidence for this assumption of limited exposure: in a French corpus of 1.6 billion word tokens, which is roughly comparable to the number of tokens an adult speaker is likely to have been exposed to, 'the average number of forms per lexemes attested in the corpus is well below the maximal number of such forms'.

Moreover, as Wilmoth & Mansfield (2021, p. 356) put it, "inflectional systems often exhibit 'unnecessary' complexity: the set of possible formatives may outnumber the set of grammatical features which can be expressed". This occurs whenever there are different allomorphs expressing the same lexical or grammatical meaning; the allomorphy may be phonologically conditioned, specific to an individual lexeme, or the lexicon may be organised into a number of inflectional classes. This complexity is 'unnecessary' in the sense that it appears to require more effort to learn and use, but without any obvious communicative benefit (Wilmoth & Mansfield 2021). Nonetheless, even in inflectional systems that show a great deal of this complexity, the complexity tends to be organised in such a way that unknown word forms are usually predictable on the basis of known forms (Ackerman & Malouf 2013, 2016; Bonami & Beniamine 2016; Lévêque & Pellard 2023; Parker & Sims 2020; Wilmoth & Mansfield 2021). We can then ask "how paradigms achieve this organisation, and what information is leveraged by speakers to facilitate prediction of unknown wordforms" (Wilmoth & Mansfield 2021). For the matter at hand, we can ask how empty morphs may serve as one source of this paradigmatic information.

Aronoff's (1994) classic analysis of theme vowels in Latin, the parade example of empty morphs, anticipates this line of reasoning. The role of theme vowels is not to contribute morphosyntactic features or semantic content — indeed, they have none — but to signal the stem's membership in one of several morphological inflection classes, which determine the set of inflectional suffixal allomorphs and stem alternations appropriate to that stem. For instance, encountering a theme vowel *ā* in an inflected word form identifies this word as belonging to the First Conjugation; once the inflection class has been identified, one can deduce all other inflected forms of that verb. Thus, 'knowing the theme vowel of a given verb will allow one to determine uniquely (and with no exceptions of any kind in the language) the entire paradigm of the forms of that verb that are formed on the present stem'. In other words, theme vowels solve the PCFP in Latin.

An even more interesting case of empty morphs indicating class membership can be found in Pitjantjatjara, a Pama-Nyungan language of central Australia (Wilmoth & Mansfield 2021). In Pitjantjatjara, verbs are obligatorily marked for seven TAM categories and two types of nominalisation. Two of the TAM categories are expressed by suffixation to the verb stem, while the other TAM inflections and nominalisations require a stem augment to be added before the suffix. There are different, suppletive allomorphs for both the suffixes and the stem augments; the choice of which allomorphs a given stem takes is determined by membership in one of four inflection classes. Examples of forms inflected for past and present tenses are shown in (113).

(113) Pitjantjatjara verb inflection (Wilmoth & Mansfield 2021)

a. Past tense:

	Ø-class	(pirci)-**ŋu**	wriggle-PST
	l-class	(paca)-ɳu	bite-PST
	ng-class	(citi)(ti-**ŋu**)	shiver-PST
	n-class	(u:)(li-nu)	irritate-PST
b.	Present tense:		
	Ø-class	(pirci)-**ɲi**	wriggle-PRS
	l-class	(paca)-ɳi	bite-PRS
	ng-class	(citi)(ti-ŋa)-**ɲi**	shiver-AUG.IPFV-PRS
	n-class	(u:)(li-na)-**ɲi**	irritate-AUG.IPFV-PRS

Inflection class membership is partly arbitrary, partly prosodically determined: two of the classes (Ø-class and l-class) only contain stems ending in a complete bimoraic foot, while the other two (ng-class and n-class) only contain stems that do not (in the above examples, footing is indicated with round brackets). In the past tense (113a), the Ø-class and the ng-class use the same suffixal allomorph, *-ŋu*, while the others have allomorphs unique to their class. However, these two classes are nonetheless distinguished in the past tense by the metrical structure of the base: the Ø-class only admits stems ending in a complete foot, while the ng-class only contains stems that do not. Therefore a single verb form inflected for the past tense is sufficient to identify the class it belongs to, despite the suffix allomorph on its own providing ambiguous information.

For the present tense (113b), three of the classes share the same suffix allomorph, *-ɲi*. The Ø-class can already be identified, as before, by the footing of the stem. But stem footing cannot distinguish the remaining classes — ng-class and n-class — from one another, as both of these apply to stems ending in an incomplete foot. However, for just these classes, stem augments are required before the present tense suffix, and each of the classes requires a different augment allomorph, keeping the inflectional form appropriate to each class distinct as a result. In other words, the stem augments distinguish classes where the suffixes do not, and the suffixes distinguish classes where the augments do not. Pitjantjatjara therefore illustrates a striking solution to the PCFP, where the "combination of suffix allomorphy, prosodically conditioned stem augmentation, and the prosodic structure of verbal roots all conspire to achieve a paradigm that is totally interpredictable: hearing one inflected verb enables a speaker to produce with certainty any other form of that verb" (Wilmoth & Mansfield 2021, p. 355). This insight is only possible by considering the complete inflected word form as the primary unit of analysis.

Of additional relevance to the present chapter is that the stem augments in Pitjantjatjara appear to be phonologically optimising: they could be construed as repairing prosodically defective bases by supplying an additional syllable, thereby creating a complete foot at the right edge of the base. For example, the final syllable in the root *cititi* 'shiver' cannot be assigned to a

foot, since feet are minimally bimoraic; in the present tense inflection, the stem augment *-na-* supplies an additional mora to meet that minimum, and thus ensures that the present tense suffix attaches to a stem ending in a complete foot.

This cannot be explained as phonological optimisation, however, for at least two reasons. For one, the choice of which stem augment allomorph to use (e.g. *-ŋa-* vs *-na-* in the present tense) is determined by inflection class membership — which is still lexically arbitrary — and not by considerations of phonological markedness. Further, only some suffixes require a fully footed base, despite others having comparable prosodic properties: all the suffixes in (113) are open monosyllables, but only those in (113b) require attachment to a complete foot. It is instead specifically where the suffix allomorphs fail to distinguish the classes within a particular inflectional category that the metrical structure becomes subject to 'repair'. Indeed, Wilmoth & Mansfield (2021, p. 365) point to synchronic deletion processes and work on historical-comparative reconstruction to show that these stem augments are the relics of historical TAM suffixes once found throughout the entire verbal inflection paradigm. These historical suffixes underwent deletion when not assigned to a foot — but, crucially, deletion only applied in those parts of the paradigm where it would not introduce uncertainty in predicting one inflected form from another. It is therefore the phonological conditioning which is in service to the morphology here, and not the other way around.

Both Latin and Pitjantjatjara therefore make use of empty morphs — theme vowels in the former, stem augments in the latter — to solve the PCFP, by enabling the identification of the morphological class to which an inflected verb belongs where the affixal allomorphy alone would be insufficient. In the next section, I apply this line of reasoning to opaque phonological alternations.

3.4.2 Phonological classes

In this section, I argue that Wubuy *ŋu-*, just like the empty morphs in Latin and Pitjantjatjara examined above, provides the speaker/hearer with crucial information for predicting the related forms of a given inflected word. In the Wubuy case, however, instead of morphological inflection *per se*, the empty morph provides information about the phonological behaviour that a root should obey in a related word form. I will begin with prefixal *ŋu-*, since the analysis there will extend to linking *ŋu-*.

Recall that our examples of 'underlyingly' stop-initial and continuant-initial roots in (48–61) show surface behaviour that varies along two dimensions in tandem. Underlying stops are realised as stops in all environments, while underlying continuants are realised as stops only after [–CONT] segments, and as continuants everywhere else. At the same time, stop-initial and continuant-initial morphs require different pronominal prefix allomorphs: the pronominal

prefixes that surface as /N/-final before underlying continuants (which then must surface as stops), are instead realised as /ŋu/-final before underlying stops (the assimilation of the prefix-final nasal will be addressed later). So, rather than separating this /ŋu/ from the segments before it with a hyphen, I'll treat it here as a portmanteau.

I will regard these two categories of roots — what previous description has analysed as underlyingly stop-initial and underlyingly continuant-initial — as belonging to different phonological classes, which I will term the *stop class* and the *alternating class*, respectively. I will develop a fuller account of this notion of a 'phonological class' in the pages to follow; for now, it can be taken to mean something like a morphological inflection class organised on the basis of phonological features as well as any phonological alternations that apply from one word form to another related form. These phonological classes are entirely independent of the verbal inflection classes, which determine TAM suffixal allomorphy.

First, let us consider some representative examples of inflected word forms from each phonological class, as in (114). The paradigms are far from complete — there are hundreds more vowel-final prefixes, and several more TAM categories — but these forms will be enough to illustrate the present argument. Also, in order to clarify the patterns of interest, which all occur at the left edge of the stem, I have restricted the examples to roots from the A2 morphological inflection class, inflected for the NPST2 TAM category; for that reason, the suffixes are the same in every cell, even though the suffixes have no bearing on the phonology at the left edge of the stem.

In the first three rows of (114), we see that the prefix allomorphs are the same in each column, while the stems differ. In the final three rows, the prefix allomorphs are different, while the stem forms are the same. As a result, each full word form here is completely unique.

(114)

	Alternating class: /juɻ-i:/ 'transport-NPST2'		Stop class: /cuɻ-i:/ 'push-NPST2'	
1sg.R	ŋa-	juɻ-i:	ŋa-	cuɻ-i:
12du.R	na-	juɻ-i:	na-	cuɻ-i:
3pl.R	wuɾu-	juɻ-i:	wuɾu-	cuɻ-i:
1sg.IRR	ŋaɲ-	cuɻ-i:	ŋaŋu-	cuɻ-i:
12du.IRR	naɲ-	cuɻ-i:	naŋu-	cuɻ-i:
2sg.R	nuɲ-	cuɻ-i:	nuŋu-	cuɻ-i:

With unique forms in every cell, one need only encounter a single inflected form of a given lexeme in order to correctly identify the class it belongs to, and thereby predict the other forms in the paradigm — just as we saw for Pitjantjatjara above. The appearance of *ŋu-* in an inflected

word form is diagnostic of membership in the stop class, and indeed, it only appears in those inflections where the manner of the initial consonant fails to provide that information (i.e., in those cells of the paradigm where the stems of both classes are stop-initial). In this, we see the crucial similarity with Pitjantjatjara stem augments, where the identity of the augment is diagnostic of inflection class, and the augment only appears at all where other diagnostic criteria — such as suffix allomorphy or metrical structure — fail to differentiate the classes.

In (114), the roots happen to be minimal pairs. But, importantly, the point being made here is not limited to minimal pairs, where total homophony is at stake. In (115), we see that, once again, knowledge of any one cell is sufficient to identify the class to which that word belongs, and this information is sufficient to predict all inflectional relatives of that word. If one's first exposure to the root *kaɾaca* 'jump' were in the form of the inflected word *ŋaŋu-kaɾac-iː* 'I will jump', one could immediately deduce that it belongs to the stop class. The stem allomorph on its own is insufficient for this, as both classes contain stop-initial stems in at least part of their paradigm. Rather, it is through the combination of stop-initial stem plus /ŋu/-final prefix allomorph that the assignment to the stop class becomes inevitable, and one can predict all of its related forms with total certainty. The role of *ŋu-* here is not to avoid confusion with some nonexistent root **waɾaca*, but to avoid uncertainty in predicting the rest of the *kaɾaca* paradigm.

(115)

	Alternating class: /wulṯ̠-iː/ 'cut-NPST2'		Stop class: /kaɾac-iː/ 'jump-NPST2'	
1sg.R	ŋa-	wulṯ̠-iː	ŋa-	kaɾac-iː
12du.R	na-	wulṯ̠-iː	na-	kaɾac-iː
3pl.R	wuɾu-	wulṯ̠-iː	wuɾu-	kaɾac-iː
1sg.IRR	ŋaŋ-	kulṯ̠-iː	ŋaŋu-	kaɾac-iː
12du.IRR	naŋ-	kulṯ̠-iː	naŋu-	kaɾac-iː
2sg.R	nuŋ-	kulṯ̠-iː	nuŋu-	kaɾac-iː

If *ŋu-* did not appear in the contexts that it does — if each prefix in the final three rows of (114) or (115) had only the single /N/-final allomorph — it would still be possible to predict, say, the 1sg.IRR form of either 'transport' or 'push' once the 1sg.R form was known. For either class, the result would be **ŋaɲcuɻiː*. But predicting in the other direction, from a word with a nasal-final prefix to a form with a vowel-final prefix, would effectively be reduced to a coin toss. Prefixal *ŋu-* reduces the uncertainty in predicting one cell on the basis of another to zero. This is not an entirely abstract question — the three pronominal categories in the lower half of (115) are among the most frequent in conversation, so it is not hard to imagine speakers encountering a previously unheard verb in this way.

There remains an important limitation on the otherwise total interpredictability between paradigm cells, however. Recall that /w/ alternates with either /p/ and /k/, depending on the lexeme. A word with a nasal-final prefix allomorph and either a /k/ or /p/-initial stem allomorph can be used to predict all the other cells of its paradigm, with total certainty: the continuant-initial stem allomorph in either case will be /w/-initial. But a word with a vowel-final prefix and a /w/-initial stem is, on its own, ambiguous in the predictions it makes regarding the other words in its paradigm. That is, without knowing any other members of the paradigm, *ŋaŋ-kulṯ-iː* predicts *ŋa-wulṯ-iː*, whereas *ŋa-wulṯ-iː* on its own can be taken to imply either *ŋaŋ-kulṯ-iː* or *ŋam-pulṯ-iː*. As it happens, this particular root has two stop-initial variants (though the /p/-initial variant is less common), and there are a handful of other roots like this. But in general, a given /w/-initial stem allomorph must be lexically specified to alternate with just one stop-initial allomorph. Nonetheless, without *ŋu-* there to distinguish the alternating and stop classes, the predictive ambiguity in /w~k/ or /w~p/ stems would run in both directions: from the post-nasal context to the post-vowel context, and vice versa. Thus, *ŋu-* substantially reduces paradigmatic uncertainty, even if it only eliminates it completely for stem-initial stops other than /k/ and /p/.

Identifying the class to which an inflected word belongs thus requires two pieces of information: the segments at the right edge of the pronominal prefix, and the continuancy of the stem-initial consonant, as spelled out in (116). Each of these two pieces of information, on its own, is necessary but not sufficient. Through the combination of the two, however, an inflected word form can be uniquely identified as a member of one or the other class, and no other phonological or morphological factors elsewhere in the word are relevant.

(116)

Vowel-final prefix + continuant-initial stem	⊃	*Alternating class*
Nasal-final prefix + stop-initial stem	⊃	*Alternating class*
Vowel-final prefix + stop-initial stem	⊃	*Stop class*
ŋu-final prefix + stop-initial stem	⊃	*Stop class*

These phonological classes run entirely orthogonal to the (morphological or lexical) classes that govern the set of TAM suffix allomorphs a verb stem must take. While (114) and (115) only contained NPST2 stems from the A2 inflection class — and therefore ended in the suffix *-iː* — in truth, stems in both the alternating and stop phonological classes can be of any inflection class, and can be inflected for any combination of realis/irrealis (encoded in the prefix) and tense-aspect category (encoded by the suffix), so long as the combination is morphosyntactically grammatical. Indeed, the generalisation in (116) is just as applicable to adjectival predicates, which take no tense-aspect suffix, or to N+V and N+Adj compounds; in every case, it is only the phonology at the left edge of the stem that determines phonological classhood. My adoption of the term 'phonological class', then, arises partly as a result of the need to distinguish the

phonologically-determined prefix and stem allomorphy from the lexically-conditioned suffix allomorphy (e.g., TAM suffix inflection classes A1, A2, A3, I1, I2, and so on).

By considering what is held in common and what varies across the inflected word forms of (114) and (115), as well as the other examples of prefixal *ŋu-* shown throughout this chapter, we can begin to generalise beyond these particular examples. A more thorough formalisation will be developed over the remaining parts of this chapter; for now, we can suggest a preliminary schematisation as in (117). The schema can be seen as a general statement about the phonological properties of words in the alternating class and stop class. Words in the alternating class have a continuant-initial stem when they follow a vowel-final prefix, and a stop-initial stem when they follow a nasal-final prefix. Words in the stop class have a stop-initial stem following a vowel-final prefix, and they also have a stop-initial stem following a /ŋu/-final prefix.

(117)

Alternating class		*Stop class*	
Prefix	Stem	Prefix	Stem
[...V-	[[−SYLL, +CONT, +SON] ...]]	[...V-	[[−SYLL, −CONT, −SON] ...]]
[ŋaju.ɻi:] 'I am transporting it'		[ŋacu.ɻi:] 'I am pushing it'	
[ŋawulʈ̪i:] 'I am cutting it'		[ŋakaɾaci:] 'I am jumping'	
[...N-	[[−SYLL, −CONT, −SON] ...]]	[...ŋu-	[[−SYLL, −CONT, −SON] ...]]
[ŋaɲcu.ɻi:] 'I will transport it'		[ŋaŋucu.ɻi:] 'I will push it'	
[ŋaŋkulʈ̪i:] 'I will cut it'		[ŋaŋukaɾaci:] 'I will jump'	

In other words, while /ŋu-/ cannot be straightforwardly associated with any semantic meaning or morphosyntactic category realisation, it is nonetheless *informative* in the context of paradigm structure, since it makes the cells of the paradigm almost totally interpredictable — what Finkel & Stump (2009) describe as a state of 'maximal paradigmatic transparency'. Indeed, this reveals the whole notion of the 'empty morph' to be the peculiar preserve of morphotactically constructive models of morphology, since they assume that each identifiable part must make an independent contribution to the meaning of the whole word, rather than making a contribution to identifying the word's place in a system of oppositions — i.e., a paradigm.

Contained in the perspective being advanced here is Heath's intuitive point that *ŋu-* allows the language to maintain contrasts in the lexicon such as *cu.ɻa* 'push' vs *ju.ɻa* 'carry, transport', which would otherwise be neutralised by Hardening: what is contrasted by the [±CONT] feature in the stem in one part of the paradigm is instead contrasted by the form of the prefix before the

stem in another part of the paradigm. But the emphasis has shifted in a small but significant way in our conception of contrast. For one, it is no longer necessary to show that *ŋu-* averts strict homophony. Absolute homophony is seldom a realistic threat — both because other factors like (suffixal) inflection class and discursive context are usually available to distinguish lexical minimal pairs, and because it is clear that some amount of homophony is tolerable both in Wubuy in particular and in languages in general. Heath's insight that ŋu-Epenthesis serves to support contrasts in the lexicon threatened by neutralisation is fundamentally an insight into the paradigmatic nature of phonological alternations. Namely, /ŋu-/ allows the speaker/listener to solve the phonological dimension of the PCFP, which is useful for any stop-initial or continuant consonant-initial stem, not just for those few such stems which happen to form minimal pairs.

3.4.3 Uncertainty reduction and linking *ŋu-*

How can we apply the same reasoning to linking *ŋu-*? Unlike prefixal *ŋu-*, it is not sensitive to the phonology of the left-hand context. It applies before underlyingly stop-initial morphs — that is, those of the stop class — when they appear following a derivational element of any kind. We saw in §3.2.6 that its distribution seems to overdetermine the stop/continuant contrast: it appears not only where the distinction between continuant and stop is threatened with neutralisation (following non-continuants), but also where it is distinguished on the surface. In (118), for example, why does it appear following the vowel-final incorporated noun *jaɾka* 'flipper', when it would still keep the classes entirely distinct if it appeared only after a stop or nasal, as when following *tan* 'guts'?

(118)

After underlying	*Alternating class*	*Stop class*
/[+CONT]/	…-jaɾk**a-j**aɾt-aŋ	…-jaɾk**a-ŋu-c**ulup-aŋ
	…-flipper-touch-NPST1	…-flipper-conceal-NPST1
/[−CONT]/	…ta**n-c**aɾt-aŋ	…-ta**n-ŋu-c**ulup-aŋ
	…-guts-touch-NPST1	…-guts-conceal-NPST1

I would suggest that this apparently excessive distribution of *ŋu-* following even continuant-final derivational affixes is fundamentally related to a species of opacity that frequently arises at this morphological juncture. Two of the most frequent derivational elements are the APPLicative prefix *aK-* and MULTiple *ŋaraK-*, both of which, under Heath's analysis, end in the archiphoneme /K/; likewise, several of the most frequent incorporable nouns end in /k/, which in this context functions identically to /K/ most of the time. This /K/ is a highly abstract segment: it is always deleted by a relatively late rule (Stop Deletion), but is detectable on the surface by its effects,

the relevant one here being Hardening (see Chapter 5 for an alternative view on /K/, based on phonological subcategorisation). On the surface, Hardening applies to continuant-initial roots, yielding a stop, but the motivation for the stop alternant is not present, due to deletion of the /K/. We therefore need to flesh out our table from (118) with stems illustrating these surface effects.

(119)

After underlying	*Alternating class*	*Stop class*
/[+CONT]/	[+CONT]-[+CONT]	[+CONT]-ŋu-[−CONT]
	…-jaɾk**a-j**aɾt-aŋ	…-jaɾk**a-ŋu-c**ulup-aŋ
	…-flipper-touch-NPST1	…-flipper-conceal-NPST1
/[−CONT]/	[−CONT]-[−CONT]	[−CONT]-ŋu-[−CONT]
	…ta**n-c**aɾt-aŋ	…-ta**n-ŋu-c**ulup-aŋ
	…-guts-touch-NPST1	…-guts-conceal-NPST1
/K,k/	[+CONT]-[−CONT]	[+CONT]-ŋu-[−CONT]
	…ŋaɾ**a-c**aɾt-aŋ	…-ŋaɾ**a-ŋu-c**ulup-aŋ
	…-MULT-touch-NPST1	…-MULT-conceal-NPST1

Again, imagine there were no *ŋu-*. Knowing one of the stems in the top left cell, such as *jaɾka-jaɾt-aŋ*, would be enough to identify the root, *jaɾt-aŋ*, as belonging to the alternating class, since it is the only one that can ever be continuant-initial. On the other hand, knowing a stem whose derivational element had a final continuant, and whose root had an initial stop — such as *ŋaɾa-caɾt-aŋ* — would present some difficulty. Assuming one did not already know other words with this derivational element, what could one deduce from the fact that the initial segment in the verb root is a stop? This root-initial stop may plausibly be either an invariant (aka underlying) stop, or one derived by the opaque workings of Hardening following velar stops, as is the case in the stems of the bottom row of (119). That is, surface strings of a continuant followed by a stop could be taken as evidence of membership in the stop class (as in the top right cell), or the alternating class (as in the bottom left cell). But owing to the appearance of linking *ŋu-* throughout the paradigm, this ambiguity does not arise.

In light of the foregoing, we can suggest a schematic representation of the linking *ŋu-* phonological paradigm based on (117), though with an additional row to account for the opaque situation with erstwhile underlying velar stops. Again, we see that the classes are distinct in each row. By removing *ŋu-*, any of the cells in the right column become identical with one of the cells on the left (not necessarily in the same row). A single word form is sufficient

to immediately identify its class, and from there to identify the other words to which it is paradigmatically related. So again, *ŋu-* is revealed to reduce paradigmatic uncertainty in a context of widespread opacity. Note also that, although we are dealing with derivational morphology here, the categories that structure the paradigm are still phonological.

(120)

Alternating class			*Stop class*		
Infl	Deriv	Root	Infl	Deriv	Root
[...-	[...[+CONT]-	[[−SYLL, +CONT, +SON] ...]]]	[...-	[...[+CONT]- ŋu-	[[−SYLL, −CONT, −SON] ...]]]
[ampurujaɾt̪ajaɾtaŋ]			[ampurujaɾtaŋucululupaŋ]		
'They will touch the flipper'			'They will hide the flipper'		
[...-	[...[−CONT]-	[[−SYLL, −CONT, −SON] ...]]]	[...-	[...[−CONT]- ŋu-	[[−SYLL, −CONT, −SON] ...]]]
[ampuɾutancaɾtaŋ]			[ampuɾutanŋucululupaŋ]		
'They will touch the guts'			'They will hide the guts'		
[...-	[...[+CONT]-	[[−SYLL, −CONT, −SON] ...]]]	[...-	[...[+CONT]- ŋu-	[[−SYLL, −CONT, −SON] ...]]]
[ampuɾuŋaracaɾtaŋ]			[ampuɾuŋaraŋucululupaŋ]		
'They will all touch it'			'They will all hide it'		

As a declarative statement, the schemas in (117) and (120) are valid generalisations. But crucial information is missing before they can encode how one *predicts* one form based on another. For the prefixal *ŋu-* schema in (117), we need to be able to express, for instance, not only that a word with a vowel-final prefix and a continuant-initial stem must be in the alternating class, but that such a word implies the existence (or possibility) of a related word whose pronominal prefix ends in a nasal, and whose stem has an initial stop, rather than a continuant, at the same place of articulation as the continuant; the rest of the stem, meanwhile, remains the same across realisations. We also need to express the fact that there is a predictable relationship between the prefix allomorphs — that which is nasal-final before a stem in the alternating class is /ŋu/-final in the stop class. The place of articulation of the prefix-final nasal is also predictable — it assimilates with the following stop — and we need a way to express this generalisation. And, further, we need to express in some way the relation between these two schemas: how can we use the knowledge of a root's phonological class as per (117) to predict its class when appearing inside a derived stem as per (120), or vice versa?

These issues are taken up in detail over the next two sections, which discuss the place of phonological generalisations in schema representations, a topic that has been seriously neglected in abstractive models of morphology. I introduce the Relational Morphology formalism (Jackendoff & Audring 2020) in §3.5, showing how it captures phonological generalisations, developing a novel account of opacity in the process. I then apply the formalism to our *ŋu-* problem in §3.6.

A final point. We have so far been considering whole words as static units, and have sought to bring out some of the contrastive value of sub-parts of these words by examining how they function as nodes in a static web of relatedness. But speech unfolds in time, so we could also ask how word-parts reduce uncertainty about the identity of a word progressively over the course of a word's production, as has been extensively studied in the psycholinguistic literature on lexical access (e.g. Baayen, Chuang et al. 2019; Norris 1994; Norris & McQueen 2008; van Son & Pols 2003). Indeed, one finding is that, all else being equal, segments earlier in a word contribute more to disambiguation from other lexical items (Cohen Priva 2015; van Son & Pols 2003). This greater informativity has been used to explain why contrasts at the left word edge are more likely to be retained than ones at the right (Hall et al. 2016, pp. 40–41). One has to wonder about the fact that this entropy-reducing element *ŋu-* is inserted immediately *before* the segment which would otherwise neutralise, and that the neutralising segment is at the left edge of the stem. This, we can assume, is a position of high cognitive salience, and we should expect it to be more decisive during the process of lexical retrieval than segments later in the word. Although at this stage it can only be fairly speculative, I presume that these two facts are related. It seems reasonable to suppose that *ŋu-* is more useful to listeners before the stem than after it, since as a prefix it effectively prepares the listener by signalling that the following stop is to be interpreted as a 'true' stop rather than one that alternates with a continuant realisation, presumably making the process of lexical retrieval less error-prone or effortful (for studies on related phenomena see Baker & Bundgaard-Nielsen 2018; Milin et al. 2009; Ussishkin et al. 2017). I leave this issue to future research.

3.5 Relational Morphology

3.5.1 Introduction

In the previous section, I showed that our Wubuy empty morph *ŋu-* could profitably be viewed in terms of reducing the uncertainty associated with predicting one inflected word form on the basis of another — that is, as drastically reducing the difficulty of solving the Paradigm Cell Filling Problem — in a context where lexical contrasts are threatened by phonological neutralisation.

To date, however, those advocating abstractive morphology have had little to say about the place of phonology in their models, being mostly concerned with lexically-determined inflection classes. It is not immediately clear, then, how to represent paradigms that are structured by genuinely *phonological* generalisations, as Round (2011) points out. What we do know is that by taking full words and their organisation into paradigms as basic, rather than building words atom by atom as in morphemic theories, we strike at one of the most fundamental assumptions of mainstream phonological theory. By discarding morphemes, we also discard the ontological distinction between underlying representations and surface forms which has, by and large, served as the core business of phonology as commonly understood. This is recognised by Booij, who states:

> Since Construction Morphology assumes that regular morphologically complex words can be stored in the mental lexicon, we do not need the concept of 'restructuring of the underlying form of a word' [i.e. Phonological Readjustment —PN]. Instead, we need a theory of how to select the proper phonological form of a base word that serves as a basis for the computation of the phonological form of words derived from that base word (Booij 2010, p. 238).

Part of the purpose of this chapter — and indeed, this thesis — is to develop this latter theory.

In this section, I build upon the Relational Morphology formalism (henceforth RM; Jackendoff & Audring 2020), as introduced in §1.3.3, and show how the approach it takes to morphological paradigms can be extended to encode phonological alternations. RM models alternations as mappings between surface words — it makes no use of underlying representations, nor any intermediate or otherwise hidden levels of phonology. Of particular interest here is that RM is particularly insightful when it comes to opacity — a phenomenon that remained out of reach in predecessors such as Construction Morphology (Booij 2010) — which it handles with the same basic mechanisms that it requires for less exotic purposes. Jackendoff and Audring (2020) do not consider this potential of their theory, but I show here that we can realise that potential without introducing any additional machinery.

3.5.2 Words, schemas, and subcategorisation

Recall from the introductory chapter that, at a very basic level, Relational Morphology proposes, like Construction Morphology (Booij 2010), that words are a kind of 'construction'. Prototypically, a *construction* is a pairing of meaning and form, stored in the lexicon. Patterns among those words are represented as schemas. A schema is also a type of construction, and is also stored in the lexicon. It differs from a word in having some part left underspecified — a variable 'slot' that must be 'filled in' in order to form a new word (or to interpret one, as in perception). Thus

the lexicon, on this view, subsumes what would in typical generative approaches be separated out into the 'grammar', whether the latter is formalised as rules or constraints. A construction is more word-like to the extent that its content is fully filled out, or more rule-like to the extent that it contains variables (Jackendoff & Audring 2020, p. 23).

To make this more concrete, we can begin with some very simple examples, without phonological conditioning, from English. I assume here an adult, competent speaker.[23] Let's suppose this speaker acquires some words — that is to say, experiences them and remembers them — such as *attentive*, *attentiveness*, *fearful*, *fearfulness*. The basic mechanism for establishing patterns over the lexicon in RM, like other approaches with storage of regular complex words (e.g., Albright 2009; Arndt-Lappe & Ernestus 2020; Baayen, Chuang et al. 2019; Blevins 2006; Booij 2014; Bybee 2006; Skousen 2009), is analogical comparison. At its simplest, this can be expressed as a classical four-part analogy, as follows:

(121) Four-part analogy (Albright 2009)

A:B :: X:Y
'Whatever the relationship is between A and B, it should also hold between X and Y'

The four-part analogy tells us that there is something which is the same between A and B — the *attentive* part — and something which is different — the additional *-ness* part in B. It also tells us that the same kind of same-different relation holding between A and B also holds between X and Y — they each have the *fearful* part, but Y additionally has *-ness*. In this example, the analogical relation holding between the forms of the words comes with a corresponding relation in the words' meaning (their morphosyntax and semantics): the first member of each pair is an adjective, while the second is a noun denoting the quality of that adjective. Thus, in our toy lexicon, *attentive* is to *attentiveness* as *fearful* is to *fearfulness*.

These four words and the relations between them can be represented as lexical entries. In the framework of the Parallel Architecture (Jackendoff 2002), which RM adopts, a lexical entry (i.e., a construction) is a set of linked structures in semantics, (morpho)syntax and phonology. We can thus represent our two pairs of words as in (122) below (though I have left the semantic level fairly vague here as it has no bearing on the discussion to follow).

[23] I thus set aside issues of how phonological categories are learned from phonetically variable speech tokens. See Archangeli & Pulleyblank 2017, 2022; Davis & Redford 2019; Frisch 2017; Olejarczuk et al. 2018; Pierrehumbert 2003; Schweitzer 2019 for relevant proposals on acquisition from a whole-word storage perspective. Relatedly, see Hall et al. 2016; Wedel 2006; Winter & Wedel 2016 for proposals on how phonological categories evolve diachronically.

(122) English deadjectival nominalising suffixation

Semantics:	a.	ATTENTIVE$_1$	c.	FEARFUL$_2$
Morphosyntax:		Adj$_1$		Adj$_2$
Phonology:		[əˈtentɪv]$_1$		[ˈfɪəfʊl]$_2$
Semantics:	b.	[QUALITY OF(ATTENTIVE$_1$)]$_4$	d.	[QUALITY OF(FEARFUL$_2$)]$_5$
Morphosyntax:		[$_N$ Adj$_1$]$_4$		[$_N$ Adj$_2$]$_5$
Phonology:		[[əˈtentɪv]$_1$ nəs$_3$]$_4$		[[ˈfɪəfʊl]$_2$ nəs$_3$]$_5$

Within a lexical entry, numerical coindices link pieces of structure on different levels. These are known as *interface links*. In (122a), coindex 1 links the phonological specification of the English word *attentive*, namely [əˈtentɪv], to the Adj category head on the morphosyntactic level and the corresponding concept on the semantic level. But there are also relations of correspondence holding between pieces of structure across lexical entries. So coindex 1 links the string [əˈtentɪv] on the phonological level of (122a) to the equivalent string [əˈtentɪv] on the phonological level of (122b), and it links the Adj head in the morphosyntax of (122a) to the Adj in (122b); the corresponding semantic concepts are linked in similar fashion. Coindex 2 performs the same function across (122c) and (122d), connecting elements of form to elements of meaning. Where coindices relate structures across lexical entries like this, they are called *relational links*.

But what is our motivation for asserting that the [əˈtentɪv] string in (122a) is the same as that in (122b)? Indeed, what is our motivation for asserting that the word in (122b) has internal structure at all? The answer is to be found in the regular form-meaning correspondence between the nouns ending *-ness* and the adjectives without it, as suggested by the four-part analogy. That is, we posit a word-internal piece of structure only if it distinguishes that word from another one. Nonetheless, we do not require any piece of structure so identified to be independently meaningful, outside of the word (or construction) it appears in.

Indeed, the nominalising suffix [nəs] in the above schemas is not held to be independently meaningful: it is relationally linked between the phonological levels of (122b) and (122d) with coindex 3, but this does not link to anything on the morphosyntax or semantics levels.[24] Rather, the change in word class, from adjective to noun, is attributed to the outer layer of constituency: coindicies 4 and 5. In this sense, the suffix is diagnostic of a certain type of construction: a deadjectival noun.

[24] This is in slight contrast to Jackendoff & Audring (2020), who link the phonology of derivational affixes to an element of the morphosyntax they label ‘aff’ (for ‘affix’). But an affix is not really a morphosyntactic category (meaning), rather it is a description of how a meaning is expressed (form). There also does not appear to be any need for other schemas — whether other word-level schemas or higher phrase-level ones — to refer to an ‘Aff’ category. As far as I can tell, it makes no difference to omit it altogether, as I do here.

Of course, we must be able to generalise from our concrete examples *attentiveness* and *fearfulness* to any noun consisting of an adjective root followed by *-ness*. In RM, as in other construction-based approaches, generalisations of this kind are expressed via schemas. Schemas are established by submitting fully-specified lexical entries — i.e., words — to a procedure Jackendoff & Audring (2020) term the Structural Intersection (SI):

> "This operation can be thought of as like Boolean intersection, except that it is defined over structures rather than sets. The input to SI is two or more structures; its output is a structure that preserves the parts of the input structures that are the same, and which leaves unspecified (or variable) the parts where the input structures do not correspond. For a simple case, the Structural Intersection of ABCD and ABCE would be ABCx." (Jackendoff & Audring 2020, p. 224)

The Structural Intersection of *attentiveness* and *fearfulness* in (122b,d) thus reveals the following structures in common: the semantic function [QUALITY OF(X)], the morphosyntax [$_{N}$ Adj], the phonology [nəs], and the arrangement of the interface links between them. These are left intact in the schema. What differs is the semantic and phonological structures that are linked to the Adj head in the morphosyntax; these are replaced with variables. And, although these structures are coindexed in the same way in (122b) and (122d), the numbers are item-specific — 1 in the former and 2 in the latter. These must be replaced by (alphabetic) *variable interface links*, which I explain further below. The structural intersection thus yields (123), a general schema for deadjectival nouns ending in *-ness*.

(123) Simple morphological generalisation: English deadjectival noun schema

Semantics:	$[\text{QUALITY OF}(X_x)]_y$
Morphosyntax:	$[_N\ \text{Adj}_x]_y$
Phonology:	$[[\ldots]_x\ \text{nəs}_3]_y$

As can be seen, schemas have the same format as the lexical entries of full surface word forms, but with one crucial difference: at least one piece of structure is not fully specified. This underspecified 'slot' can be filled in by — can unify with — a piece of structure from another schema. But it cannot be filled by simply *any* piece of structure. Just as the roots in the phonology of (122b,d) were coindexed with interface links to an Adj head in the morphosyntax, the variable slot in the phonology of (123) bears a *variable* interface link, *x*, which is also linked to an Adj head in the morphosyntax. This imposes a condition on what kind of material can fill the ellipses in the phonology — that is, it must be a string that is associated with an Adj. Otherwise, whatever unifies with this slot can be of any phonological shape, since there are

no further specifications on the phonological level (and note that there is nothing in common phonologically between the roots of *attentive* and *fearful*, nor any other adjective that can appear here). The semantics associated with coindex *x* is an empty variable, just like the phonology, so the schema does not impose any restrictions on the semantic properties of the Adj that it unifies with.

The schema in (123) can therefore unify with the adjectives (122a) or (122c) to form their nominalised counterparts in (122b) and (122d), respectively. But it can also unify with any new adjective that enters the lexicon, unifying with, say, *dizzy* to form *dizziness*, or *shallow* to form *shallowness*, and so on. Schemas can thus function either declaratively — they describe latent patterns in the lexicon — or generatively — they provide a model for creating new words on the basis of known ones (in production), or interpreting new words that fit the pattern (in perception). Links are, moreover, nondirectional: a noun ending in [nəs] can be used to derive (predict the form and meaning of) a corresponding adjective just as much as the other way around.

Note that anything that is specified within a variable functions as a subcategorisation requirement on what kind of structure can 'fill in' that slot, and therefore imposes a limitation on what the schema can unify with, and this is a direct consequence of the Structural Intersection operation. In (123), the phonology and semantics of the variable slot is left entirely unspecified (as it changes between the entries compared by SI), while the morphosyntax specifies some explicit structure (as the linkage to the Adj head does not change between entries). In this case, the subcategorisation is morphosyntactic: *-ness* can only attach to adjectives. But it is equally possible for subcategorisation requirements to be imposed at any level of representation, or even multiple levels. For instance, plurals in English are only possible for countable nouns (i.e. individuals). Expressing this fact requires a schema that specifies structure at both the morphosyntactic level (it can unify with a noun) and the semantic level (it must be an individual), but that leaves the phonological slot open.

Of course, a schema's subcategorisation may also apply at the level of phonology, which brings us to the heart of the issues of this chapter. Examples of phonologically-conditioned morphology are by no means hard to find. The English comparative adjective suffix *-er* (simplifying somewhat) only occurs on adjectives of either one syllable, or two syllables if the second is weak; otherwise a syntactic paraphrase is necessary. In German, past participles feature a prefix /gə-/ if the stem takes initial stress, and no prefix otherwise (though they also feature a suffix and/or stem vowel ablaut). In Armenian, the definite suffix has two allomorphs, /-n/ and /-ə/. The former appears following vowel-final stems, while the latter appears on consonant-final stems; some examples are shown in (124).

(124) Definite suffix allomorphy in Armenian (Vaux 1998, cited in Inkelas 2014)

	gloss	plain	definite
a.	'tongue'	lezu	lezu-n
	'cat'	gadu	gadu-n
	'wine'	kini	kini-n
b.	'chair'	atorr	atorr-ə
	'book'	kirk	kirk-ə
	'piece'	hat	hat-ə

Just as we saw for (123), we produce schemas for each allomorph by considering what changes and what stays constant between each full word: all the Definite words in (124a) feature a noun ending in a vowel, followed by [n]; all those in (124b) feature a noun ending in a consonant, and are followed by [ə]. This is illustrated below, where the variable slot for the base noun in the phonology imposes a subcategorisation requirement on what it will unify with, namely that it must be something ending in a vowel in (125a), and something ending in a consonant in (125b). The part linking to the definite meaning — the suffixes in each schema — are notated with constants.

(125) Armenian definite suffix allomorph schemas

	a.	b.
Semantics:	$[\text{DEF}(X_x)]_y$	$[\text{DEF}(X_x)]_y$
Morphosyntax:	$\{N_x, \text{DEF}_6\}_y$	$\{N_x, \text{DEF}_7\}_y$
Phonology:	$[[\ldots[+\text{SYLL}]]_x\ n_6]_y$	$[[\ldots[-\text{SYLL}]]_x\ ə_7]_y$

Formally, a subcategorisation requirement is a constant — an explicit piece of structure — inside what is otherwise an open variable. In (125), the constant within the phonological variable notated with ellipses is the [+/−SYLL] feature at the right edge of the root. This emerges directly from the comparison of the actual word forms in (124) — looking at the range of bases that occur with the /-ə/ suffix, we can see that they are completely heterogeneous *except for* the fact that they all end in a consonant.[25] In that way, both the explicit subcategorisation requirement — that the final segment of the part linked to N in the morphosyntax is a consonant — and the suffix itself — that it has the form /-ə/ and is linked to the definite feature in the

[25] There is another generalisation, namely that all the plain nouns in (125a) have a CVCV structure, but this is a tiny data set. Nonetheless, one of the predictions of RM is that any property held in common among a set of words that take the same allomorph can be seized on by speakers as the basis of formulating a generalisation that adequately describes the allomorphic pattern, and that can then be extended to new items. This therefore allows for etymologically inaccurate analyses and spurious back-formations.

morphosyntax — are objects of the same type, namely constants. Constants are simply that which is held in common between analogical sets. With the constants so identified, the residue — whatever changes between the compared items — is expressed as a variable. The notion of a variable is therefore intrinsically connected to underspecification — that which changes cannot be specified. (Note that the role of underspecification presented here differs from those theories that propose underspecification in order to remove all redundancy from the grammar.)

However, there is a range of situations in which it is necessary to relate the variables of one schema to the variables of another schema. This comes up especially in situations where two morphologically complex words share a base, but that base does not exist as a free form. This is commonplace in inflection, where it is not unusual to find some overt inflectional material in every cell of a paradigm. To take a very simple example, in Spanish, the largest class of adjectives must be inflected for gender (and number, which I set aside here). Some examples of singular masculine and feminine forms are shown below.

(126) Spanish singular adjectives

M.Sg	F.Sg	
feo	*fea*	'ugly'
lindo	*linda*	'pretty'
amarillo	*amarilla*	'yellow'

The obvious generalisation to make here is that the masculine forms all end in *-o*, while the feminine forms end in *-a*. From one pair to another, there is a consistent relationship between an aspect of the meaning — the gender agreement — and the form — the final vowel. Meanwhile, the forms meaning 'ugly' all contain the sequence *fe-*, the forms meaning 'pretty' contain *lind-*, and the forms meaning 'yellow' contain *amarill-*. The part preceding this final vowel remains constant between the members of each pair but changes from one pair to another. Let us suppose a masculine-feminine pair of lexical entries like those in (127).[26]

(127) Lexical entries for some Spanish adjectives

Semantics:	a. PRETTY_8	b. PRETTY_8
Morphosyntax:	$\{\text{Adj}_8, \text{MASC}_9\}_{10}$	$\{\text{Adj}_8, \text{FEM}_{11}\}_{12}$
Phonology:	$[\text{lind}_8\ \text{o}_9]_{10}$	$[\text{lind}_8\ \text{a}_{11}]_{12}$

[26] Note that, following Jackendoff & Audring (2020, p. 137), I have not included any explicit reference to gender on the semantic level, as this is a grammatical gender system — that is, a matter for morphosyntax only. The issue is not pertinent to the discussion here, however.

In each, an interface link relates the final vowel to the gender feature it realises in the morphosyntax, and another interface link coindexes the stem in the phonology to the Adj head in the morphosyntax and the PRETTY concept in the semantics. But the interface links for the stem also serve as relational links between the two entries: the *lind-* string in *lindo* is claimed to be the same as the *lind-* string in *linda*. In order to generalise, it would seem straightforward enough to set up a schema for the masculine singular with a phonology like [… o], where the *-o* would be linked to the MASC feature in the morphosyntax, and the ellipses would link to the Adj head with a variable interface link, much as we saw in the English *-ness* schema above. But unlike in the English examples in (122) or the Armenian examples in (124), the Spanish words in (127) are obviously related to one another without either being the morphological base of the other: there is no free form **lind*. We need a way of expressing the fact that the existence of a masculine singular form ending in *-o* predicts the existence of a feminine singular form ending in *-a*, where the remaining part (the stem) is held constant between forms. In the context of the formalism presented here, we require a way for variables to refer to each other in the same way that relational links can establish correspondences between constants across different lexical entries.

In RM, this is achieved by *variable relational links*, notated with Greek coindices. In the sister schemas below, coindex α effectively states that whenever one finds some phonological sequence before *-o* in the masculine singular form, one should also expect to find that same phonological sequence before *-a* in the feminine singular form. Since the links are nondirectional, the feminine form can just as easily predict the masculine.

(128) Spanish singular adjective agreement sister schemas

Semantics:	a.	X_{α}	b.	X_{α}
Morphosyntax:		$\{Adj_{\alpha}, MASC_{9}\}_{x}$		$\{Adj_{\alpha}, FEM_{11}\}_{x}$
Phonology:		$[\ldots_{\alpha}\ o_{9}]_{x}$		$[\ldots_{\alpha}\ a_{11}]_{x}$

Variable relational links essentially state that a piece of structure can only instantiate a variable slot in a schema if it can also instantiate variables marked with the same coindex in any other schema. In this way, variable relational links can model class behaviour. The schema above carries no implication that all feminine singular adjectives must end in *-a*, much less — as in a more derivational framework — that to form a feminine singular adjective, one must take some abstract adjectival stem and attach *-a* to it. Indeed, a masculine singular form not ending in *-o* does not predict a feminine singular ending in *-a* — which is appropriate, given the existence of pairs that do not show this *-o/-a* alternation, like *verde* 'green.M.sg' ~ *verde* 'green.F.sg' and *capaz* 'capable.M.sg' ~ *capaz* 'capable.F.sg', which clearly belong to a different inflection class. If the feminine schema in (128b) instead simply had a variable interface link coindexing the

root slot, the prediction would be that one could fill the variable with an adjective root like *capaz*, on the basis of knowing *capaz* 'capable.M.sg', producing **capaza*. The variable relational links allow us to leverage knowledge of paradigms in constructing schemas, providing us with a means to specify what characteristics about a given word allow one to predict a related word (in either direction). In other words, variable relational links come into play whenever the applicability of one schema depends on the applicability of another.

Recall from Chapter 1 that we can also address non-affixational morphology. For processes like vowel ablaut in English verbs, Jackendoff & Audring (2020) introduce another device in RM, the 'same-except' relation (see also Culicover & Jackendoff 2012), notated by placing asterisks around any part of a constituent that differs in its realisation in different contexts. In this case, the part that differs across realisations is the vowel, on the phonological level.[27] In (10), the phonological forms are indicated to be the same in the present and past forms (via coindex 13), *except for* the vowel, as denoted by the asterisks.

(129) English verbal ablaut in *sing-sang*

Semantics:	a.	$[\text{SING(X)}]_{13}$	b.	$[\text{PAST}([\text{SING(X)}]_{13})]_{14}$
Morphosyntax:		V_{13}		$\{V_{13}, \text{PAST}\}_{14}$
Phonology:		$[s\ *i*\ ŋ]_{13}$		$[s\ *æ*\ ŋ]_{13,14}$

We have seen that the basic unit in RM is the construction, which consists of a set of linked structures on parallel levels — typically phonology, (morpho)syntax and semantics — each with their own conditions on well-formedness. Pieces of structure are linked between levels by *interface links*, and across entries by *relational links*, both of which are notated with numerical coindices. In schemas (as opposed to fully-specified word entries), at least some structure is made into a variable, as are any links to that variable. Interface links thereby become *variable interface links*, notated with alphabetic coindexation, while relational links become *variable relational links*, coindexed with Greek letters. If there is any part that changes between two or more structures that are linked relationally, the change is surrounded by asterisks in each entry. As we will see, these basic properties make an entirely surface-based analysis of phonological alternations possible, including those with opaque interactions.

3.5.3 Phonologically-conditioned allomorphy, transparent and opaque

Variable relational links (Greek subscripts) are able to subsume the functions of variable interface links (Latin subscripts): a Greek subscript marking some piece of structure on each of the

[27] Conversely, leaving the phonological form between two entries the same, but specifying a difference in the morphosyntax with the same-except asterisks, yields 'morphomic' phenomena.

semantic, morphosyntactic and phonological levels will link those structures together (i.e., will function as a variable interface link), while also linking those structures to any structures with the same coindex in another schema. If such is the case, one may well ask whether we need variable interface links at all. After all, variable relational links can perform the same work and more. It is true that the schemas in (123) and (125) could just as easily be modelled as sister schemas, with one schema for the simplex form and one for the complex form, as below. In this formulation, variable relational links relate the free form to the base of the complex form, at the same time that they interface the relevant structures on each of the levels (as also suggested by Booij 2010, p. 254).

(130) English noun and deadjectival nominal sister schemas (alternative)

	a.		b.	
Semantics:		X_{β}		$[\text{QUALITY OF}(X_{\beta})]_{x}$
Morphosyntax:		Adj_{β}		$[_{N}\ \text{Adj}_{\beta}]_{x}$
Phonology:		$[\ldots]_{\beta}$		$[[\ldots]_{\beta}\ \text{nəs}_{3}]_{x}$

Note the difference from the original schema in (123), where we used only a variable interface link for the stem, allowing it to unify with any lexical item headed by Adj in the morphosyntax, such as (122a,c). Instead, (130) depicts paired schemas that specify the expected composition of both the complex form and the free form whose base it shares, and explicitly connects the shared parts with variable relational links. This tells us that whatever the form of an adjective in its free form, this will also be its form when functioning as the base of a deadjectival noun construction ending in *-ness*.

For examples like our English deadjectival nouns, very little hangs on whether we analyse them with variable interface links as in (123) or with variable relational links as in (130). But this is not true of every situation, and keeping the variable relational and variable interface links formally separate affords us greater flexibility. One of the advantages of this separation of functions — perhaps unexpectedly — is that it allows us to distinguish between patterns of allomorphy that would be described as sensitive to *surface* versus *underlying* phonology in more traditional models.

To see this, let us imagine here a hypothetical language where noun forms can appear with one of two plural suffix allomorphs, /-foo/ and /-bar/. Suppose that the appearance of these allomorphs were conditioned by the surface phonology of the stem they attach to, such that /-foo/ appears with vowel-final stems and /-bar/ appears with consonant-final stems (equivalent to our Armenian definite case above). Schemas for each of the allomorphs can be represented as below (I omit the semantics here, since it is not relevant):

(131) Allomorphy conditioned by **surface** phonology (Hypothetical language L):

a. Morphosyntax: $\{N_x, PL_{15}\}_y$
   Phonology: $[[...V]_x\ foo_{15}]_y$

b. Morphosyntax: $\{N_x, PL_{16}\}$
   Phonology: $[[...C]_x\ bar_{16}]_y$

In both schemas, there is a phonological variable slot connected to the N head in the morphosyntax by a variable interface link, *x*. This variable slot is not entirely open: in (131a) it is specified to end with a vowel, and in (131b) it is specified to end with a consonant. These are pure constraints on output form: any noun that meets the surface constraints of one of these schemas can unify with it. As indicated by the use of variable interface (not relational) links on the variable slot, the form that a root takes in other contexts is irrelevant. Therefore, even if there are phonological processes that can delete or add consonants in suffixing contexts, the allomorphy here will be conditioned only by the result of those processes.

Allomorphy may, however, be sensitive to a base's underlying phonology (Paster 2006). However, it is important to clarify what is meant by 'underlying'. In generative phonology, the underlying representation is a device for stating that two morphological constituents are 'the same' despite observed variability in form across different contexts. That variability in form is accounted for by a separate device — the rules/constraints — which furnish the language user with a surface realisation (an output) that may turn out subtantially different from the underlying representation (the input).

Constructionist approaches, as discussed in Chapter 1, reject the generative split between storage and computation — i.e. a lexicon of minimal (Platonic) forms and a grammar that combines them. That being so, the work that underlying representations and rules perform must be achieved by other means. In RM, relational links perform this work. Relational links assert that a set of linked structures in one lexical entry are 'the same' as a set of linked structures in another lexical entry. They assert identity — or more precisely, equivalence — in the face of change. The change itself is stated in these schemas too, by means of the 'same-except' asterisks.

Phonologically-conditioned allomorphy that is sensitive to the *underlying form* of the base can therefore be recast in RM in terms of a sensitivity to *the form that the base takes in some other context*. Suppose a variation on our hypothetical language. In this variation, the plural allomorphs /-foo/ and /-bar/ subcategorise, following traditional analysis, not for bases that are vowel- and consonant-final in their surface form, but rather in their underlying form. Further, let's suppose that 'underlying' root-final consonants appear intact when no suffix is present, but are deleted before suffixes. What this means in RM is that the roots that are consonant-final in an unsuffixed context are the same in a suffixing context *except that* they lack this consonant.

This is illustrated in (132), where roots are indicated by their Greek coindex to be the same across different words, except for any part between asterisks. As we see, the conditioning for the different suffixes is not distinct in the surface forms (132a,c), as both plural suffixes appear following a noun ending in a vowel. But in the forms that these nouns are relationally linked to, (132b,d), we see that these allomorphs are phonologically conditioned after all: /-foo/ is constrained to appear following bases that have a final vowel in their unsuffixed form, while /-bar/ is constrained to appear following bases that have a final consonant in their unsuffixed form.

(132) Allomorphy conditioned by **underlying** phonology, rendered opaque by a phonological alternation (Hypothetical Language L′):

Morphosyntax:	a.	$\{N_{\gamma}, PL_{15}\}_{x}$	b.	N_{γ}
Phonology:		$[[...V]_{\gamma}\ foo_{15}]_{x}$		$[...V]_{\gamma}$
Morphosyntax:	c.	$\{N_{\delta}, PL_{16}\}_{x}$	d.	N_{δ}
Phonology:		$[[...V]_{\delta}\ bar_{16}]_{x}$		$[...V\ *C*]_{\delta}$

In rule-based serialism, the choice of allomorph would precede the phonological rule deleting consonants before other consonants. Because of its early position in the ordering, the allomorphy rule has access to the underlying representation, and is therefore unaffected by the later deletion process. This makes the interaction a species of *overapplication* opacity: there is a statement that can be made about the distribution of the /-bar/ allomorph, but in (132c), it appears outside of its stated distribution. In RM, the choice of plural allomorph is conditioned by the phonology of the noun *as it appears elsewhere*. The relational link δ effectively tells us that the /-bar/ allomorph appears following a noun that is consonant-final when it appears without a suffix. The asterisks tell us that the consonant that appears in the unsuffixed form is absent from the suffixing form. In that way, the deletion and the allomorphy are handled simultaneously, by the same mechanism: relationally-linked schemas.

As we have seen in these hypothetical examples, phonological alternations that appear to be sensitive to underlying phonological representations can be restated in terms of relational links between surface word forms. Opaque phonological interactions therefore emerge as a phenomenon very similar to classic inflection classes. In fact, the only difference between phonologically-conditioned allomorphy and arbitrary, lexically-conditioned inflection classes is the phonological generalisation regarding the base.

Suppose for our /-foo/~/-bar/ example that the alternation were lexically conditioned, and that it induced no changes in the base. To capture this in RM notation, we must simply remove the generalisations concerning the phonology, as in (133). This is just a necessary consequence of

the way that constants and variables emerge from the lexicon — to say that inflection classes are lexically conditioned is really to say that there is no constant element in the meaning or form of the bases they attach to, thus the base is an open variable. In (133), the variable relational links simply tell us that if there is some plural word ending in /-foo/, then we should also expect to find a word without the /-foo/ ending and without the plural meaning. Likewise, the existence of a plural ending in /-bar/ predicts a related unsuffixed form. But in this case, the prediction no longer works in the other direction. If all we knew was the simple unsuffixed noun, we would have no way of knowing which plural suffix allomorph it should take, as it could just as easily match the subcategorisation of (133b) as (133d). This undecidability is in fact precisely what the language learner is faced with when classes are lexically conditioned.

(133) Allomorphs associated with arbitrary lexical classes (Hypothetical Language L″):

Morphosyntax:	a.	$\{N_{\gamma}, PL_{15}\}_{x}$	b.	N_{γ}
Phonology:		$[[\ldots]_{\gamma}\ foo_{15}]_{x}$		$[\ldots]_{\gamma}$
Morphosyntax:	c.	$\{N_{\gamma}, PL_{16}\}_{x}$	d.	N_{γ}
Phonology:		$[[\ldots]_{\delta}\ bar_{16}]_{x}$		$[\ldots]_{\delta}$

We have seen that RM can capture a wide variety of morphological and phonological generalisations, from surface-oriented phonologically-conditioned allomorphy, to opaque phonologically-conditioned allomorphy, to arbitary class behaviour. It does so using only statements of similarity and difference — variables and constants — derived by analogical comparison of surface word forms.

This is a striking result. One of the most widespread criticisms of standard parallel OT is that it cannot handle opacity, and this is usually diagnosed as stemming from it only allowing one 'stage' of evaluation, whereas rule-based serialism allows the output of one process to feed another. The surface-oriented character of the EVAL component is also held to be responsible here, because opacity involves phonological generalisations that may not be true or apparent at the surface. It is therefore rather surprising that we should be able to model opaque behaviour in a model that posits just one level of phonological representation, namely the surface form.

I would therefore argue that OT's inability to capture opaque interactions is not so much a consequence of its parallel evaluation mechanism but rather of its assumption that word forms are derived in isolation. It therefore cannot model any kind of class behaviour, of which opacity is just one, because that necessarily involves making generalisations (or perhaps, expectations) about the form of a word based on the knowledge of another word.

3.5.4 Summary

In this section, we have seen that Relational Morphology expresses phonological alternations as predictable relations of similarity and difference between related word forms (i.e. constructions). Shared parts of form within these words are linked with relational links; if any element within one of these linked parts differs from one word to another, it is formally notated with asterisks, indicating an exception from the sameness claimed by the relational link.

These shared sub-word parts (i.e. stems and affixes) can subcategorise for the elements they unify with. Formally, subcategorisation is expressed as some constant within a schema's variable. Where this constant is on the Phonology level, we are in the presence of phonologically-conditioned allomorphy, since it entails that the schema will only unify with structures that have certain phonological properties. Opaque phonologically-conditioned allomorphy emerges when a variable subcategorises for phonological properties of a morph not as it would appear in the local context of the given schema, but rather as it appears in other contexts (i.e., in other constructions to which it is relationally linked).

It should be noted that this approach to opacity is the logical product of combining three of RM's key theoretical claims. The first of these claims is simply that word forms are related through shared parts of structure. The second is that those parts may differ from one form to another in predictable ways. The last claim is the notion that subcategorisation can apply to any level of representation (Phonology, Morphosyntax, Semantics). The inevitable consequence of combining these three claims is that a morph can subcategorise for a phonological structure in another, related lexical entry. Each of these claims is argued for on independent grounds by Jackendoff & Audring (2020), and the authors do not consider questions of phonological opacity at all.

The analysis of opacity in RM is therefore not an ad-hoc stipulation, but the necessary consequence of theoretical claims that would be needed for a theory of morphology even if opacity did not exist. Interestingly, when looking to the future of OT, Kager (1999) speculated that a surface-based account of allomorphy would finally lead to that long hoped-for surface-based account of opacity:

> "One development that can be foreseen is that an even greater role will be assigned to *surface forms* in explaining alternations, moving away from abstract underlying forms… More on a speculative note, allomorphy may also offer an alternative account of opacity, a phenomenon for which current OT has no genuinely surface-based analysis" (Kager 1999, p. 420)

Although the means are different from those of OT, the analysis of opacity proposed here begins at much the same starting point, and achieves the desired end. In the next section, I develop a

formal analysis of ŋu-Epenthesis that draws together all the tools of RM described here.

3.6 Analysing ŋu- in Relational Morphology

As the previous section has illustrated, RM offers several advantages that address our requirements for a compelling account of ŋu-Epenthesis, insofar as we understand *ŋu-* to be a phonologically conditioned empty morph sensitive to a two-sided context, in a manner that is both non-optimising and opaque. For one, the abstractive architecture frees us from the requirement, generally assumed in morphemic theories, that each piece of form must be independently associated with some meaning. Second, with phonological conditioning treated as subcategorisation, the theory predicts the existence of non-optimising patterns of phonologically-conditioned allomorphy. That phonological subcategorisation can refer to features not only within a lexical entry but also across linked entries means that this phonological conditioning can be opaque. Moreover, since bound morphs are only represented in the lexicon in the context of the constructions they appear in — that is, words — subcategorisation can target any part of that construction, not just 'inner' morphological layers; this means that RM does not rule out outwards-sensitive allomorphy. As an additional benefit, relational links specify shared structure between surface word forms directly, and thereby obviate the need for underlying representations. By abandoning URs, we also partly sidestep the w_1/w_2 conundrum, which is irresolvable under morphemic approaches. But we also move closer to a goal held by many, especially those of the OT persuasion, of achieving a truly surface-based analysis of phonological generalisations.

Thus, we are now in a position to present a fully surface-oriented analysis of Wubuy that can encompasses prefixal and linking *ŋu-*, including all of its opaque interactions with Hardening, Nasal Deletion, and Nasal Assimilation, simultaneously. I begin in §3.6.1 with an RM analysis of prefixal *ŋu-* and Hardening, which also gives a treatment of the assimilation of prefix-final nasals. Then, in 3.6.2, I extend this to linking *ŋu-*. Finally, I consider the residual issue of non-alternating sonorants in §3.6.3.

3.6.1 Prefixal ŋu-

Recall from the discussion of the PCFP in §3.4.2 that, by comparing simple Wubuy verbs to one another, we established the existence of two 'phonological classes' — the *stop class* and the *alternating class* — on the basis of the behaviour of the stem-initial segment and the inflectional pronominal prefix. These classes are independent of the TAM inflection classes, which govern the choice of suppletive suffixal allomorph at the right edge of the verb root. We then proposed some simple schemas that would account for the form, but these were not

linked to any meaning, nor to each other. Having now addressed the RM approach to opaque phonologically-conditioned allomorphy, we can develop those schemas further.

To determine what is required for our schemas, let us first begin with some lexical entries for some of the inflected verbs that we saw in (115). Lexical entries for [ŋawulṯi:] 'I am cutting it' and [ŋaŋkulṯi:] 'I will cut it' can be presented in RM notation as in (134).

(134) Lexical entries: alternating class (verbs with *wulṯa* 'cut')

a. Semantics [PRES([CUT(Agent:1sg$_5$, Patient:3.RESID)]$_1$)$_{2,4}$]$_6$
Morphosyntax { [V$_1$, NPST2$_2$]$_3$, REALIS$_4$, [$_{\text{Pro}}$ P:1, NUM:Sg]$_5$ }$_6$
Phonology [ŋa$_{4,5}$ [*w*ulṯ$_1$ i:$_2$]$_3$]$_6$

b. Semantics [FUT([CUT(Agent:1sg$_8$, Patient:3.RESID)]$_1$)$_{2,7}$]$_9$
Morphosyntax { [V$_1$, NPST2$_2$]$_3$, IRREALIS$_7$, [$_{\text{Pro}}$ P:1, NUM:Sg]$_8$ }$_9$
Phonology [ŋa*ŋ*$_{7,8}$ [*k*ulṯ$_1$i:$_2$]$_3$]$_9$

Clearly, given Wubuy's polysynthetic verb morphology, these entries require considerable unpacking. Note first of all the nested constituent structure on the Phonology level. The verb root and its TAM inflectional suffix form a constituent — i.e. a stem — which is nested inside a larger constituent containing the pronominal prefix. The [wulṯ] string on the phonological level in the stem of (134a) has an interface link to the V category head in the morphosyntax. The [kulṯ] string in (134) also has an interface link to the V in its morphosyntax. Additionally, these interface links serve as relational links to each other, meaning that the two phonological forms [wulṯ] and [kulṯ] are to be taken as instantiating the same morphosyntactic and semantic elements, namely a verb root with the meaning 'cut'. Their phonological form is also the same, *except for* the part delimited with asterisks. (The prefix in (134b) also features an asterisked part; this distinguishes its form from the prefix it is relationally indexed to in (135b), as we will see.) The root and its tense-aspect suffix are grouped into a constituent also, and this constituent is the one that will mostly be referred to in the discussion of stem classes below.

The notation of the pronominal prefixes also requires some comment. My analyses here take the pronominal prefix as an undifferentiated block, and index it to various morphosyntactic features: the person, number and gender/noun class features of the subject and object pronominal arguments, and realis/irrealis mood. Also, in order to avoid cluttering the page with a needless profusion of coindices, I group the morphosyntactic features associated with each pronominal argument into a constituent and coindex that to the prefix in the phonology.

Finally, the semantics is not important to the questions in this chapter, but the interested reader may note a couple of matters of detail. For one, there is a patient argument in the semantics that is not connected to a pronominal argument in the morphosyntax. This can be

taken to reflect the fact that any intrasitive prefix can be interpreted as transitive with a RESID object; to explicitly include this argument would be just as acceptable. (For discussion of the potential virtues of decoupling semantic interpretation and morphosyntax, see Jackendoff & Audring 2020, pp. 164–166.) Secondly, I take the tense reading to be associated not with any one affix but with the combination of the choice of realis/irrealis prefix and the suffix on the verb root (see Verstraete 2005). This is reflected in the coindexation of the PRES function in the semantics of (134a), which simply links it to the outer interface link surrounding the whole word — i.e., the combination of prefix and inflected verb root. An alternative would be to index the tense function in the semantics to the REALIS and NPST2 categories in the morphosyntax, and from there to those categories' corresponding structures on the phonological level. But since the embedding of functions in the semantics does not neatly follow the embedding of morphological constituents, and I am only interested here in the latter, I leave a more carefully articulated formalisation of the semantics and its interfaces to future research. To make sense of the lexical entries here and in subsequent sections, readers would suffer no great loss by ignoring the semantic level altogether.

Now let us consider some examples of verb stems from the stop class. In (135), we have lexical entries for [ŋakaɾaci:] 'I am jumping' and [ŋaɲukaɾaci:] 'I will jump'. This time, the stem form is the same across both the realis and the irrealis entries, but the prefix is different in both form and meaning. In (135a) the prefix has the form [ŋa] and is linked to the REALIS mood category in the morphosyntax, while in (135b) it has the form [ŋaɲu] and is linked to the IRREALIS. But note that the indexation on each of these prefixes is also linked relationally to the prefixes in (134a) and (134b), respectively. The REALIS-realising prefixes in (134a) and (135a) have the same form: [ŋa]. The form of the IRrealis prefix differs between classes, however: it is [ŋaŋ-] in (134b) and [ŋaɲu-] in (135b). In other words, these two prefixes are allomorphs of one other — they appear in different words but are associated with the same morphosyntactic categories and semantics. Their phonological form in each context is the same *except for* the part in asterisks. (Meanwhile, the suffixal TAM inflection on the verb root bears the same coindex as the inflection on the root of (134a,b), since both verbs are in the A2 inflection class, though this has no bearing on the phonology.)

(135) Lexical entries: stop class (verbs with *karaca* 'jump')

a.	Semantics	$[\text{PRES}([\text{JUMP}(\text{Agent:1sg}_8)]_{10})_{2,4}]_{12}$
	Morphosyntax	$\{[\text{V}_{10}, \text{NPST2}_2]_{11}, \text{REALIS}_4, [_{\text{Pro}}\ \text{P:1}, \text{NUM:Sg}]_5\ \}_{12}$
	Phonology	$[\text{ŋa}_{4,5}\ [\text{kaɾac}_{10}\ \text{i:}_2]_{11}\ ]_{12}$

b. Semantics $[\text{FUT}([\text{JUMP}(\text{Agent:1sg}_8)]_{10})_{2,7}]_{13}$
Morphosyntax $\{[\text{V}_{10}, \text{NPST2}_2]_{11}, \text{IRREALIS}_7, [_{\text{Pro}}\ \text{P:1}, \text{NUM:Sg}]_8\ \}_{13}$
Phonology $[\text{ŋa*ŋu*}_{7,8}\ [\text{kaɾac}_{10}\ \text{i:}_2]_{11}\ ]_{13}$

Taking the lexical entries for our *wulṯa* and *kaɾaca* verbs as our point of departure, while also comparing across all the examples of alternating and stop class stems provided in this chapter, we can propose abstract schemas to describe the prefix and stem allomorphy. To do that, we apply the Structural Intersection operation as described in §3.5.3, replacing any elements of structure that vary from one instance to another with variables, and leaving in place any elements that remain constant. This process yields the alternating class schema in (136) and the stop class schema in (137).

(136) Alternating class: inflectional prefix-stem schemas

a. Semantics $[\text{F}([\text{PRED}(\ldots_\gamma)\ ]_\alpha\ )_\delta\ ]_\text{x}$
Morphosyntax $\{\text{X}_\alpha, \text{MOOD}_\delta, \text{Pro}_\gamma\ \}_\text{x}$
Phonology $[\ldots\text{V}_{\delta,\gamma}\ [\ \begin{bmatrix} \text{*+SON*} \\ \text{*+CONT*} \\ \text{PLACE}_\varepsilon \end{bmatrix} \ldots]_\alpha]_\text{x}$

b. Semantics $[\text{F}([\text{PRED}(\ldots_\zeta)\ ]_\alpha\ )_\eta\ ]_\text{x}$
Morphosyntax $\{\text{X}_\alpha, \text{MOOD}_\eta, \text{Pro}_\zeta\ \}_\text{x}$
Phonology $[\ldots^*\ \begin{bmatrix} \text{+NASAL} \\ \text{PLACE}_\varepsilon \end{bmatrix}^*_{\eta,\zeta}\ [\ \begin{bmatrix} \text{*−SON*} \\ \text{*−CONT*} \\ \text{PLACE}_\varepsilon \end{bmatrix} \ldots]_\alpha]_\text{x}$

(137) Stop class: inflectional prefix-stem schemas

a. Semantics $[\text{F}([\text{PRED}(\ldots_\gamma)\ ]_\beta\ )_\delta\ ]_\text{x}$
Morphosyntax $\{\text{X}_\beta, \text{MOOD}_\delta, \text{Pro}_\gamma\ \}_\text{x}$
Phonology $[\ldots\text{V}_{\delta,\gamma}\ [\ \big[\text{−SON}\big]\ldots]_\beta]_\text{x}$

b. Semantics $[\text{F}([\text{PRED}(\ldots_\zeta)\ ]_\beta\ )_\eta\ ]_\text{x}$
Morphosyntax $\{\text{X}_\beta, \text{MOOD}_\eta, \text{Pro}_\zeta\ \}_\text{x}$
Phonology $[\ldots^*\text{ŋu*}_{\eta,\zeta}\ [\ \big[\text{−SON}\big]\ldots]_\beta]_\text{x}$

Setting aside the phonological aspects for one moment, we can consider the features that need to be referenced on the Morphosyntax and Semantics levels, which are identical across the alternating class and stop class schemas. Firstly, we note that the basic constituent structure of a prefix and stem remains constant across the words in question. The prefixes need to encode

at least one pronominal argument and a mood category, as in our lexical entries above. While the stems we have seen so far have been verbs of the A2 class, we know from other examples in this chapter that the stem here can be of any verbal inflectional class, and indeed may not be a verb at all, since adjectival stems take the same prefixes. For that reason, our schemas must demand only that the stem be linked to a category head in the morphosyntax which is linked to a predicate function in the semantics; this condition can be satisfied by verbs as well as adjectives. We must also abstract away from any internal structure in the stem, as it does not matter whether the stem in question is just an inflected root, or is derivationally complex.

The crux of the analysis of prefixal *ŋu-* and its interactions with Hardening, Nasal Deletion and Nasal Assimilation can be grasped by examining the relational variable links in each of these schemas, especially those on the phonological level. As a visual aid, the relational links connecting the phonology levels of (136a,b) and (137a,b) are depicted as association lines in the diagram in (138). The association lines are also annotated with the constituent they link (either 'prefix' or 'stem').

At the heart of this analysis is the idea that prefix allomorphs and stem allomorphs *sub-categorise for each other*. The schema of (136a) features a relational variable (α) joining its continuant-initial stem to the stop-initial stem of (136b). This performs the same work as Hardening, but the continuant-initial allomorph and the stop-initial one are on equal footing; neither is basic to the other. In like fashion, the stop class stem in (137a) is linked to the stem in (137b), and is specified to be stop-initial in both cases.

Meanwhile, the prefix in (136a) is relationally linked to the prefix in (137a), with no change in meaning, and no change in form: in both contexts, this prefix is vowel-final. But the prefixes in (136b) and (137b), also relationally linked to one another, show a predictable discrepancy in form: that prefix which ends in a nasal in (136b) is the same in (137b) except that it ends in [ŋu]; in each schema, it is linked to the same morphosyntactic categories. So, rather than Nasal Deletion, we simply have replacement of the prefix-final nasal with [ŋu]. Moreover, the final nasal in the prefix of (136b) must have the same place features as the stop that follows it, which is the same as the place features of the stem-initial continuant in (136a), as indicated by the relational variable indexed to the PLACE feature; this encodes the same generalisation as the Nasal Assimilation rule.

(138) Relational variable links as bidirectional association lines

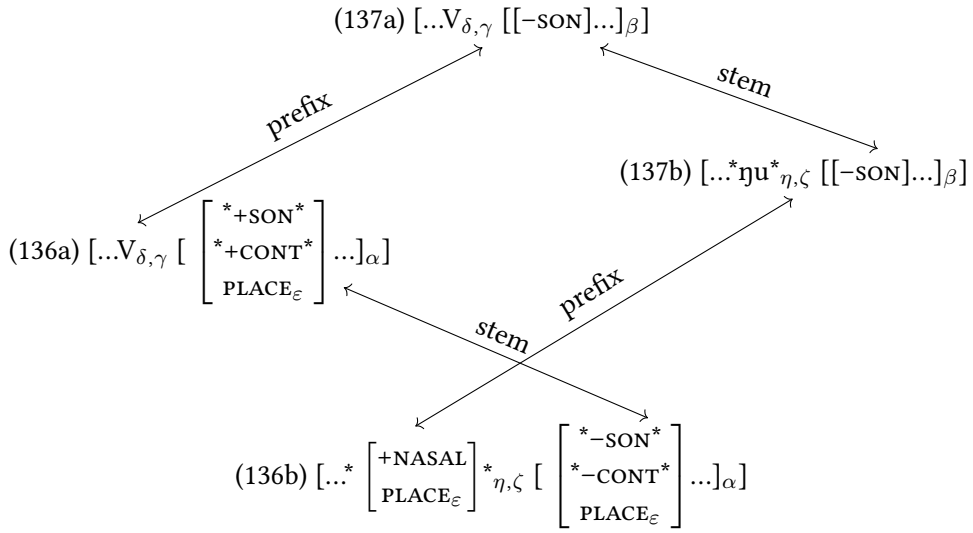

The schemas in (136a,b) and (137a,b) thus serve to identify a fully inflected word as a member of either the stop class or the alternating class, on the basis of two key pieces of information in conjunction: the type of prefix and the type of stem. Neither piece of information on its own is sufficient. In this mode, schemas are interpreted declaratively, abstracting a generalisation over the lexicon. But schemas can also be interpreted procedurally, using knowledge of existing words in the lexicon to form new or unknown words.

By identifying the phonological class of a given word, we can predict the form of either the prefix or the stem in some other context. The schemas effectively tell us that if a prefix is vowel-final before a continuant (136a), it will have that same form before a stop (137a). Meanwhile, if a stem is continuant-initial after one of those vowel-final prefixes (136a), it will be stop-initial after a nasal-final prefix (136b). And if a prefix is nasal-final preceding a stop that is elsewhere a continuant (136b), it will be instead *ŋu*-final before a stop that is elsewhere a stop (137b). Schemas, then, encompass both representations and rules.

The empty slots in these schemas can be filled in by ('can unify with'), for example, the prefix and stem constituents from our lexical entries (134) and (135). The prefix /ŋa-/ from our (134a) *ŋa-wulṯi:* entry can instantiate the prefix slot in (136a), which means that it can also instantiate the prefix slot it is relationally linked to in (137a) — as can /ŋa-/ from (135a) *ŋa-karaci:*, since it bears the same coindex. The stem with the form [wulṯi:] from (134a) *ŋa-wulṯi:* alternates with [kulṯi:] in (134b) *ŋaŋ-kulṯi:*. Thus [wulṯi:] and [kulṯi:], bearing coindex 3, can instantiate the constituent coindexed with the relational variable α in (136a) and (136b), respectively.

The opaque behaviour of *ŋu-* emerges from the phonological subcategorisation requirements

of these relationally linked stems and prefixes. Consider again the underapplication opacity problem, where ŋu-Epenthesis only applies to 'underlying' stops, but fails to apply to stops derived by Hardening. Here, the prefixal allomorphs ending in [ŋu] — schematised in (137b) — subcategorise for a stop-initial stem. But only a stop-initial stem of a particular kind: namely, one that is also stop-initial when following a vowel-final prefix, as in (137a). And note that this is not just any vowel-final prefix either, as its relational link indicates that it must be vowel-final preceding both stop-initial and continuant-initial stems. On the other hand, a stem that is continuant-initial following one of these vowel-final prefixes, as in (136a), is linked to a stop-initial stem allomorph in a different context, (136b). But even this stop-initial allomorph cannot ever appear with prefixal *ŋu-*, as it does not meet the subcategorisation requirement indicated by the relational variable link β in (137b). The contrast in the behaviour of 'underlying' versus 'derived' stops is here transformed into a contrast between phonological classes, formally expressed as sets of sister schemas related by relational variables, each set bearing different phonological subcategorisation features. The *ŋu*-final prefixes appear with one class, but not the other.

The overapplication (self-destructive feeding) problem is addressed by exactly the same elements. Instead of an invariant underlying prefix form ending in /N/, which 'triggers' application of ŋu-Epenthesis only to be subsequently deleted, our schemas only represent the surface result of both 'rules'. The string which ends in [ŋu] in (137b) is linked to that which ends in a nasal (homorganic to the following stop) in (136b). Rather than locating the trigger for ŋu-Epenthesis in an abstract underlying segment, we locate it in a word of a different phonological class: indeed it was owing to the existence of words with nasal-final prefixes, after all, that Heath found motivation for positing nasals in the underlying representation, despite their absence at the surface. Here, we keep the generalisation that the abstract segment enables intact, but without requiring underlying representations or any other hidden levels.

The relational links can be considered to perform an important blocking function as well. Recall from (105) that there are various prefixes — such as *ŋaŋu-* '1sg>3Fem' — that happen to end in [ŋu], but which are not implicated in ŋu-Epenthesis. These do not alternate: they appear in the same form before stems of both the stop class and the alternating class. This means that a word with one of these prefixes can only be an instantiation of regular vowel-final prefixes, as in either (136a) or (137a).

Finally, note that we also have a simple account of transparent phonology here, as in our treatment of Nasal Assimilation, which again makes use of variable relational links. The final segment of the prefix in (136b) is underspecified: it is a nasal without a constant place of articulation. Again, this is a consequence of comparing what is the same and what is different among all its realisations. While the prefix realising first singular irrealis in (134b) *ŋaŋ-kulṯi:*

has the form [ŋaŋ], the prefix with the same meaning in *ŋaŋ-cuɻi:* or *ŋam-paɭaɾ̰ɭaŋ* has different place features. But crucially, in all such words, the prefix-final nasal and the stem-initial stop always have *the same place features*. This is indicated by a variable relational coindex on the PLACE feature: it simply means that whatever unifies with one of these must unify with both — that is, unification must result in a homorganic nasal-stop sequence.

The place of articulation in the initial segment of the stem, meanwhile, does not vary across the words it appears in, even though it does vary in manner (e.g. *ŋa-**w**ulʈ̰i:, ŋaŋ-**k**ulʈ̰i:; pa-**w**ulʈ̰i:, nuŋ-**k**ulʈ̰i:*, etc.). Therefore, the stem must be specified to have place features. If these verb stems necessarily have place features, then they can 'fill in' the empty place node in the stem-initial segment of (136b). Due to the variable relational coindex, this unification simultaneously propagates leftwards, filling in the variable place feature of the prefix-final nasal. Although Jackendoff & Audring (2020) do not discuss the affordances of relational variables for addressing classically 'phonological' processes like assimilation, it falls out naturally from their definition. Nor is my treatment here entirely new. Similar proposals for assimilation through feature sharing, attained via the logical unification of partial statements about the phonological form of surface words, is one of the leading ideas of Declarative Phonology (Bird 1990; Bird & Ellison 1994; Coleman 1990, 1998; Ogden 1999; Scobbie 1991). Various kinds of featural agreement by correspondence between surface segments have also been proposed in OT work (Hansson 2007; Inkelas & Shih 2014; Rose & Walker 2004; Walker 2000).

We have seen that the phonological classes we established more informally in §3.4 are expressed formally in RM as sets of sister schemas related through variable relational links, where each class subcategorises for different phonological features, and the phonology of linked constituents may differ across contexts (via the same-except relation). By these mechanisms, we can finally address the systematic differences in phonological behaviour between the stop class and the alternating class, without invoking underlying representations, intermediate derivational stages, or any other hidden constructs. That is, we reach a fully surface-based analysis of opacity — indeed, both the underapplication and overapplication dimensions of the opacity problem are dealt with simultaneously, and using the same mechanisms. As I will show in the next section, the approach taken here for prefixal *ŋu-* also extends quite straightforwardly to linking *ŋu-*.

3.6.2 Linking ŋu-

We can now extend the analysis of prefixal *ŋu-* in terms of phonological classes to linking *ŋu-*. We are concerned here with words in which some element — a derivational prefix or an incorporated noun — intervenes between the pronominal prefix and the root. The schema in (139) states that, in a word where the root is of the stop class, the string [ŋu] must appear on the

phonological level between a derivational element and the root. While the prefix, derivational element and root bear interface links to their corresponding morphosyntactic and semantic structures, *ŋu-* is not linked to anything. This follows Jackendoff & Audring's (2020, p. 101) approach to linking elements in Dutch compounds. The derivational element is here a completely open variable, since there is no subcategorisation involved, phonological or otherwise: the only feature of consequence is its linear position, which distinguishes it from the inflectional pronominal prefix.

(139) Derivationally complex words: *stop class*

a.	Semantics	$[[\text{PRED}(\ldots_x)_\beta; \ldots_y]_z]_a$
	Morphosyntax	$\{ [X_y\ X_\beta]_z, \text{Pro}_x \}$
	Phonology	$[\ldots_x [\ldots_y$ ŋu $[[-\text{SON}]\ldots]_\beta]_z]_a$

Where the root is of the alternating class, as in (140), the initial segment once again alternates between a continuant and a stop. This time, however, we only know that the continuant-initial allomorph must follow a continuant (140a), whereas the stop-initial allomorph can follow either a non-continuant or a continuant (140b). The latter situation arises in the case of morphs ending in /k/ and /K/, which must be deleted in this context, as was discussed in §3.4.3. I analyse these cases in more detail in Chapter 5, and argue that such 'underlyingly' velar-final morphs subcategorise for a following non-continuant. If we allow that the schemas for a given derivational element subcategorise for the element to their right, then they will only be able to unify with either (140a) or (140b), not both.

(140) Derivationally complex words: *alternating class*

a.	Semantics	$[[\text{PRED}(\ldots_x)_\alpha; \ldots_y]_z]_a$
	Morphosyntax	$\{ [X_y\ X_\alpha]_z, \text{Pro}_x \}$
	Phonology	$[\ldots_x [\ldots[+\text{CONT}]_y [\begin{bmatrix} *+\text{CONT}* \\ *+\text{SON}* \end{bmatrix} \ldots]_\alpha]_z]_a$
b.	Semantics	$[[\text{PRED}(\ldots_x)_\alpha; \ldots_y]_z]_a$
	Morphosyntax	$\{ [X_y\ X_\alpha]_z, \text{Pro}_x \}$
	Phonology	$[\ldots_x [\ldots_y [\begin{bmatrix} *-\text{CONT}* \\ *-\text{SON}* \end{bmatrix} \ldots]_\alpha]_z]_a$

Again, variable relational links ensure that only roots of the right kind can instantiate the schemas in (140) and (139). The rightmost element of alternating schema (140) — the root — is coindexed α. Thus, that which can fill this slot must also be able to fill the likewise coindexed slots in the prefix-stem schemas of (136). In other words, if a root appears continuant-initial

after a vowel-final prefix and stop-initial after a nasal-final prefix, it is also predicted to have these forms when appearing in more complex constructions.

On the other hand, roots of the stop class can only unify with the appropriate variable of (139). They cannot unify with the schema for the stop-initial allomorph of the alternating class in (140b), since they cannot also unify with (140a), nor with the inflectional schemas (136), as required by the variable relational links. Note that, in referring to the stop and alternating classes here, we are talking only of the inner (right-hand) constituent of the complex stem here. The class of the stem as a whole is still determined by its leftmost element, as per the prefix-stem schemas in (136) and (137) above.

In fact, the constituent that instantiates the right-hand slot in the complex stems of (139) and (140) need not be a simple one like a root, but may itself be complex — it may, for instance, consist of an incorporated noun and a verb root. This complex constituent may then be fed back into one of (139) or (140) to yield yet another complex constituent. This nesting of linked morphological constituents thus implies a degree of recursion: a complex constituent can unify with a complex constituent, and so on. However, as is standard in models based on constituent structure and unification, I assume that a construction cannot unify with another construction containing a (constant) constituent of the same type, and that therefore, this recursiveness is not infinite. For example, a stem containing a verb root cannot unify with a stem already containing a verb root. Rather, a verb can unify with a variable slot that subcategorises for a V head in the morphosyntax and a predicate in the semantics. (More generally, we can say that a variable can unify with a constant of the same type, and a variable can unify with a variable of the same type, but a constant cannot unify with a constant of the same type.)

3.6.3 Nasals and other non-alternating sonorants

The RM treatment of stems in the alternating class developed here has drawn on the feature specification [+CONT,+SON] to target the segments that alternate with a stop. As noted in §3.2.2, there are in fact two continuants that escape Hardening: /l/ and /l̠/. In stem-initial position, these instead pattern with the nasals. Before addressing the problem of the non-alternating laterals, then, let us briefly formalise the nasals.

Nasal-initial stems have not figured prominently in this chapter, since they do not alternate. But they do implicate a phonological process we have seen already: Nasal Deletion. That is, those pronominal prefixes previously analysed as underlyingly ending in /N/ undergo deletion of the final nasal when appearing before a stem beginning in a nasal (or one of the non-alternating laterals). The first singular irrealis, for instance, which is elsewhere /ŋaN-/ or /ŋaŋu-/, is in this context just /ŋa-/.

A pair of sister schemas can be set up as in (141). Although the stem form does not alternate,

the prefix in (141b) is without the part between asterisks in the prefixes it is relationally linked to in (136b) and (137b). This therefore does the same work as Nasal Deletion: a prefix-final nasal that appears in some other context is absent in this context.

(141) Sonorant class: inflectional prefix-stem schemas

a. Semantics $[\mathrm{F}([\mathrm{PRED}(\ldots_{\gamma})\ ]_{\vartheta}\ )_{\delta}\ ]_{\mathrm{x}}$
Morphosyntax $\{\ \mathrm{X}_{\vartheta}, \mathrm{MOOD}_{\delta}, \mathrm{Pro}_{\gamma}\ \}_{\mathrm{x}}$
Phonology $[\ldots \mathrm{V}_{\delta,\gamma}\ [\ \begin{bmatrix}+\mathrm{SON}\\ -\mathrm{CONT}\end{bmatrix} \ldots]_{\vartheta}]_{\mathrm{x}}$

b. Semantics $[\mathrm{F}([\mathrm{PRED}(\ldots_{\zeta})\ ]_{\vartheta}\ )_{\eta}\ ]_{\mathrm{x}}$
Morphosyntax $\{\ \mathrm{X}_{\vartheta}, \mathrm{MOOD}_{\eta}, \mathrm{Pro}_{\zeta}\ \}_{\mathrm{x}}$
Phonology $[\ldots_{\eta,\zeta}\ [\ \begin{bmatrix}+\mathrm{SON}\\ -\mathrm{CONT}\end{bmatrix} \ldots]_{\vartheta}]_{\mathrm{x}}$

There are two potential avenues for addressing the patterning of the non-alternating laterals within the formalism. One could either redefine the feature specification of these segments, or specify exceptional schemas for them at a lower level of abstraction.

The first option would define /l/ and /l̠/ as [−CONT], rather than [+CONT] like the dental lateral. This would make it meet the subcategorisation requirement of the schemas in (141). Typological evidence suggests that this would not be entirely unwarranted: Mielke's (2008) survey finds that laterals are conspicuously prone to ambiguity with respect to continuancy, patterning with continuants in some languages and non-continuants in others, a fact he attributes to their ambiguous phonetic attributes. In his Emergent Feature Theory, features are assigned to segments on a language-particular basis, according to the alternations they participate in.

However, /l/ and /l̠/ only pattern with the sonorants in morph-initial position. In morph-final position, they do not trigger Hardening in following continuants, unlike the nasals; in this respect they pattern with the continuants. Reclassifying these two segments as [−CONT] therefore does not resolve but rather displaces the problem of their exceptionality.

The second option is to define schemas for these segments at a lower level of abstraction, working at the level of the individual phoneme rather than the feature, as in (142) and (143). Since RM assumes that a more specific schema takes precedence over a more general one, these schemas will be the ones to generalise to new contexts involving stems with an initial /l/ or /l̠/, rather than the alternating class schemas above.

(142) /l/-initial stems: inflectional prefix-stem schemas

a. Semantics [F([PRED(...$_\gamma$)]$_\iota$)$_\delta$]$_x$
Morphosyntax { X$_\iota$, Mood$_\delta$, Pro$_\gamma$ }$_x$
Phonology [...V$_{\delta,\gamma}$ [l...]$_\iota$]$_x$

b. Semantics [F([PRED(...$_\zeta$)]$_\iota$)$_\eta$]$_x$
Morphosyntax { X$_\iota$, Mood$_\eta$, Pro$_\zeta$ }$_x$
Phonology [...$_{\eta,\zeta}$ [l...]$_\iota$]$_x$

(143) /l̩/-initial stems: inflectional prefix-stem schemas

a. Semantics [F([PRED(...$_\gamma$)]$_\kappa$)$_\delta$]$_x$
Morphosyntax { X$_\kappa$, Mood$_\delta$, Pro$_\gamma$ }$_x$
Phonology [...V$_{\delta,\gamma}$ [l̩...]$_\kappa$]$_x$

b. Semantics [F([PRED(...$_\zeta$)]$_\kappa$)$_\eta$]$_x$
Morphosyntax { X$_\kappa$, Mood$_\eta$, Pro$_\zeta$ }$_x$
Phonology [...$_{\eta,\zeta}$ [l̩...]$_\kappa$]$_x$

While I expect that not everyone will find this solution entirely satisfying, it is certainly possible in the theory. Moreover, the ability to generalise at multiple levels of abstraction is taken to be a key strength of lexicon-based approaches (Albright & Hayes 2003; Booij 2010; Ernestus & Baayen 2004; Wedel 2009).

But it should also be noted that we already have empirical reasons to regard phoneme-level generalisations as playing a role in speakers' knowledge of continuant-stop alternations anyway. For one, the continuants that alternate with stops are fairly heterogeneous. Secondly, there are two overlapping mappings for /w/: for /w/-initial stems, feature-level generalisations do not completely reduce uncertainty in predicting what the related stop form will be.

3.7 Conclusion

The present chapter has provided a detailed examination of what is a truly puzzling conundrum: ŋu-Epenthesis. As I demonstrated, *ŋu-* is an empty morph whose conditioning is phonologically conditioned, in a way that is phonologically non-optimising, opaque and outwards-sensitive. In implicating all these properties simultaneously, it tests the limits of our theories of the phonology-morphology interface. Indeed, an analysis within mainstream generative theory was shown here to be unachievable, owing to its basic assumption that morphology is 'constructive'

in nature (Blevins 2006) — that is, its task is to build words out of morphemes, which are then passed to the phonological component to be 'readjusted' according to well-formedness criteria. Our *ŋu-* empty morph is troubling under such a view: it is not really a morpheme, and it is not a phonological repair.

As an alternative, I showed that *ŋu-* is better analysed from the perspective that morphology is 'abstractive': whole words are stored in the mental lexicon, while recurrent parts of words emerge as abstractions over these stored word forms. Where a language has several lexically-assigned inflection classes, empty morphs have been analysed under this perspective as recurrent parts that identify the inflection class to which a word belongs; such has been argued for Latin and Pitjantjatjara, for instance. As I argued here, the *ŋu-* empty morph in Wubuy reduces uncertainty about a word's phonological class — which, owing to the neutralising effects of Hardening, would not otherwise be apparent. That is, it identifies a stop-initial stem as being of the kind that is a stop in all environments, and not the kind that alternates with a continuant. In doing so, *ŋu-* significantly reduces uncertainty in predicting an unknown word form based on a known related form.

I then proposed a formalisation in terms of the relatively new theory of Relational Morphology. Although the theory was developed in order to answer more properly morphological questions, I showed that it could be deployed to address phonological alternations as well. Rather surprisingly, it turned out that, if we follow the assumptions of the model to their natural conclusion, we arrive at a novel, fully surface-based account of opacity that rests on neither underlying representations nor serial derivations. That is, it turns out we can capture the *effects* of underlying representations — such as the alternations that are sensitive to 'underlying' phonology in Paster's (2006) typological survey — without having to subscribe to the *existence* of underlying representations.

Nonetheless, some lingering questions remain. For one, while the RM account of opacity holds a great deal of promise, it will obviously be necessary to test it on a much wider sample of cases. Second, the schemas developed in the account here work at a high level of abstraction. This largely reflects the tradition of formal linguistic analysis that I am in dialogue with, which has tended to seek the smallest and simplest set of grammatical statements that captures the data. However, as a model of how speakers actually generalise their knowledge to a new context in a concrete situation, this is probably incomplete. Rather than generalising on the basis of schemas that capture, say, the entire alternating class, they may instead generalise on the basis of just the subset of stems that alternate, for instance, /l̪/ and /t̪/ (see Albright 2002). But keep in mind that a crucial aspect of RM, as with other lexicon-based approaches, is its assumption that analogical generalisations can apply at any level of abstraction. A model that can analyse *ŋu-* at the most abstract level and at more concrete ones is still, to my mind, preferable to the

available generative alternatives, which fail to capture it at any level. Still, a fuller model of speaker knowledge, addressing aspects of both competence and performance, must surely be stochastic and probabilistic (e.g. Albright & Hayes 2003; Hall 2012; O'Donnell 2015; Rácz et al. 2020). I leave this to future research.

In the next chapter, I turn to a detailed examination of Wubuy's typologically unique reduplication pattern, whereby the size of the reduplicant alternates according to the sonority of the initial segment of the stem it copies. Again, this alternation is phonologically conditioned but non-optimising, and it once again hinges on the difference between stop-initial stems and stems of other kinds. In order to address it, I develop an RM analysis that builds on the analysis of ŋu-Epenthesis and Hardening here. Again, the generative offerings are not able to provide a similarly successful solution, for reasons that are related to those discussed in this chapter, but are nonetheless quite distinct from them.

4 Sonority-conditioned reduplication

4.1 Introduction

An extremely pervasive phenomenon of Wubuy discourse, apparent on virtually any page of Heath's (1980) collection of texts, is the repetition of part of a word in order to indicate plural entities or events, or some greater extension in time or space. This pattern of repetition, known as *reduplication*, is a feature of nouns, verbs, demonstratives, and sometimes adjectives. The first text in Heath (1980), 'Olive Python and the Two Boys', consists of 288 words, of which 51 (17.7%) show reduplication. Clearly, then, an important part of Wubuy speakers' linguistic competence is the ability to produce and interpret reduplicative words, including novel ones.

This chapter aims at a detailed description and analysis of reduplicative word-formation processes in Wubuy, situating the discussion within current debates in the literature on the phonology and morphology of reduplication, and comparing the Wubuy patterns with those of neighbouring languages. Although the semantic, syntactic and discursive characteristics of reduplication constructions would be a fertile area of study, the discussion in this chapter is limited to issues of reduplicative form.[1]

Of particular interest is the fact that in the only productive reduplication pattern, the size of the reduplicant is conditioned by the sonority of the initial segment of the base, which is typically the entire stem (everything that follows the inflectional pronominal prefix). Stop-initial stems take a monosyllabic reduplicant, while sonorant-initial stems take a disyllabic reduplicant. Thus sonorant-initial *wuɾuc* '(Aboriginal) person' reduplicates as *wuɾu-wuɾuc* '(Aboriginal) people', while stop-initial *ṯuɾapata* 'non-Aboriginal person' reduplicates as *ṯu-ṯuɾapata* 'non-Aboriginal people'. Other than Wubuy's closest genetic neighbour, Anindilyakwa, I know of no other language with a reduplicant size alternation conditioned by the featural makeup of the input.

Although the pattern is fairly straightforward to describe, it does not admit a ready explanation in terms of either Base-Reduplicant Correspondence Theory (BRCT; McCarthy & Prince 1993b, 1995) or Morphological Doubling Theory (MDT; Inkelas & Zoll 2005), two major formal theories of reduplication. I instead adopt a different approach, using the basic tools of Relational

[1] See Bednall (2019) for a detailed examination of the semantics of reduplication constructions in Anindilyakwa, much of which is likely to be applicable to Wubuy.

Morphology (RM; Jackendoff & Audring 2020), none of which are specific to reduplication, and argue that the approach performs better than the alternatives in capturing the Wubuy facts.

In this chapter, then, I hope to address one of the major frustrations in the literature on reduplication, which is that there is an alarming paucity of detailed descriptive work on reduplication in any language; as a result, theories are often built on the basis of a small number of examples, often recycled from one theoretical formulation to another, and we thus lack a grasp of the full scope of possibility in the typology of reduplication (Hurch & Mattes 2005; Kiyomi 1993).

The chapter is organised as follows. Section §4.2 establishes the basic facts of Wubuy reduplication, covering its basic functions before examining reduplicant form, where I demonstrate that the shape of the reduplicant is conditioned by the sonority of the base-initial segment. I then discuss the interactions between reduplication and other phonological processes like Hardening, ŋu-Epenthesis, and vowel hiatus resolution. I also note some exceptions to the main generalisations.

Having established the empirical phenomena at issue, the next part of the chapter, §4.3, compares the Wubuy pattern with reduplicative structures in other languages of the region. The sonority-conditioned shape alternation is shown to exist in almost identical fashion in nearby Anindilyakwa, though it is absent from related languages Ngandi, Ngalakgan and Bininj Kunwok. I then seek an explanation for Wubuy and Anindilyakwa's uniqueness in terms of lexical access, making crucial reference to the comparative rarity of stops in these two languages vis-à-vis their genetic relatives.

The remainder of the chapter is more theoretical in orientation. In §4.4, I develop an analysis in RM, building on the foundations that were set out in Chapter 3. RM can capture reduplication by means of analogical comparison between surface words, expressed as relations of sameness and difference. These are the same mechanisms assumed by the theory for expressing the relationship between any lexical entries, and are not specific to reduplication. The defining characteristic of reduplication is instead about how the tools are applied: reduplication constructions feature relations of correspondence between phonological sequences within the same lexical entry — that is, between reduplicant and base — as well as between a reduplicated word and its unreduplicated counterpart. A key plank of the model is the ability for a reduplicant shape to subcategorise for the phonological class of its base, even in the face of opaque interactions between reduplication and other processes. Since reduplication constructions have not otherwise been given any serious attention in RM, one contribution of this chapter is to broaden the theory's empirical coverage.

I then consider alternative approaches in §4.5. I first attempt an analysis in BRCT, showing that, while each reduplicant shape is analysable on its own in BRCT, capturing the alternation

between them leads to an intractable ranking paradox. Echoing some of the discussion of empty morphology in Chapter 3, I show that BRCT — couched as it is in strictly parallel OT — lacks a sufficiently articulated notion of phonological subcategorisation, and therefore cannot capture alternations that are not in service of either markedness or identity, such as the one under discussion. Following this, I suggest an analysis in MDT. Insofar as the model allows for phonological subcategorisation of the input, it can generate the different reduplicant shapes by assigning stop-initial and sonorant-initial bases to different cophonologies. However, MDT's basic procedure of generating a reduplicant by copying an entire stem (of any length) and truncating it down to the required size is laden with dubious cognitive and historical implications. Finally, the conclusion in §4.6 draws together the findings and suggests directions for future research.

4.2 Descriptive overview

In this section, I provide the necessary empirical background to the theoretical discussion to follow, which is concerned with the formal characteristics of reduplication. As Baker & Harvey (in press) point out, examples from Australian languages have often figured prominently in these theoretical debates, but the descriptive materials the data come from typically contain only scant discussion and exemplification of reduplication, and therefore tend to be inadequate for answering even basic empirical questions, such as whether the pattern is productive. While the main reduplication pattern under consideration here is of a type unheard of in the literature, it is by no means marginal in Wubuy, and I take effort here to amass a large number of examples and situate them within the broader system of the grammar.

To begin, some definitions. In labelling a particular word as an example of reduplication, I assert that some substring within the grammatical word is repeated, and that this repetition is the only formal property distinguishing the word from another word to which it is semantically related (its unreduplicated counterpart). Throughout this section, I follow more traditional treatments of reduplication in referring to the part common to both the reduplicated and unreduplicated words as the *base*, and the additional material as the *reduplicant*, and will say that the reduplicant *copies* material from the base. These terms, while convenient for descriptive purposes, come with theoretical baggage, and are not uncontested (e.g. Inkelas & Zoll 2005); I set those issues to the side.

Not all kinds of repetition count as reduplication under this definition. For one, it excludes those roots that happen to contain a repeated substring in their basic form, but that do not have an unreduplicated counterpart, such as the verb root *kuckuci* in (146) below (see Kiyomi 1993). So too does it exclude syntactic repetition: entire grammatical words are often repeated

in Wubuy discourse — as in the repetition of *ka.ɻa* 'hole' and *maʈa* 'grass' in (147) below — and this fact no doubt bears some connection to the phenomenon of reduplication, but I will not be analysing such repetitions here (see Heath 1984, pp. 340–343).

The next section, §4.2.1, presents some of the functions of reduplication, in broad strokes. The following sections all concern the form of reduplication constructions: reduplicant shapes and their conditioning in §4.2.2, interactions with segmental phonology in §4.2.4, issues with vowel-initial bases in §4.2.5, and delimiting the base within the morphological template of the verb in §4.2.3. I summarise the main points in §4.2.6.

4.2.1 Functions of reduplication

Semantically, reduplicated verbs typically indicate a greater distribution in time and/or space than their unreduplicated counterparts, or a greater quantity of object referents. Thus their interpretation ranges from repetitive or prolonged action to actions being carried out in many places ('here and there', 'all over the place', 'everywhere', etc.). Speakers tend to give English translations like, '(just) keeps doing [verb]', or '[verb] all day' For that reason, reduplication is restricted to verbs inflected for continuous aspect, not punctual aspect. In narrative discourse, reduplicated verbs may (also) serve as a kind of backgrounded clause for introducing a new event (Heath 1984, p. 343). None of these usages appear to be obligatory, though they are extremely common in texts. Given the wide variety of functions, the free translations of reduplicated verbs I provide here should be taken as very loose approximations. Note also that Wubuy differs from some other Gunwinyguan languages such as Ngalakgan, where (Root-level) reduplication is obligatory to fill out certain cells in the verbal tense-aspect inflection paradigm for some monosyllabic verb roots. In Wubuy, reduplication is never used to realise verbal inflectional categories.

These usages are in evidence in the following textual examples. In each, the Reduplicant is highlighted in bold.[2] In (144), the specification of where the action takes place — a place called Wugayuwanyung — suggests that the verb reduplication in the clause should be interpreted to indicate greater extension in time rather than space. The characters sit (i.e. stay), eat and sleep there over several nights. In (145), on the other hand, the reduplication in *niwu-maŋa-ma-ŋi* 'he was getting it' and *niwu-ŋunu-ŋu-ni* 'he was eating it' could plausibly be interpreted as indicating that the events occurred over a longer time and/or that the subject had to travel around a large area to collect the kurrajong pods. In (146), we see the extremely frequent reduplicated stem *jaŋka-jaŋki:* 'Red-go.Past2' being used as a backgrounded event to introduce

[2] As elsewhere in this thesis, I observe Heath's convention of transcribing deleted or elided vowels with an apostrophe, which makes for easier glossing; problems in the formal treatment of such vowels will be discussed in this chapter.

the verb *nu-kuckuci:ni-wukic* 'he was still holding him up'. Note that the form of this latter verb happens to involve repetition, but is not an instance of reduplication proper (it is not possible to remove the initial /kuc/). Similar examples are presented in (147)–(149), and more can be found on practically any page of the texts in Heath (1980).

(144) wini-**pu**-puɾi ana-wukajuwaɲuŋ anupani wini-**pu**-puɾi aʈapa
3duM.R-Rdp-sit.Past2 (place name) this 3duM.R-Rdp-sit.Past2 then

winiwi:-'caŋa-ŋ, wini-**ŋunu**-ŋu-ni ana-ļaku
3duM>3Neut.R-think-Past1 3sgM>3Resid-Rdp-eat-Past2 Top.Resid-honey

wini-**ja**-ja-j
3duM.R-Rdp-sleep-Past2

'The two of them (Moon and Native Cat) were staying at that place Wugayuwanyung. They thought about honey and ate some. They spent some nights there' (NMET 34.2.2)

(145) atica aʈapa nika niwu-mi-ɲ, na-ŋala:liki,
kurrajong then 3sgM 3sgM>3Neut-get-Past1 Top.Masc-green.turtle

niwu-**maŋa**-ma-ŋi(:::) aʈapa niwu-**ŋunu**-ŋu-ni
3sgM>3Neut-Rdp-get-Past2 then 3sgM>3Neut-Rdp-eat-Past2

'For his part, Green Turtle collected (edible pods of) kurrajong tree (*Brachychiton diversifolium*). He was collecting and eating them.' (NMET 42.2.2)

(146) ni-jaŋki, ni-**jaŋka**-jaŋki ni-**jaŋka**-jaŋki ni-**jaŋka**-jaŋki
3sgM.R-go.Past2 3sgM.R-Rdp-go.Past2 3sgM.R-Rdp-go.Past2 3sgM.R-Rdp-go.Past2

nu-kuckuci:ni-wukic ŋa
3sgM>3sgM.R-lift.Past2-still and

'He (the rescuer, with the dog) went along. He was still holding him up.'
(NMET 43.12.1)

(147) waː-l̪ɪrak, kaɻa wɪri-maːɳʈ̪-aŋi kaɻa kaɻa, maʈa
TOP.COLL-bluetongue hole 3COLL>3NEUT.R-make-PAST2 hole hole grass

maʈa jiː-cuŋ, maʈa, wɪri-**ka**-kaɻa-maːɳʈ̪-aŋi
grass this sort of thing grass 3COLL>3NEUT.R-RDP-hole-make-PAST2

wuː-**ɻ̪ulku**-ɻ̪ulkuɻ̪-aŋi
3COLL>3RESID.R-RDP-make.shelter-PAST2

'Blue-tongue was making a hole. It was making a hole with this kind of thing, grass. It was making a shelter.' (NMET 36.1.1)

(148) paku-miri apa wuɾaː-kamp-aː, wuɾaː-**'ka**ː-'kamp-aː
ANAPH.LOC-INSTR now 3pl>3RESID.R-roast-PAST2 3pl>3RESID.R-RED-roast-PAST2

ŋa, wuɾu-**ta**-tan-ma-ŋiːːː ŋa, wuɾu-jɪɾɪl̪ipalaw-aː,
then 3pl>3RESID.R-RED-guts-get-PAST2 then 3pl>3RESID.R-cut.down.middle.PAST2

wuɾaː-'kamp-aː wuɾu-**muntu**-muntuwa-ɲciː-niːːː ŋa
3pl>3RESID.R-RED-roast-PAST2 3pl.R-RED-gather-RECIP-PAST2 and

'Right there now, they roasted it. They roasted it, they got the guts and cut it down the middle. They all came together.'[3] (NMET 95.10)

(149) pukuni mana-waːŋ, mana-wulaŋ, ma-**wulu**-wulaŋ-cɪrjɪr-aŋi
there TOP.VEG-whatchamacallit TOP.VEG-blood 3VEG.R-RED-blood-drip-PAST2

apa wu-ku-ɾu an-uwa-kaː-'la waː-**'ɻ̪ɻ̪a**ː-'ɻ-l̪uʈ̪ʈuɭw-aː
now it (NEUT) NEUT-DIST-LOC-CENTRIP 3NEUT.R-RED-water-rise-PAST2

anaː-kuku wu-l̪**a**-l̪aɭakiː-'-ni, wu-l̪aŋaʈp-aŋi-makaː
TOP.NEUT-water 3NEUT.R-RED-lift-REFL-PAST2 3NEUT.R-emerge-PAST2-Ø[4]

wu-wajam-aŋi-makaː ana-wiɲik wu-l̪aŋaʈp-aŋi
3NEUT.R-proceed-PAST2-Ø TOP.NEUT-small 3NEUT.R-emerge-PAST2

waŋku-na-ni waŋku-mal-na-ni wu-ku-ɳi jaː-ci
3NEUT>3pl.R-see-PAST2 3NEUT>3pl.R-sky-see-PAST2 they (MASC.DU) here

wini-**wuʈ̪u**-wuʈ̪i aɾwaɻ, a-muŋkala
3Mdu.R-RED-perch.PAST2 above OBL.RESID-tree.fork

'The blood was still dripping (into the water of the billabong). (The blood aroused the pythons.) The water began to rise up like a flood. A small python came up to the surface

[3] My more literal translation, based on Heath's.

and poked its head out. It swam along, and looking up it saw them (the two boys), who (by then) were up in the fork of a tree (hanging over the billabong).' (NMET 1.1)

Nouns can be reduplicated to indicate plurality. Morphosyntactic plurality is typically restricted to nouns denoting humans (150–152), or named non-human animals such as domestic dogs and mythological beings (153). In such uses, reduplication is not obligatory, and plurality can also be indicated simply by using a human plural noun class prefix, as in (152a). Reduplication can also be used where non-human noun roots are semantically extended to humans. The standard Wubuy term for a non-Aboriginal (i.e. European) person is a semantic extension of the word for a type of spear, *ṯuɾapata* (152a); the plural is formed by reduplication, as *ṯu-ṯuɾapata* (152b). This usage appears to be quite productive: I have also encountered the word *muɾaɻpu* 'crab' reduplicated as *muɾu-muɾaɻpu*, to refer to police officers. Less commonly, nouns denoting geographical features, or adjectives describing them, can be reduplicated to indicate semantic (not morphosyntactic) plurality (154), as Heath (1984, p. 194) notes. Note that there are certain human nouns that do not reduplicate, being lexically specified to take a special human plural prefix of the form *mic-* or *miɲ-* instead.

(150) a. wuɾuc
human
'(Aboriginal) person' (NMET 20.15.4)

b. **wuɾu**-wuɾuc
Red-human
'(Aboriginal) people' (NMET 34.7.4)

(151) a. muɾuŋun-paː
(clan n.)-Du
'two Murrungun clansmen' (NMET 20.12.5)

b. **muɾu**-muɾuŋun
Red-(clan n.)
'people (Pl) of Murrungun clan' (NMET 20.11.5)

[4]The postposition /-makaː/ has many different uses; in this example it is a kind of discourse particle that attaches to a backgrounded clause in the narrative. See Heath (1984, pp. 449–450)

(152) a. waː-ʈuɾapata-wuj
PL-European-ALL
'to the Europeans/white people' (NMET 34.8.4)

b. waɾa-**ʈu**-tuɾapata
PL-RED-European
'the Europeans/white people' (NMET 165.1.3)

(153) waɾ-upa-wi-juŋ waː-ɭanʈuɾk, waː-ɭ**anʈa**-ɭanʈuɾk waɾa-maja-na,
that (COLL) TOP.COLL-dog TOP.COLL-RED-dog 3pl>3COLL.R-call-NPST2

waɾa-maja-na, waɾ-**umu**ː-'-muwac ama-kunapipi-jala,
3pl>3COLL.R-call-NPST2 TOP.PL-RED-PL-name OBL.VEG-(ritual)-ABL

waɾ-upa-wi-juŋ, ŋuɲcu a-maʈajin-ciː
that (COLL) similar OBL.NEUT-(ritual)-SIMIL

'They (people) call those dogs by names. (Names) from the Gunabibi ritual, or sometimes the Maḏayin ritual' (NMET 88.1.1)

(154) wuɾu-jaŋki wiɾi-jaɻat-aŋi, ana-**wuɻu**-wuɻuku,
3pl.R-go.PAST2 3pl>3NEUT.R-look.for-PAST2 TOP.NEUT-RED-billabong

ana-wuɻuku-**ḻama**-ḻamunkuɾ
TOP.NEUT-billabong-RED-short

'They (people) went along heading for billabongs, short (small) billabongs' (NMET 95.1.1)

There is one situation in which nominal reduplication is obligatory, namely in a part-whole possession where the possessor is of the PLural noun class. This is formed by prefixing the 'derivational' noun class prefix *uK-* and reduplicating the resulting stem; the PL inflectional noun class prefix can then be attached to this prefixed and reduplicated stem. An example is *waɾ-**umu**ː-'-muwac* 'their names' in (153): this is derived from the root *muwac* 'name', which together with the derivational prefix is *u-muwac*, and this prefixed form is the base to reduplication; vowel coalesence then applies over the reduplicant-base boundary. (Although dogs are generally of the COLL class, here they are being treated as the human PL class, as is common for named animals.)

A further function of reduplication is as a word-formation process. Verb stems can be productively reduplicated to coin new nouns. Less productively, verbs also sometimes appear

reduplicated and suffixed with a (no longer productive) nominalising suffix /-c/ or /-uc/, to produce a noun or adjective (see the discussion in Heath 1984, pp. 484–486).

(155) waɾa-**naja**-na-jiː-wuy pa-jaːɾiː ucpil-wuy
PL-RED-see-NPST2-ALL (= doctors) 2sg.IRR-go.NPST2 hospital-ALL
'Go to the clinic and see the doctors [those who look and look]' (HN; PN2020-11-08.4, –)

(156) waɾa-ŋumic dawu aː-kuku wuɾaː-**'l̪ŋa**ː-'l̪ŋu-jiː-jiɲuŋ
TOP.COLL-ant there(IMM.COLL) OBL.NEUT-freshwater 3pl.R-RED-drink-NPST2-REL
'The ant is there, in the grog' [that which they drink and drink]
(DM; PN2020-11-12, 00:25:18)[5]

Where a reduplicated verb becomes a relative clause by taking the REL suffix, and so functions as a kind of nominal expression — equivalent to 'those who…' or 'the one that…' etc. — it can be difficult to discern if the reduplication indicates plurality of referents, pluractionality or spatial distribution of events, or perhaps temporal distribution (indicating that the referents died at different times, rather than all at once).[6]

(157) wuː-**l̪aka**-l̪akaɾa-ŋaw-iː-ni-jiɲuŋ
3pl.R-RED-MULT-die-REFL-PAST2-REL
'They who all died' (NMET 47.9.2)

We have seen examples in this section featuring reduplicants of various shapes — some monosyllabic, some disyllabic — but, importantly for our purposes, there is no indication that these different formal patterns are associated with different functions. Rather, whatever interpretations are available to reduplicated words of a given word class generally will also be available to reduplicated words with any given reduplicant shape, subject to the semantic properties of the base itself. As we shall see in the next section, the basic formal alternation between monosyllabic and disyllabic reduplicants is conditioned not by semantics but by initial segment sonority.

[5] The speaker explained immediately after this that [wuɾaːl̪ŋaːl̪ŋujiːjiɲuŋ] was equivalent to [ana-gˌɻog] (Eng. *grog*) — /kuːku/ 'water' is sometimes used euphemistically to mean 'alcohol'.

[6] Evans (2003) notes that reduplication of intransitives in Bininj Kunwok can be used to indicate plurality of subjects (or perhaps emphasise plurality, if the plural is otherwise marked). The same may perhaps be in evidence here.

4.2.2 Sonorant-initial disyllabic reduplication and stop-initial monosyllabic reduplication

As a way in to the formal characteristics of the data, it will be helpful to observe the range of variation in the *shape* of the reduplicant and its *position* with respect to the base, and what factors condition the variation.

The position of the reduplicant is relatively invariant. In the most productive pattern, the reduplicant is positioned immediately following any inflectional prefix, and copies the material to its right, starting at the leftmost segment. There are some exceptions to this, however, as will be discussed in §4.2.3.

The leftmost segment of the base, in nearly all cases, determines the shape of the reduplicant. Where the initial segment of the stem is a sonorant consonant, the reduplicant is a minimal disyllable — that is, a $C_1V_1C_2{}^*V_2$ sequence, where the medial consonant can be a cluster of any length. Both vowels of the reduplicant are short, regardless of the length of the corresponding vowels in the base, and the second vowel must have the same quality as the first: $V_1 = V_2$. Thus the stem in (158a), *wulṯ-aŋi* 'cut-PAST2', reduplicates as (158b) *wulṯu-wulṯ-aŋi*, where the reduplicant is *wulṯu-*. In the base, the first and second vowels are /u/ and /a/, respectively; in the reduplicant they are both /u/.

(158) a. wini-wulṯ-aŋi
3duM>RESID.R-cut-PAST2
'They two (men) cut it (RESID)' (NMET 17.1.3)

b. wini-**wulṯu**-wulṯ-aŋi
3duM>RESID.R-RED-cut-PAST2
'They two (men) cut it (RESID)' (NMET 17.1.4)

In (159b), it is instead a difference in vowel quantity that is ignored in the reduplicant. The second vowel of the stem *wanŋ-a:* 'danced' is long, but in the reduplicant, *wanŋa-*, its corresponding vowel is short.

(159) a. wuɾu-wanŋ-a:
3pl.R-dance-PAST2
'They danced' (NMET 9.8.1)

b. wuɾu-**wanŋa**-wanŋ-a:
3pl.R-RED-dance-PAST2
'They danced' (NMET 9.6.1)

The fact that the two vowels in the reduplicant must be identical, even where the corresponding vowels in the base may differ in quality or quantity, provides justification for the claim that reduplication is prefixing rather than infixing — we should expect that, absent some convincing evidence to the contrary, the base is that part which stays the same between the reduplicated and unreduplicated words.

In order to demonstrate that the disyllabic reduplication pattern is available to sonorant-initial stems generally, and is not an accidental artifact of the few examples so far presented, further examples are presented below. Stems beginning with all sonorant consonants are exemplified, except /ɾ/, /l/, /ɳ/ and /ɲ/, which are very rare stem-initially.

(160) a. ni-wajam-aŋi
3sgM.R-proceed-Past2
'He went along' (NMET 1.5.5)

b. ni-**waja**-wajam-aŋi
3sgM.R-Red-proceed-Past2
'He went (kept going) along' (NMET 1.5.5)

(161) a. ŋu-waɭaɾl̪-ick-aɲ
3sgM>3Fem.R-lie-Caus-Past1
'He flipped her over [the mother python]' (NMET 2.6.4)

b. ŋu-**waɭa**-waɭaɾl̪-ick-a:
3sgM>3Fem.R-Red-lie-Caus-Past2
'He went (kept going) along' (NMET 1.6.10)

(162) a. waɾa-juɻ-aŋi
3pl>3pl.R-transport-Past2
'They transported them' (NMET 50.5.6)

b. waɾa-**juɻu**-juɻ-aŋi
3pl>3pl.R-Red-transport-Past2
'They transported them' (NMET 40.1.3)

(163) a. ŋi-jamp-i:na
3sgF.R-Red-speak-Npst2
'She speaks' (NMET 10.3.3)

b. ŋi-**jampa**-jamp-iːna
3sgF.R-Red-speak.Npst2
'She keeps talking' (HN; PN2020-11-04 00:12:27)

(164) a. ŋunu-jaɾpi-ni
3Fem>3sgM.R-wash-Past2
'It (Fem) washed him' (NMET 13.17.1)

b. ŋunu-**jaɾpa**-jaɾpi-ni
3Fem>3sgM.R-Red-wash-Past2
'It (Fem) washed him' (NMET 13.17.2)

(165) a. na-jiwaŋku-waː
M-old.man-Du
'the two old men' (NMET 63.1.5)

b. waɾa-**jiwi**-jiwaŋku
Pl-Red-old.man
'the old men' (NMET 167.10.4)

(166) a. wiː-ji-w-ini
3pl>3Neut-scorch-Aug-Past2
'They scorched it' (NMET 53.1.6)

b. wiː-**jiwi**-ji-w-ini
3pl>3Neut-Red-scorch-Aug-Past2
'They scorched it' (NMET 53.1.7)

(167) a. ma-jiwuɲ-maː
3Veg.R-cool-Inch.Past2
'It (Veg) became cool' (118.3.2)

b. wuː-**jiwi**-jiwuɲ-maː
3Coll.R-cool-Inch.Past2
'It (Veg) became cool' (125.2.1)

(168) a. waɲi-ɻa-ni
3Fem>3pl.R-spear-Past2
'They (mosquitoes) speared them' (NMET 33.4.2)

b. waɲi-**ɻana**-ɻa-ni(:::)[8]
3Fem>3pl.R-Red-spear-Past2
'They (mosquitoes) speared (all of) them' (NMET 33.4.3)

(169) a. wu:-ɻa-ji:
3pl>3Resid.R-spear-Past2
'They spear them (marine mammals)' (NMET 122.1.4)

b. wu:-**ɻaja**-ɻa-ji:
3pl>3Resid.R-Red-spear-Past2
'They keep spearing them (small rays)' (NMET 109.1.3)

(170) a. wunu-ɻaku-ni
3pl>3sgM.R-seek-Past2
'They looked for him' (NMET 43.8.2)

b. wunu-**ɻaka**-ɻaku-ni
3pl>3sgM.R-Red-seek-Past2
'They were looking around for him' (NMET 43.8.1)

(171) a. wu:-ɻuku-ni
1pl>3Resid.R-weep-Past2
'They wept' (NMET 41.17.1)

b. ŋi-**ɻuku**-ɻuku-ni
1pl>3Resid.R-Red-weep-Past2
'She wept' (NMET 71.5.3)

[8]Here and elsewhere, IPA length symbols inside parentheses — e.g., '(:::)' — indicate stylistic, not phonemic, lengthening of the phrase-final syllable with a flat, high pitch, as transcribed by Heath (1980), and described in Heath (1984, p. 602). This is widely reported in Australian languages; see Mailhammer & Caudal (2019), and Chapter 6 of this thesis.

(172) a. wiː-ɻuc-aŋi
3pl>3Neut.R-scrape-Past2
'They scraped it [wood]' (NMET 127.1.2)

b. wiː-**ɻucu**-ɻuc-aŋi
3pl>3Neut.R-scrape-Past2
'They used to scrape it [wood]' (NMET 127.1.2)

(173) a. ma-ɻiɾ-aː
3Veg.R-dry-Npst2
'It dries off' (NMET 115.4.1)

b. ma-**ɻiɾi**-ɻiɾ-aː(ːː)
3Veg.R-Red-dry-Npst2
'It keeps drying off' (NMET 115.4.1)

(174) a. nuː-ɻulkuḻ-aŋi
1pl>3Resid.R-make.shelter-Past2
'We made a shelter' (NMET 145.1.5)

b. nuː-**ɻulku**-ɻulkuḻ-aŋi
1pl>3Resid.R-Red-make.shelter-Past2
'We made a shelter' (NMET 145.1.6)

(175) a. nuː-ḻ-aŋi
1pl>3Resid.R-chop-Past2
'We chopped it (wood)' (NMET 143.1.3)

b. nuɾu-**ḻaŋa**-ḻ-aŋi
1pl>3Resid.R-Red-chop-Past2
'We chopped it (wood)' (NMET 143.1.3)

(176) a. wini-ḻaɾiː-ni
3duM>3Resid.R-untie-Past2
'They untied it (fish trap)' (NMET 108.6.6)

b. wini-**ḻaɾa**-ḻaɾiː-ni
3duM>3Resid.R-Red-untie-Past2
'They untied it (fish trap)' (NMET 108.6.7)

(177) a. ŋu-l̪am-aɲ
3sgM>3sgF.R-beget-Past1
'He begat her' (NMET 14.3.3)

b. nampi-**l̪ama**-l̪am-aː-jiɲuŋ
3pl>1pl.R-Red-beget-Past2-Rel
'We whom they begat' (NMET 114.11.6)

(178) a. wiɾima-l̪alwuɭ̪-iː
3pl>3Veg.R-soak-Npst2
'They soak it (Veg)' (NMET 156.2.3)

b. wiɾima-**l̪alwa**-l̪alwuɭ̪-iː[10]
3pl>3Veg.R-Red-soak-Npst2
'They soak it (Veg)' (NMET 156.2.3)

(179) a. wu-ɭaɭm-aɲ
3Resid.R-dawn-Past1
'It dawned' (NMET 4.5.4)

b. wu-ɭ**a**ɭ**ma**-ɭaɭm-aː
3Resid.R-Red-dawn-Past2
'It was dawning' (NMET 4.5.5)

(180) a. waːɻi ampiː-ɭiɾpi-j-aː waɾa-miɲ-ŋampaɻa
not 3pl>3Neut.Irr-fun-Fact-Past2 the elders
'The elders did not take [such words] lightly' (NMET 47.7.3)

b. ŋiŋku-ɭ**iɾpi**-ɭiɾpi-ma-ŋi
3Resid>3sgF.R-Red-fun-Inch-Past2
'It (buffalo) had fun with her' (NMET 44.1.6)

[10] This reduplicated form is also used in text 127 and 132, each by a different speaker from 156. In text 115 (yet a different speaker), there are two tokens of an alternative, *ma-l̪ala-l̪alwuɭ̪-aː* (NMET 115.5.2), with /w/ deleted just in the reduplicant. This verb is in a different inflection class, however, and is interpreted as intransitive (unaccusative): 'it (Veg) soaked'. It is unclear whether /w/ can be deleted from the reduplicant when reduplicating the stems in (178) as well, or indeed from reduplicant-medial clusters generally.

(181) a. wini-maṯant-iŋ
3duM.R-paddle-PAST1
'They (two men) paddled' (NMET 14.17.2)

b. wini-**maṯa**-maṯant-aŋi
3duM.R-RED-paddle-PAST2
'They (two men) paddled' (NMET 14.17.1)

(182) a. wani-na-ni
3sgM>3pl.R-see-PAST2
'He saw them' (NMET 12.10.1)

b. wani-**nana**-na-ni
3sgM>3pl.R-RED-see-PAST2
'He saw them' (NMET 12.10.3)

(183) a. wiɾiŋa-na-ji:
3pl>3FEM.R-see-NPST2
'They see it (Fem)' (NMET 106.1.3)

b. wiɾima-**naja**-na-ji:
3pl>3VEG.R-RED-see-NPST2
'They see it (Veg)' (NMET 115.2.1)

(184) a. ma-nak-i:ni
3VEG.R-burn-PAST2
'It burned' (NMET 119.3.6)

b. ma-**naka**-nak-i:ni(:::)
3VEG.R-RED-burn-PAST2
'It burned' (NMET 118.2.2)

(185) a. wa:-naɾ-aŋi
3pl>3pl.R-wait-PAST2
'They waited for them' (NMET 55.11.1)

b. wa:-**naɾa**-naɾ-aŋi
3pl>3pl.R-RED-wait-PAST2
'They waited for them' (NMET 55.11.2)

(186) a. ni-ŋaʈukump-iːni
3sgM.R-fish-Past2
'He fished (with line and hook)' (NMET 15.5.7)

b. ni-**ŋaʈa**-ŋaʈukump-iːni(:::)
3sgM.R-Red-fish-Past2
'He fished (with line and hook)' (NMET 15.5.3)

(187) a. niɾima-ŋaɾkiw-aː
1pl>3Veg.R-pull.out-Past2
'We pulled it out' (NMET 157.5.5)

b. niɾima-**ŋaɾka**-ŋaɾkiw-aː
1pl>3Veg.R-Red-pull.out-Past2
'We pulled it out' (NMET 157.5.6)

If the stem-initial segment is instead a stop, the reduplicant will be a light (CV) monosyllable. In (188), the stem *puɾi* 'sat' reduplicates as *pu-puɾi*. As with the disyllabic reduplicant, the vowel in the stop-initial reduplicant must be short even if the corresponding vowel in the base is long, as in (189) and (190). Further examples of the stop-initial, disyllabic pattern can be observed below; each stop phoneme is represented with at least one example.

(188) a. wini-puɾ-i
3duM.R-sit-Past2
'Those two (men) sat' (NMET 34.1.1)

b. wini-**pu**-puɾ-i
3duM.R-Red-sit-Past2
'Those two (men) sat' (NMET 34.2.2)

(189) a. niwu-ʈuːɭw-iːni
3sgM>3Neut.R-sing-Past2
'He sang for it' (NMET 19.8.6)

b. niwu-**ʈu**-ʈuːɭw-iːni
3sgM>3Neut.R-Red-sing-Past2
'He sang for it' (NMET 19.8.7)

(190) a. ni-ʈaːl̪-aɲi
3sgM.R-sing-Past2
'He sang (ritual song)' (NMET 29.2.4)

b. ni-**ʈa**-ʈaːl̪-aɲi
3sgM.R-Red-sing-Past2
'He sang (ritual song)' (NMET 29.1.3)

(191) a. ŋaŋu-paɻawut-aɲi(:::)
1sg>3Fem.R-throw-Past2
'I threw it (Fem) away' (NMET 111.4.3)

b. ŋaŋu-**pa**-paɻawut-aɲi
1sg>3Fem.R-Red-throw-Past2
'I threw it (Fem) away' (NMET 111.4.3)

(192) a. ŋaŋu-paḻ̪-uː
1sg>3Fem.R-cut.up-Past2
'I cut it (Fem) up' (NMET 111.4.4)

b. ŋaŋu-**pa**-paḻ̪-uː
1sg>3Fem.R-Red-cut.up-Past2
'I cut it (Fem) up' (NMET 111.4.4)

(193) a. ma-pilja-ɻa
3Veg.R-tilt-Npst2
'(The boat) is on its side' (YN, PN2017-07-07.2 00:01:06)

b. ma-**pi**-pilja-ɻa
3Veg.R-Red-tilt-Npst2
'(The boat) is on its side' (YN, PN2017-07-07.2 00:01:07)

(194) a. pa-puŋuɾuʈ̪a-ŋ
2sg.Irr-get.drunk-Npst1
'You'll get drunk' (HN, PN2020-11-04.2 00:21:16)

b. ni-**pu**-puŋuɾuʈ̪-iːni-wala ŋa ni-ŋaw-iɲ
3sgM.R-Red-get.drunk-Past2-Abl and 3sgM.R-die-Past1
'He drank too much, and he died' (HN, PN2020-11-04.2 00:20:36)

(195) a. wuɾu-t̪uɾma-na
3pl>3Resid.R-mould-Npst2
'They mould it (Resid) with their hands' (NMET 106.3.2)

b. wuɾu-**t̪u**-t̪uɾma-na:::
3pl>3Resid.R-Red-mould-Npst2
'They mould it (Resid) with their hands' (NMET 106.3.2)

(196) a. ni-t̪aːɾa-kajiki-ni
3sgM.R-beard-play-Past2
'He was playing with his beard' (DM; PN2020-11-12, 00:30:59)

b. ni-**t̪a**-t̪aːɾa-kajiki-ni
3sgM.R-Red-beard-play-Past2
'He's playing with his beard' (DM; PN2020-11-12, 00:31:59)

(197) a. wuɾu-ta-ŋi
3pl.R-defecate-Past2
'They defecated' (NMET 132.4.1)

b. wiŋi-**ta**-ta-ŋi-maka:::
3duF.R-Red-defecate-Past2
'They two (women) defecated' (NMET 11.3.2)

(198) a. pawu-cuɻa-ŋ
2sg>3Neut.Irr-push-Npst3
'Push it! (Resid)' (NMET 162.8.2)

b. nima-**cu**-cuɻ-aŋi
3sgM>3Veg.R-Red-push-Past2
'He was pushing it (Veg)' (NMET 12.8.4)

(199) a. nuɾu-kuʈ-aŋi
1pl>3Resid.R-catch.fish-Past2
'We caught fish' (NMET 166.15.3)

b. nuɾu-**ku**-kuʈ-aŋi
1pl>3Resid.R-Red-catch.fish-Past2
'We caught fish' (NMET 162.17.1)

(200) a. ŋu-kulmuŋ-kaːj-uː
3sgM>3Fem.R-belly-cut-Past2
'He cut up her belly' (NMET 1.7.1)

b. ŋu-**ku**-kulmuŋ-kaːj-uː
3sgM>3Fem.R-Red-belly-cut-Past2
'He cut up (all) her belly' (NMET 1.7.1)

We can already see from these examples that the internal composition of the base has no effect on reduplicant size. As Table 4.1 shows, the base can contain long or short vowels in any position, it can be any number of syllables in length, and these syllables can be simple CV syllables or complex CVCC ones. None of these factors has any influence on the shape of the reduplicant: only the sonority of the initial segment is relevant. This is an important fact to note, as it shows that, although the effect of the alternation is prosodic, the factors that condition the alternation cannot be explained in terms of prosody.

The characterisation provided so far accounts for almost all consonant-initial stems (vowel-initial stems are treated in §4.2.5 below). It also extends to loanwords, as the following examples attest. In (201), the base, derived from Kriol *.ɻiŋim*, is continuant-initial, and therefore conditions a disyllabic reduplicant.[11] The base in (202) is, by contrast, a stop-initial borrowed verb, and conditions a monosyllabic reduplicant. Examples of nouns follow in (203) and (204). The latter of these, strikingly, has an initial cluster /kɻ/ which, presumably owing to the initial stop, is reduplicated with a CCV monosyllable (complex onsets do not occur outside of loanwords in Wubuy).

(201) a. naɳ-ʈiŋim-ʈ̪-ika-ɲc-ina
/naN-ɻiŋim-ʈ̪-icka-ɲc-ina/
12du.Irr-phone-Aux-Caus-Recip-Npst2
'You and I are going to ring each other up' (HN; PN2020-11-08.4 00:23:20)

b. waɾa-**ɻiŋi**-ɻiŋim-ʈ̪-iː (< Kriol *.ɻiŋim* 'ring up (on phone)')
/waɾa-ɻiŋi-ɻiŋim-ʈ̪-iː/
3pl>3pl.R-Red-phone-Aux-Npst2
'They keep ringing/are always ringing them up' (Baker & Harvey in press, p. 256)

[11] The stem in (201a) shows Hardening also, as befits stems that have an initial continuant when also word-initial. Note that here and elsewhere in this chapter, forms given between slashes represent not the raw 'underlying representation' (which would contain no segmental content for the reduplicant) but rather the immediate output of reduplication, before undergoing further changes such as Hardening.

ex. no.	reduplicant	base	syllables in base	structure
(188)	pu-	puɾi	2	CV.CV
(192)	pa-	paɭ̈uː	2	CV.CVː
(199)	cu-	cuɻaɲi	3	CV.CV.CV
(195)	ʈ̈u-	ʈ̈uɾmana	3	CVC.CV.CV
(190)	ʈa-	ʈaːɭaɲi	3	CVː.CV.CV
(189)	ʈu-	ʈuːlwiːni	3	CV:C.CVː.CV
(191)	pa-	paɻawutaɲi	5	CV.CV.CV.CV.CV
(196)	ʈ̈a-	ʈ̈aːɾakajikini	6	CVː.CV.CV.CV.CV.CV
(168)	ɻana-	ɻani	2	CV.CV
(150)	wuɾu-	wuɾuc	2	CV.CVC
(169)	ɻaja-	ɻajiː	2	CV.CVː
(179)	ɭaɭma-	ɭaɭmaː	2	CVC.CVː
(159)	wanŋa-	wanŋaː	2	CVC.CVː
(153)	ɭanʈ̈a-	ɭanʈ̈uɾk	2	CVC.CVCC
(162)	juɻu-	juɻaɲi	3	CV.CV.CV
(163)	jampa-	jampina	3	CVC.CV.CV
(164)	jaɾpa-	jaɾpini	3	CVC.CV.CV
(184)	naka-	nakiːni	3	CV.CVː.CV
(176)	ɭ̈aɾa-	ɭ̈aɾiːni	3	CV.CVː.CV
(187)	ŋaɾka-	ŋaɾkiwaː	3	CVC.CV.CVː
(160)	waja-	wajamaɲi	4	CV.CV.CV.CV
(181)	maʈ̈a-	maʈ̈antaɲi	4	CV.CVC.CV.CV
(180)	ɭiɾpi-	ɭiɾpimaɲi	4	CVC.CV.CV.CV
(186)	ŋaʈa-	ŋaʈukumpiːni	5	CV.CV.CVC.CVː.CV
(229)	ɭ̈acŋa-	ɭ̈acŋupaɻawutaɲi	7	CVC.CV.CV.CV.CV.CV.CV

Table 4.1: Reduplicant shape, sonority of the base-initial segment, and skeletal composition of the base.

(202) a. nuɾu-pajan-t̪-aŋi (< Kriol *pajan* 'buy')
1pl>Resid.R–buy-Aux-Past2
'We bought it' (NMET 166.2)

b. ŋa-**pa**-pajan-t̪-iː
1sg>Resid.R-Red-buy-Aux-Npst2
'I keep buying (things)/I'm always buying (things)'
(Baker & Harvey in press, p. 256)

(203) waɾa-**ci**-cicta (< English *sister*)
Pl-Red-nurse
'the nurses' (NMET 139.1.4)

(204) waɾa-**kɻi**-kɻictan
Pl-Red-Christian
'the Christians' (NMET 162.16.4)

There are, nonetheless, a handful of exceptions to the main pattern, as also noted in Heath (1984). Some of these involve a simple stem in which the first two syllables of the root are already identical; reduplication involves copying only one of these syllables, not both, even if the initial segment is a sonorant. The reduplicated stem *-l̪a-l̪al̪ak-iːni* appeared in (149) above; I show it separately here, with its unreduplicated counterpart for comparison. As we see, it takes a CV reduplicant despite its initial sonorant. Another example of a sonorant-initial but CV reduplicant is given in (206). In both of these examples, the two identical syllables in the base have a CV shape, and so does the reduplicant. The reduplicated stem *wuɾu-mal̪-mal̪mal̪-aŋi* in (207b) shows that if the identical syllables in the base have a CVC shape, the reduplicant has that shape also. The examples in (208) and (209) also have an unexpected CVC reduplicant, but the bases are instead stop-initial. While it is notable that each of these exceptions have a repeated initial syllable in their unreduplicated form, the number of roots falling into this category is very small, so it is difficult to say whether this minor pattern is truly general.

(205) a. ni-l̪al̪ak-iɲ
3Neut.R-Red-lift.Refl-Past1
'He got up' (NMET 5.1.4)

b. wu-l̪**a**-l̪al̪ak-iːni
3Neut.R-Red-lift.Refl-Past2
'It kept rising up' (NMET 1.1.5)

(206) a. wuŋunu-mamaɾ-iɲ
3duF>3sgM.R-hide-PAST1
'They hid it from him' (NMET 13.8.3)

b. wuŋunu-**ma**-mamaɾ-aŋi
3duF>3sgM.R-RED-hide-PAST2
'They kept it hidden from him' (NMET 13.8.4)

(207) a. wuɾu-maɭmaɭ-aŋi
3pl.R-dance.circumcision-PAST2
'They danced the circumcision dance' (NMET 55.4.5)

b. wuɾu-**ma**ɭ-maɭmaɭ-aŋi
3pl.R-RED-dance.circumcision-PAST2
'They danced the circumcision dance' (NMET 55.4.5)

(208) a. wunu-kuckuc-iːni
3pl>3sgM.R-lift-PAST2
'They lifted him up' (NMET 44.1.3)

b. ŋiŋku-**kuc**-kuckuc-iːni
3RESID>3sgF.R-RED-lift-PAST2
'It (buffalo) lifted her up [playing with her]' (NMET 44.1.3)

(209) a. wuɾu-cuɾcuɾk-iːni
3pl>3RESID.R-push-PAST2
'They pushed it' (NMET 110.2.3)

b. wuɾu-**cuɾ**-cuɾcuɾk-iːni
3pl>3RESID.R-RED-push-PAST2
'They pushed it' (NMET 110.2.5)

Interestingly, two formally very similar synonyms for 'make' exceptionally take a monosyllabic reduplicant, although beginning with the sonorant /m/.

(210) a. nima-maːɳʈ-aɲi
3sgM>VEG.R-make-PAST2
'He made it [rope]' (NMET 4.4.3)

b. nima-**ma**-maːɳʈ-aɲi
3sgM>VEG.R-RED-make-PAST2
'He was making it [rope]' (NMET 4.4.3)

(211) a. ŋani-maːcka-na
1sg>MASC.R-make-PAST2
'I will make it [his camp]' (NMET 62.2.2)

b. ŋani-**ma**-maːcka-na
1sg>MASC.R-RED-make-PAST2
'I will make it [his camp]' (NMET 62.2.2)

There are a handful of other exceptions that appear in the texts. (212a) has an unexpected CV reduplicant, despite the initial segment of the base being /ɻ/; however, the expected form (212b) also occurs in the texts.

(212) a. nuː-**ɻi**-ɻi-ɲc-iːni
1pl.R-RED-spear-REFL-PAST2
'We used to spear each other' (NMET 69.1.2)

b. nuɾu-**ɻiɲci**-ɻi-ɲc-iːni
1pl.R-RED-spear-REFL-PAST2
'We used to spear each other' (NMET 161.4.3)

Heath (1984) also documents some stop-initial roots with exceptional reduplication patterns, one of them being *ʈamunkuɾ* 'short', which in his work is recorded with a special reduplicated form *ʈama-ɭamunkuɾ* 'short.PL'. In this exceptional pattern, the reduplicant is stop-initial but disyllabic, and the base-initial segment has been lenited to a continuant. My consultants, however, told me the forms in (213a,b) immediately upon being asked to translate 'short', without me trying to elicit a plural; they recognised and were happy to produce (213c) when I prompted them with it, but told me it meant 'dwarves' (*ŋuɲcu a*-dwarf-*jiː* 'same like a dwarf'). I presume this more lexicalised meaning is not the only possible one, since Heath documents examples where the referent is *wuɻuku* 'billabong' and *ɭalpic* 'fan palm'. Examples like (213b) may suggest that the productive pattern described above is strong enough to overgenerate, at

least among this somewhat younger cohort of speakers. Similar patterns to the *ṯ̪ama-l̪̈amunkuɾ* example occur in Ngandi (see §4.3.2), and Yolŋu languages. There are just two other Wubuy roots claimed to have this behaviour in Heath's documentation, both of them verbs: *paɾa* 'to become bent/crooked' and *pilwila* 'to sway'. They are not really comparable to *ṯ̪amunkuɾ*, however. The 'reduplicated' forms Heath has in mind seem to be the adjectival forms *paɾawaɾuc* 'crooked' and *pilwilwiluc* 'weak, soft' — the effect of reduplication in these forms is not morphosyntactically or semantically equivalent to the productive verbal reduplication pattern.

(213) a. ṯ̪amunkuɾ-juŋ
short-SG
'short (person)' (MM; PN2017-07-06.1, 00:46:48)

b. **ṯ̪a**-ṯ̪amunkuɾ
RED-short
'short (people)' (MM; PN2017-07-06.1, 00:46:50)

c. **ṯ̪ama**-l̪̈amunkuɾ
RED-short
'short (people); dwarves' (MM; PN2017-07-06.1, 00:48:32)

All told, the exceptions to the basic generalisation that stop-initial bases condition a monosyllabic reduplicant shape, and sonorant-initial bases a disyllabic one, are surprisingly few. In the next section, I consider the position of the reduplicant, which is normally immediately to the right of the inflectional prefix, but does occasionally appear inside complex stems.

4.2.3 Reduplicant position

In the canonical reduplication pattern, the reduplicant is positioned to the left of the stem, and copies from the left edge of the stem rightwards, as we've seen. Just as the stem can be morphologically complex, the base to reduplication can be also. In (214), each of the verbs contains the same inflected root *paɾk-a:*. When this is the only element of the stem, as in (214a), the reduplicated form copies from this edge to form *pa-paɾk-a:* in (214b) — note the monosyllabic reduplicant, conditioned by the initial stop. When the stem is complex, as with *jika-ŋu-paɾk-a:* in (214c), reduplication takes this entire stem as the base for copying, attaching to its left, as in (214d), *jiki-jika-ŋu-paɾk-a:* — and note that this time, the reduplicant must be disyllabic due to the initial sonorant.

(214) a. ŋawu-paɾk-a: ana-ŋuɻa
/niwu-paɾk-a:/ /ana-ŋuɻa/
3sgM>NEUT.R-blow.fire-PAST2 NCANA-fire
'I blew on the fire (to make it bigger)' (JM, PN2020-11-03, 00:07:27)

b. ŋawu-**pa**-paɾk-a: ana-ŋuɻa
/ŋawu-**pa**-paɾk-a:/ /ana-ŋuɻa/
1sg>3NEUT.R-RED-blow.fire-PAST2 TOP.NEUT-fire
'I kept blowing on the fire' (MM, PN2017-07-07.2, 00.03.12)

c. ŋawu-jika-ŋu-paɾk-a:
/ŋawu-jika-ŋu-paɾk-a:/
1sg>3NEUT.R-fire-blow.fire-PAST2
'I blew on the fire (to make it bigger)' (YN, PN2017-07-07.2, 00:00:15)

d. ŋawu-**jiki**-jika-ŋu-paɾk-a:
/ŋawu-**jiki**-jika-ŋu-paɾk-a:/
1sg>3NEUT.R-RED-fire-Ø-blow.fire-PAST2
'I kept blowing on the fire' (YN, PN2017-07-07.2, 00:01:08)

Morphological boundaries are ignored in this process. In (215a), the stem consists of the inflected verb root *waɭk-a:* 'pounded', reduplicated as *waɭka-waɭk-a:* in (215b). The stem in (215c) is headed by the same verb root but features a monosyllabic incorporated noun. In order to be disyllabic, the reduplicant in (215d) must copy across the boundary between incorporated noun and verb root. Further examples of copying across boundaries follow.

(215) a. wiŋiwu-waɭk-a:
/wiŋiwu-w$_2$aɭk-a:/
3duF>3NEUT.R-pound-PAST2
'The two women were pounding it (yam; NEUT)' (NMET 126.2.3)

b. wiŋima-**waɭka**-waɭk-a:
/wiŋima-w$_2$aɭka-w$_2$aɭk-a:/
3duF>3VEG.R-RED-pound-PAST2
'The two women were pounding it (sedge sp.; VEG)' (NMET 126.2.3)

c. wiŋiwu-ɻaŋ-paɭk-aː
/wiŋiwu-ɻaŋ-w$_2$aɭk-aː/
3duF>3Neut.R-mush-pound-Past2
'The two women were pounding the mush (of yam; Neut)' (NMET 126.2.3)

d. wiŋiwu-**ɻaŋpa**-ɻaŋ-paɭk-aː
/wiŋiwu-ɻaŋw$_2$a-ɻaŋ-w$_2$aɭk-aː/
3duF>3Neut.R-Red-mush-pound-Past2
'The two women were pounding the mush (of yam; Neut)' (NMET 126.2.4)

(216) a. wuɾaː-'jacaɾ-aŋi[13]
/wuɾu-w$_1$ajacaɾ-aŋi/
3pl>3Resid.R-lay-Past2
'They spread things out' (NMET 163.4.4)

b. wuɾaː-**'ja**ː-'jacaɾ-aŋi
/wuɾu-w$_1$aja-w$_1$ajacaɾ-aŋi/
3pl>3Resid.R-Resid-lay-Past2
'They spread (many) things out' (NMET 163.4.3)

c. wiɾima-mun-kajacaɾ-iŋ
/wiɾima-mun-w$_1$ajacaɾ-iŋ/
3pl>3Veg.R-foot-lay-Past1
'They spread things out under the tyre [i.e. foot, of the car = Veg]' (NMET 163.4.5)

d. wiɾima-**munku**-mun-kajacaɾ-aŋi
/wiɾima-munw$_1$u-mun-w$_1$ajacaɾ-aŋi/
3pl>3Resid.R-Resid-lay-Past2
'They spread (many) things out under the tyre' (NMET 163.4.4)

(217) a. ŋaŋku-l̪ac-kulʈ̪-iː
/ŋaŋku-l̪ac-w$_1$ulʈ̪-iː/
1sg>3Neut.Irr-firewood-chop-Npst2
'I'm going to chop firewood' (HN, PN2020-11-11.1, 00:27:53)

[13]The deletion of w$_1$ is discussed below in §4.2.5.

b. ŋawu-**l̪acka**-l̪ac-kulṯ-aŋi
/ŋawu-l̪acw$_1$a-l̪ac-w$_1$ulṯ-aŋi/
1sg>3Neut.Irr-Red-firewood-chop-Past2
'I was chopping firewood all day' (HN, PN2020-11-11.1, 00:28:25)

(218) a. wunu-ɻa-watcaŋi
/wunu-ɻa-w$_2$atcaŋi/
3pl>3M.R-tooth-hit.Past2
'They hit him in the teeth' (HN, PN2020-10-30.3, 00:28:46)

b. wunu-**ɻawa**-ɻa-watcaŋi
/wunu-ɻawa-ɻa-w$_2$atcaŋi/
3pl>3M.R-Red-tooth-hit.Past2
'They hit him in the teeth'[14] (HN, PN2020-10-30.3, 00:28:46)

The reduplicant appears in this position in the vast majority of cases, but not without exception. The forms in (219) were offered spontaneously by two speakers in a single elicitation session, after I had elicited a different word with a similar meaning. In (219), reduplication begins at the left edge of the root *anṯuɾk* 'straight', rather than the entire stem *jinak-anṯuɾk* 'head-straight', as would normally be expected. Heath (1984) also reports cases of 'internal' reduplication, and suggests that it appears to be more common for incorporation constructions headed by adjectives, as in (154) above. (The internally reduplicated (219c) example also contains an adjective root, but it is used there to derive a Factitive/Reflexive verb.)

(219) a. waːɻi amaː-'anṯuɾ-k-iːni
/waːɻi/ /ama-anṯuɾk-w$_1$a-i-ni/
Neg 3Veg.Irr-straight-Fact-Refl-Past2
'It (the canoe) didn't go straight' (MM, PN2017-07-07, 00:15:03)

b. ama-jina-kanṯuɾ-k-iːni[16] waːɻi
/ama-jinak-anṯuɾk-w$_1$a-i-ni/ /waːɻi/
3Veg.Irr-head-straight-Fact-Refl-Past2 Neg
'The prow (of the canoe) wasn't straight' (YN, PN2017-07-07, 00:15:09)

[16]The *kanṯuɾk* allomorph is due to the workings of Stop Insertion, which conditions deletion of the /k/ in *jinak*, which is why it appears instead as *jina*; see Chapter 5.

c. ma-jina-**kanʈa**:-’nʈuɾ-k-iːni
/ma-jinak-**anʈa**-anʈuɾk-w$_1$a-i-ni/
3Veg.R-head-Red-straight-Fact-Refl-Past2
‘The prow (of the canoe) was going along straight’ (YN, PN2017-07-07, 00:19:20)

(220) a. niː-’mpat-aɲi
/ni-ampat-aɲi/
3Masc.R-burst-Past2
‘It [Masc] burst’ (NMET 98.2.1)

b. wuɾaː-’ɾaː-**’mpa**ː-’mpat-aɲi
/wuɾu-aɾa-**ampa**-ampat-aɲi/
3pl.R-Mult-Red-burst-Past2
‘They all burst’ (NMET 16.22.1)

(221) a. waɾa-juɭu-juɭ-aːni
/waɾa-juɭu-juɭ-aːni/
3pl>3Coll.R-Red-stir-Past2
‘They stirred it’ (NMET 125.9.1)

b. waɾa-ŋaɾaɳ-**cu**ɭ**u**-juɭ-aːni
/waɾa-ŋaɾaɳ-**ju**ɭ**u**-juɭ-aːni/
3pl>3Coll.R-Mult-Red-stir-Past2
‘They stirred it all’ (NMET 125.9.2)

(222) waː-ɻaː-**’la**ː-’laːti-jiɲuŋ
/waɾa-ɻa-**ala**-alaːti-jiɲuŋ/
Top.Pl-tooth-Red-bad-Rel
‘bad-toothed ones’ (NMET 122.4.3)

(223) mana-pa-**ʈuŋku**-ɻuŋkal-iɲuŋ pakaɾaːk
/mana-paK-**ɻuŋku**-ɻuŋkal-iɲuŋ/ pakaɾaːk
Top.Veg-eye-Red-big-Rel cycad nut
‘the big nuts (from among the cycad nuts)’ (NMET 116.2.4)

In a few cases, both a stem-internal constituent and the stem itself are reduplicated. Such examples are extremely rare, and virtually impossible to elicit directly, in my experience.

It is therefore unclear whether there is any meaningful difference between ordinary stem reduplication and reduplication of the stem and a stem-internal constituent. The root in (226) is reduplicated at the root in most of its attested compounds, so it may be somewhat frozen in that form.

(224) wuɾaː-**ʼɾa**ː-ʼɾaː-**ʼpa**ː-ʼ-put-i-jiɲuŋ
/wuɾu-**w$_1$aɾa**-w$_1$aɾa-**aKw$_2$a**-aK-w$_2$ut-i-jiɲuŋ/[18]
3pl.R-Red-Mult-Red-ground-swollen-Past2-Rel
'They who were all swollen up' (NMET 139.9.1)

(225) aŋku-**jiki**-jika-**maɲca**-ma-ɲc-iːna
/aŋku-**jiki**-jika-**maɲca**-ma-aɲci-iːna/
3Neut.Irr-Red-fire-Red-get-Recip-Npst2
'Fires will get each other [join up]' (NMET 97.4.2)

(226) wini-**ŋaɾa**-ŋaɾaɳ-**pulu**-wula-ka-na
/wini-**ŋaɾa**-ŋaɾaɳ-**w$_2$ulu**-w$_2$ula-ka-na/
3pl>3Masc.R-Red-heap-Red-make.two-Fact-Npst2
'They keep multiplying it [making piles of two]' (NMET 121.3.1)

As we have seen, reduplication can target either the stem or, less often, a constituent inside the stem. This shows some apparent similarity to the closely related language Anindilyakwa. According to Bednall's (2019, p. 237) analysis, among verb stems consisting of a quantifying derivational prefix /məɳʈakə-/ — the equivalent of the Wubuy Mult prefix — and an inflected verb root, reduplication can target either the root or the quantifying prefix, with different consequences for interpretation. Reduplicating the quantifier emphasises the plurality of the subject or object, whereas reduplicating the inflected root emphasises 'the intensity of the situation'; reduplicating both is also possible, and its interpretation is more or less compositional. To the best of my knowledge, we do not yet have any evidence bearing on the question of whether reduplication of roots vis-à-vis stems in Wubuy has semantic consequences. I therefore make the conservative assumption that the position of the reduplicant has no bearing on the semantic interpretation of reduplicated words.

[18] Heath (1980) claims the root is /w$_2$uta/ as shown, but the status of this incorporated noun /aK/ 'ground' is a little dubious, and is clearly not recoverable from the semantics of the whole here. An alternative would be to represent the root as simply /aputa/.

4.2.4 Reduplication with Hardening and ŋu-Epenthesis

Where a reduplicated stem begins in a continuant consonant subject to Hardening, it will have a stop-initial — but still disyllabic — realisation after any of the nasal-final prefixes. In rule-ordering terms, reduplication feeds Hardening (and not vice versa). Thus the stem which has the form *jampa-jamp-iːni* in (227b), is instead *campa-jamp-iːni* in (227d), as it follows the consonant-final prefix *ŋaN-* '1sg.Irr' and must therefore undergo Hardening.

(227) a. ŋa-jamp-iːna-jiɲuŋ
/ŋa-jamp-iːna-jiɲuŋ/
1sg.R-speak-Npst2-Rel
'What I'm talking about' (NMET 47.23.7)

b. ŋa-**jampa**-jamp-iːni-jiɲuŋ
/ŋa-jampa-jamp-iːni-jiɲuŋ/
1sg.R-Red-speak-Past2-Rel
'What I was talking about' (NMET 47.23.2)

c. ŋaɲ-camp-iːna
/ŋaN-jamp-iːna/
1sg.Irr-speak-Npst2
'I will speak' (NMET 47.22.2)

d. waːɻi ŋaɲ-**campa**-jamp-iːni
/waːɻi/ /ŋaN-jampa-jamp-iːni/
not 1sg.Irr-Red-speak-Past2
'I didn't speak [about such things]' (NMET 47.22.4)

Only the reduplicant-initial segment is affected by this alternation — the stop realisation is not copied back onto the base, i.e. as **ŋaɲ-campa-camp-iːni*. A stop-initial monosyllabic reduplicant cannot be used here either: e.g. **ŋaɲ-ca-jamp-iːni* or **ŋaɲ-ca-camp-iːni*. In that sense, the interaction between Hardening and reduplication is opaque: the generalisation that stop-initial reduplicants should be monosyllabic and sonorant-initial reduplicants should be disyllabic is not surface-true, as we do in fact encounter stop-initial disyllabic reduplicants on the surface — namely, those derived by Hardening. This point will be relevant to the theoretical discussion in 4.5.

As for the stop-initial monosyllabic reduplication pattern, it feeds the insertion of prefixal *ŋu-* following consonant-final prefixes, as befits stop-initial stems generally — or, in the terms

of Chapter 3, stems of the *stop class*. So, just as the initial /p/ of the stem *puɾ-aː* 'sit-NPST2' conditions prefixal *ŋu-* in (228a), the initial /p/ of the entire reduplicated stem *pu-puɾ-aː* does likewise in (228b). One cannot insert *ŋu-* between reduplicant and base.

(228) a. na-**ŋu**-puɾ-aː
/naN-puɾ-aː/
12du.IRR-sit-NPST2
'We will sit' (NMET 20.18.4)

b. na-**ŋu**-pu-puɾ-aː
/naN-pu-puɾ-aː/
12du.IRR-RED-sit-NPST2
'We will sit (for some time)' (NMET 20.18.5)

Linking *ŋu-*, on the other hand, can feed reduplication, such that /ŋu-/ will appear in the reduplicant if it is close enough to the left edge of the stem to come within the domain for copying. In all cases, this requires that the reduplicant is disyllabic and that whatever derivational element precedes linking /ŋu-/ is monosyllabic. Thus the domain for reduplication in (229) encompasses the monosyllabic incorporated noun *l̪ac-* 'firewood' and the linking /ŋu-/ that follows it. (In the glosses, I have included linking *ŋu-* in the forms between slashes, to make this feeding relationship more apparent.)

(229) a. wiɾi-l̪ac-**ŋu**-paɻawut-aɲi
/wiɾi-l̪ac-ŋu-paɻawut-aɲi/
3pl>3NEUT.R-firewood-Ø-discard-PAST2
'They threw away the firewood' (PN2020-11-03)

b. wiɾi-l̪ac**ŋa**-l̪ac-**ŋu**-paɻawut-aɲi
/wiɾi-l̪acŋa-l̪ac-ŋu-paɻawut-aɲi/
3pl>3NEUT.R-RED-firewood-Ø-discard-PAST2
'They threw away the firewood' (PN2020-11-03)

Although it is possible for a reduplicated word to contain both a prefixal /ŋu-/ *and* a linking /ŋu-/, the linking /ŋu-/ cannot be copied into the reduplicant in such a case. This is because the necessary conditions for, on the one hand, prefixal /ŋu-/ appearing before the reduplicated stem, and on the other, a copy of linking /ŋu-/ appearing in the reduplicant, do not overlap.

Consider that in order for a linking /ŋu-/ to come within the domain for copying, it must be as close as possible to the left edge of the stem. But by definition, linking /ŋu-/ cannot ever appear

	Prefixal /ŋu-/	Reduplicated linking /ŋu-/
Reduplicant must be:	monosyllabic	disyllabic
Stem must be:	stop-initial	sonorant-initial
Example:	nu-ŋu-**pa**-pa-ŋu-t̪it-iːna 2sg.R-Ø-Red-eye-Ø-close-Npst2	wiɾi-**l̪acŋa**-l̪ac-ŋu-paɻawut-aɲi 3pl>Neut.R-Red-firewood-Ø-discard-Past2

Table 4.2: The necessary conditions for prefixal /ŋu-/ applying and for a linking /ŋu-/ appearing inside the reduplicant are mutually exclusive.

as the leftmost element of the stem (hence 'linking'), so the closest it can be to the beginning of the stem is the second syllable. If the reduplicant is then disyllabic, this *ŋu-* will be copied. But if the reduplicant is monosyllabic, it will not. From this it follows that the stem must be sonorant-initial — therefore conditioning a disyllabic reduplicant — in order to copy linking *ŋu-*. However, a sonorant-initial, disyllabic stem cannot condition prefixal *ŋu-*, which requires an underlyingly stop-initial stem. Thus, the conditions for prefixal *ŋu-* and the conditions for a reduplicated linking *ŋu-* logically preclude each other, as Table 4.2 illustrates. In the examples, the reduplicant is shown in bold, and the corresponding portion of the base is underlined.

We might say, then, that the processes we find at the left edge of a reduplicated stem are just those we find at the left edge of unreduplicated stems as well. Stop-initial reduplicated stems behave like other stop-initial stems in being subject to ŋu-Epenthesis, and continuant-initial reduplicated stems are exactly like ordinary continuant-initial stems in being subject to Hardening. The reduplication process — and its shape allomorphy — does not 'care' about what processes apply at the prefix-stem boundary, nor about the internal composition of the base beyond the leftmost segment and what processes may have applied there (e.g. linking *ŋu*). The interactions between reduplication and other processes discussed in this section are illustrated with some sample derivations in Table 4.3, which broadly assumes the underlying representations, transformational rules, and their respective ordering, as in Heath's (1984) analysis — the only departure is in splitting ŋu-Epenthesis into two rules (for its two domains of application) rather than employing a cyclic mechanism. In the next section, we consider what happens when the base is vowel-initial.

4.2.5 Reduplication of vowel-initial stems

In Heath's (1984) analysis, underlyingly vowel-initial stems also copy an open disyllable, like sonorant consonant-initial stems; the disyllabic character of the reduplicant is then partially obscured by vowel coalescence rules. His reduplication rule takes a stem like *akamp-aː* 'roast-Past2' and produces the intermediate form *aka-akamp-aː*, with vowels on either side of the boundary

between base and reduplicant. His 'VV-Contraction' rules must repair the vowel hiatus between reduplicant and base, yielding *aka:kamp-a:*, as in (230b). The same rules also repair the hiatus between the stem-initial vowel and the prefix-final vowel, if there is one, as in (230a,b). I present several further examples of vowel-initial stems below. This time, both surface and underlying forms[19] are shown, the former without morphological boundaries — as we shall see, the vowel hiatus resolution makes mophological segmentation difficult.

(230) a. wiɾiŋaːkampaː
/wiɾiŋa-akamp-aː/
3pl>3Fem.R-roast-Past2
'They roasted it (body of Emu; Fem)' (NMET 95.7.1)

b. wiɾiŋaːkaːkampaː
/wiɾiŋa-aka-akamp-aː/
3pl>3Fem.R-Red-roast-Past2
'They roasted it' (NMET 95.7.2)

(231) a. wuɾaːkampaː
/wuɾu-akamp-aː/
3pl>3Resid.R-roast-Past2
'They roasted it (Resid)' (NMET 95.10.1)

[19] Again, the forms are not quite underlying, since reduplication has already applied, to spell out the segmental content of the reduplicant.

	/ŋaN-(Red-)jamp-iːni/	/ŋaN-(Red-)puɾ-aː/	/wiɾi-(Red-)l̪ac-paɻawut-aɲi/	/nuN-(Red-)pa-ʈ̪it-iːna/
Linking ŋu-	—	—	wiɾi-l̪ac-ŋu-paɻawut-aɲi	nuN-pa-ŋu-ʈ̪it-iːna
Reduplication	ŋaN-jampa-jamp-iːni	ŋaN-pu-puɾ-aː	wiɾi-l̪acŋa-l̪ac-ŋu-paɻawut-aɲi	nuN-pa-pa-ŋu-ʈ̪it-iːna
Prefixal ŋu-	—	ŋaN-ŋu-pu-puɾ-aː	—	nuN-ŋu-pa-pa-ŋu-ʈ̪it-iːna
Hardening	ŋaN-campa-jamp-iːni	—	—	—
Dorsal del.	—	ŋa-ŋu-pu-puɾ-aː	—	nu-ŋu-pa-pa-ŋu-ʈ̪it-iːna
Nasal assim.	ŋaɲ-campa-jamp-iːni	—	—	—

Table 4.3: Sample derivations illustrating interactions between reduplication and other processes, according to the underlying forms and ordered rules in Heath (1984).

b. wuɾaːkaːkampaː
/wuɾu-aka-akamp-aː/
3pl>3Resid.R-Red-roast-Past2
'They roasted it' (NMET 95.10.1)

(232) a. nimaːɭɲawanŋaː
/nima-aɭɲa-wanŋ-aː/
3sgM>3Veg.R-mud-step-Past2
'He stepped in the mud (Veg)' (NMET 10.15.1)

b. nimaːɭɲaːɭɲawanŋaː
/nima-aɭɲa-aɭɲa-wanŋ-aː/
1pl>3Resid.R-Red-mud-step-Past2
'He stepped in the mud (Veg)' (NMET 10.15.3)

(233) a. wiɾiŋaːɾawulkulṯaŋi
/wiɾiŋa-(w$_1$)aɾa-wulkulṯ-aŋi/
3pl>3Fem.R-Mult-chop-Past2
'They chopped it all up (Fem)' (NMET 45.11.8)

b. wiɾiŋaːɾaːɾawulkulṯaŋi
/wiɾiŋa-aɾa-(w$_1$)aɾa-wulkulṯ-aŋi/
3pl>3Fem.R-Red-Mult-chop-Past2
'They chopped it all up (Fem)' (NMET 45.11.7)

(234) a. wuɾaːwaŋk-iːni(:::)
/wuɾu-awaŋk-iːni/
3pl.R-listen-Past2
'They listened' (NMET 41.9.3)

b. wuɾaːwaːwaŋkiːni
/wuɾu-awa-awaŋk-iːni/
3pl.R-Red-listen-Past2
'They listened' (NMET 95.2.4)

(235) a. juka nimpa:ŋuṯitaŋi
/juka/ /nimpa-aK-ŋu-ṯit-aŋi/
Q 2sg>1sg.R-Appl-Ø-close-Past2
'Did you shut [the doors] up for me?' (MM; PN2017-07-07.1 00:22:01)

b. juka nimpa:ŋa:ŋuṯitaŋi anu-ṯawaŋ
/juka/ /nimpa-aKŋa-aK-ŋu-ṯit-aŋi/ /ana-uK-l̪awaŋ/
Q 2sg>1sg.R-Red-Appl-Ø-close-Past2 Top.Neut-Neut-door
'Did you shut the doors up for me?' (MM; PN2017-07-07.1 00:23:10)

(236) a. ŋuna:nṯarmana
/ŋunu-anṯarma-na/[21]
1sg>2sg.R-imitate-Npst2
'I'm copying you' (HN; PN2020-10-30; 00:03:43)

b. ŋuna:nṯa:nṯarmana
/ŋunu-anṯa-anṯarma-na/
1sg>2sg.R-Red-imitate-Npst2
'I'm copying you' (HN; PN2020-10-30; 00:06:42)

(237) a. ŋaŋa:ŋupakal̪anaŋana
/ŋaŋu-aK-ŋu-pakal̪a-naŋa-na/
1sg>3Fem.Irr-Appl-Ø-eye-burn-Npst2
'I'm going to burn its eye for her' (BB2017-12-08)

b. ŋampa:ŋa:ŋupakal̪anaŋana
/ŋampa-aKŋa-aK-ŋu-pakal̪a-naŋa-na/
1sg>3pl.Irr-Red-Appl-Ø-eye-burn-Npst2
'I'm going to keep burning its eye for them' (BB2017-12-08)

[21]Heath (1982) regards this stem as composed of the verb root /l̪arma/ 'to chase, follow' along with compound initial -an-, which does not appear with any other root. This complexity is omitted here in order to simplify the presentation. Note also that Heath lists two other free variants of this construction, /w$_1$an-ṯarma/ and /w$_2$an-ṯarma/, but this speaker only accepted the vowel-initial variant /an-ṯarma/ (including all its predicted allomorphs).

(238) a. wiɾima:ɲciŋucuɻaŋi
/wiɾima-aɲci-ŋu-cuɻ-aŋi/
3pl>3Veg.R-Comit-Ø-push-Past2
'They pushed [the canoe] together' (PN2017-07-07)

b. wiɾima:ɲca:ɲciŋucuɻaŋi
/wiɾima-aɲca-aɲci-ŋu-cuɻ-aŋi/
3pl>3Veg.R-Red-Comit-Ø-push-Past2
'They pushed [the canoe] together' (PN2017-07-07)

(239) a. ŋaɾa:ŋupuɾi
/ŋaɾa-aK-ŋu-puɾ-i/
1sg>3pl.R-Appl-Ø-sit-Past2
'I sat down for them' (HN, PN2017-06-27 00:49:45)

b. ŋaɾa:ŋa:ŋupuɾi
/ŋaɾa-aKŋa-aK-ŋu-puɾ-i/
1sg>3pl.R-Red-Appl-Ø-sit-Past2
'I had to sit down for/because of them'[22] (HN, PN2017-06-27 00:50:25)

The vowel coalescence process has several important consequences for a formal analysis. First of all, it obscures the generalisation that sonorant-initial stems take a disyllabic reduplicant, since it is only disyllabic at an abstract level of representation, and not on the surface. Whereas in the examples of sonorant consonant-initial stems above, reduplication results in a word with two more syllables than its unreduplicated counterpart, vowel-initial stems gain just one extra syllable by reduplication. And, whereas the vowel in this extra syllable is long, the vowels in the sonorant consonant-initial stems must be short. It is therefore not the case that long vowels are banned from the reduplicant generally, only that they must be the result of resolving hiatus. One is faced with a choice between either treating the vowel-initial stems as a class distinct from the stop-initial and sonorant consonant-initial stems, or assuming some sort of serial mechanism (as Heath does), such that reduplication can feed more surface-level processes like vowel coalescence. Later in this chapter, I defend an analysis that takes the former approach.

A further consequence of vowel coalescence is that actually distinguishing the reduplicant from the rest of the word is problematic if one only has access to the surface forms. Comparing

[22]Speaker explanation: "sitting down there for them and you didn't get away from them, just sat there for them till daybreak or whatever"

the surface forms *wiɾiŋa:kampa:* in (230a) and *wiɾiŋa:ka:kampa:* in (230b), there are really three possibilities: either reduplication adds a V:C(C) prefix (*a:k-*), or a C(C)V: infix (*-ka:-*) inserted after the first vowel in the base, or a V:C(C) infix inserted after the first consonant of the base. These possibilities are illustrated, respectively, as (240b), (240c) and (240d). If one were to subtract the reduplicant identified by the box in any of these three reduplicative word forms, one would be left with (240a), namely its unreduplicated counterpart. Note also that, no matter which substring of the word is singled out as the reduplicant in (240b,c,d), the vowel is always long. Consequently, a hallmark of the V-initial reduplication pattern is that both the first and second vowels of the reduplicated stem are long.

(240) a. wiɾiŋa:kampa: *Unreduplicated word*
b. wiɾiŋ a:k a:kampa: *Reduplicant as V:C(C) prefix*
c. wiɾiŋa: ka: kampa: *Reduplicant as C(C)V: infix*
d. wiɾiŋa:k a:k ampa: *Reduplicant as V:C(C) infix*

This problem of morphological segmentation is not unique to reduplication. The heteromorphemic VV sequence is resolved by the same phonological process here as it is in Wubuy generally, and this process, by realising two underlying segments that are associated with two different morphological constituents as one segment on the surface — albeit with potentially two timing units — serves to obscure the morphological boundary between the two.[23]

To add a further complication, the surface long vowel straddling the boundary between inflectional prefix and stem — whether the stem is reduplicated or not — changes in quality depending on the underlying vowels it is derived from. This creates the potential for discrepancies between the surface reduplicant — however delimited — and the corresponding portion of the base.

Consider the following partial paradigm for the verb stem *api-na* 'jump-NPST2', shown in (241) in its reduplicated and unreduplicated forms, following each of the three vowels.

(241) Consequences of reduplicating stem *api-na* 'jump-NPST2', following pronominal prefixes with different underlying final vowels

	/ŋa-/ '1sg.R'	/wu-/ '3RESID.R'	/ŋi-/ 3sgF.R
Unredup.	a. ŋa:pina	c. wa:pina	e. ŋi:pina
Redup.	b. ŋa:pa:pina	d. wa:pa:pina	f. ŋi:pa:pina

Among the reduplicated words in (241), note that the first vowel, which is the surface result of underlying vowels on either side of the prefix-stem boundary, alternates between /a:/ and

[23]For what it's worth, this is also reflected in Wubuy orthography: major morpheme boundaries are indicated with hyphens, except where VV coalescence has applied to produce a long vowel.

/iː/ (but not /uː/) depending on the underlying vowels involved. Whatever quality this first vowel has in the unreduplicated form, so we predict it to be in the reduplicated form, and vice versa. The first vowel associated with the stem in (241a,c) is /aː/, as is the first vowel in their reduplicated counterparts (241b,d); the first vowel in (241e) is instead /iː/, just as it is in (241f). However, the second (long) vowel in the reduplicated forms does not change, no matter the surface quality of the first vowel: it is always /aː/.

The stems in the V-initial examples presented so far happen to all be /a/-initial, which reflects their overwhelming predominance in the Wubuy lexicon. Among the vowel-initial verb roots listed in Heath (1982), 110 are /a/-initial, six are /i/-initial, and only one is /u/-initial. The Appl *aK-*, Comit *aɲci-* and Mult *(w$_1$)aɾa-* derivational prefixes, which must come immediately after the inflectional pronominal prefix, boost the number of /a/-initial stems even further, as do the many /a/-initial adjective roots and incorporable nouns.[24] Paradigms like (241) are therefore the overwhelming norm among vowel-initial stems.

On the other hand, there are only 15 attested roots or compounds in the dictionary beginning in /i/, five of them featuring the incorporated noun *-ic-* 'words, truth', and four featuring *-iɲkaɭa-* 'shortcut', none of them appearing frequently in the texts. The compound *i-caŋa* 'hear' (analysable as /ic/ 'words' + /jaŋa/ 'to hear') is the only one of these in common use — it has six tokens in the texts, one of which is reduplicated (242b). The four other compound constructions in the same corpus with the incorporated noun *-ic-* 'words' at their left edge have only one token each, and none are reduplicated. None of the compounds with *-iɲkaɭa-* 'shortcut' listed in the dictionary appear in the text corpus at all.

(242) a. waːɻi aɲiwiːcaŋani
/waːɻi/ /aɲiwu-ic-jaŋa-ni/
not 3sgF>Neut.R-word-hear-Past2
'She didn't think' (NMET 7.6.2)

b. ŋiwiːciːcaŋani
/ŋiwu-icji-ic-jaŋa-ni/
3sgF>Neut.R-Red-word-hear-Past2
'She thought' (NMET 45.5.3)

The only other example of a reduplicated /i/-initial stem in the entire text corpus is that shown in (243), where the stem is, according to Heath, underlyingly /il-aɲi/ 'slice-Past2'. It

[24] The Mult prefix is not strictly vowel-initial; it is listed in Heath's work as *(w$_1$)aɾa-*. As will be explained below, /w$_1$/ is optionally deleted in morph-initial position, in which case the morph behaves as if it were vowel-initial. For this prefix in particular, deletion applies amost all the time.

seems telling that the speaker — no less than the founder of the settlement at Numbulwar, and a Wubuy storyteller of the highest standing — seems rather unsure what the vowels should be. His first use of this verb in the story is unreduplicated (*waɾa:laŋi*), where the long vowel over the prefix-stem boundary is /a:/, as would be predicted if the stem-initial vowel were underlyingly /a/. Immediately following, he produces the reduplicated counterpart of that word, once again as if the stem were /a/-initial: *waɾa:la:laŋi*. Then he instead offers, perhaps as a self-correction, a reduplicated form where the first two vowels of the stem are /i:/ — *waɾi:li:laŋi* — instead suggesting an underlyingly /i/-initial root. Later in the same text, the speaker goes back to the first form, with long /a:/. It seems probable that his ambiguity around which vowel quality is appropriate here reflects the vastly higher frequency of /a/-initial reduplicated forms, leading him to regularise in that direction. The four different forms in (243) are set out for easier comparison in (244).

(243) ŋa aʈapa **waɾa:laŋi** **waɾa:la:laŋi(:::)** **waɾi:li:laŋi(:::)**
/ŋa/ /aʈapa/ /waɾa-al-aŋi/? /waɾa-ala-al-aŋi/? /waɾa-ili-il-aŋi/
and now 3pl>3Coll.R-slice-Past2 3pl>3Coll.R-Red-slice-Past2 3pl>3Coll.R-Red-slice-Past2

waɾa:ɾaḻuḻuwi:ni, **waɾi:li:laŋi(:::)** waɾa:ɾaḻuḻuwi:ni,
/waɾa-aɾa-ḻuḻuw-i:ni/ /waɾa-ili-il-aŋi/ /waɾa-w$_1$aɾa-ḻuḻuw-i:ni/
3pl>3Coll.R-Mult-soak-Past2 3pl>3Coll.R-Red-slice-Past2 3pl>3Coll.R-Mult-soak-Past2

waɾaŋaɲpantaŋi jaku, **waɾi:laŋi**
/waɾa-ŋaɲpant-aŋi/ /jaku/ /waɾa-il-aŋi/
3pl>3Coll.R-want-Past2 or 3pl>3Coll.R-slice-Past2

'Then they sliced the yams (with the blade). They wanted them, they sliced them up.'
(NMET 125.6)

(244) a. waɾa:laŋi
b. waɾa:la:laŋi
c. waɾi:laŋi
d. waɾi:li:laŋi

I have also found it very difficult to obtain reduplicated /i/-initial stems under elicitation with current speakers, despite their familiarity with the unreduplicated base forms. One of the very few examples is shown in (245).

(245) a. ŋuɾi:ɲkaḻa:pa:paɲcina
/ŋuru-iɲkaḻa-apa:pa-ɲci-na/
12pl.R-shortcut-sneak[26]-Recip-Npst2
'We cut each other off [on the path]' (HN; PN2020-10-28)

b. ŋuɾiːɲkiːɲkaɭaːpaːpaɲcinawukic
/ŋuɾu-iɲki-iɲkaɭa-apaːpa-ɲci-na-wukic/
12pl.R-Red-shortcut-sneak-Recip-Npst2-still
'We kept cutting each other off [on the path]' (HN; PN2020-10-28)

There is only one /u/-initial verb stem — the Npst2 inflected root *u-ma-na* 'hit'. This is quite frequent in both its reduplicated and unreduplicated form, with 6 and 11 tokens respectively in Heath (1980). One of each is shown in (246), taken from the same text. Other tense-aspect inflections of this verb instead behave like /w_2/-initial roots, including by hardening to /p/. The only other instance of a /u/-initial root mentioned in Heath's work is a special suppletive form of /ji/~/i/ 'give' — namely /u/ — used as the right-hand member of a compound, but by being constrained to this position, it cannot appear at the left edge of the base to reduplication. It also appears not to be productive: the Wubuy speakers I consulted would not accept any novel compounds using this form.

(246) a. wiɾimuːmana
/wiɾima-u-ma-na/
3pl>3Veg.R-hit-Ø-Npst2
'They were hitting it (Veg)' (NMET 155.1.5)

b. wiɾimuːmuːmana
/wiɾima-umu-u-ma-na/
3pl>3Veg.R-Red-hit-Ø-Npst2
'They were hitting it (Veg)' (NMET 155.1.5)

In addition to underlyingly vowel-initial stems, there are also derived vowel-initial stems, owing to a productive but variable process deleting intervocalic /w_1/ at the beginning of a stem or root (rule P9 w_1 deletion in Heath 1984). Baker & Harvey (2022) calculate that, between low vowels, /w_1/-deletion applies 96% of the time in Heath's (1980) text corpus; in other environments it does not apply. Where deletion leaves a V-V sequence, the result is either two vowels in hiatus (247b), or a long vowel (248b). Following /a/, then, a /w_1a/-initial stem is indistinguishable from an /a/-initial one; but both of these are unlike /w_2/-initial stems, as underlying /aw_2a/ is realised [awa], not [aː]. The availability of a /w/-initial stem for deletion would therefore appear to provide morphophonemic evidence for the initial labiovelar being /w_1/, not /w_2/.

[26] In Heath (1982), the root is listed as *iɲkaɭaːpaːpa* 'to go past (someone, direct object) by taking a shortcut to cause (someone) to detour or swerve'. My gloss reflects the assumption that it is a lexicalised compound of *iɲkaɭak* 'shortcut' and verb root *apaːpi* 'to sneak along', which occur independently.

(247) a. wunu-wap-a:
wunu-$w_1$ap-a:
3pl>3sgM.R-wrap-PAST2
'They wrapped him up' (NMET 51.2.1)

b. wunu-wapa-ap-a:
wunu-$w_1$apa-$w_1$ap-a:
3pl>3sgM.R-RED-wrap-PAST2
'They wrapped him up' (NMET 51.2.3)

c. waɾa:-'pa:-'p-a:
waɾa-$w_1$apa-$w_1$ap-a:
3pl>3pl.R-RED-wrap-PAST2
'They wrapped them up' (NMET 50.5.2)

(248) a. ŋu-waɾa-na-ŋa:
ŋu-w_1aɾa-na-ŋa:
3sgM>3sgF.R-MULT-burn-PAST2
'He burned her all over' (NMET 5.20.3)

b. ŋu-waɾa:-'ɾa-na-ŋa:
ŋu-w_1aɾa-w_1aɾa-na-ŋa:
3sgM>3sgF.R-RED-MULT-burn-PAST2
'He kept burning her all over' (NMET 5.20.6)

Although it is rare, I also have an instance of /w_2/ undergoing deletion in this way, as in (249). However, in this example, the identity of the initial labiovelar is somewhat ambiguous. The root showing deletion here is the bound root /$w_2$al/ 'stem, trunk', which appears in the texts in some part-whole constructions where it hardens to [p], as expected of /w_2/. However, Heath (1982) notes that when incorporated into a verb headed by /w_2aɾaɻa/ 'to spear at tree/bush (to shake off fruit)' — itself a special compounding form of /ɻa/ 'to spear' — the initial segment can harden to either [p] or [k], which would suggest underlying /w_1/. Indeed, this speaker also accepted *nuŋ-kal-waɾaɻaŋ* as an alternative to *num-pal-waɾaɻaŋ* in (249b). It is therefore notable that in (249d), the initial /w/ of the reduplicant undergoes Hardening to [p] (as expected of /w_2/), but the initial /w/ of the base undergoes deletion (as expected of /w_1/).

(249) a. na-waʎji-ɲuŋ nataku ni-jal-waɾaɻa-ŋ
/na-waʎja-ɲuŋ/ /nataku/ /ni-w$_2$al-waɾaɻa-ŋ/
Top.M-man-Sg Prox.M 3sgM.R-vegetation-spear.Past1
'That man threw a spear at the tree' (HN; PN2020-10-30.3, 00:00:22)

b. juka num-pal-waɾaɻa-ŋ
/juka/ /nuN-w$_2$al-w$_2$aɾaɻa-ŋ/
Q 3sgM.R-vegetation-spear.Past1
'Did you throw a spear at the tree?' (HN; PN2020-10-30.3, 00:00:24)

c. naː-**'lwa**ː-'l-waɾaɻa-ni-jiɲuŋ
/na-**w$_2$alwa**-w$_2$al-w$_2$aɾaɻa-ni-jiɲuŋ/
Top.M-Red-vegetation-spear.Past2-Rel
'About that man throwing spears into the tree' (HN; PN2020-10-30.3, 00:01:01)

d. wuɾu-cuŋ na-jiwaŋku-ɲuŋ-maciː ani-janṯawiw-ana-maciː juka
/wuɾuc-juŋ/ /na-jiwaŋku-ɲuŋ-maciː/ /ani-janṯawiw-ana-maciː/ /juka/
person-Sg Top.M-old.man-Sg-if 3sgM>2sg.Irr-ask-Npst2-if Q

num-**palwa**ː-'l-waɾaɻa-ni ani-jam-aː-maciː nukawi-wuj
/num-**w$_2$alwa**-w$_2$al-w$_2$aɾaɻa-ni/ /ani-jam-aː-maciː/ /nukawi-wuj/
2sg.R-Red-vegetation-spear.Past2 3sgM>2sg.Irr-do-Npst2-if 2sg-All

'A person, an old man perhaps, he might ask you, "Were you throwing spears at the tree?" He might say that to you.' (HN; PN2020-10-30.3, 00:01:45)

4.2.6 Summary: reduplicant shape alternations in Wubuy

The preceding sections have set out the key empirical dimensions of reduplication in Wubuy. We have established that reduplication is partial, and that the partial reduplicant attaches to the left of the base it copies from. The target of reduplication is almost always the stem (rather than the root or the word). Where the stem is morphologically complex — involving, for instance, a derivational prefix or an incorporated noun — reduplication copies across any morphological boundaries within its domain, including linking *ŋu-* if present. A reduplicated stem is subject to all the same phonological alternations at the prefix-stem boundary as unreduplicated stems, including ŋu-Epenthesis and Hardening. Where the latter applies, there is a surface discrepancy between the manner of the reduplicant-initial segment (a stop), and the base-initial segment it corresponds to (a continuant).

I have taken particular care here to demonstrate that the partial copy in reduplication constructions varies in shape according to the sonority of the base-initial segment. It is striking that the crucial split is between stops and everything else, just as we saw for ŋu-Epenthesis in Chapter 3: stop-initial bases condition a monosyllabic (CV) reduplicant, while other kinds of bases condition — at some level of analysis — a disyllabic one. For sonorant consonant-initial bases, the reduplicant is disyllabic (CVC(C)(C)V) at the surface; the vowels are short, and the second vowel harmonises with the first. For vowel-initial bases, or /w_1/-initial bases where the initial consonant has been deleted, the reduplicant is only disyllabic (VC(C)V) at an abstract level, namely at an intermediate stage in a serial derivation. I presented evidence from speakers' uncertainty around /i/-initial bases that suggested no such abstract level is utilised by speakers, and that they instead analogise from the surface forms of reduplicated words with /a/-initial stems, which are overwhelmingly more common. Though there are exceptions to almost every one of these generalisations — sometimes a monosyllabic reduplicant appears with a sonorant-initial base, sometimes the reduplicant appears inside the stem, etc. — they are notably rare, and the dominant pattern described here is extremely well attested, and a pervasive feature of narrative texts.

In the next section, I compare these features of reduplication in Wubuy with those of its neighbours. Doing so suggests some avenues by which sonority-conditioned reduplicant shape allomorphy may have emerged, which helps us understand how it should be analysed synchronically.

4.3 The emergence of sonority-conditioned reduplication

The reduplication patterns described here for Wubuy are puzzling for the theory and typology of reduplication in a number of respects. In particular, the dependency of the reduplicant size on the sonority of the initial segment of the base is, to the best of my knowledge, unattested in the literature on reduplication anywhere in the world, with one single exception: an almost identical pattern occurs in Wubuy's genetically closest neighbour, Anindilyakwa (van Egmond 2012). In other Gunwinyguan languages, both monosyllabic and disyllabic reduplicants occur, but are they are not conditioned by base-initial sonority.

In this section, I consider some possible pressures that may have exerted an influence on the diachronic development of sonority-conditioned reduplication in Wubuy. These pressures make reference to the feature that distinguishes Wubuy and Anindilyakwa from the rest of the family and establishes their status as a genetic subgroup — their loss of the historical fortis stop series and lenition of singleton stops to continuants, which resulted in a relative underrepresentation of stops in the lexicon. This underrepresentation makes stem-initial stops

more informative on average than sonorants, with the effect that any segments following a stem-initial stop contribute less to the identification of the stem than they would after a sonorant. If these segments cease to be sufficiently informative, they should be more prone to diachronic processes of reduction. Thus, monosyllabic reduplicants may have emerged from disyllabic reduplicants just where the initial consonant is a stop. The synchronic result of this is a system in which reduplicant size subcategorises for base-initial sonority. And again, the classes involved for reduplication are precisely those that we saw for the alternations analysed in Chapter §3: stops, alternating stops/continuants, and non-alternating sonorants, as well as a vowel-initial class which is addressed for the first time in this chapter.

4.3.1 Reduplication in Anindilyakwa

Like in Wubuy, reduplication in Anindilyakwa targets the stem (in most cases), and the shape of the reduplicant is conditioned by the sonority of the stem-initial consonant. The examples in (250) show that stop-initial bases condition a monosyllabic reduplicant, while the sonorant-initial bases in (251) condition a disyllabic reduplicant.[27]

(250) Stop-initial monosyllabic reduplication in Anindilyakwa (van Egmond 2012, p. 77)

a. -ku-kuɻa-
-RED-hook.fish

b. wurə-pu-puŋkawa
3a-RED-boss

c. a-ci-ciɾku-wiʎaɾa
NEUT-RED-river-middle

d. -t̪ə-t̪ərənʈa-
-RED-descend

(251) Sonorant-initial disyllabic reduplication in Anindilyakwa

a. mə-l̪əkə-l̪əkaɾku-wiʎaɾa
VEG-RDP-river-middle

b. -mici-micikeji-
-RED-search

c. -manʈə-manʈarka-
-RED-point

[27] All examples taken from van Egmond (2012) have been converted to IPA transcription.

d. -ŋurkpə-ŋurk+paḻa-
-Red-mouth+wide-

(Anindilyakwa: van Egmond 2012, p. 77)

Van Egmond (2012) argues that, in both patterns, the quality of the reduplicant-final vowel is predictable on the basis of the surrounding consonants, just like the epenthetic vowels that appear quite prolifically elsewhere in Anindilyakwa (see Mansfield, Billington et al. 2022). For that reason, she analyses the stop-initial template as underlyingly C, and the sonorant-initial one as underlyingly CVC(C)(C); vowel epenthesis then breaks up the consonant cluster at the reduplicant-base boundary. This does not substantially diminish the comparison with Wubuy: on the surface, the reduplicant is still either monosyllabic or disyllabic, and the shape allomorphy depends on base-initial sonority. (For its part, a case could be made for the second vowel in Wubuy disyllabic reduplicants being epenthetic also; my proposal in §4.4.4 could be interpreted in such terms.)

The two languages differ more sharply in their treatment of vowel-initial stems. In Anindilyakwa, when the stem begins in a vowel, two different reduplicative patterns are in evidence. In one pattern, the initial vowel is skipped over and a CV or CVC(C)V reduplicant is infixed after it (252); in the other, the initial vowel is included in the reduplicant, which is in this case either VC or VCVC(C), and is prefixed to the stem (253). In both of these patterns, the sonority of the first consonant still determines reduplicant size: monosyllabic if the first consonant is a stop, and disyllabic otherwise. At this stage, it is unclear what determines whether the vowel-skipping or vowel-retaining pattern applies for any given vowel-initial stem, or indeed whether both patterns are always available.[28]

(252) Vowel-initial base, CV∼CVC(C)V reduplicant infixed after vowel

a. -apaɻumə- -a-pə-paɻəmə-[30]
search Red.search

b. aɻakpa=wija a-ɻakpə-rakpa=wija
compl.act=Perg Red.compl.act=Perg

[28] The difference between the two strategies essentially boils down to the quality of the vowel at the boundary between the reduplicant and the base. If it is a schwa, as in /-a-p**ə**-paɻəmə-/, it is characterised as adhering to the first vowel-initial reduplicative pattern. If it is not, as in /-ak-**a**kuma-/, it is characterised as belonging to the second pattern. If /-ak-akuma-/ instead followed the first pattern, we would derive */a-k**ə**-kuma-/. If underlyingly non-schwa vowels can variably be realised as schwa (cf. Mansfield, Billington et al. 2022), it may be that there is no real difference between the two patterns at all.

[30] This is a vowel change here between base and reduplicant, from /-abaɻ**u**mə-/ to /-a-bə-baɻ**ə**mə-/. Van Egmond does not comment on this.

Initial segment of base	Redup. shape	Redup. placement	Example
Stop	C_1V	Left edge of stem	-ku-ku.ɻa-
Sonorant consonant	$C_1V_1C_2(C_3)V$	Left edge of stem	-mici-micikeji
Vowel followed by stop C	C_1V	After first V of stem	-a-pə-pa.ɻəmə-
	$V_1C_1(C_2)$	Left edge of stem	-ak-akuma-
Vowel followed by sonorant C	$C_1V_1C_2V$	After first V of stem	a-ɻakpə-ɻakpa=wija
	$V_1C_1V_2C_2(C_3)$	Left edge of stem	amənɳḍ-amənɳḍəra

Table 4.4: Summary of main reduplication patterns for Anindilyakwa, based on examples and discussion in van Egmond (2012, pp. 76–79).

(253) Vowel-initial base, VC∼VCVC(C) reduplicant prefixed before vowel

a. -akuma- -ak-akuma-
put RED-put

b. amənɻəra amənɻ-amənɻəra
slowly RED-slowly

Anindilyakwa also differs from Wubuy in terms of what counts as a stop: stems whose initial consonant is a homorganic nasal-stop cluster or heterorganic dorsal+labial stop cluster behave as if they were single stops, conditioning a monosyllabic reduplicant. Van Egmond (2012) also notes some marginal patterns and lexical exceptions, though the discussion and exemplification of reduplication in her grammar is in any event fairly brief. These details notwithstanding, Wubuy and Anindilyakwa have the most striking features of their reduplication patterns in common. In both languages, reduplication targets the stem, and the shape of the reduplicant is conditioned by sonority: monosyllabic for stop-initial bases and disyllabic otherwise. The main generalisations are summarised in Table 4.4.

4.3.2 Reduplication in other Gunwinyguan languages

The sonority-conditioned reduplicant shape allomorphy shared by Wubuy and Anindilyakwa does not seem to occur anywhere else in the family, nor indeed further afield. In other Gunwinyguan languages, different reduplicant shapes may occur, but they are not primarily conditioned by phonological features of the base.

For Ngandi,[31] Wubuy's next closest genetic neighbour after Anindilyakwa (van Egmond & Baker 2021), Heath (1978b) identifies three types of reduplicant template: CVCV- (Type A), CVCVʔ- (Type B), and reduplication of the whole root (Type C). The reduplicant, whatever its

[31]When presenting examples from Ngandi and other languages with a fortis-lenis stop contrast, I use voiced symbols for lenis, and voiceless for fortis.

shape, is always prefixed to the root, not the stem or the word. While nominal reduplication is rare in Ngandi, verbal reduplication is quite common and apparently productive, and is associated with a similar range of meanings to verbal reduplication in Wubuy — repetition, distribution, and continuity — and is likewise largely restricted to verbs inflected for Continuous aspect. As far as I can glean from Heath's materials, the different reduplication patterns are not associated with different semantics. The choice of reduplicant template is in part determined by the verb's inflection class, such that each class tends to favour one type of reduplication pattern: verbs of classes 1 and 2 usually reduplicate following the Type C pattern, as in (256), while other verb classes usually follow Type A (254), with the occasional example of Type B (255). In general, root reduplication (Type C) is only available where the root is clearly segmentable, as in the large number of verbs that require a following thematic augment (256). Many of these roots can appear independently in an uninflected 'root form' (Heath 1978b, p. 87), a phenomenon which is almost entirely absent from Wubuy. However, in the very few cases where that segmentable root is longer than two syllables, the disyllabic reduplication pattern (Type A) occurs instead, as in (254d).[32] The Type A CVCV- reduplicant, then, acts as a kind of default shape, applicable whenever the root cannot be targeted as required by Type C; the Type B reduplicant is a much rarer variant of Type A (Heath 1978b, pp. 14–15).

(254) Type A: CVC(C)V- reduplication (Ngandi: Heath 1978b, pp. 14–15)

a. ɖaː-bu-
try.out-Aux-
'to try out'

b. **ɖabu**-ɖa-bu-
Red-try.out-Aux-
'to try out repeatedly'

c. ɟapaɖaʔ-ʈu-
go.again-Aug-
'to go to the same place'

d. **ɟapa**-capaɖaʔ-ʈu-
Red-go.again-Aug-
'to go to the same place repeatedly'

(255) Type B: CVC(C)Vʔ- reduplication (Ngandi: Heath 1978b, p. 99)

a. ju-ɻi
put.on-PCont
'put on'

b. **juɻiʔ**-ju-ɻi
Red-put.on-PCont
'put on repeatedly'

[32] The change from /c/ to /ɟ/ is due to a type of hardening rule that applies in reduplication.

(256) Type C: root reduplication (Ngandi: Heath 1978b, p. 93)

a. ɖak-ɖu-
cut-AUG-
'to cut'

b. **ɖak**-ɖak-ɖu-
RED-cut-AUG-
'to cut repeatedly'

Reduplicative constructions in Ngandi are subject to certain phonological alternations, particularly fortis-lenis alternations under certain conditions (254d), and shortening of vowel length (254b). However, there is no indication that segmental phonology plays a role in determining the choice of reduplicant shape.

A little further afield, in Ngalakgan, there are two productive reduplication patterns: CV?- and CVC(C)V- reduplication (Baker 2008b, p. 272).[33] For many words, speakers appear to have some degree of freedom in using one pattern or another; in other cases it appears to be lexically restricted (Brett Baker p.c.). Below, (257) exemplifies CV?- reduplication, while (258) is an example of CVC(C)V- reduplication. As in Ngandi, the reduplicant prefixes to the verb root, rather than the entire stem.[34] Reduplication, with either pattern, is associated with either a *distributive* or *intensive* reading, depending on the meaning of the base (Baker 2008b, p. 105).

(257) CV?- reduplication in Ngalakgan

a. maŋiŋ
get.PCONT
'got'

b. **ma?**-maŋiŋ
DISTR-get.PCONT
'got here and there' (Ngalakgan: Baker 2008b, p. 272)

(258) CV(C)CV- reduplication in Ngalakgan

a. worowk-mi+ŋ
jump.along-AUX+PPUNC

b. puru-**woro**-worowk-mi+ŋ
3pl-DISTR-jump.along-AUX+PPUNC
'they jumped along' (Ngalakgan: Baker 2008b, p. 273)

[33] Phonetically, the glottal coda is realised as creaky phonation during the preceding vowel.
[34] There is also total reduplication of nouns, which is interpreted as a literal count (Baker 2008b, p. 273.)

Reduplication in Bininj Kunwok is substantially more complex, with numerous patterns and sub-patterns, which I only very briefly cover here. Among nouns, Evans (2003, p. 168) notes that 'reduplication for plural is lexeme-specific and, unlike with verbs, one cannot apply the various formal types productively throughout the nominal lexicon'. Among verbs, reduplication is productive. There are three basic functions of verbal reduplication — termed *inceptive*, *iterative* and *extensive* — each associated with one or more reduplicant templates. For instance, the inceptive is formed with a CV?- reduplicant, as in (259a). The iterative has a range of different templates associated with it — too numerous to detail exhaustively here — which are conditioned partly by morphology, partly by phonology.[35] For instance, verbs with thematic augments *-g/ke* and *-ma* (among others) reduplicate by copying the entire syllable (including coda) preceding the augment, and, if that syllable ends in a nasal or lateral, adding a glottal stop (Evans 2003, p. 382); as in (259c). Inflected verbs that are longer than two syllables typically take a glottal-closed open disyllable (CV(C)CV?-) reduplicant, as in (259b). Disyllabic inflected forms of monosyllabic roots also reduplicate an open disyllable, but without the final glottal stop, as in (259d). There is also a special iterative reduplicant that applies when the inflected root is monosyllabic, which has the form $C_1V_1NV_1$-, as in (259e). The nasal is /ŋ/ if the root-initial consonant is coronal, and /n/ otherwise. CV- reduplicants have a very restricted distribution, only appearing on monosyllabic inflected roots in the Kune and Manyallaluk Mayali dialects, with iterative reading. Finally, 'extensive' reduplication has its own patterns and its own subcategorisation requirements, which I will not go into here. The point to take from all of this is that reduplicant shape in Bininj Kunwok is determined in the first instance by its function. Within each of these functional categories, certain inflection classes (particularly those with thematic augments) have their own preferred pattern. Where the pattern is not determined by inflection class, the prosodic makeup of the base plays a further deciding role. Evans also makes note of some dialectal variation in reduplication.

(259) Reduplication in Bininj Kunwok (Evans 2003, pp. 381–385)

a. barı-**ja?**-jame-ŋ guɲ [INCEP = CV?-]
3a>3P-INCEP-spear-PPUNC kangaroo
'They tried to spear the kangaroo…'

[35] In truth, several of the examples from Evans (2003) are also consistent with an analysis that treats the coverb root (which his glossing does not segment) as the target for copying, just like root reduplication in Ngandi. If that is the case, several of these patterns, with their putatively distinct phonological targets, may be collapsed into one basic pattern with a morphological target.

b. barɪ-**waɾeʔ**-waɾewo-ŋ [ITER = CV(C)CVʔ-]
3a>3P-ITER-miss-PPUNC
'…but they kept missing it'

c. ga-ɾulk-**ɟob**-ɟobge mungui? [ITER = σ-]
3-tree-ITER-cut.NPST always
'He keeps cutting trees down all the time'

d. bi-jau-**ŋune**-ŋuneŋ giŋa [ITER = CV(C)CV-]
3>3hP-child-ITER-eat.PPUNC crocodile
'The crocodile ate the baby right up'

e. aban-bo-**wono**-wo-ŋ goba-goʔbaɲ [ITER = $C_1V_1NV_1$-]
1>3plP-liquid-ITER-give-PPUNC RED-old.person
'I gave all my grog to the old people'

Segmental features of the base only influence reduplicant form in two ways, each of them in the iterative. In verbs that copy the whole syllable preceding the thematic augment, the presence of a nasal or lateral determines whether the reduplicant ends in a glottal;[36] in the nasal-inserting pattern, the place of articulation in the root-initial consonant determines the place of the nasal. As interesting as all of these reduplication patterns are, they are obviously not the same as the sonority-conditioned shape alternation in Wubuy and Anindilyakwa.

Indeed, while I have not examined the reduplication strategies of every Gunwinyguan language here, nor those of the unrelated languages of the region, it should be clear that the Wubuy/Anindilyakwa pattern is unique. This claim for uniqueness is explicitly supported in Baker & Harvey's (in press) recent large-scale survey of the phonology and morphology of Australian languages; further, they single out these two languages as among only six across the continent that have CV- reduplication at all. As we saw, Ngandi typically reduplicates an open disyllable (sometimes also inserting a final glottal stop), and otherwise reduplicates the entire root — the latter target is not a prosodic constituent but a morphological one. Ngalakgan productively reduplicates either an open disyllable or a glottal-closed monosyllable (CVʔ-). Both of these are considered bimoraic in Ngalakgan, and therefore licit prosodic words; CV- does not reach this bimoraic minimum. Bininj Kunwok has a much broader array of reduplicant shapes than any of the others, but CV- reduplicants have an extremely narrow distribution and only occur within two of the dialects, and there is no evidence of conditioning by sonority among any of the attested patterns.

[36] This is probably of little significance in any case, as glottals aren't licensed after obstruents in BKW generally.

4.3.3 Reduplication, lenition, and informativity

Wubuy and Anindilyakwa stand out from other languages of the region for their unusual sonority-dependent reduplicant shape alternations. In searching for an explanation of this fact, it seems sensible to look at what else these two languages have in common, especially if it can be linked to consonant sonority.

Fortunately, the very feature that classifies these two languages as genetic neighbours is of this character. In van Egmond & Baker's (2021) reconstruction of the common ancestor of Wubuy and Anindilyakwa, which they dub proto-Insular-Gunwinyguan, they identify four main phonological innovations from proto-Gunwinyguan, of which two are relevant here: the elimination of the proto-Gunwinyguan fortis-lenis contrast (with just the lenis series remaining), and the lenition of proto-Gunwinyguan lenis stops to continuants. These two developments may be described as together constituting a leniting chain shift (as also mentioned in the previous chapter) — though van Egmond & Baker (2021) suggest that the lenition of stops to continuants applied somewhat variably in proto-Insular-Gunwinyguan, since in each of the modern languages there are certain lexical items that retain inherited stops as stops, but the items in question often differ from one language to another. The result, as they point out, is that the lexicons of these languages contain a substantially smaller proportion of stops than the cognate lexicon of proto-Gunwiynguan, and indeed that of other Gunwinyguan languages like Ngandi, which have not undergone historical stop lenition.

In light of this, Baker & Bundgaard-Nielsen (2018) suggest an account of the Wubuy reduplicant shape alternation in terms of lexical access. That is, the authors ground the alternation in how listeners recognise a string of discrete words from continuous speech input. Psycholinguistic models such as Shortlist (Norris 1994), Shortlist B (Norris & McQueen 2008), TRACE (McClelland & Elman 1986), or the Neighbourhood Activation Model (Luce & Pisoni 1998), posit that recognising words in speech requires evaluating multiple lexical candidates in parallel, continuously updating one's predictions about what word is being uttered as speech unfolds in time. These lexical candidates compete with each other for recognition: upon hearing a word in speech, all other words sharing initial strings are 'activated' in the listener's mind until the set of possible words reduces to one (known as the 'uniqueness point'). That is, when an English-speaking listener hears [kʰæ], their lexical representations for *cat, category, carriage, cattle* and any other words they know that also share the initial string /kæ/ will be activated in their mind. If successful word recognition depends on narrowing down the list of competing possible words as speech unfolds, then phones that are less frequent will, on average, be more useful in ruling out competitors than more frequent ones (Balling & Baayen 2012; Wurm et al. 2006). Moreover, phones at the beginning of a word will be more informative than ones toward the end, as each successive phone cumulatively contributes to word identification (Geertzen

et al. 2016).

Recall that in Wubuy, as in Anindilyakwa, there is a significant asymmetry in the lexicon such that sonorants vastly outnumber stops. This would appear to predict that stops will hold greater discriminatory value during lexical access. In information-theoretic terms, we should expect stops in these languages to be more *informative*: less predictable in context, averaged across all the contexts they appear in (Baese-Berk et al. 2018; Balling & Baayen 2012; Cohen Priva 2017; Seyfarth 2014). Baker & Bundgaard-Nielsen (2018) suggest that a partial reduplicant need only 'minimally evoke' the base it copies, pointing the listener in the general direction of its location in the mental lexicon during lexical retrieval. The reduplicant therefore needs to be just long enough to activate the lexical representation of the base, without also activating too many other lexical neighbours (other lexemes sharing the initial string). They point out that in the case of a language like Wubuy, upon hearing a root beginning with a stop, the listener can already significantly narrow down the range of possible competitors since there are relatively few roots beginning with a stop in the Wubuy lexicon. Upon hearing a sonorant-vowel sequence however, the field of competitors is still too large to identify the correct lexeme, and the listener must wait for more information (in the form of the phones following this initial sonorant-vowel sequence). Thus, a sonorant-initial reduplicant meets this minimum informativity threshold when it is disyllabic, while a stop-initial reduplicant need only be monosyllabic to reach it.

In truth, we know almost nothing about the processing of reduplicative words, as reduplication has barely figured at all in the psycholinguistic literature.[37] However, the general thrust of their argument is anticipated by Blevins (2005), who also seeks to establish a relationship between reduplicative phonology and predictability. She presents comparative data from a range of Austronesian languages demonstrating that reduplicants are more prone to leniting sound change than ordinary morphs, precisely because, in the context of the base they attach to, reduplicants are under less pressure to be informative.

> If, quite generally, hypoarticulation correlates with the degree to which speakers or listeners are able to supplement acoustic information with information from other sources, then a strong prediction is made for reduplicated forms: all else being equal, phonetic reduction should occur more often in reduplicants than elsewhere in words, because the phonological information present in the reduplicated substring is predictable from the base. (Blevins 2005, p. 518)

In recent years, a considerable body of evidence has shown that low informativity of a linguistic unit correlates with the probability of it being phonetically reduced or deleted, and that this

[37] The only exception I know of is Gow et al (2023), though this is within an Artificial Language Learning experimental paradigm; it does not investigate the processing of a natural language that uses reduplication.

effect holds even in those contexts where it is in fact informative (Cohen Priva 2015, 2017; Cohen Priva & Gleason 2020; Hall et al. 2016; Seyfarth 2014). Blevins' contribution is to point out that reduplicants are *necessarily* of low informativity, and therefore especially susceptible to reduction, explaining why they so often end up as partial copies.

Her argument can be taken to suggest a possible diachronic pathway for sonority-conditioned reduplicant shape in Wubuy and Anindilyakwa. In this scenario, all reduplicants are initially disyllabic, and then the segments following the initial CV portion of reduplicants undergo deletion just when the reduplicant is stop-initial, yielding an alternation. For this, suppose that at an early stage of the ancestor language, before the emergence of the lenition chain shift, it had open disyllable reduplication as its dominant pattern, as in modern Ngandi. As stop lenition becomes entrenched in lexical representations, stops become comparatively underrepresented in the lexicon, and therefore more informative. If a stop-vowel sequence is more informative than a sonorant-vowel sequence, then any reduplicative segments after this stop-vowel sequence do comparatively little work in lexical identification, making them ripe for deletion. A reduplicant-initial sonorant-vowel sequence, however, is not sufficiently informative on its own, so there is greater usefulness in retaining the following portion. An attractive feature of this line of analysis is that it accounts simultaneously for both Wubuy and Anindilyakwa, and it does so by binding their unique reduplication pattern to the other feature that distinguishes them from the rest of the family: their widespread lenition of stops.

However, not all stops are alike in this regard, at least for Wubuy. The stops that condition a monosyllabic reduplicant are those of what I analysed in the previous chapter as the *stop class* (aka underlying stops). As we saw above in §4.2.4, a disyllabic reduplicant, prefixed to a continuant-initial base, will be stop-initial if it occurs in a Hardening context, as in (227d) *ŋaɲ-campa-jamp-iːni*. If stop-vowel sequences are informative enough on their own that the following reduplicative segments can be deleted, why does this not occur in these 'derived' stops (which are in any case historically antecedant)? Why is the reduplicated stem in (227d) *campa-jamp-iːni* and not simply *ca-jamp-iːni*?

The answer would seem to be that, as I argued in Chapter 3, the stop alternant of the alternating class is disambiguated from a stop of the stop class by the prefix-final nasal. Listeners may thus treat it as a kind of 'continuant in disguise', meaning that it would activate the same set of lexical competitors as a continuant following a vowel-final prefix — namely, those of the alternating class. To that extent, listeners still need to 'supplement acoustic information [in the reduplicant] with information from other sources', beyond what is provided by the initial CV sequence. This line of reasoning assumes that stop-initial and continuant-initial alternants are linked together in the mental lexicon, as Relational Morphology would claim. Given the critical endangerment of Wubuy, this assumption can scarcely be investigated directly, but support for

it comes from a study on lexical access and consonant mutation in Scots Gaelic (Ussishkin et al. 2017), a phenomenon not unlike Hardening in Wubuy. The study finds evidence of priming between mutated and unmutated alternants, independent of phonological overlap, which the authors interpret to suggest that speakers link these alternants to each other in the mental lexicon.

Finally, I note here that the disyllabic reduplicant has not entirely escaped leniting sound change either, in either Wubuy or Anindilyakwa. Open disyllable reduplication in the other Gunwinyguan languages exemplified here retains in every case the quality of the second vowel, as in (259d) *bi-jau-ŋun**e**-ŋun**e**ŋ*. In Wubuy, by contrast, the second vowel of the reduplicant must always be identical to the first, as in (186b) *ni-ŋaʈ**a**-ŋaʈ**u**kump-i:ni*. This can be understood as lenition in the sense that it minimises the effort involved in moving to a distinct vowel target (at V_2) between two identical ones (V_1 in the reduplicant and base), allowing the tongue dorsum and jaw configuration to remain comparatively uninterrupted throughout the reduplicant and into the first syllable of the base. In Anindilyakwa, this second vowel of the disyllabic reduplicant is the schwa, which assimilates in various ways to the consonants on either side, and in that sense can be considered less effortful than a full vowel.

4.4 Sonority-conditioned reduplicant shape in Relational Morphology

4.4.1 Orientations

The remainder of this chapter is concerned with developing a satisfactory theoretical treatment of the Wubuy reduplication data. The present section elaborates an account within Relational Morphology (RM; Jackendoff & Audring 2020), extending the theoretical machinery of the previous chapter. Section 4.5 compares this account against the affordances of two better-known theories of reduplication in the literature, Base-Reduplicant Correspondence Theory (BRCT; McCarthy & Prince 1993b), and Morphological Doubling Theory (MDT; Inkelas & Zoll 2005), both of which come with a number of unwelcome predictions. Each of the three theories of reduplication make different assumptions about the relationship between words that are reduplicated and those that are not. Before examining the finer details, it will be worth orienting ourselves by laying out some of the key differences between their general approaches.

At the outset, we can observe that both BRCT and MDT formalise this relationship as one between a surface reduplicative construction and some form of abstract underlying representation. BRCT assumes that this underlying representation contains an abstract RED morpheme whose phonological content is not supplied by the lexicon. The form this morpheme takes on

the surface — the reduplicant — is the product of competing demands to be more like the base to which it corresponds, and to satisfy general constraints on phonological well-formedness, modelled as an OT grammar. BRCT therefore characterises reduplication as a phenomenon of *phonological* identity between reduplicant and base. Partial reduplication obtains wherever markedness constraints requiring the reduplicant to be a particular prosodic shape — such as Red=σ or Red=PrWd — outrank base-reduplicant identity, which requires the reduplicant to be a full copy of the base. BRCT thus seeks to explain the surface form *wini-pu-puɾi* in (188b) as the product of an underlying form /wini-Red-puɾi/ and an OT grammar which ranks Red=σ over Max-BR.[38]

MDT instead sees reduplication as a phenomenon of *morphosemantic* identity: a reduplicative construction calls for two morphological units with (near-)identical semantics. There is thus no Red morpheme. Both 'reduplicant' and 'base' — though the theory seeks to problematise both of these concepts — are fully specified in the input underlying representation, and each is underlyingly a full copy of the targeted constituent. So in Wubuy, where reduplication targets the stem, *wini-pu-puɾi* would be underlyingly /wini-puɾi-puɾi/, with two complete copies of the stem. The theory assumes a model of grammar in which word patterns are modelled as constructional schemas which unify with one another. Phonological modifications can apply to the inputs at each morphological layer, and the grammar governing those modifications can differ from one construction to another: these subgrammars are known as 'cophonologies', and are modelled as OT grammars with differing constraint rankings. Partial reduplication obtains wherever the cophonology for one copy in the reduplicative construction ranks a restriction on output size, such as Output=σ, over input-output faithfulness constraints, such as Max-IO. Because it places no restrictions on what constraint rankings are possible within a given cophonology, MDT predicts that many other kinds of divergence are possible between the full copy in the input and its surface manifestation.

Although RM, as we have seen in previous chapters, also adopts a unification-based model of grammar, it assumes a monostratal phonology, and therefore does not permit underlying representations at all. Instead of relating a surface reduplicative form to some abstract underlying one, it treats the identity relationship at the heart of reduplication to be one holding between pairs of surface word forms, where one member of a pair is reduplicated and the other is not.[39] As with any other pair of morphologically related word forms, relational links establish correspondences between the constituents that these forms share. This means, first of all, that the base in a reduplicated word is linked to the corresponding constituent in an unreduplicated

[38] This presentation summarises only the 'Basic Model' of McCarthy & Prince (1993b, p. 25). The 'Full Model' differs by the addition of input-reduplicant correspondence constraints. In general, however, the Full Model has rarely been taken up in the literature, and the Basic Model is therefore more representative of typical BRCT analyses.

word (typically, the stem). The distinguishing feature of reduplicative constructions is that they contain two word-internal constituents with the same relational coindex: the reduplicant and the domain for copying within the base. In *wini-pu-puɾi*, both the reduplicant and the portion of the base which it resembles — *pu* and *pu* — are relationally coindexed to the corresponding portion of its unreduplicated counterpart: the *pu* in (188a) *wini-**pu**ɾi*. That is, both 'reduplicant' and 'base' are linked by means of relational coindices to the corresponding phonological constituent in the reduplicative word's unreduplicated counterpart.

Each of the three theories thus seeks to explain reduplicative word forms in relation to some other form: an underlying representation with a RED morpheme in BRCT, an underlying representation with two complete copies of the targeted constituent in MDT, and a non-reduplicative but otherwise equivalent surface form in RM. This is illustrated in (260). Dotted lines connect the elements in correspondence.

(260)

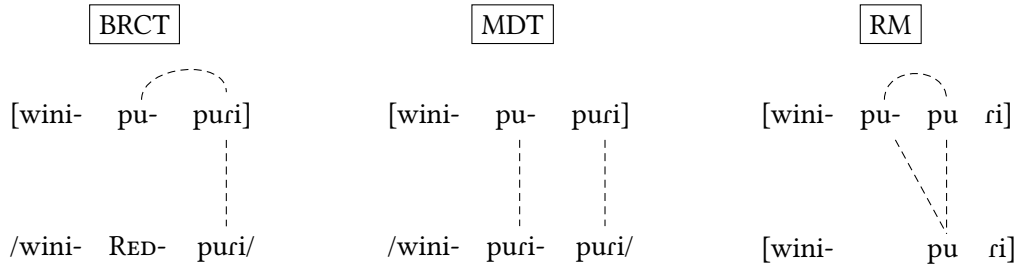

Any of these three theories can give an account of most of the familiar cases of reduplication in the literature. The unique challenge presented by Wubuy lies in its sonority-conditioned reduplicant shape alternation. As we will see, each of these theories has different means of answering this challenge, not all of them successful. And further, the fact that this reduplicant shape alternation is situated within what is already a complex array of phonological alternations, such as those examined in the previous chapter, sets a particularly high bar for what any theory of reduplication must be capable of.

4.4.2 Reduplication in Relational Morphology

In this section, I present a detailed RM analysis of reduplication in Wubuy. Since reduplication, in any language, has not yet been given detailed treatment in RM, the discussion here significantly extends the empirical coverage of the theory. I show that no additional mechanisms beyond the basic tools of the formalism — abstraction, coindexation, and unification — are necessary to

[39] Although I will not discuss it in any detail, note that Saperstein (1997) also pursues a theory of reduplication eschewing morphemes and underlying forms, couched within the Word and Paradigm tradition represented by Anderson (1992) and Stump (2001).

capture reduplication patterns. Relational coindices, otherwise used to relate structures between lexical entries, are here used also to relate phonological structures *within* lexical entries, such as between reduplicant and base. The otherwise vexing issue of sonority-conditioned reduplicant shape alternation is formally characterised in terms of phonological subcategorisation, just as we saw for the opaque interactions between ŋu-Epenthesis and Hardening in Chapter 3.

4.4.3 Existing analyses of reduplication in RM

Jackendoff & Audring (2020) offer only a brief and fairly tentative discussion of how reduplication can be captured in Relational Morphology. They illustrate their general approach with examples from the Australian languages Warlpiri and Mangarrayi, which mark the plural by, respectively, total reduplication of the entire nominal root, and partial (VC(C)) reduplication infixed before the first vowel of the root. Curiously, they do not demonstrate the application of their model to prefixing partial reduplication, the cross-linguistically most common type by a fairly wide margin (Inkelas & Downing 2015; Parncutt 2015). In this section, I summarise their approach to these two fairly simple examples, and attempt to draw out their implications for reduplication generally, including prefixing partial reduplication.

In Warlpiri, reduplication is total, and is used on human nouns to indicate plurality, as in *kurdu* 'child' and *kurdukurdu* 'children', which are given lexical entries below.

(261) Warlpiri sister lexical entries for nominal reduplication

a.

Semantics	CHILD_1
Morphosyntax	N_1
Phonology	$[\text{kuɽu}_1]$

b.

Semantics	$[\text{PLUR}(\text{CHILD}_1)]_2$
Morphosyntax	$\{\text{N}_1, \text{Pl}\}_2$
Phonology	$[\text{kuɽu}_1\ \text{kuɽu}_1]_2$

(Warlpiri; reconstructed from Jackendoff & Audring 2020, p. 122, citing Nash 1986)

The key formal device here is the relational coindex. Whereas relational coindices are typically used to mark shared pieces of structure on the same level between two or more lexical entries, in reduplication constructions they have the additional function of relating shared pieces of structure on the same level within a single lexical entry.

Consider coindex 1. In (261a) it functions on the one hand as an interface link, relating the phonological string /kuɽu/ to the N category head in the morphosyntax, and to the 'child' concept in the semantics. But it also serves as a relational link to the equivalent structures on each level of (261b). However, on the phonological level of (261b) there are two /kuɽu/ strings,

and coindex 1 appears on both of them. Thus, the plural word is characterised by a noun root appearing twice in the same word — a noun root that otherwise appears on its own (in the singular).

True to the declarative nature of the Parallel Architecture, there is no copying procedure at work here, but rather a statement that there are two pieces of phonological structure which are the same as each other, and which are the same as a piece of phonological structure in another word. Nor is there any requirement to make a distinction here between which copy is the 'base' and which the 'reduplicant'. Neither copy has priority, and neither copy introduces the plural meaning on its own: the plural function is a property of the construction as a whole, denoted with the outer interface link 2.

Other plurals can be formed in the same way — for example, *kamina* 'girl' ~ *kaminakamina* 'girls', or *mardukuja* 'woman' ~ *mardukujamardukuja* 'women'. Note that, as indicated in the source of these examples, human nouns are the only nouns to form plurals, through reduplication or otherwise; some other nominal stems can reduplicate to give a spatially distributed sense, but this latter use is less productive (Nash 1986, p. 130). We can therefore motivate general schemas for singular and plural human nouns in Warlpiri by replacing the coindices with variables, as in (262). The relational variable coindex α links the two schemas, effectively stating that for any word consisting of a given phonological string linked to a noun category head in the morphosyntax and some human entity in the semantics, one can also expect to find a word with two such strings concatenated together, and it will have a plural reading in both morphosyntax and semantics. As always, the statement can be interpreted declaratively, to motivate relationships between entries in the lexicon, but can also be used generatively, to form plurals from known singulars, or even form unknown singulars from known plurals — so long as such a plural can unify with (262b) in every particular.

(262) Warlpiri sister schemas for nominal reduplication (adapted from Jackendoff & Audring 2020, p. 123)

a.

Semantics	HUMAN ENTITY$_\alpha$
Morphosyntax	N$_\alpha$
Phonology	[...$_\alpha$]

b.

Semantics	[PLUR(HUMAN ENTITY$_\alpha$)]$_x$
Morphosyntax	{N$_\alpha$, PL}$_x$
Phonology	[...$_\alpha$ ...$_\alpha$]$_x$

Of course, since reduplication is total in Warlpiri, representing it formally is relatively straightforward. In order to analyse cases of partial reduplication, the basic approach to the

linkage between levels remains largely the same, but the representation of the phonological level requires a more complex elaboration. Jackendoff & Audring discuss some well-known examples from Mangarrayi, which forms (some) plurals by infixing a VC(C) reduplicant after the first consonant of the base, and copies the material to its right: e.g. *gabuji ~g**ab**abuji* 'old person(s)', *waŋgij ~w**aŋg**aŋgij* 'child(ren)'. Lexical entries for the first of these pairs are shown in (263) below. The semantic level has been omitted for clarity.

(263) Mangarrayi nominal lexical entries (Jackendoff & Audring 2020, p. 123)

a. Morphosyntax N_1
   Phonology $[g\ [ab]_2\ uji]_1$

b. Morphosyntax $\{_N\ N_1, PL\}_3$
   Phonology $[g\ {}^*[ab]^*_2\ [ab]_2\ uji]_{1,3}$

As an interface link, coindex 1 associates the phonological string with the N category head in the morphosyntax, just as we saw in the Warlpiri examples. In its relational function, coindex 1 associates the whole phonological string in (263a) with the whole phonological string in (263b). In (263b) there is additionally the coindex 3, which associates the entire string, here understood as the reduplicative construction, with the set of morphosyntactic features that includes the plural. Again, plurality is associated with the construction as a whole, not with the reduplicant or the part it 'copies'.

The two identical parts in the reduplication construction (263b) are again indexed to one another (coindex 2). But where in the Warlpiri example the two parts are complete instances of an otherwise existing lexical item, here the resemblance is between *substrings*: /gabuji/ exists independently (i.e. as the lexical entry 263a), but /-ab-/ does not. Since this substring does not contribute any meaning on its own, it is not indexed to anything in the morphosyntax (unlike the copied parts in the Warlpiri example). In (263b) this substring /-ab-/ — i.e., the part normally called the reduplicant — is relationally coindexed to its corresponding string within the same lexical entry (the 'base'), but it is also indexed to the corresponding part of its unreduplicated counterpart (263a).

The stars around the reduplicant indicate the 'same-except' relation: the two strings indexed 1 in the phonology of (263a) and (263b) are the same, *except for* the part delimited with stars. This star notation is part of the authors' general approach to formalising infixation, and is not otherwise necessary to the analysis of reduplication. Moreover, these two strings /gabuji/ and /gababuji/ interface with equivalent pieces of structure in the morphosyntax (and semantics, though it is not shown), namely the N category head.

Again, abstract schemas can be formulated by turning the constants into variables, as below.

The relational variable α links the stems to each other; meanwhile β links the corresponding strings in the reduplicant, the base, and the unreduplicated form, asserting that they must be the same. This additional identical substring in (264b) is what distinguishes it from (264a) – linked to it via coindex α — and is therefore surrounded with same-except asterisks.

(264) Mangarrayi reduplication sister schemas (Jackendoff & Audring 2020, p. 123)[40]

a. Morphosyntax N_{α}
   Phonology $[C\ [VC(C)]_{\beta}\ ...]_{\alpha}$

b. Morphosyntax $\{_{N}\ N_{\alpha},\ PL\}_{x}$
   Phonology $[C\ {}^{*}[VC(C)]_{\beta}{}^{*}\ [VC(C)]_{\beta}\ ...]_{\alpha,x}$

In the analyses of both Warlpiri and Mangarrayi, it is the presence of the two phonologically identical strings together in the same word that licenses the plural meaning in the morphosyntax and the semantics, but neither string contributes this meaning on its own. That is, the semantic contribution of reduplication is a property of the construction as a whole, associated with the outer interface link. In partial reduplication, the partial copy is not independently meaningful, but is relationally linked to something that is — the base within the same word, or the base in the corresponding unreduplicated word. This partial copy is not linked by an interface link to any element of meaning. In total reduplication, each copy is meaningful, and linked (by an interface link) to the lexical semantics, but neither copy contributes the plural meaning on its own.

A further point to draw out concerning reduplication in RM is that it involves correspondence — represented with relational links — of three kinds: between the reduplicant and the part of the base within the same word that it is identical to, between the reduplicant and the part of the unreduplicated word that it is identical to, and between the identical parts of the base in the reduplicated word and the unreduplicated word. Reduplication sister schemas may therefore be characterised as having a triangular set of relational links holding between the unreduplicated and reduplicated words, and within the reduplicated word.[41]

With these points established, we can begin to develop an account of prefixing partial reduplication in RM. In order to keep things simple, I will first treat prefixing partial reduplication in general terms, illustrated by way of a hypothetical language, before moving to the much

[40] I have added additional bracketing to Jackendoff & Audring's (2020) partial reduplication schema in order to clarify the scope of β, which links the (sub)strings in correspondence between base and reduplicant.

[41] Note here the parallel with McCarthy & Prince's (1993b) proposed but largely unexplored 'Full Model' in Base-Reduplicant Correspondence Theory, which adds constraints on correspondence between Input and Reduplicant alongside the more familiar Base-Reduplicant and Input-Output identity constraints.

more complicated Wubuy problem. Let us suppose a language that prefixes a CV- reduplicant in order to form the plural and that, like in Warlpiri and Mangarrayi, this is restricted to human nouns. We represent prefixation here with a nested constituent structure, just as we did in Chapter 3, such that the base is nested inside the constituent containing the reduplicant prefix. There are also brackets around the corresponding CV substrings.

(265) Hypothetical CV- prefixing reduplication in human plurals

a. Semantics HUMAN_{γ}
Morphosyntax N_{γ}
Phonology $[[\text{CV}]_{\delta}\ \ldots]_{\gamma}$

b. Semantics $[\text{PLURAL}(\text{HUMAN}_{\gamma})]_{x}$
Morphosyntax $\{\text{N}_{\gamma}, \text{PL}\}_{x}$
Phonology $[[\text{CV}]_{\delta}\ [[\text{CV}]_{\delta}\ \ldots]_{\gamma}]_{x}$

The sister schemas for monosyllabic reduplication in (265) can essentially be read in the following way: for any noun denoting a human and beginning in a CV sequence when appearing on its own (265a), we can expect there to be a related form (265b) which is the same except that the CV sequence will appear twice, and the word will have a plural interpretation in semantics and morphosyntax. But it can also be read in the opposite direction, such that the existence of a plural noun with an initial CV sequence that is repeated implies the existence of a non-plural noun where the initial CV sequence only appears once. Keep in mind that such interpredictability depends on every part of the schema being satisfied: a *non*-plural noun whose initial CV sequence is repeated cannot be used as the basis for deriving an unreduplicated form — at least, the schemas in (265) cannot be used for this.

The prefix in (265b) subcategorises for phonological form: it must consist of a consonant followed by a vowel, and this CV sequence must be the same as the initial CV sequence of the base it attaches to. This relation of identity between the two CV strings is expressed by the relational link δ: whatever the initial CV portion of the base is, so will be the CV of the reduplicant. These strings are only linked on the phonological level. They are not linked to any structure in the semantics or morphosyntax, as neither of them independently introduce any meaning. But they are linked relationally to the identical CV string in the phonology of (265a), which is likewise meaningless. Thus, the reduplicant and the copying domain of the base are coindexed to each other within the same lexical entry as well as to the equivalent string in (265a). Again, the plural reading in (265b) is associated with the general gestalt of the entire phonological form — namely, that it contains a repeated part — but neither of the CV copies contributes this meaning independently.

In order to restrict our reduplication pattern in (265) to a more specific phonological context — so that it only applies to, for instance, stop-initial nouns — we need only further specify the initial C to be [–SON]. Of course, this is precisely what is required to analyse Wubuy reduplication, since different reduplicant shapes subcategorise for different phonological features in the initial segment of the base. In the next section, I develop an analysis of Wubuy reduplication along such lines.

4.4.4 Reduplication of consonant-initial verb stems in Wubuy

To formalise Wubuy verbs, let us first consider a mini-paradigm, illustrating all the crucial relations of similarity and difference between reduplicated and unreduplicated word forms that we need to account for, as in (266). The consonant-initial stems will be treated here, while vowel-initial ones are dealt with in §4.4.5. Just as we saw in Chapter 3 for prefix allomorphy and the empty morph *ŋu-*, differences in reduplicant allomorphy can be understood in terms of phonological classes, where the different classes are established by subcategorising for phonological features across relationally linked word forms. Although we have so far treated stems with initial sonorant non-continuants and those with sonorant continuants as largely the same for the purposes of reduplication — both require a disyllabic reduplicant — there are additional complexities in the latter class introduced by Hardening, as the table shows. We therefore have the same three classes for reduplication as we did in Chapter 3: a stop class, a sonorant class, and an alternating class.

(266) Partial paradigm for reduplication of consonant-initial verb stems

	Stop class *puɾi* 'sit.PAST2'	Sonorant class *ŋuni* 'eat.PAST2'	Alternating class *jampiːni* 'speak.PAST2'
1sg.R	ŋa-puɾi	ŋa-ŋuni	ŋa-jampiːni
1sg.IRR	ŋaŋu-puɾi	ŋa-ŋuni	ŋaɲ-campiːni
1sg.R-RED	ŋa-pu-puɾi	ŋa-ŋunu-ŋuni	ŋa-jampa-jampiːni
1sg.IRR-RED	ŋaŋu-pu-puɾi	ŋa-ŋunu-ŋuni	ŋaɲ-campa-jampiːni

As we can see from (266), the stop class differs from the others in taking a monosyllabic reduplicant, and requiring the /ŋu/-final allomorph for the 1sg.IRR inflectional prefix. The sonorant class and the alternating class also each have their own distinct prefix allomorphs (for the 1sg.IRR), though both take a disyllabic reduplicant, in which the second vowel is the same as the first. Where Hardening applies to the stems of the alternating class, the reduplicant

will be stop-initial, while the corresponding base-initial segment is continuant-initial, as in *ŋaɲ-campa-jampi:ni*.

To capture these differences, we can elaborate sets of sister schemas for each of the classes of stems. I present those schemas here; they are explained in the discussion to follow.

(267) Sister schemas for verbal reduplication: stop class (monosyllabic)

a. Sem [PRED(Participant$_x$)$_\alpha$]$_y$
Morph { V$_\alpha$, Pro$_x$ }$_y$
Phon [...V$_x$ [[[−SON]V(*:*)]$_\beta$ ...]$_\alpha$]$_y$

b. Sem [DISTR([PRED(Participant$_x$)]$_\alpha$)$_y$]$_{??}$
Morph { V$_\alpha$, Pro$_x$, DISTR$_y$ }$_z$
Phon [...V$_x$ [[[−SON]V]$_\beta$ [[[−SON]V(*:*)]$_\beta$ ...]$_\alpha$]$_y$]$_z$

(268) Sister schemas for verbal reduplication: sonorant class (disyllabic)

a. Sem [PRED(Participant$_x$)$_\gamma$]$_y$
Morph { V$_\gamma$, Pro$_x$ }$_y$
Phon [...V$_x$ [[$\begin{bmatrix}-\text{CONT}\\+\text{SON}\end{bmatrix}$VC(C)(C)]$_\delta$ ...]$_\gamma$]$_y$

b. Sem [DISTR([PRED(Participant$_x$)]$_\gamma$)$_y$]$_z$
Morph { V$_\gamma$, Pro$_x$, DISTR$_y$ }$_z$
Phon [...V$_x$ [[$\begin{bmatrix}-\text{CONT}\\+\text{SON}\end{bmatrix}$[V$_\varepsilon$]C(C)(C)]$_\delta$ V$_\varepsilon$ [[$\begin{bmatrix}-\text{CONT}\\+\text{SON}\end{bmatrix}$V(:)C(C)(C)]$_\delta$ ...]$_\gamma$]$_y$]$_z$

(269) Sister schemas for verbal reduplication: alternating class (disyllabic)

a. Sem [PRED(Participant$_x$)$_\zeta$]$_y$
Morph { V$_\zeta$, Pro$_x$ }$_y$
Phon [...V$_x$ [[$\begin{bmatrix}-\text{SYLL}\\ *+\text{CONT}*\\ *+\text{SON}*\end{bmatrix}$VC(C)(C)]$_\eta$ ...]$_\zeta$]$_y$

b. Sem [PRED(Participant$_x$)$_\zeta$]$_y$
Morph { V$_\zeta$, Pro$_x$ }$_y$
Phon [...C$_x$ [[$\begin{bmatrix}-\text{SYLL}\\ *-\text{CONT}*\\ *-\text{SON}*\end{bmatrix}$VC(C)(C)]$_\eta$ ...]$_\zeta$]$_y$

c. Sem [DISTR([PRED(Participant$_x$)]$_\zeta$)$_y$]$_z$
Morph { V$_\zeta$, Pro$_x$, DISTR$_y$ }$_z$
Phon [...V$_x$ [[$\begin{bmatrix} -\text{SYLL} \\ *+\text{CONT}* \\ *+\text{SON}* \end{bmatrix}$ [V$_\vartheta$]C(C)(C)]$_\eta$ V$_\vartheta$ [[$\begin{bmatrix} -\text{SYLL} \\ *+\text{CONT}* \\ *+\text{SON}* \end{bmatrix}$ V(:)C(C)(C)]$_\eta$ ...]$_\zeta$]$_y$]$_z$

d. Sem [DISTR([PRED(Participant$_x$)]$_\zeta$)$_y$]$_z$
Morph { V$_\zeta$, Pro$_x$, DISTR$_y$ }$_z$
Phon [...C$_x$ [[$\begin{bmatrix} -\text{SYLL} \\ *-\text{CONT}* \\ *-\text{SON}* \end{bmatrix}$ [V$_\vartheta$]C(C)(C)]$_\eta$ V$_\vartheta$ [[$\begin{bmatrix} -\text{SYLL} \\ *+\text{CONT}* \\ *+\text{SON}* \end{bmatrix}$ V(:)C(C)(C)]$_\eta$ ...]$_\zeta$]$_y$]$_z$

In each of these schemas, the base to reduplication is the stem, which is — as in Chapter 3 — simply that which follows the pronominal prefix, the latter being some morphological constituent at the left edge of the word that has an interface link to a Pro (pronoun) in the morphosyntax and a participant argument in the semantics. There is no need to explicitly identify any unit here as a 'stem' other than for descriptive convenience: the unit has no special status in the RM formalism. I also assume here, like in many unification-based theories, that the V may itself have internal structure (Jackendoff & Audring 2020, p. 139), such that the stem may contain derivational morphology like the MULT prefix or an incorporated noun, or may be headed by a transparently derived deadjectival verb. So long as the category of the stem as a whole is V, it should be able to unify with one of these schemas (whichever one conforms to the phonological feature subcategorisation).

The difference in meaning between the unreduplicated and reduplicated forms is captured here by a morphosyntactic feature DISTR(ibuted) and a corresponding semantic function DISTR: reduplicated forms should be interpreted as having a greater distribution in time or space. This is obviously vague, but it will serve our purposes here since we are principally focused on the reduplicant shape alternation.

Phonological subcategorisation here works much as it did in Chapter 3. To see this, let us begin with the stop class in (267), which indicates that the pronominal prefix must end in a vowel, and the stem must begin in a stop. From our mini-paradigm in (266), only the verbs of the left-hand column have a stop-initial stem when following a vowel-final prefix. The word *ŋa-puri*, for instance, meets this definition, and indeed it matches the phonological subcategorisation requirements of (267a) only, and not those of any of the other schemas presented here. The stop-initial alternants of the alternating stem class, such as *ŋaɲ-campi:ni*, meanwhile, are not stop-initial when they follow a vowel-final prefix; words like *ŋaɲ-campi:ni* therefore cannot instantiate (267a), and instead meet the requirements of the alternating class schema in (269b).

The two stop class schemas in (267) express the formation of a reduplicated word on the basis

of its unreduplicated counterpart (and vice versa) by means of relational links. There is, first of all, a relational link (α) between the stem in (267a) and the base in the reduplicated stem of (267b), both of which are subcategorised as stop-initial. Reduplication is expressed in (267b) by means of a prefix attaching to this relationally-linked base, consisting of a CV string that is the same as the first CV string of the base it attaches to. This relation of correspondence or identity between reduplicant and base is expressed by the relational link β; this link also relates the corresponding CV portions of the reduplicated stem to the first CV portion of the unreduplicated stem in (267a). Note also that the vowel in the base may be long, but the corresponding vowel in the reduplicant is always short. This pair of sister schemas therefore tells us that, for any verb that has a vowel-final pronominal prefix and a stop-initial stem, where the first vowel is potentially long (267a), we can expect there to be a related form (267b) that also has a vowel-final prefix and the same stop-initial stem, but also repeats the first CV portion of that stem; the first copy of that repeated part has a short vowel, even if the second copy is long. This word with a doubled CV string is interpreted as the same predicate as the one without doubling, but more distributed (in time or space).

The sonorant class functions similarly, but the reduplicant prefix is in this case an open disyllable — that is, it has a CVC(C)(C)V structure — and the initial segment of the base, the reduplicant, and the linked constituent in the unreduplicated counterpart is instead specified to be a sonorant non-continuant. There are, however, some additional complexities introduced by the vowel harmony seen in the reduplicant.

There are a couple of ways the vowel harmony could be addressed. The approach I take here is to index the initial CVC(C)(C) strings in the reduplicant and the base to each other, and to treat the second vowel of the reduplicant as a copy of the first vowel — the second vowel is thus akin to an epenthetic vowel with predictable quality. This vowel is not linked to anything in the morphosyntax or semantics. It is linked relationally to the first vowel in the reduplicant, which is itself linked relationally the first vowel of the base. This forms a line of feature inheritance: whatever the quality of first vowel in the unreduplicated stem is, such will be the quality of the first vowel in the reduplicant, and whatever the first vowel of the reduplicant is, such will be the following vowel. We find thus two relations of identity between word parts, each linked with variable relational links: between the initial CV(C)C strings in base and reduplicant (δ), and between the first vowel of these CVC(C)(C) strings and the vowel sandwiched between them (ε).

Phonological subcategorisation ensures that words participating in this class are those with a stem-initial [+SON,−CONT] segment when following a vowel-final prefix. A form like *ŋa-ŋuni* can therefore unify with the schema in (268a). As with the stop class, the phonological form of the two schemas in (268) is the same, except that (268b) has a copy of the first part of the stem

that it shares with (268a) prefixed to it: the reduplicant. The reduplicant in (268b) consists of the same CVC(C)(C) string as the base it attaches to, plus a V that is identical to the V in that CVC(C)(C) string; in other words, the reduplicant is disyllabic for sonorant-initial bases. Thus, *ŋa-ŋunu-ŋuni* from (266) can instantiate this schema, since it obeys all of its subcategorisation requirements. In the morphosyntax and semantics, the schemas of (268) are exactly parallel to those of (267) — the difference between the classes is strictly phonological, and has no bearing on function.

The alternating class also reduplicates a disyllable, but the segments in the reduplicant and the base are not always identical, due to the regular alternation between stops and continuants that affects stems of this class anyway, as in (269a,b). In (269c), the initial CVC(C)(C) interval is identical between reduplicant and base, and the initial consonants of each are continuants. In (269d), however, the initial consonant of the reduplicant is a stop, while the corresponding consonant of the base is a continuant, as indicated by the 'same-except' asterisks around the [CONT] and [SON] features of each. The pronominal prefix in this context, moreover, is consonant-final. Any one of these four schemas in (269) can be used to predict the others, and none of them overlaps completely with any of the schemas in (267) and (268).

In fact, we can be more specific here. According to the subcategorisation stated in (269b,d), the pronominal prefix is only required to be consonant-final, but in fact that we know that the set of prefixes that actually appear here can be captured with a narrower feature specification: all of them are nasal-final, where the place of the nasal assimilates to the place of the following stem-initial stop. To include this information in (269b,d), we need only unify them with the nasal-final prefix allomorph schemas in (136) of Chapter 3. I have declined to do so here only in order to avoid cluttering the page.

Another alternation from Chapter 3 that I have not explicitly referred to in these schemas is ŋu-Epenthesis. But again, note that the stop class schemas shown here, (267a) and (267b), can unify with any vowel-final prefix. This would include prefixes that are always vowel final, like underlying /ŋa-/ '1sg.R' as well as those that with prefixal *ŋu-*, like '1sg.IRR' when it appears as the *ŋaŋu-* prefix allomorph, as in the bottommost row of the stop class paradigm in (266). As for linking *ŋu-*, it is effectively irrelevant to reduplication. Like any other stem-internal constituent, it will appear in the reduplicant if it is close enough to the left edge to be copied. Since reduplication is blind to stem-internal structure generally, we have no need to specify the presence or absence of linking *ŋu-* in our reduplication schemas; the requirement for linking *ŋu-* in any given stem is, as always, a matter of following the derivational schemas elaborated in §3.6.2 of the previous chapter.

The schemas shown here approach the sonority-conditioned reduplicant shape alternation problem as a problem of predicting an unknown surface word form on the basis of a known one.

If we take any of the words from the paradigm in (266), it will be able to unify with only one of the schemas in (267)–(269), and that schema can then be used to predict the rest of the stem's paradigm. This is because a lexical entry that can unify with one of the above verb schemas will be able to unify with any of the paradigmatically related schemas bearing the same (Greek) relational variable coindex (except where blocked by a more specific entry, in the case of lexical exceptions). A bare stem can therefore be used to predict a reduplicated stem, and vice versa. In the alternating class, a stop-initial stem allomorph can be used to predict a continuant-initial stem allomorph, or a reduplicated continuant-initial stem with a stop-initial allomorph, or one with a continuant-initial allomorph; all these predictions run bidirectionally between any given schema and each other one within the alternating class, which is to say any other schema with a stem bearing the same relational coindex.

The above schemas account for the productive pattern of reduplication among consonant-initial stems, but as always they may be blocked by a more specific lexical entry, such as an idiosyncratic lexical exception. Vowel-initial stems consitute a different phonological class altogether, and raise a different set of issues for an RM analysis. I turn to those issues in the next section.

4.4.5 Reduplication of vowel-initial verb stems

As discussed in §4.2.2, Heath (1984) proposes a reduplication rule that treats underlyingly vowel-initial stems in the same way as other sonorant-initial stems. This rule produces an intermediate output that spells out the content of the reduplicant morpheme as a minimal disyllable: VC(C)(C)V. At this stage of the derivation, an input like /ni-apina/ '3sgM.R-jump.NPST2' produces the output /ni-apa-apina/. Later rules — notably VV-Contraction, which resolves vowel hiatus — obscure the disyllabic origins of this reduplicant, producing the surface form [ni:pa:pina], where the underlying /i-a/ sequence has surfaced as [i:], while /a-a/ is realised [a:].

This is undoubtedly the most challenging characteristic of Wubuy reduplication for an RM account. We must begin with an account of how vowel coalescence should be analysed in RM generally, before combining it with our approach to reduplicative structures specifically.

Part of the problem here is identifying the constituents. Where the coalesced vowel is the product of the underlying final vowel of the prefix and the underlying initial vowel of the stem, the question arises as to whether we regard the surface vowel as 'belonging' to the prefix or the stem. On the grounds that the identities of both the prefix and the stem determine the quality of this surface long vowel, it seems reasonable to regard it as belonging to both constituents. In derivational terms, in any underlying heteromorphemic V1-V2 sequence, one must know the underlying quality of each vowel in order to know what kind of vowel it will yield on the surface. In the terms of RM, one must know what phonological class each morphological

constituent belongs to in order to know what form they take when they appear adjacent to one another. And in order to know that, one must know how the constituents appear in contexts other than this one.

The RM formalism provides a variant of the usual coindexation for situations where phonological constituents partially overlap, for instance in blends (Jackendoff & Audring 2020, p. 112). In this variant, not only the right edge but also the left edge bears a subscript coindex. The general approach is illustrated in (270). The element of the morphosyntax notated X and bearing the variable interface coindex *x* is associated with a phonological string that extends from the left edge of the word to a vowel; the element Y is associated with a string which also includes this vowel, and then extends to the right edge of the word. Since links are already assumed to be potentially many-to-many in RM, the novelty here is purely notational.

(270) Simple overlap schema, where a medial vowel is associated with two larger morphological constituents to its left and right

Morphosyntax	$\{ X_x, Y_y \}_z$
Phonology	$[_x \ldots {}_y V_x \ldots_y]_z$

Vowel coalescence in Wubuy can be analysed in like manner. In what follows, I only elaborate schemas for vowel-initial stems with an initial /a/, and not those with an initial /i/ or /u/. This is partly for reasons of space, and partly because, as discussed in §4.2.5, the set of vowel-initial stems in the lexicon is made up almost exhaustively of /a/-initial ones, with only a few exceptions. There are, then, many opportunities to generalise the /a/-initial stems to new contexts. On the other hand, the set of /u/-initial stems includes only the NPST2 and NPST3 inflections of /(w_2)u/ 'hit' (which are nonetheless very frequent) — once a learner has acquired each possible inflection of these two stems, there are no new stems to generalise to. As for the /i/-initial stems, we also saw that speakers may be unsure as to what the quality of the coalesced vowel should be, and treat them as if they were /a/-initial. This could be taken to suggest that /i/-initial stems must be memorised, but are infrequent enough that those memories are unstable, and may be subject to regularisation in the direction of the more well-established /a/-initial pattern. In any case, the logic of the approach taken here for vowel coalescence in /a/-initial stems should apply equally well for /i/ and /u/-initial stems. The alternations that will be addressed here are therefore those captured in earlier documentation by the rules in (271).

(271) Vowel coalescence rules (Heath 1984, p. 94)

a. /a-a/ → [a:]

b. /u-a/ → [a:]

c. /i-a/ → [iː]

We may begin with the form of the prefixes. The sister schemas in (272) state that any pronominal prefix which is /a/-final preceding a consonant-initial stem will be /aː/-final preceding a vowel-initial stem. However, the initial vowel of the stem and the final vowel of the prefix are the same: they overlap. There is a relational coindex on the vowel-initial stem of (272b), which we will return to. Similar sets of sister schemas for erstwhile underlying /i/-final and /u/-final prefixes follow in (273) and (274), respectively. In the latter case, the quality of the prefix-final vowel changes across contexts: /u/ with consonant-initial stems alternates with /aː/ with vowel-initial stems.

(272) Sister schemas for pronominal prefixes: /a/-final class

	Morphosyntax	$\{ V_x, Pro_\iota \}_y$
a.	Phonology	$[...{}^*a{}^*_\iota\ [C...]_x\ ]_y$
b.	Phonology	$[_\iota... \ [_\nu {}^*aː{}^*_\iota \ ...]_\nu\ ]_y$

(273) Sister schemas for pronominal prefixes: /i/-final class

	Morphosyntax	$\{ V_x, Pro_\kappa \}_y$
a.	Phonology	$[...{}^*i{}^*_\kappa\ [C...]_x\ ]_y$
b.	Phonology	$[_\kappa... \ [_\nu {}^*iː{}^*_\kappa \ ...]_\nu\ ]_y$

(274) Sister schemas for pronominal prefixes: /u/-final class

	Morphosyntax	$\{ V_x, Pro_\lambda \}_y$
a.	Phonology	$[...{}^*u{}^*_\lambda\ [C...]_x\ ]_y$
b.	Phonology	$[_\lambda... \ [_\nu {}^*aː{}^*_\lambda \ ...]_\nu\ ]_y$

The alternations in the form of /a/-initial stems, including under reduplication, is depicted in (275). The semantic level has been omitted for clarity. For the moment, we can set aside the reduplicated stems in (275c,d) to concentrate just on the vowel quality alternations in (275a,b).

(275) Sister schemas for verbal reduplication: vowel class

a. Morphosyntax $\{ V_{\nu}, Pro_{x} \}_{y}$
   Phonology $[_{x}\ldots\ [_{\mu}{}^{*}a{:}^{*}{}_{x}\ C(C)(C)_{\mu}\ \ldots]_{\nu}\ ]_{y}$

b. Morphosyntax $\{ V_{\nu}, Pro_{\kappa} \}_{y}$
   Phonology $[_{\kappa}\ldots\ [_{\mu}{}^{*}i{:}^{*}{}_{\kappa}\ C(C)(C)_{\mu}\ \ldots]_{\nu}\ ]_{y}$

c. Morphosyntax $\{ V_{\nu}, \textsc{Distr}_{\xi}, Pro_{x} \}_{y}$
   Phonology $[_{x}\ldots\ [_{\mu}{}^{*}a{:}^{*}{}_{x}\ C(C)(C)_{\mu}\ [_{\mu}{}^{*}a{:}{}^{*}C(C)(C)_{\mu}\ \ldots]_{\nu}\ ]_{\xi}\ ]_{y}$

d. Morphosyntax $\{ V_{\nu}, \textsc{Distr}_{\xi}, Pro_{\kappa} \}_{y}$
   Phonology $[_{\kappa}\ldots\ [_{\mu}{}^{*}i{:}^{*}{}_{\kappa}\ C(C)(C)_{\mu}\ [_{\mu}{}^{*}a{:}{}^{*}C(C)(C)_{\mu}\ \ldots]_{\nu}\ ]_{\xi}\ ]_{y}$

The analysis here exploits an assumption in RM that a more specific schema takes logical precedence over a more general one, and that a relational variable is more specific than an interface variable (since, in referring to a sister schema, it specifies more). In the phonology of (275a), there is a pronominal prefix at the left edge of the word, notated with the interface variable *x*, and there is a stem at the right edge of the word, notated with the relational variable ν. These constituents overlap: the final segment of the prefix, /a:/, is also the first segment of the stem. This segment is asterisked, as it alternates across its realisations. In (275b), another realisation of the stem indexed with ν, the overlapping vowel is instead /i:/.

The prefix in (275b) is more restricted than in (275a): it is notated with the relational variable κ, and can consequently only be instantiated by a prefix that is able to instantiate all such variables in all the schemas they appear in, such as those in our /i/-final prefix schema (273) above. So, the kind of prefix that can unify with (275b) is that which can unify with (273b), which must also be able to unify with our /i/-final (273a) — these are the prefixes that would be underlyingly /i/-final in a derivational model. Since the prefix slot in (275b) is more specific than (275a), it takes precedence. That is, if the prefix is one which is otherwise (before consonants) /i/-final, it must unify with (275b) rather than (275a), and the coalesced vowel that appears over the prefix-stem boundary must therefore be /i:/. If the prefix has any other kind of final vowel, like /a/ or /u/, then it will unify with the less specific (275a), and the coalesced vowel will be /a:/.

Finally, we are at the point where we can examine the reduplicative schemas (275c) and (275d). The reduplicant is understood here to be a V:C string, which attaches to the base as a prefix, and is a copy of the initial V:C string of the base, as suggested in §4.2.5. Like in our consonant-initial reduplication schemas above, the copying relation is notated with a relational variable. The relational variable μ establishes the link between the corresponding V:C string in the reduplicant and the base, and between each of these and the corresponding string in the non-reduplicative forms (275a) and (275b).

As the asterisks indicate, the quality of that V: alternates across contexts. The reduplicant-initial vowel is /i:/ in exactly the same conditions as any stem-initial vowel — namely, when it follows a prefix that would be /i/-final before a consonant-initial stem. Following other kinds of prefixes, it is /a:/, and the initial vowel it is linked to in the base is always /a:/.

4.4.6 Summary of RM analysis

The consonant-initial verb schemas in (267–269) and the vowel-initial verb schemas in (275) are able to capture the sonority-conditioned reduplicant shape alternation as an unmediated relation between full surface word forms in the lexicon. There is no suggestion here that speakers must draw on such schemas any time they utter a reduplicative word. In RM, we assume that most of the time, speakers simply retrieve reduplicated words from memory, and only draw on schematic knowledge when generalising to new items. The theory also assumes that speakers may generalise patterns in the lexicon at any level of granularity. The schemas here are developed at a very high level of abstraction, organised into broad features like [SON] or [CONT], but it may be that speakers in fact draw on narrower generalisations, such as at the level of the individual phoneme (Albright 2002; Albright & Hayes 2003; Ernestus & Baayen 2004), as suggested might be required in the previous chapter.

However, although these schemas are abstract, they do not involve any hidden (concrete, phonological) structure, and mappings between lexical entries are direct: there are no underlying or intermediate forms, and no special RED morpheme. Neither do we need any devices to manage the interactions between different layers of phonological representation: there is no need to postulate separate sub-grammars, cophonologies, alternative constraint rankings or any other generative device. The devices used to relate an unreduplicated word to a reduplicated one — variables and indices — are just those required for morphological relations generally. As I will show in the next section, both BRCT and MDT require a range of reduplication-specific devices, and various hidden levels of structure.

Whatever processes may have shaped the Wubuy lexicon historically, such as those explored in §4.3, are not invoked here as a synchronically active force. That is, we assume the speaker is producing monosyllabic reduplicants for stop-initial stems and disyllabic reduplicants for sonorant-initial verbs for no other reason, synchronically, than the fact that the speaker's mental lexicon already contains reduplicated words with stop-initial stems and monosyllabic reduplicants, and reduplicated words with sonorant-initial stems and disyllabic reduplicants. If there is any process of 'optimisation' here, it is a diachronic one only. As we will see in the following section, this stands in stark contrast to BRCT, according to which the pattern *must* be analysed a reflection of synchronic optimisation. This theoretical commitment makes the Wubuy shape alternation intractable in a BRCT analysis.

4.5 Comparison with other models

4.5.1 Base-Reduplicant Correspondence Theory

Reduplication has long been central to the development of phonological theory. Nowhere is this more apparent than in the early development of Optimality Theory, where insights gleaned from the analysis of reduplication led to significant refinements in the machinery available for analysing phonological patterns outside of reduplication.[42] Reduplication remains one of the central empirical domains of OT research, and continues to drive further developments of the theory (e.g. McCarthy, Kimper et al. 2012; Saba Kirchner 2013; Zukoff 2020)

Base-Reduplicant Correspondence Theory began as McCarthy & Prince's attempt to incorporate the central insights of their Prosodic Morphology research program (McCarthy & Prince 1986) into the then-new violable constraint architecture of Optimality Theory (McCarthy & Prince 1993b, p. 110). Prosodic Morphology sought to understand a range of phenomena in which morphological form is sensitive to prosodic structure. A key device is the 'prosodic template', which constrains a morphological unit to have a certain shape — for instance, a heavy syllable, or a disyllabic foot, and so on. Its core claims are summarised in the following three principles:

(276) Prosodic Morphology within Optimality Theory (McCarthy & Prince 1993b, p. 110)

a. *Prosodic Morphology Hypothesis*
Templates are constraints on the prosody/morphology interface, asserting the coincidence of morphological and prosodic constituents.

b. *Template Satisfaction Condition*
Templatic constraints may be undominated, in which case they are satisfied fully, [or] they may be dominated, in which case they are violated minimally, in accordance with general principles of Optimality Theory.

c. *Ranking Schema*
P » M [constraints on prosodic structure take precedence over constraints on morphological structure]

[42] Take, for example, the now-ubiquitous constraint MAximality, which penalises phonological deletion: what began as a device to penalise any kind of mismatch between the form of the reduplicant and the form of the base (McCarthy & Prince 1993b), and was then refined to only penalise elements of the base that are absent from the reduplicant (McCarthy & Prince 1994), is eventually generalised to penalise any segments in the underlying form (the input) that do not appear in the output (McCarthy & Prince 1995), and has therefore effectively become a standard component of OT analysis inside and outside of reduplication.

For reduplication, the crucial morphological constituent referred to by the first of these principles is the morpheme RED, whose exponent in the output is termed the *reduplicant*. RED differs from ordinary fixed-segment (i.e. non-reduplicative) morphemes in that the lexicon does not provide it with any phonetic content — its content must instead be supplied by constraints governing the output (McCarthy & Prince 1993b, p. 66). Consequently, whereas fixed-segment morphemes are subject to constraints demanding faithfulness to their input underlying representations, prosodic morphemes like RED have no underlying representation to be faithful to. Therefore, input-output correspondence constraints have no scope over the content of the reduplicant.

Instead, the reduplicant must be faithful to the *base*; this is what gives reduplication its 'copying' character. This faithfulness requirement is formalised as the family of base-reduplicant correspondence constraints (MAX-BR, DEP-BR and IDENT-BR), which penalise any deviation from total identity between base and reduplicant (McCarthy & Prince 1995). As such, only total reduplication obeys these constraints absolutely.

These correspondence constraints are in tension with templatic alignment constraints (McCarthy & Prince 1993a), which require that the reduplicant in the output be of a certain prosodic shape, preferably an 'authentic unit of prosody': e.g. a light syllable, a heavy syllable, foot, prosodic word. Partial reduplication results when a language ranks such a prosodic alignment constraint higher than MAX-BR, which demands full copying of the base, and the prosodic category in question is smaller than the base.

As we shall see in the following pages, there is no implication here that only one reduplication pattern is possible in a given language. Because templatic alignment constraints are violable, as per (276b) and the general constraint architecture of OT, the actual shape that the reduplicant takes within a language may vary. But by (276c), BRCT makes the prediction that such variability should be optimising, where what is 'optimal' is determined by language-particular rankings of universal prosodic constraints. As I will demonstrate in this section of the chapter, the reduplicant shape alternation we find in Wubuy is not possible to analyse in BRCT precisely because it is non-optimising.

We can begin building up a BRCT account of Wubuy reduplication by defining the following constraints. (The linear position of the exponent of RED and the delimitation of the domain that it copies from are formalised with constraints from the Alignment and Anchoring families. Since they are never violated, they are undominated, and have no bearing on the analysis. I therefore omit these constraints from the tableaux to follow, and I do not consider any candidates that would violate them.)

(277) **Align(Red,R,Stem,L)**
The reduplicant is at the left edge of the stem. (Red is a prefix.)

(278) **Anchor(Red,L,Base,L)**
The leftmost element of the reduplicant corresponds to the leftmost element of the base. (Copying proceeds rightward)

Accounting for the size of the partial reduplicant instead depends on the relative ranking of the following constraints (the first two are of the type sometimes referred to as 'size restrictors'):

(279) **Red=σσ**
Align the left and right edges of the reduplicant with the left edge of a syllable and the right edge of a following syllable. (The reduplicant must consist of two syllables.)

(280) **Red=σ**
Align both edges of the reduplicant with both edges of a single syllable. (The reduplicant must consist of one syllable.)

(281) **Max-BR**
Every element of the base has a correspondent element in the reduplicant. (No partial reduplication.)

Red=σσ and Red=σ are templatic constraints on reduplicant shape. When ranked highly, they have the effect of restricting the size of the reduplicant to the specified prosodic unit. Technically, they are alignment constraints, understood in terms of McCarthy & Prince's (1993a) theory of Generalised Alignment, and are satisfied when a designated morphological constituent edge and a designated prosodic constituent edge align. The reduplicant template constraints each subsume constraints on the alignment of left edges and right edges — e.g. Red=σ subsumes Align(Red,L,σ,L) and Align(Red,R,σ,R). This simplification does not affect the analysis: the left and right constraints are not in conflict with one another and must occupy the same position in the ranking.[43] Where misalignment occurs, I count violations by the number of syllables intervening between the designated morphological boundary and the designated prosodic boundary. Red=σσ (called Disyll in McCarthy & Prince 1993b, p. 87) is somewhat non-standard, as 'two syllables' is not strictly an 'authentic unit of prosody' as required in much work on Prosodic Morphology (e.g., McCarthy & Prince 1986, 1990); the reader may prefer

[43] Although for criticisms of equating prosodic templates with alignment constraints, see Baker (2008b, 159ff).

an alternative like RED=Ft or RED=PrWd. Ultimately, however we characterise the prosodic template responsible for the disyllabic reduplication pattern — σσ, Ft, PrWd — has little bearing on the principal issue under examination here, which is how it interacts with the other templatic constraint, RED=σ. Note also that the universalising assumptions of the requirement to express prosodic morphological constituents exclusively in terms of 'authentic units of prosody' is questioned by typological work in this area, which finds many cases of reduplicants with prosodic characteristics which are neither authentic nor optimising (e.g. Baker 2014b; Blevins 2003; Yu 2007).

MAX-BR is the identity constraint demanding a full copy of the reduplicant. The templatic constraints are clearly in direct conflict with MAX-BR: where MAX-BR demands the reduplicant copy as much material as possible from the base (which may be many syllables long), the size restrictors demand that the reduplicant be no more than a syllable or a foot in size. Partial reduplication results when templatic constraints on the size of the reduplicant outrank MAX-BR.

Temporarily setting aside the problem of the conflict between the monosyllabic and disyllabic reduplication patterns, we can see the work of ranking the templatic constraint RED=σσ above MAX-BR in the tableau (282), for the earlier example (181b) *wini-maṯa-maṯant-aŋi* 'those two men kept paddling'. The reduplicant is the underlined portion in each output candidate. Candidate *b* loses to *a* because, although *b* satisfies MAX-BR, it violates the highly-ranked templatic constraint RED=σσ by being too large: the right edge of the reduplicant occurs two syllables too late. RED=σσ incurs a violation mark for each of these syllables. Candidate *c* also violates RED=σσ, as its reduplicant is too small. This time, the right edge of the second syllable occurs one syllable *later* than the edge of the reduplicant, and therefore incurs one violation. Note that I do not provide an analysis of syllabification here; it is not crucial to the concerns of the chapter.

(282) Partial reduplication: RED=σσ » MAX-BR

/wini, RED, maṯantaŋi/	RED=σσ	MAX-BR
☞ a. wi.ni.-ma.ṯa.-ma.ṯan.ta.ŋi		ntaŋi
b. wi.ni.-ma.ṯan.ta.ŋi.-ma.ṯan.ta.ŋi	∗!∗	
c. wi.ni.-ma.-ma.ṯan.ta.ŋi	∗!	ṯantaŋi

This then raises the question of why the reduplicant does not copy more segments from the base while remaining disyllabic, as in **wini-maṯan-maṯant-aŋi*, since a candidate like this would satisfy RED=σσ and incur fewer violations of MAX-BR. In other words, what ensures that the reduplicant copies the minimum amount of material required by the templatic constraint, and not the maximum? We can restrict the winning candidate to have a vowel-final disyllabic

reduplicant with the standard OT constraints NoCoda, which penalises codas, and Max-IO, which penalises deletion of input segments (McCarthy & Prince 1995).

(283) **NoCoda**
Assign one violation mark for every output segment syllabified as a coda.

(284) **Max-IO**
Every element of the input has a correspondent element in the output. (No deletion.)

There is an implied conflict between the two constraints: a candidate can satisfy NoCoda by deleting consonants found in the input, but such a candidate will violate Max-IO. However, Max-IO can only penalise deletion from lexically specified morphemes, not from the reduplicant, which is the domain of Max-BR — recall that, according to the theory, Red has no input segmental specification, so Max-IO has no scope over it. The ranking Max-IO»NoCoda ensures that deletion does not occur outside the reduplicant in order to appease NoCoda. Ranking NoCoda over Max-BR ensures that an output that keeps the reduplicant short to avoid a coda is preferred over copying more of the base. Thus, candidate *b* in (285) loses to *a* due to its unnecessary violation of NoCoda. Though *c* causes no violations of NoCoda, it loses to both the other candidates by violating the higher-ranked Max-IO.

(285) Avoidance of final coda in reduplicant: Max-IO » Red=σσ » NoCoda » Max-BR

/wini, Red, maṯantaŋi/	Max-IO	Red=σσ	NoCoda	Max-BR
☞ a. wi.ni.-ma.ṯa.-ma.ṯan.ta.ŋi			n	ntaŋi
b. wi.ni.-ma.ṯan.-ma.ṯan.ta.ŋi			nn!	taŋi
c. wi.ni.-ma.ṯa.-ma.ṯa.ta.ŋi	n!			taŋi

Note also that a candidate which avoids violations of NoCoda by deleting coda segments from the reduplicant but not the base, such as **ŋaŋu-muɾu-muɾkuʈi-l̪a-ŋi* in (286c), can be ruled out by Contig-BR (a 'contiguity' constraint), which demands that the corresponding strings in base and reduplicant be contiguous strings (McCarthy & Prince 1995, p. 123). Since no violations of this are attested, it is presumably undominated. I therefore omit it from all tableaux to follow.

(286) Contig-BR protects coda in reduplicant if it has a correspondent in the base:

/ŋaŋu, Red, muɾkuʈi, l̪aŋi/	Contig-BR	Max-IO	Red=σσ	NoCoda	Max-BR
☞ a. ŋa.ŋu.-muɾ.ku.-muɾ.ku.ʈi-l̪a.ŋi				ɾɾ	ʈil̪aŋi
b. ŋa.ŋu.-mu.ɾu.-mu.ɾu.ʈi.-l̪a.ŋi		k!			ʈil̪aŋi
c. ŋa.ŋu.-mu.ɾu.-muɾ.ku.ʈi.-l̪a.ŋi	k!			ɾ	kʈil̪aŋi
d. ŋa.ŋu.-mu.ku.-muɾ.ku.ʈi.-l̪a.ŋi	ɾ!			ɾ	ɾʈil̪aŋi

The particular ranking we arrive at here is an instance of a more general one: Faith-IO » Markedness » Faith-BR. This is the ranking responsible for the phenomenon termed 'The Emergence of the Unmarked', whereby the effects of a markedness constraint can be seen in the reduplicant but not in those parts of the word that have an input specification, where higher-ranking faithfulness constraints can take effect (Gafos 1998; McCarthy & Prince 1994; Spaelti 1996; Urbanczyk 2001; Ussishkin 2005; cf. Blevins 2003; Inkelas & Zoll 2005).

In the specific examples of *wini-maṯa-maṯant-aŋi* and *ŋaŋu-muɾku-muɾkuṯi-l̪-aŋi*, it so happens that the first two vowels of the base are the same in quality and quantity, as are the first two vowels of the reduplicant. But as discussed above, the vowels of the disyllabic reduplicant will have these properties even where the base vowels differ in quality, or are long. I set aside those issues here, since they do not illuminate the main problem, but note that they would also amenable to an analysis in terms of the Emergence of the Unmarked: a constraint enforcing vowel harmony or penalising long vowels would be defeated by input-output faithfulness in lexically specified morphemes but not in the reduplicant.

On the face of it, monosyllabic reduplication should be derived the same way as we have seen for disyllabic reduplication. The ranking Red=σ » Max-BR generates partial reduplication with a monosyllabic reduplicant, while Max-IO » NoCoda ensures that codas are avoided in the reduplicant (but not elsewhere), thus preferring CV over CVC reduplicants.

(287) CV reduplication with Red=σ

/waɾa, RED, ṯuɾapata/	Max-IO	Red=σ	NoCoda	Max-BR
☞ a. wa.ɾa-ṯu.-ṯu.ɾa.pa.ta				ɾapata
b. wa.ɾa-ṯuɾ.-ṯu.ɾa.pa.ta			ɾ!	apata
c. wa.ɾa-ṯu.ɾu.-ṯu.ɾa.pa.ta		∗!		pata

4.5.2 Unifying the rankings

The problem, of course, is that there is no straightforward way to unify the analyses of monosyllabic and disyllabic reduplication without ensuring that one or the other reduplicant size always wins, irrespective of the sonority of the base-initial segment. Returning to the examples analysed above, consider what happens if we attempt to unify the rankings such that the size-restricting constraints are ranked equally, as in (288). Candidate *b*, with a heavy monosyllable as the reduplicant, is ruled out by NoCoda, as before. Max-BR will emerge to break the tie between the monosyllabic and disyllabic candidates that remain (*a* and *c*, respectively). Since Max-BR favours copying every segment of the base, it will always prefer the candidate with the largest reduplicant, which is the disyllabic one (288c).

(288) Max-BR breaks the tie between Red=σ and Red=σσ, favouring disyllabic reduplicant:

/wara, RED, ʈurapata/	Max-IO	Red=σ	Red=σσ	NoCoda	Max-BR
☞ a. wa.ɾa-ʈu.-ʈu.ɾa.pa.ta			*		ɾapat!a
b. wa.ɾa-ʈuɾ.-ʈu.ɾa.pa.ta			*	ɾ!	apata
☹ c. wa.ɾa-ʈu.ɾu.-ʈu.ɾa.pa.ta		*			pata

An additional side effect of this unified ranking — which does not occur with the separate monosyllabic and disyllabic rankings — is that disyllabic reduplicants with a medial coda will always lose. This is because NoCoda becomes decisive: a disyllabic reduplicant (289a) and a monosyllabic reduplicant (289d) incur the same number of violations of the templatic constraints, but the disyllabic reduplicant, by copying more material, contains an extra coda.

(289) When the domain for copying contains a medial coda, monosyllabic reduplicant incurs fewer violations of NoCoda:

/ŋaŋu, Red, muɾkuʈi, ɭaŋi/	Max-IO	Red=σ	Red=σσ	NoCoda	Max-BR
☞ a. ŋa.ŋu.-muɾ.ku.-muɾ.ku.ʈi.-ɭa.ŋi		*		ɾɾ!	ʈiɭaŋi
b. ŋa.ŋu.-mu.ɾu.-mu.ɾu.ʈi.-ɭa.ŋi	k!	*			ʈiɭaŋi
c. ŋa.ŋu.-muɾ.-muɾ.ku.ʈi.-ɭa.ŋi			*	ɾɾ!	kuʈiɭaŋi
☹ d. ŋa.ŋu.-mu.-muɾ.ku.ʈi.-ɭa.ŋi			*	ɾ	ɾkuʈiɭaŋi

We get no relief from promoting one size restrictor constraint to crucially dominate the other. If Red=σ dominates Red=σσ, as in (290), the effect is to uniformly prefer monosyllabic reduplicants irrespective of the base. If Red=σσ dominates Red=σ, as in (291), the inverse problem holds: the ranking uniformly prefers disyllabic reduplicants irrespective of the base.

(290) Red=σ » Red=σσ ensures disyllabic reduplicant can never win:

/wini, Red, maʈantaŋi/	Max-IO	Red=σ	Red=σσ	NoCoda	Max-BR
☞ a. wi.ni.-ma.ʈa.-ma.ʈan.ta.ŋi		*!		n	ntaŋi
b. wi.ni.-maʈ.-ma.ʈan.ta.ŋi			*	ʈn!	uŋun
☹ c. wi.ni.-ma.-ma.ʈan.ta.ŋi			*	n	ʈantaŋi

(291) Red=σσ » Red=σ ensures monosyllabic reduplicant can never win:

/wara, RED, ʈurapata/	Max-IO	Red=σσ	Red=σ	NoCoda	Max-BR
☞ a. wa.ɾa-ʈu.-ʈu.ɾa.pa.ta		*!			ɾapata
b. wa.ɾa-ʈuɾ.-ʈu.ɾa.pa.ta		*!		*	apata
☹ c. wa.ɾa-ʈu.ɾu.-ʈu.ɾa.pa.ta			*		pata

The different constraint rankings and their respective predictions for the stems /maṯant-aŋi/ ‘paddle-Past2’, /ṯurapata/ ‘white person’ and /muɾkuʈi-ḻa-ŋi/ ‘ribs-chop-Past2’ are summarised below.

(292) a. Max-IO » Red=σ » Red=σσ » NoCoda » Max-BR
A monosyllabic reduplicant is always more harmonic.
e.g., /muɾuŋun/ ⟶ *mu-muɾuŋun
/muɾkuʈi-ḻa-ŋi/ ⟶ *mu-muɾkuʈi-ḻa-ŋi
/ṯurapata/ ⟶ ṯu-ṯurapata

b. Max-IO » Red=σσ » Red=σ» NoCoda » Max-BR
A disyllabic reduplicant is always more harmonic.
e.g., /muɾuŋun/ ⟶ muɾu-muɾuŋun
/muɾkuʈi-ḻa-ŋi/ ⟶ muɾku-muɾkuʈi-ḻa-ŋi
/ṯurapata/ ⟶ *ṯuru-ṯurapata

c. Max-IO » {Red=σ, Red=σσ}» NoCoda » Max-BR
A disyllabic reduplicant is more harmonic, unless the first syllable of the base has a coda, in which case a monosyllabic reduplicant is more harmonic.
e.g., /muɾuŋun/ ⟶ muɾu-muɾuŋun
/muɾkuʈi-ḻa-ŋi/ ⟶ *mu-muɾkuʈi-ḻa-ŋi
/ṯurapata/ ⟶ *ṯuru-ṯurapata

As can be seen, we are undoubtedly faced with a ranking paradox. Although a standard BRCT analysis of partial reduplication for each shape allomorph works separately (with just one operative templatic constraint in each case), as soon as we try to unify the analyses, we are forced to always prioritise one shape at the expense of another. The crux of the problem, then, is how to relate each shape allomorph to its conditioning environment: the sonority of the initial segment of the base. As we will see, BRCT and its offshoots have suggested various approaches to reduplicant shape allomorphy, but none of them is applicable to the present conundrum.

4.5.3 Morphologically indexed constraints

Finding multiple reduplicant shapes within the same language is hardly unusual, and much of BRCT’s theoretical machinery was developed precisely in order to address shape alternations. But the alternations analysed in this extensive literature generally pursue one of two strategies, neither of which appear to be available in the present case. One option is to attribute the different reduplicant shapes to distinct instances of Red supplied by the lexicon, where each instance differs in its grammatical function, or morphological category (e.g. root vs affix).

Another option is to argue that, despite appearances, the alternation really is phonologically optimising. Finally, when these two options fail to yield a compelling result, one may be forced to accept parochial constraints that apply only to this alternation.

Urbanczyk's (2001, 2006) analysis of Lushootseed reduplication can serve as an example of the first approach. In this language, CVC- and CV- reduplicants are associated with distributive and diminutive functions, respectively, as in [bəd-bəda?] 'children' and [ča-čaləs] 'little hand'. Since they are associated with different functions, they can be subject to different instances of MAX-BR, ranked differently with respect to markedness constraints. MAX-BR_{Dist} is ranked above NOCODA, predicting a larger reduplicant, while MAX-BR_{Dim} is ranked below it.

(293) Two morphologically-indexed reduplicant shapes in Lushootseed (adapted from Urbanczyk 2001, 2006)

a. CVC Distributive reduplication: MAX-BR_{Dist} » NOCODA

/MAX-BR_{Dist}, bəda?/	MAX-BR_{Dist}	NOCODA
☞ a. bəd-bəda?	**	**
b. bə-bəda?	***!	*

b. CV Diminutive reduplication: NOCODA » MAX-BR_{DIM}

/MAX-BR_{DIM}, čaləs/	NOCODA	MAX-BR_{DIM}
a. čal-čaləs	**!	**
☞ b. ča-čaləs	*	***

The two partial rankings in (293) can be conflated to MAX-BR_{Dist} » NOCODA » MAX-BR_{Dim} without any change in predictions, since the domains over which the MAX constraints have scope do not overlap. Note that in this analysis, reduplicative templates are done away with altogether, and differences between reduplicant shapes are attributed to different coexisting — and differently ranked — correspondence constraints. In fact, Urbanczyk ultimately takes this even further, associating the different patterns to the morphological category of the respective RED morphemes. The distributive RED is classified as a root, and is subject to all correspondence constraints indexed to roots, while diminutive RED is analysed as an affix. This, she argues, accords with the observation that fixed-segment (i.e. lexical, as opposed to reduplicative) roots in Lushootseed permit codas (and marked structures generally), while fixed-segment affixes do not. Eliminating templatic constraints from OT prosodic morphology constitutes the main aim of Generalised Template Theory (GTT; other examples include Alderete & MacMillan 2015; Gafos 1998; Spaelti 1996; cf. Caballero 2006).[44]

[44]Relatedly, in Ngalakgan, there are two CV?- reduplicants that are both prefixes, but, according to Baker's (2008b,

There are two key ingredients in the Lushootseed data that make this sort of analysis possible. First of all, there are genuine differences in the grammatical functions of the CV- and CVC-reduplicants. The same base can be submitted to each of the reduplicative patterns, and produce distinct meanings which are a function of the base and the reduplicant shape: e.g. /dʒə́səd/ 'foot', /dʒə́s-dʒəsəd/ 'feet' (DIST-foot), /dʒí-dʒəsəd/ 'little foot' (DIM-foot). If each shape is associated with a different meaning, it is reasonable to suppose that they are lexically distinct.

Secondly, roots and affixes in Lushootseed are associated with different phonological characteristics outside of reduplication — i.e. codas appear in roots but not in affixes. These same characteristics differentiate the reduplicant shapes — the distributive reduplicant has a coda, while the diminutive does not. This allows Urbanczyk to make the stronger claim that the CVC- distributive is not only formally similar to other roots, but really *is* a root, just as the CV-diminutive reduplicant really is an affix. If, as she assumes, differences in the phonological properties of roots and affixes are the product of general constraints on the morphology-phonology interface in an optimising grammar, she can suppose that the reduplicant shapes follow from this grammar also.

Either of these ingredients might be considered sufficient grounds for asserting that two reduplicant shapes correspond to two RED morphemes, but neither of them happens to be true in Wubuy. The monosyllabic and disyllabic reduplicants in Wubuy are, as was amply demonstrated above, not associated with different functions. The range of interpretations available to reduplicated words is the same whether it features a monosyllabic or disyllabic reduplicant, and a given base has only one possible reduplicative pattern that can apply to it, which is (notwithstanding a few exceptions) always determined by the sonority of its initial consonant.

We cannot establish any general association between prosodic shape and morphological category in Wubuy either. Both roots and affixes can be monosyllabic, disyllabic, or indeed longer; they can contain syllables with long or short vowels, with or without a coda; they can begin with a consonant or a vowel, a sonorant or a stop. While it is true that there are no underlyingly stop-initial suffixes, there are stop-initial prefixes (e.g. *pa-* '2sg.IRR', *para-* '2sg>3pl.R', *pampa-* '2sg>3pl.IRR', and six others), and the reduplicant appears to the left of its base, like a prefix. There are of course also sonorant-initial prefixes with equivalent prosodic shapes (e.g. *na-* '12du.R', *nara-* '12du>3pl.R', *nampa-* '12du>3pl.IRR', etc.). One major difference between verbal prefixes and verbal roots is that the former have much more constrained phonotactics (no heteroganic clusters, for instance). The reduplicant cannot be either a root or

pp. 104–108) level-ordered analysis, attach to different morphological domains — the ROOT level and the WORD level. The phonological characteristics associated with each of these domains generally is also reflected in the CV?- reduplicants (e.g. only the WORD-level reduplicant constitutes its own prosodic word, and undergoes monosyllabic lengthening).

an affix, though: if it were, reduplicative constructions would be equivalent to compounding or derivational constructions, and should therefore trigger ŋu-Epenthesis preceding stop-initial bases, incorrectly predicting *ʈuɾapata* 'white person' to reduplicate as **ʈu-ŋu-ʈuɾapata* instead of *ʈu-ʈuɾapata*.

In Lushootseed, the semantic and phonological differences between distributive and diminutive reduplication patterns offer independent but converging evidence that these patterns are associated with different morphemes in the lexicon. A grammar that can index its constraints to morphological category can produce different outcomes for roots versus affixes. Since these semantic and phonological properties are not apparent in Wubuy, we cannot solve the shape alternation problem by means of morphologically-indexed constraints.

4.5.4 Making the alternation optimising

This is leaves us with the second option: deriving the shape alternation by general constraints on phonological well-formedness. For this to work, we would have to show that the alternation is the result of a repair strategy — that is, that the output would be more marked if there were no reduplicant shape alternation.

Given that the alternants differ in their prosodic templates, we might assume that they repair some dispreferred prosodic structure — the alternation might, for instance, avoid a stress clash, or restore a degenerate foot. But we simply have no evidence that this is the case. We have seen that the base can be of any number of syllables, and consist of any segmental string which is possible in stems generally.

But also, crucially, we would have to show that it is something about the sonority of the initial segment that demands repair in this way. We would have to show that a morphological constituent (i.e. reduplicant) of two syllables is somehow generally less marked *when the constituent it attaches to is sonorant-initial*, and a morphological constituent of one syllable is less marked *when the constituent it attaches to is stop-initial*.[45]

We need to discover some high-ranking markedness constraint — let's call it M — that penalises any candidate with a stop-initial affix of more than one syllable, or a sonorant-initial candidate of less than two syllables. Since non-reduplicative affixes do not similarly alternate in size according to the sonority of the initial segment of their host, M must be dominated by Max-IO, as in (294).

[45]We could frame this instead as the relative markedness of the initial segment of the reduplicant, rather than the base, but this can be violated. Recall that the disyllabic reduplicant can surface as stop-initial as a result of Hardening — that is, in order to satisfy the undominated constraint on sonority sequencing.

(294) RED=σσ violated in order to satisfy some unidentified constraint, M:

/RED, {muɾuŋun ~ t̪uɾapata}/	MAX-IO	M	RED=σ	RED=σσ	NOCODA	MAX-BR
☞ a. mu.ɾu.-mu.ɾu.ŋun			**		*	ŋun
b. muɾ.-mu.ɾu.ŋun		*!		**	**	uŋun
c. mu.-mu.ɾu.ŋun		*!		**	*	ɾuŋun
d. t̪u.ɾu.-t̪u.ɾa.pa.ta		*!	**			pata
e. t̪uɾ.-t̪u.ɾa.pa.ta				**	*!	apata
☞ f. t̪u.-t̪u.ɾa.pa.ta				**		ɾapata

Of course, by this point, we are really stretching the theory of markedness beyond the limits of credibility. Constraint M is nothing more than an ad-hoc attempt to stipulate the resolution of a ranking paradox; it has nothing to do with markedness per se. Note also that a solution of this kind would be unavoidable if we were to jettison templatic constraints or even all reduplication-specific constraints altogether, as many OT theories subsequent to BRCT sought to do, as these all typically subscribe to the P»M postulate of the Prosodic Morphology Hypothesis (e.g. Hendricks 1999; Saba Kirchner 2010; Spaelti 1997).

Ryan & Heath (2021) — to date, the only other analysis of the Wubuy reduplicant shape alternation to appear in the literature since Heath (1984) — end up in an analogous dilemma. They suggest that the appearance of a size alternation belies an underlying preference for a minimally bimoraic reduplicant, and that low-sonority onsets (i.e. stops) contribute weight in the reduplicant (and only in the reduplicant). TV- reduplicants (where T is any stop) and NVCV- reduplicants (where N is any sonorant) are just different ways of satisfying a single constraint demanding a bimoraic minimum in the reduplicant. So the monosyllabic reduplicant in *t̪i-t̪ilila* 'RED-drip' and the disyllabic reduplicant in *l̪iri-l̪iɾipala* 'RED-inside' equally satisfy this single bimoraic template: the stop onset in the former reduplicant contributes one mora in addition to the mora of the vowel, while the two vowels in the latter contribute one mora each. Thus, rather than a ranking in which one templatic constraint must be violated in favour of a second templatic constraint in some context, they suggest that both outcomes can emerge from a single templatic constraint. But although the technical implementation differs from that in (294), the logic is largely the same: they seek to argue that the alternation is, despite appearances, optimising, so long as we assume some overarching markedness consideration which is not otherwise relevant outside the alternation in question.

The fundamental empirical claim, it should be noted, is based on the observation that patterns of onset weight sensitivity are claimed to exist in other languages — a claim with which I have no quarrel (Gordon 2005; Ryan 2014). We have no evidence that Wubuy prosody is sensitive to onset sonority — in stress assignment, for instance — and Ryan & Heath (2021) do not provide

any, as they acknowledge. Indeed, as I show in Chapter 6, the evidence we do have for weight sensitivity in Wubuy is in the rime, affecting the placement of phrase-final pitch accent. Ryan & Heath's solution therefore only works by means of a hopeful but empirically unsupported redefinition of the two reduplicant shapes, such that they can be regarded as realising one invariant shape: a bimoraic one.

4.5.5 Resorting to parochial constraints

Given that we cannot index the different templates to different RED morphemes, and attempting to cast the alternation in optimising terms remains stipulative, we may feel it necessary to finally resign to an even more stipulative solution, proposing parochial constraints that combine the sonority of the initial segment of the reduplicant with its prosodic shape, as in (295). The more general RED=σσ can then be ranked either equal to or immediately below it, thus functioning as a kind of default which can be satisfied when no initial stop is available to satisfy the monosyllabic templatic constraint.

(295) **RED=TV**
'Align both edges of the reduplicant with both edges of a stop-initial syllable'.

It is difficult to determine how violations of a constraint like (295) should be counted; as a simplifying assumption, I will treat it here as a binary choice between one violation mark or none. And, in order to compare the templatic constraints fairly, I will also regard violation of RED=σσ as binary. This does not affect the central argument.

Now consider a tableau like (296). As it stands, the constraint performs its role well enough. The constraint is decisive in selecting monosyllabic reduplicants as optimal when they are stop-initial, while it rejects all sonorant-initial reduplicants no matter their shape, and 'passes them on' to be adjudicated by the following constraint, RED=σσ.

(296) RED=σσ employed as the default shape when no stop is available to satisfy RED=TV:

/RED, {muɾuŋun ~ t̪uɾapata}/	MAX-IO	RED=TV	RED=σσ	NOCODA	MAX-BR
☞ a. mu.ɾu.-mu.ɾu.ŋun		*		*	ŋun
b. mu.-mu.ɾu.ŋun		*	*!	*	ɾuŋun
c. t̪u.ɾu.-t̪u.ɾa.pa.ta		*!			pata
☞ d. t̪u.-t̪u.ɾa.pa.ta			*		ɾapata

However, recall that underlyingly continuant-initial stems — i.e. not those in (296), which are both non-continuant-initial — take a disyllabic reduplicant that hardens following a prefix-final nasal. This gives us disyllabic reduplicants that alternate between a continuant-initial

realisation and a stop-initial surface realisation, as in (227b) *ŋa-**j**ampa-jamp-iːni* and (227d) *ŋaɲ-**c**ampa-jamp-iːni*.

The constraint ranking in (296) fails to predict the correct result for either of these forms, because it tolerates gratuitous application of Hardening for the sake of appeasing our parochial constraint, RED=TV, which is satisfied when the reduplicant is both monosyllabic and begins with a [–SON] segment. With our current constraint ranking, as we see in (297), the unattested candidate *d* satisfies RED=TV with its stop-initial monosyllabic reduplicant *ca-*, even though the base, *jamp-iːni*, is continuant-initial. (Note that it would make no difference to make the two templatic constraints of equal ranking, as in that case NOCODA would break the tie between candidate *a* and candidate *d*, still favouring the latter; reversing their ranking would universally prefer disyllabic reduplicants even where a monosyllabic one is licensed.)

(297) RED=TV satisfied by changing the reduplicant-initial continuant to a stop:

/ŋa-, RED, jampiː, -ni/	MAX-IO	RED=TV	RED=σσ	NOCODA	MAX-BR
☞ a. ŋa.-jam.pa.-jam.piː.-ni		∗!		∗∗	na
b. ŋa.-ja.-jam.piː.-ni		∗!	∗	∗	mpina
c. ŋa.-cam.pa.-jam.piː.-ni		∗!		∗∗	na
☹ d. ŋa.-ca.-jam.piː.-ni			∗	∗	mpina

Without some countervailing constraint banning changes in sonority between input and output, it should be possible for any candidate to emerge as optimal if it merely changes the reduplicant-initial segment into a stop. Suppose, then, that we have faithfulness constraints enforcing sonority identity between (298) input and output and (299) reduplicant and base.[46]

(298) **IDENT-IO[SON]**
'Corresponding segments in the input and output must be identical with respect to the feature [SON]'.

(299) **IDENT-BR[SON]**
'Corresponding segments in the reduplicant and base must be identical with respect to the feature [SON]'.

Recall that, according to the claims of Correspondence Theory, the reduplicant in the output corresponds to a RED morpheme in the input, which has no segmental specification. Therefore,

[46] Constraints enforcing identity of continuancy, i.e. IDENT-IO[CONT], would be necessary in a fuller analysis, since this feature is also implicated in continuant-stop alternations. However, these matters are largely irrelevant to the issues discussed here. See further discussion in Chong (2011) on the implications of treating either [CONT] or [SON] as driving these alternations.

the sonority of reduplicant segments can only be governed by base-reduplicant faithfulness (IDENT-BR[SON]), not input-output faithfulness (IDENT-IO[SON]). In order to block candidates like **ŋa-ca-jamp-iːna*, both must be ranked above the templatic constraints, as we can see in (300). Given that Hardening applies to reduplicants without affecting the base, base-reduplicant sonority faithfulness violations are tolerated in such contexts: thus IDENT-IO[SON] must be ranked above IDENT-BR[SON]. (For considerations of space, I omit MAX-IO from the tableaux to follow — for present purposes we can assume it occupies the top position in the ranking.)

(300) RED=TV cannot be satisfied by gratuitous changes in sonority:

/ŋa-, RED, jampiː, -ni/	ID-IO[SON]	ID-BR[SON]	RED=TV	RED=σσ	NOCODA	MAX-BR
☞ a. ŋa.-jam.pa.-jam.piː.-ni			*		**	ni
b. ŋa.-cam.pa.-jam.piː.-ni		*!	*		**	ni
c. ŋa.-cam.pa.-cam.piː.-ni	*!		*		**	ni
d. ŋa.-ja.-jam.piː.-ni			*	*!	*	mpini
e. ŋa.-ca.-jam.piː.-ni		*!		*	*	mpini
f. ŋa.-ca.-cam.piː.-ni	*!			*	*	mpini

In order to capture Hardening of disyllabic reduplicants just in the correct context — following stops or nasals — IDENT-BR[SON] must be violable in that context. Following Chong (2011), we can analyse the pressure to realise underlying continuants as stops as emerging from a constraint on sonority sequencing, as below. Post-nasal stops will satisfy (301), while post-nasal continuants will violate it.

(301) **Syllable Contact Law** (SCL; simplified from Chong 2011)
'For a heterosyllabic sequence A.B, assign one violation mark for every segment, B, which is higher in sonority than A'.
(Effectively, no [−CONT][+CONT] sequences.)

SCL should be undominated, as no violations of it are attested. However, as we can see in (302), adding this to the tableau fails to select the attested output as optimal. Whereas both the attested candidate (302c), *ŋaɲ-campa-jamp-iːni*, and the unattested optimal candidate (302d), *ŋaɲ-ca-jamp-iːni*, both violate IDENT-BR[SON] in order to satisfy SCL, the former also violates our parochial constraint RED=TV, and thereby loses to the latter.

(302) RED=TV becomes decisive where SCL forces violation of IDENT-BR[SON]:

/ŋaN-, RED, jampiː, -ni/	SCL	ID-IO[SON]	ID-BR[SON]	RED=TV	RED=σσ	NOCODA	MAX-BR
a. ŋaɲ.-jam.pa.-jam.piː.-ni	∗!			∗		∗∗	ni
☞ b. ŋaɲ.-cam.pa.-jam.piː.-ni			∗	∗!		∗∗	ni
c. ŋaɲ.-cam.pa.-cam.piː.-ni		∗!		∗		∗∗	ni
d. ŋaɲ.-ja.-jam.piː.-ni	∗!			∗	∗!	∗	mpini
☹ e. ŋaɲ.-ca.-jam.piː.-ni			∗		∗	∗	mpini
f. ŋaɲ.-ca.-cam.piː.-ni		∗!			∗	∗	mpini

As we have seen, the fact that a stop-initial monosyllabic reduplicant is required in some contexts means it is difficult to rule it out in the narrow context in which Hardening applies. It would obviously not be helpful to further specify RED=σσ to require an initial sonorant either — i.e. RED=NVXV, where N is any sonorant and X is any consonant or consonant cluster. This would have to be violated where Hardening applies (i.e. to avoid violation of SCL), thus prioritising candidates that can satisfy RED=σ (or RED=TV), just like candidate (302d).

It should also be pointed out that, given the SCL constraint is never violated within Wubuy words, any structure which obeys it should be regarded as quintessentially unmarked. As such, one may be tempted to attribute the appearance of derived stops in disyllabic reduplicants to the workings of the so-called Emergence of the Unmarked (TETU). But TETU is only really relevant to where an alternation is the product of input-output correspondence dominating a markedness constraint, which then dominates base-reduplicant correspondence; this causes markedness effects to obtain in the reduplicant but not in the base. It is therefore only implicated in the preference of (302e) over (302f); both the attested candidate and the winning candidate incur the same violations of the sonority faithfulness constraints.

Since we cannot index reduplicant size to the sonority of the reduplicant-initial segment, we may instead seek a solution that indexes it instead to the base-initial segment, as this reliably remains faithful to its input underlying specification. It is difficult to see how anyone committed to the idea of a universal constraint set could accept a constraint so uniquely attuned to the idiosyncracies of Wubuy reduplicative morphophonology. However, one could imagine that a conjoined constraint (in the sense of Smolensky 1995) could offer at least a technical fix: where a templatic constraint on reduplicant size, i.e. RED=σσ, is conjoined to a phonotactic constraint requiring, say, the left edge of the base to be aligned to a sonorant. This is nonetheless an extremely 'brute force' patch, and gives no greater insight into why prosodic shape might have any relation to initial segment sonority.

It is important to note here that the problem of the non-optimising character of the shape alternation, and the difficulty of generating it with templatic constraints in BRCT (or any

other parallel OT model), would exist even without the problem of opacity, as I showed with the RED=σ/RED=σσ ranking paradox above. However, the opacity problem is conceptually dependent upon the existence of the shape alternation problem: disyllabic reduplicants appearing stop-initial on the surface is troubling because it renders the generalisation regarding the distribution of monosyllabic and disyllabic reduplicants non-surface-true. If only disyllabic reduplicants were possible, forms like *ŋaɲ-campa-jamp-iːni* would present no great difficulty, as IDENT-BR[SON] would be violated for a principled reason: satisfaction of SCL. And, even if one had a compelling BRCT (or other parallel OT) account of the basic alternation, the opacity problem would nevertheless persist, as it does for parallel OT generally.

4.5.6 Summary of BRCT

Base-Reduplicant Correspondence Theory makes the claim that reduplication involves a special kind of morpheme, a reduplicant, which is unspecified for phonological content. Its content is instead supplied by the language-particular ranking of universal but violable constraints, expressed as an OT grammar. The relevant constraints for reduplicative words are templatic constraints on the prosodic shape of the reduplicant, faithfulness constraints demanding identity between the reduplicant and the base, as well as the usual OT markedness and input-output identity constraints — though the latter have no scope over the reduplicant, as it has no specified input phonology. This makes the crucial prediction that reduplicant shape, like productive phonological behaviour generally, should be optimising.

In Wubuy, a monosyllabic reduplicant occurs with stop-initial bases, and a disyllabic reduplicant occurs with sonorant-initial bases. As I have attempted to show, BRCT is inadequate to the task of describing a reduplicant shape alternation such as this. While we have no great issue analysing the monosyllabic and disyllabic patterns individually, we are inevitably led to a ranking paradox once we attempt to unify the analyses into one OT grammar, as the theory requires. We cannot resolve the paradox by attributing the different reduplicant shapes to different reduplicant 'morphemes' in the lexicon, and nor by positing some overarching markedness constraint that we lack independent evidence for. Even parochial constraints that bind prosodic templates to initial segment sonority are of no use once Hardening is involved.

The fundamental problem with the Wubuy reduplicant shape alternation is akin to what we saw with the *ŋu-* alternation in Chapter 3. We are again faced with an alternation which is exclusively determined by phonology, and yet which cannot be construed as optimising, since it is not a repair for the conditions that trigger it. In both the *ŋu-* problem and the reduplicant shape alternation, the generalisation involves a systematic difference between (non-alternating) stops and consonants of other kinds, and yet there is no obvious inherent phonetic relationship between occluded oral airflow and a light monosyllabic prosodic template, just as there is no

obvious inherent phonetic relationship between stops and the syllable [ŋu]. (It should also be remembered that any successful account of reduplication will ultimately have to address ŋu-Epenthesis as well, since both processes can appear in the same word, and indeed *ŋu-* can be part of the material to be copied in the reduplicant.)

Blevins (2003) has shown that, far from being the domain where unmarked structures are most likely to emerge, reduplicants can and often do feature marked structures that are nonetheless the result of leniting sound change. The discussion I have presented here can be seen as contributing a different kind of problem for analysing reduplicative patterns in terms of markedness. Here, neither the monosyllabic nor the disyllabic reduplicant are problematic for markedness on their own; the problem is that they both occur in the same language, but in different phonological environments, and their relationship to those environments is unrelated to markedness. Instead, the best explanation we have available is in terms of the relative informativity of stops and sonorants in stem-initial position, as suggested in §4.3.3. That argument located optimisation, if it can be called such, in the diachronic forces that led to the present pattern. Synchronically, it assumes that speakers learn to associate the different reduplicant shapes to the sonority of the base-initial segment because that is the dominant pattern in the particular lexicon they have been exposed to, not because there is some cross-linguistic universal that commands them to do so.

A BRCT analysis might be saved here by invoking some mechanism of phonological subcategorisation, such that each reduplicant shape can select for bases with segmental phonology of a specified kind. This possibility has sometimes been hinted at in applications of correspondence theory. Urbanczyk (2001, p. 24), for example, suggests that "because different reduplicative morphemes [with different shapes] have distinct lexical entries, the phonological differences should follow from some distinguishing feature in a lexical entry". These lexical entries can have different types of subcategorisation restrictions — syntactic, semantic, and phonological — and these restrictions organise morphemes into classes. But for reduplicant morphemes, she only acknowledges prosody as a relevant phonological subcategorisation feature. In any case, allowing morphological constituents to select for phonological properties of their hosts is antithetical to the basic claims of the BRCT. In particular, the universal P»M ranking requires that all allomorphy be resolved by surface markedness constraints (Paster 2009; Yu 2007). And as a result of these assumptions, no true phonological subcategorisation is technically feasible in BRCT (or basic OT), by design.

To that end, Morphological Doubling Theory (Inkelas & Zoll 2005) requires consideration here — not only because it is at present the most significant challenge to Correspondence Theory in the theoretical literature on reduplication, but also because its central claims are so different that comparing its predictions with those of Correspondence Theory and those of Relational

Morphology illuminates what is at stake in any theoretical account of reduplication in Wubuy. The next section takes up MDT in detail.

4.5.7 Morphological Doubling Theory

On the basis of the preceding discussion, it may be said that a satisfying OT treatment of the reduplicant shape alternation in Wubuy seems to require some ability to reorder the templatic constraints RED=σ and RED=σσ according to the sonority of the base-initial segment. Morphological Doubling Theory (hereafter MDT; Inkelas & Zoll 2005), probably the most significant competitor to BRCT in recent theoretical literature on reduplication, makes constraint reordering a central plank of its model. In this section, I present the main features of MDT and show that, so long as we allow a few assumptions regarding subcategorisation, it has the formal machinery to handle the Wubuy reduplicant size alternation. Nonetheless, I will also ultimately argue that the historical and psycholinguistic implications of this approach are unconvincing.

MDT adopts the general approach to the morphology-phonology interface known as Sign-Based Morphology, a declarative, constituent structure-based framework (Orgun 1996, 1998). Like some other such frameworks, it treats linguistic expressions — words as well as phrases — as signs: pairings of meaning (syntax/semantics) and form (phonology). A morphologically complex word is a hierarchical constituent structure (aka construction) containing multiple signs: a 'mother' sign at the upper node of the hierarchy encompasses 'daughter' signs for roots and affixes at the lower nodes.[47] Each node specifies a phonological function which takes the phonology of its daughters as input and produces an output; terminal nodes take underlying representations as input. Crucially, these functions can differ from one node to another, or from one construction to another. The functions are termed 'cophonologies', and can be formalised as OT grammars with different constraint rankings.

A simple example depicting English plural affixation is shown below. The construction on the left is the general schema for the English regular plural; the construction on the right is one concrete instance of the regular plural, the word *books*. As can be seen, the plural suffix is associated with an underlying representation, /z/, which concatenates to the root by means of the phonological function *g* (which presumably has a high-ranking constraint ensuring the final sibilant agrees with the preceding segment in voicing). As the *books* example shows, this function can result in an output form that differs from the underlying representation (/z/ → [s]).

[47] Sign-Based Morphology is intended to be a flexible framework that can accommodate different theories of morphology and phonology. A morphemic interpretation, whereby roots and affixes have the same formal status as terminal nodes, is just one interpretation possible under the framework, but it is generally the one assumed by Inkelas & Zoll (2005). See Orgun 1996, pp. 68–92 for more discussion of this, and note that he ultimately prefers an interpretation in which affixes (unlike roots) are not signs at all, and have "no life outside of the construction" they occur in (p.94), which would bring SBM closer to RM.

(303) An English affixation construction in Sign-Based Morphology (Inkelas & Zoll 2005, pp. 13, 72).

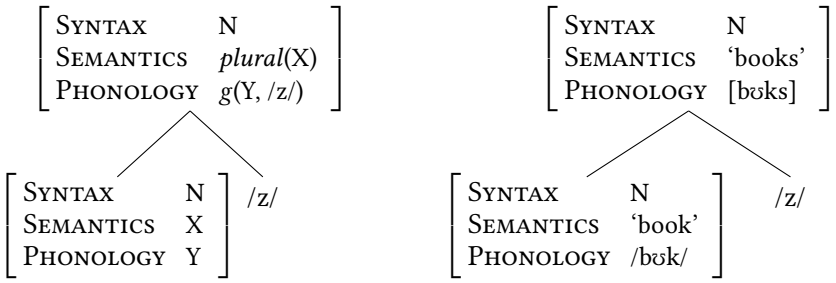

A reduplication construction in MDT is one whose input contains two morphological constituents with the same (or similar) semantics. Whatever morphological constituent is targeted by the reduplication construction (root, affix, stem or word) is present at the input in its entirety, in both form and function — there is no RED morpheme. Like in other constructionist approaches, however, the semantics of the whole need not be a simple function of the composition of its parts. If we submit the sign *wuɾuc* '(Aboriginal) person' to a reduplication construction, the sign will be doubly present at the input, i.e. 'person'+'person', but the meaning of the whole is the plural '(Aboriginal) people', rather than, say, 'two people' or 'one person and another person'.[48] Far more radical divergences from compositionality are of course possible, as we should expect from the cross-linguistic typology of reduplicative semantics (Inkelas & Downing 2015). Whatever meanings reduplication might have in a particular language, that meaning is associated with the reduplication construction as a whole: the meaning is not associated with a RED morpheme, nor is it present in the input at all.

(304)

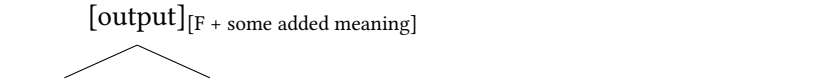

Because different cophonologies can be associated with each node of a construction, each of the copies in a reduplication construction can apply different phonology, as illustrated in (305). Partial reduplication arises when one of the daughters — what we would elsewhere call the 'reduplicant' — is subject to a *truncating* cophonology, though various sorts of phonological modification other than truncation are predicted to be possible. Whatever phonological modification does occur, however, must be analysed in MDT in terms of the relationship between input and output, not between 'reduplicant' and 'base'. The daughters are phonologically independent

[48] The DUAL, of course, is one of the possible values of the NUMBER feature and therefore contrasts with PLURAL. It takes an overt suffix.

of one another. In (305), then, Cophonology X truncates the input to a specified size, while Y and Z are faithful renderings of their inputs.

(305) Simplified diagram of a partial reduplication construction in MDT

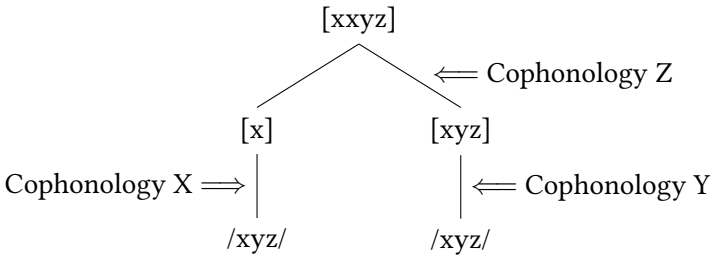

Recall that in BRCT, a reduplicated word is one whose input contains a RED morpheme. Identity between reduplicant and base is essentially the result of phonological repair: because RED does not come with a segmental specification, the grammar must supply it with one in order to be pronounceable. In MDT, phonological resemblances between the two copies in a reduplication construction (erstwhile reduplicant and base) do not arise out of a drive to maintain identity, but are rather an epiphenomenon of the fact that their inputs are the same. Thus in (305), the output of X partially resembles the output of Y because both have identical inputs.

Unsurprisingly, given its typological rarity, Inkelas & Zoll (2005) do not discuss cases of phonologically-conditioned shape allomorphy, but they do argue for constructions that subcategorise for features of the daughters they unify with. Stem type selection in Sye reduplication is an example of these features, as is vocabulary stratum in Hindi synonym compounding (Inkelas & Zoll 2005, p. 61). So long as phonology is among the features a construction can subcategorise for, then we can posit two basic types of reduplicative constructions for Wubuy. One construction (306a) takes two identical stop-initial stems as daughters, and the left-hand daughter has a cophonology that truncates to a light syllable; in the other (306b), the daughters are sonorant-initial, while the left hand daughter truncates to a disyllabic unit. The numbered boxes indicate variables, much like relational variables in RM. I will explain how truncation works for each of these schemas in the paragraphs to follow.

(306) Wubuy reduplication constructions as MDT/SBM signs

a. Monosyllabic reduplication of stop-initial bases.

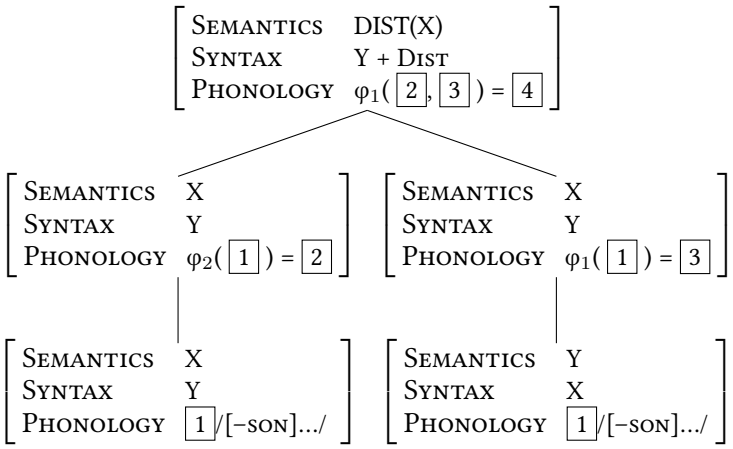

b. Disyllabic reduplication with sonorant-initial bases.

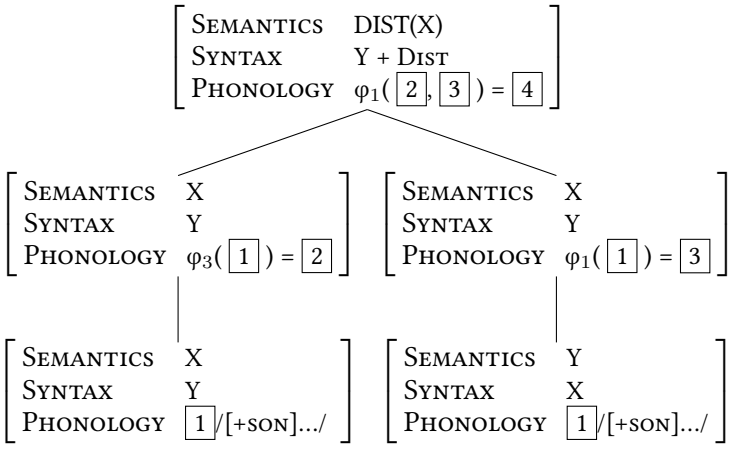

The essential difference between the reduplicant template constraints we saw in the BRCT analysis and the templatic constraints required for truncation is in the morphological category targeted for alignment: instead of requiring the *reduplicant* to be the size of some unit from the prosodic hierarchy, we need to restrict the entire *output* to be a certain size. With that in mind, we can make use of the following constraints, in (307) and (308). Presumably, other constraints would be required — for instance, alignment constraints requiring that the left edge of the output aligns to the left edge of the input, so that the left side is not truncated — but these constraints are sufficient to demonstrate the fundamental workings of the approach. I count violations by the number of syllables in excess of the template.

(307) **Output=σσ**
Align the left and right edges of the output with the left edge of a syllable and the right edge of a following syllable.
(The output must consist of two syllables.)

(308) **Output=σ**
Align both edges of the output with both edges of a single syllable.
(The output must consist of one syllable.)

Since the input to both copies in the reduplication construction is fully specified, Max-IO has scope over the entire constituent, not only the parts outside the reduplicant as in BRCT. It therefore takes over the functions of Max-BR, and, similarly, ranking it lower than templatic constraints will result in truncation to a specified size. Indeed, there is no longer any need for reduplication-specific constraints at all; cophonologies obviate the need for morphologically indexed constraints generally.

We require three cophonologies: one that favours the retention of all segments, one that favours truncation to a syllable, and one that truncates to a disyllable. Ranking either of the output size constraints above Max-IO is what will ensure truncation; NoCoda ensures that the winning candidate has only the minimum number of segments necessary to meet the template. Thus, we have constraint rankings for each of the cophonologies as below:

(309) a. Identity (φ_1)
Max-IO » {Output=σ, Output=σσ} » NoCoda

b. Monosyllabic truncation (φ_2)
Output=σ » Output=σσ» NoCoda» Max-IO

c. Disyllabic truncation (φ_3)
Output=σσ» Output=σ» NoCoda» Max-IO

Let us return to the *ṯu-ṯurapata* 'non-Aboriginal people' example from earlier. Assume that the input consists two instances of /ṯurapata/. Recall that the two reduplication constructions in (306) subcategorise for the phonology of their daughters; /ṯurapata/, being stop-initial, can unify only with (306a). The right-hand daughter of (306a) — elsewhere termed the 'base' — is fed into cophonology φ_1, as shown in (310), which simply returns an identical output, due to high-ranking Max-IO. The left-hand daughter — the 'reduplicant' — is submitted to cophonology φ_2, which instead ranks Output=σ above Max-IO and therefore selects a truncated output as optimal, as per (311). The winning candidates of (310) and (311) are then concatenated at the mother node, which is submitted to cophonology φ_1 once again, as in (312). This finally yields the attested reduplicated word, *ṯu-ṯurapata*.

(310) Right-hand daughter: Identity (φ_1)

/t̪ɹapata/	Max-IO	Output=σ	Output=σσ	NoCoda
a. t̪u	ɾ!apata		*	
b. t̪uɹ.	a!pata		*	*
c. t̪u.ɾa.	p!ata	*		
d. t̪u.ɾap.	a!ta	*		*
☞ e. t̪u.ɾa.pa.ta		***	**	

(311) Left-hand daughter: Truncation to monosyllable (φ_2)

/t̪ɹapata/	Output=σ	Output=σσ	NoCoda	Max-IO
☞ a. t̪u		*		ɾapata
b. t̪uɹ.		*	*!	apata
c. t̪u.ɾa.	*!			pata
d. t̪u.ɾap.	*!		*	ata
e. t̪u.ɾa.pa.ta	*!**	**		

(312) Mother: Identity (φ_1)

/t̪u, t̪ɹapata/	Max-IO	Output=σ	Output=σσ	NoCoda
a. t̪u.	t̪!ɹapata		*	
b. t̪u.t̪u	ɾ!apata	*		
c. tu.t̪uɹ.	a!pata	*		*
☞ d. t̪u.t̪u.ɾa.pa.ta		****	***	

Though I will not burden the reader with further tableaux demonstrating the fact, it should be fairly obvious by now that the basic approach to the monosyllabic reduplication pattern in MDT will encounter no difficulty in accounting for the disyllabic pattern as well. The mechanism used to generate partial reduplication in general — constraint reranking at each nested morphological domain — is the same as that required to generate alternation among different partial reduplicant templates, so long as constructions can subcategorise for segmental features.

For the disyllabic pattern, which specifies cophonology φ_3 for the reduplicant (aka the left-hand daughter), instead of promoting Output=σ to dominate Max-IO as in cophonology φ_2, we promote Output=σσ instead. The result is truncation to a disyllable. Since cophonology φ_3 in the reduplication schema (306b) subcategorises for a sonorant-initial input, only sonorant-initial (not stop-initial) stems will truncate to a disyllable.

By allowing multiple constraint rankings, together with subcategorisation for features of the input to each cophonology, MDT avoids the ranking paradox that is unavoidable for BRCT when facing the Wubuy reduplicant shape alternation. Certainly, the fact that MDT can address the Wubuy problem at all is an advantage it has over BRCT, though not RM. An advantage it holds over both other theories is that the nested morphology-phonology interleaving also has a simple explanation for Hardening in continuant-initial stems: the input to truncation is continuant-initial, licensing the disyllabic shape, and only hardens to a stop when it follows a prefix-final nasal at a higher node of affixation. The same logic extends to another 'late' phonological process, vowel hiatus resolution, as in Heath's (1984) rule-based account: vowel-initial stems can be reduplicated by means of the disyllabic truncation cophonology φ_3 also, after which the mother node resolves the hiatus, erasing the syllable boundary between reduplicant and base. With its reliance on stored word forms as the basis of morphological generalisations, however, RM has a better account of speakers' extension of /a/-initial reduplication to the /i/-initial stem that we saw in (243).

Undoubtedly, cophonologies are a powerful device. The looming question, however, is whether this greater power comes with any greater insight. The basic mechanism of generating a partial reduplicant by copying a full stem and then truncating it to the required size works on a purely formal level, and little more. First of all, one has to wonder how it could be translated into processing terms. The implication seems to be that, whenever a speaker utters a predictably-derived reduplicated word, they must compute all its nested phonological functions. Since reduplication in Wubuy most typically targets the entire stem, which can be very long, one would presumably have to plan an utterance by assembling as many as four or five syllables (ten or more segments) that will never see the light of day. One might then expect to find speech errors involving reduplicants of three syllables, or involving a full copy of the base; these are not the sorts of errors or incomplete changes we find, in Wubuy or elsewhere (Blevins 2005). It seems difficult to imagine that speakers would spend so much processing time on a generalisation that could simply be memorised and retrieved. Moreover, it would suggest that learners, on hearing reduplicative words with a reduplicant of one syllable and a base of six syllables, for some reason reconstruct the missing five syllables for the reduplicant also, even though to do so adds nothing to their construal of the semantics. For all we know, the part of the input that is deleted by truncation (in the reduplicant) could even be different to the equivalent part of the input to the base, since it is ultimately unobservable.

Despite its various shortcomings, the relative formal economy of BRCT, having only one input and one output, at least recognises that reduplication constitutes a distinct morphological strategy, crucially involving similarity between parts of a single word, and it does not burden the grammar with multiple stages of evaluation involving hidden structures to produce this surface

similarity. While I remain skeptical of the idea that speakers compute every predictably-derived word in order to say it, computing one OT tableau is surely less burdensome than computing several. The processing implications of RM, meanwhile, are straightforward: speakers store whole reduplicated and unreduplicated words in memory and retrieve them when necessary. 'Rules' for reduplication, reformulated as schemas, are only required when generalising beyond this store of words, producing a new word. Since schemas express relations between full surface words, no hidden levels of structure or intermediate stages are necessary for this generalisation: the inference requires just a single step, from known word to unknown word.

The reliance on truncation raises further questions when we situate reduplication in its historical and typological context. Truncation is not a morphological strategy otherwise found in Wubuy — nor is it reported in Australian languages generally beyond the occasional case of single syllable deletion, as in Lardil (Round 2011) or Yidiny (Dixon 1977). Reduplication, on the other hand, is absolutely widespread, and partial reduplication is found in all languages of the Gunwinyguan family. That being so, MDT's derivations have contemporary Wubuy speakers reconstructing total reduplicants that have not been a feature of the family for as far back as we can reconstruct with historical-comparative methods, only to cut them down to size by means of a strategy that is attested nowhere on the continent.

Contrary to MDT's claims that partial reduplication is equivalent to compounding with truncation, and that 'there is nothing unique about the phonology of reduplication constructions' (Inkelas & Zoll 2005, p. 69), Juliette Blevins' (2005) work makes clear that partial reduplication and morphological truncation are not equivalent. The fact that the material in the reduplicant is also recoverable from the base to which it is attached means that it is necessarily predictable, unlike any other kind of morphology, and therefore especially prone to leniting sound change processes — though as she also shows, the synchronic effects of those processes are not limited to the workings of Markedness and BR-Identity, as in BRCT.

4.5.8 Summary of MDT

In sum, then, an analysis of sonority-conditioned reduplicant shape in Wubuy is possible in MDT. Reduplicative words are here conceived as constructions with a hierarchical structure, where the daughters of that structure each call for a morphological unit with the same morphosyntax and semantics. At each node of the structure, an OT grammar evaluates an input according to a constraint ranking (a cophonology) specified for that node, and selects an optimal output, which then becomes the input to the node above it in the hierarchy. There is no limit on the number of possible constraint rankings operative in a given language. Partial reduplication is analysed as truncation of one of the morphosemantically identical daughters, and is held to arise when the cophonology for that daughter node ranks a templatic constraint on output size above

maximality constraints (which penalise deletion). This same mechanism of constraint reordering at each node of the construction also means we can have different prosodic templates for different reduplication constructions. And, because constructional schemas can subcategorise for features of the elements they unify with — including, presumably, segmental phonological features — MDT predicts the possibility of reduplication schemas with different reduplicant templates indexed to the sonority of the initial segment in the input. Thus, a formalisation of the Wubuy reduplicant shape alternation is achievable in MDT.

It is a clear strength of MDT that it predicts a far wider range of reduplicative phenomena than BRCT, including various sorts of behaviour unrelated to markedness, such as the alternation under discussion in this chapter. As I have argued, however, the reliance on full copy and truncation does not necessarily make for a compelling explanation of speakers' ability to produce and understand reduplicative words, given the reliance on highly abstract levels of computation.

4.6 Conclusion

In this chapter, I have undertaken the first detailed examination of Wubuy's striking and puzzling patterns of reduplication. I have presented a range of data showing that, in the language's only productive pattern, the size of the partial reduplicant is conditioned by the sonority of the initial segment in the base it attaches to: stop-initial bases require a monosyllabic reduplicant, sonorant-initial and continuant-initial bases require a disyllabic one, and vowel-initial bases may be understood as a third type, depending on the analysis. As far as I have been able to determine, sonority-conditioned reduplicant shape has not been reported in any other language in the world, save Wubuy's closest genetic relative, Anindilyakwa. This chapter thus contributes vital data, extending the range of empirical phenomena that will need to be covered in future debates on the nature of reduplication and how it varies cross-linguistically.

While it is fairly simple to state the main features of the pattern descriptively, it raises fundamental challenges for the most well-known formal theories of reduplication to date — Base-Reduplicant Correspondence Theory and Morphological Doubling Theory. The former theory simply has no answer to this challenge, while the latter can account for it only at the cost of dubious implications for psycholinguistic processing and production. I presented an analysis in Relational Morphology, building on the analysis of empty morphology in Chapter 3, which avoids these limitations. A notable advantage over the alternatives is that, as in the previous chapter, this analysis required no new formal mechanisms that are not already required to address phonology and morphology outside of reduplication. Moreover, the phonological classes that were motivated in that chapter turned out to be adequate here. To the best of

my knowledge, this chapter represents the first attempt to provide a thorough analysis of reduplication in Relational Morphology, in any language. That the theory has withstood such a formidable challenge suggests that it holds great promise in generalising to other problematic cases of reduplication; future work will determine whether this hope is well-founded.

Naturally, there still remain many unanswered questions. At the more fine-grained end, I have largely left out any consideration of the behaviour of /w_1/-initial stems under reduplication. These stems variably but frequently show deletion of the glide in base-initial position, reduplicant-initial position, or both. In the latter case, this renders the stem (in most but not all contexts) identical to a stem of the vowel-initial class in certain contexts, inducing the same hiatus resolution alternations (Baker 2021). For the formal analysis in RM, this would appear to require a special subclass of the alternating class, just affecting the /w/ that alternates with /k/, but its relationship to the vowel-initial class remains to be worked out.

Another open question, also concerning an aspect of variation, is how often and under what conditions the reduplicant can occur inside the stem, a configuration that does appear sometimes in the text corpus. Informally, it appears that this is only likely following one of the frequent derivational prefixes like Applicative or the Multiple, but this needs to be explicitly quantified. Another important topic requiring quantification is the argument around informativity and lexical access: in particular, a more precise formulation of Baker & Bundgaard-Nielsen's (2018) notion of a minimum informativity threshold for reduplicative segments is needed in order to test their claim against other languages with partial reduplication. And of course, I have had nothing to say on the semantics and discursive functions of reduplicative constructions, an understanding of which would be essential for a complete account of reduplication in Wubuy.

5 Variability and sound change in the lexicon: velar stop deletion

5.1 Introduction

The previous two chapters examined two of Wubuy's most pervasive phonological alternations — phonologically-conditioned empty morph insertion and phonologically-conditioned reduplicant shape allomorphy — in order to probe the nature of phonological generalisations and the sorts of representations required to analyse them. This chapter is concerned with the nature and extent of Velar Stop Deletion, a process that deletes /k/ when it precedes any other consonant. Like those other processes, Velar Stop Deletion also interacts opaquely with continuant-stop alternations. As such, the crucial features of the Relational Morphology analyses in Chapters 3 and 4 can be adapted to Velar Stop Deletion also.

However, there are some additional challenges here that cannot be straightforwardly addressed by a purely formal model, since — contra the description in Heath (1984) — deletion is variable. Although the literature since Heath (1984) has occasionally noted that deletion fails to apply uniformly, this chapter is the first to quantify that variability and analyse it in detail, including its phonetic manifestation. I show that the variability is structured according to, firstly, the type of morph it appears in, and secondly, the type of word the velar-final morph appears in. I suggest an explanation of this structured variability in terms of the entrenchment of frequency-driven gestural overlap, arguing that deletion is a sound change in progress, diffusing at different rates throughout the lexicon.

The chapter is organised as follows. Firstly, in §5.2, I provide the necessary background on Velar Stop Deletion and its interactions with other processes in Wubuy phonology, with attention given to Heath's motivations for positing a 'stop archiphoneme' /K/, which in his analysis functions identically to ordinary /k/. I then provide preliminary evidence of variation which suggests that this equivalence is false: whereas /K/ undergoes deletion uniformly, /k/ does not. Moreover, morph-final velar stops in recent loanwords are not deleted, while morph-final velar stops native roots are deleted some of the time.

This variation suggests that Velar Stop Deletion is an ongoing process of sound change, rather

than a fixed, stable pattern. In §5.3, then, I review the literature on phonetic pressures driving consonant deletion. I show that while consonant codas are often subject to tight articulatory control in Australia, unlike what may be expected elsewhere, this generalisation does not appear to hold for the velar place of articulation.

The most substantive contribution of this chapter lies in §5.4, where I present a quantitative study of Velar Stop Deletion. This study quantifies the distribution of deleting and non-deleting tokens, but it also examines their phonetic properties, specifically their duration and formant transition properties. I present the research questions and detail the methodology I adopt to address them, and report on the results. The results support the tripartite division of /k/-final morphs mentioned above — native /k/, loan /k/, and archiphoneme /K/ — but they also find a clear difference in the native (non-suppletive) noun root category on the basis of whether the root is incorporated into a verb (where deletion is the more common variant) or is the head of a noun word (where non-deletion is the more common variant).

The implications of these findings are discussed in §5.5 and situated in relation to theories of sound change and the phonetics-phonology interface, particularly those from the usage-based tradition. There, I offer an explanation of the different rates of deletion in different parts of the lexicon in terms of the relative frequency with which a velar-final morph appears at the right word edge, as opposed to some nonfinal position in the word. The higher this frequency, the higher the likelihood of deleting the velar stop. I attribute this to the automatisation of gestural sequences through practice, which favours articulatory routines that yield a percept of deletion. Finally, conclusions are given in §5.6.

5.2 Background

In this section, I go over what has previously been decribed regarding velar stop deletion in Wubuy. I pay particular attention to the opaque interactions with Hardening and Stop Insertion, and show how they logically lead Heath (1984) to posit a 'stop archiphoneme' /K/. This background is crucial, since much of the chapter will be devoted to investigating the precise relationship between archiphoneme /K/ and true /k/. I also present some of the similarities between the deletion of /k/ and the deletion and assimilation of /ŋ/, which provides additional evidence to the suggestion that velar gestures in Wubuy are more susceptible to reductive processes than gestures at other places of articulation.

5.2.1 Stop Deletion, Stop insertion, and Hardening

Within native roots and affixes of Wubuy, the dorsal stop /k/ is never encountered immediately preceding any other consonant. However, morphs ending in /k/ are relatively common, as are

morphs beginning in consonants. As a result, underlying sequences of /k/ followed by another consonant often occur straddling a morphological boundary: /k-C/. According to Heath's (1984) description, these dorsal stops must be deleted from the surface: /k-C/ → [C]. In (313), the examples in the left column show unsuffixed nouns ending in /k/; the right column shows the same nouns in their suffixed forms, where the /k/ is absent.

(313) a. ļakaja**k**
saltwater (NMET 64.6.2)

b. ļakaja-miri
saltwater-INSTR (NMET 119.4.5)

c. ɻaŋa**k**
wood (NMET 50.1.3)

d. ɻaŋa-miri
wood-INSTR (NMET 40.2.2)

If the morph following this underlying /k/ is from the class that alternates between a continuant-initial and a stop-initial realisation, the stop-initial form is required here — even though the /k/ which conditions it is absent from the surface. For example, from (314a), we can see that the root *ļakajak* 'saltwater' ends in /k/ when it is word-final. In its suffixed form (314b), *ama-ļakaja-tuc*, the Locative suffix is [-tuc], not [-ɾuc]. Where the root is instead continuant-final in both its underlying form and its surface form, as with (315a) *ɳuka* 'stone', the Locative suffix instead surfaces as [-ɾuc], as in (315b) *ama-ɳuka-ɾuc*. Underlying representations in the examples below follow Heath (1984).

(314) a. ļakajak
/ļakajak/
saltwater
'saltwater; sea' (NMET 64.6.2)

b. ama-ļakaja-tuc
/ama-ļakajak-ɾuc/
VEG.OBL-saltwater-LOC
'in the sea' (NMET 41.24.5)

c. ɻaŋak
/ɻaŋak/
wood
'tree; wood' (NMET 50.1.3)

d. a-ɻaŋa-tuc
/a-ɻaŋak-ɾuc/
NEUT.OBL-wood-LOC
'in the tree' (NMET 2.1.3)

(315) a. ɳuka
/ɳuka/
stone
'stone' (NMET 15.12.4)

b. ama-ɳuka-ɾuc
/ama-ɳuka-ɾuc/
VEG.OBL-stone-LOC
'on the stone' (NMET 15.12.5)

Input	saltwater-Instr /l̪akajak-miɾi/	Veg.Obl-saltwater-Loc /ama-l̪akajak-ɾuc/	Veg.Obl-saltwater-All /ama-l̪akajak-w$_1$uj/
Hardening	—	ama-l̪akajak-tuc	ama-l̪akajak-kuj
Stop del.	l̪akajaØ-miɾi	ama-l̪akajaØ-tuc	ama-l̪akajaØ-kuj
Output	l̪akaja-miɾi	ama-l̪akaja-tuc	ama-l̪akaja-kuj

Table 5.1: Sample derivations illustrating Hardening and Velar Stop Deletion, according to the underlying forms and ordered rules in Heath (1984).

c. a-muŋkala
/a-muŋkala/
Neut.Obl-platform
'(the) platform' (NMET 43.2.3)

d. a-muŋkala-ɾuc
/a-muŋkala-ɾuc/
Neut.Obl-platform-Loc
'on the platform' (NMET 43.2.3)

Comparing the suffixed forms in the right column of (314) and (315), the choice of suffix allomorph — between *-ɾuc* or *-tuc* — does not appear phonologically motivated. In each case, the immediate phonological environment is the root-final /a/. It is only by comparing these forms with their unsuffixed counterparts that any phonological conditioning is apparent: [-tuc] if the related free form ends in /k/ (or any non-continuant), [-ɾuc] otherwise.

This makes the interaction between Deletion and Hardening a prototypical case of 'self-destructive feeding' (Baković 2007), a variety of 'overapplication' or 'counterbleeding' opacity, just like we saw for the relation between ŋu-Epenthesis and Nasal Deletion in Chapter 3. In the serial rule-based phonology assumed by Heath, the /k-ɾ/ sequence over the root-suffix boundary of (314b,d) feeds Hardening in the suffix-initial /ɾ/, yielding /k-t/, which then feeds Deletion of the root-final /k/, as illustrated with the derivation in Table 5.1. The final surface form is then simply [t], with no trace remaining of the segment that conditioned it.

When underlying /k/ is followed by /w$_1$/, the result on the surface is singleton [k], by the same rule ordering: /k-w$_1$/ → /k-k/ → /k/. This is illustrated in (316) with nouns inflected for Allative case with the suffix /-w$_1$uj/. In, (316a), the conditions for Hardening are not met, so the suffix surfaces as [-wuj]. In (316b,c), a root-final non-continuant is there to trigger Hardening, so the suffix instead surfaces as the stop-initial allomorph [-kuj]. In (316d), the suffix is [-kuj] once again, but this time its conditioning context (a root-final /k/) is absent from the surface form. If one only looked at this last example, one could reasonably ask whether it is really the root-final /k/ or the suffix-initial /k/ that is deleted here. But the fact that we know *l̪akajak* 'saltwater' has a final velar stop when it appears unsuffixed, and that this velar stop also disappears before other consonant-initial suffixes *-miɾi* and *-tuc*, suggests that it is the root-final /k/ that is deleted

here as well. Moreover, (316b,c) provides transparent motivation for a *-kuj* suffix allomorph anyway.

(316) a. a-ŋuka-wuj
/a-ŋuka-$w_1$uj/
NEUT.OBL-stone-ALL
'to stone' (NMET 116.6.2)

b. ji:-wa:jin-kuj
/ji:-wa:jin-$w_1$uj/
FEM.OBL-emu-ALL
'to emu' (NMET 7.12.3)

c. a-l̪al-aɲpac-kuj
/a-l̪al-aɲpac-$w_1$uj/
NEUT.OBL-country-other-ALL
'to another country' (NMET 29.9.4)

d. ama-l̪akaja-kuj
/ama-l̪akajak-$w_1$uj/
VEG.OBL-saltwater-ALL
'to the sea' (NMET 28.13.5)

Making matters somewhat more complicated here is the fact that velar stops are also *inserted* at morphological boundaries wherever a morph ending in a stop or nasal is followed by a morph beginning with a vowel. The process can be observed in the following examples. The adjectival root *ala:ti* 'bad' can be observed in isolation, as in (317a), where we see that it is vowel-initial. Morph-initial vowels coalesce with preceding vowels, as in (317b), where *ala:ti* follows the pronominal agreement prefix *wu-*, or in (317c), where it follows the compound initial *waɭakaɭa* 'halfway'. But note (317d): the final /c/ of the bound root *l̪amic* 'taste' triggers insertion of a velar stop before *ala:ti*, yielding *l̪amic-**k**ala:ti*. Finally, in (317e) we can observe what happens when the non-continuant to trigger Stop Insertion is a morph-final /k/. Again, this morph-final /k/ is then deleted from the surface, just as we saw with Hardening. The examples in (317f,g) provide justification for positing an underlying final /k/ in the incorporated noun *mantak* 'group'.

(317) a. ala:ti
/ala:ti/
'bad' (NMET 18.19.4)

b. ana:ni:ni:la ana-mu:n wa:-'la:ti
/ana:ni:ni-ala/ /ana-mu:n/ /wu-ala:ti/
along this side (NEUT) NEUT.TOP-foot 3NEUT.R-bad
'The foot on this side is bad' (HN; PN2017-06-27 01:09:04)

c. wuɾu-waɭakaɭa:-'la:ti-ji:-'-na
/wuɾu-waɭakaɭa-ala:ti-w$_1$a-i-na/
3pl-halfway-bad-FACT-REFL-NPST2
'They became halfway bad' (NMET 67.1.3)

d. ni-l̪amic-kala:ti
/ni-l̪amic-ala:ti/
3MASC-taste-bad
'It (fan-palm) has a bad taste' (NMET 117.4.3)

e. ŋaɾa-manta-kala:ti-ja-ŋun-maki
/ŋaɾa-mantak-ala:ti-w$_1$a-ŋun-maki/
1sg>3pl.R-group-bad-FACT-EVIT-EVIT
'Lest I destroy them all' (NMET 7.14.6)

f. wu:-manta-ʈam-a:-ac
/wuɾu-manta-ɻam-a:-w$_1$ac/
3pl.R-group-be_abundant-PAST2-RPERG
'Back when they were abundant [elders]' (NMET 139.8.5)

g. waɾ-u-mantak
PL.TOP-PL-group
'group (of people)' (NMET 17.10.4)

Velar Stop Insertion can be thought of as relating to Hardening in the sense that it repairs an underlying heteromorphemic sequence of rising sonority, just as Hardening does. In (317d), for instance, the /…c-a…/ sequence in the underlying stem /l̪amic-ala:ti/ has a stop (lowest sonority) followed by a vowel (highest sonority); the surface form repairs this to a sequence of stop followed by another stop, where the second consonant does not have higher sonority than the first.

Finally, I have mentioned that fact that the sequence /kC/ never occurs within roots or affixes of Wubuy. This is not a mere accident, but is rather the result of velar stop deletion having applied to forms that historically had a preconsonantal /k/. As van Egmond & Baker (2021) show,

Input	3NEUT.R-bad /wu-alaːti/	3MASC-taste-bad /ni-l̪amic-alaːti/	1sg>3pl.R-group-bad-FACT-EVIT-EVIT /ŋaɾa-mantak-alaːti-w$_1$a-ŋun-maki/
Stop insert.	—	ni-l̪amic-kalaːti	ŋaɾa-mantak-kalaːti-ja-ŋun-maki
Stop del.	—	—	ŋaɾa-mantaØ-kalaːti-ja-ŋun-maki
VV-coalesc.	waː-ʼlaːti	—	—
Output	waː-ʼlaːti	ni-l̪amic-kalaːti	ŋaɾa-manta-kalaːti-ja-ŋun-maki

Table 5.2: Sample derivations illustrating Insertion and Deletion of /k/, according to the underlying forms and ordered rules in Heath (1984).

Anindilyakwa roots that have cognates in Wubuy show a regular correspondence between /kC/ in the former and just /C/ in the latter. For instance, Anindilyakwa *aɻa**kp**a* 'now' corresponds to Wubuy *aɻa**p**a*, *t̪ə+ma**kp**əl̪a* 'pelican' corresponds to *maː**p**ul̪u*, and so on. This historical deletion of /k/ will become relevant to the analysis later in this chapter, since it appears to provide evidence of sound change by lexical diffusion.

5.2.2 The archiphoneme /K/

The upshot of all of the processes described here is that the final /k/ in morphs like /l̪akajak/ 'saltwater' or /ɻaŋak/ 'wood; canoe' are deleted when followed by morphological material of any phonological substance. If they are followed by any consonant, they are deleted; if they are followed by a vowel, a consonant is inserted, and then they are deleted.

How then do we decide that a /k/ is part of the representation here at all? For the examples described so far, the answer must be that we look to their realisation at the right word edge, which is the only location at which a /k/ can surface as [k] (or at least, where the [k] that is realised can be reliably associated with the morph preceding it, rather than the one following it). We know that, as indepedent nouns, /l̪akajak/ and /ɻaŋak/ are realised as [l̪akajak] and [ɻaŋak], respectively, so we can feel reasonably confident that a final /k/ is part of the representation.

However, other morphs — such as derivational prefixes, bound roots and suppletive incorporated nouns — can never appear at the right word edge, by definition. Where morphs of this type induce all the phonological alternations in the following morph that are characteristic of velar stops, they are listed in Heath's work as ending in a 'stop archiphoneme' /K/. It is so labelled in recognition of the fact that there are, as a logical consequence of the interactions of the various phonological rules he describes, no conditions under which this segment can be directly observed — i.e., as [k] — and yet it participates in phonological alternations in identical fashion to a true /k/. All of the transformational rules Heath proposes assume that /K/ and /k/ are treated identically by the phonological grammar, in that they categorically trigger Hardening and Velar

Insertion, and are categorically subject to deletion preceding another consonant. Examples are given here showing the phonological alternations induced by the /K/-final morphs (318) *ŋaraK-* 'MULT', (319) *aK-* 'APPL', and (320) *paK-* 'eye' (a suppletive incorporated noun). The expected outputs of input sequences involving the velar stop followed by other classes of segments are summarised in Table 5.3.

(318) a. wuː-jam-aː
/wuɾu-jam-aː/
3pl.R-do-PAST2
'They did that' (NMET 9.9.1)

b. wuː-ŋaɾa-**c**am-aː
/wuɾu-ŋaɾa**K**-jam-aː/
3pl.R-MULT-do-PAST2
'They all did that' (NMET 9.9.1)

c. wuɾaː-'tat-iɲ
/wuɾu-atat-iɲ/
3pl.R-go_up-PAST1
'They went up on shore' (NMET 18.13.1)

d. wuː-ŋaɾa-**k**atat-iɲ
/wuɾu-ŋaɾa**K**-atat-iɲ/
3pl.R-MULT-go_up-PAST1
'They all went up on shore' (NMET 26.2.3)

(319) a. ŋunu-ɻaŋ
/ŋunu-ɻaŋ/
1sg>2sg.R-spear.PAST1
'I speared you' (HN; PN2017-09-24.1 00:10:01)

b. ŋunaː-'-**ʈ**aŋ
/ŋunu-a**K**-ɻaŋ/
1sg>2sg.R-APPL-spear.PAST1
'I speared it for you' (HN; PN2017-09-24.1 00:12:32)

Input	Output	Example	Rule
/k/ + word edge	[k]	/ɻaŋak/ > ɻaŋak ‘wood, tree’	
/k, K/ + continuant	stop	/paK-ɻuŋkal/ > pa-ʈuŋkal ‘eye-big’	Hardening, *kC
/k, K/ + stop	[ŋu]-stop	/paK-ʈitiɲ/ > pa-ŋu-ʈitiɲ ‘eye-closed’	ŋu-Epenthesis, *kC
/k, K/ + nasal	nasal	/paK-naji:/ > pa-naji: ‘eye-look’	*kC
/k, K/ + vowel	Ø-kV	/paK-ala:ti/ > pa-kala:ti ‘eye-bad’	Stop Insertion, *kC

Table 5.3: Surface products of underlying sequences involving /k/, following rules proposed in Heath (1984)

c. nuɾa:-’jik-i:na
/nuɾu-ajik-i:na
1pl.R-play-NPST2
‘We (exc.) play’ (NMET 57.1.6)

d. ŋa:-’-**k**ajik-iɲ-kala
/ŋu-a**K**-ajik-iɲ-w$_1$ala/
3sgM>3sgF.R-APPL-play-PAST1-ABL
‘After he played for her’ (NMET 21.10.6)

(320) a. ŋu-waɾa-wuɻ̪i-ɲ
/ŋu-w$_1$aɾa-w$_2$uɻ̪i-ɲ/
3sgM>3sgF.R-MULT-rub-PAST1
‘He rubbed her all over’ (NMET 28.9.2)

b. ŋu-pa-**p**uɻ̪i-ɲ
/ŋu-pa**K**-w$_2$uɻ̪i-ɲ/
3sgM>3sgF.R-eye-rub-PAST1
‘He rubbed her all over’ (NMET 28.9.2)

c. ni-pa-**k**ala:ti
/ni-pa**K**-ala:ti/
3sgM.R-eye-bad
‘He was blind’ (NMET 40.2.2)

A table with all of the morphs analysed in Heath as /K/-final is provided in the Appendix to this chapter. Most of these morphs are basically unproductive, appearing only in a few compounds, often with fairly non-transparent semantics, but the remainder of this chapter will

focus only on those /K/-final morphs that are clearly productive. The nominal prefixes *(w)uK-*, *niK-*, *ŋaɾiK-* and *maK-* constitute the entire set of noun class derivational prefixes, typically used to indicate (part-whole) possession, where the noun class corresponds to the possessor. More obviously recent applications include *wu-ŋu-pakaɭaŋ* for 'bullet', which literally means 'eye of NEUT' (cf. *pakaɭaŋ* 'eye [VEG]' and *patili* 'gun [NEUT]'). It would seem uncontroversial to regard these prefixes as productive. The verbal/adjectival derivational prefixes *aK-* 'APPL' and *ŋaɾaK-* 'MULT' are similarly unrestricted, and this thesis contains numerous examples of each. Finally, the two most productive /K/-final suppletive incorporated nouns are *paK-* 'eye' and *ŋuɻaK-* 'penis'. The former of these is a little dubious: although it appears in many fairly transparent compounds, and speakers can recognise new compounds formed with it, it is also possible to use the transparent root *pakaɭaŋ*, and in my experience speakers always incorporate this when asked to translate from English prompts unless there is already a well-established incorporation construction using *paK-*. The latter, *ŋuɻaK-*, is very productive, since it is used in adjectival incorporation constructions to form novel insults aimed at men.

5.2.3 Formalising deletion in RM

As mentioned, the interaction between /k/-Deletion and Hardening can be understood as self-destructive feeding opacity. Since the interaction between Nasal Deletion and ŋu-Epenthesis was also of this type, all of the formal difficulties outlined for that interaction in Chapter 3 apply to this interaction as well. And again, it is possible to capture it in terms of surface correspondence between full words using RM, as the sister schemas in (321) demonstrate. A noun root that ends in a /k/ when appearing on its own (321a) lacks that /k/ when appearing word-medially (321b), with the change between the linked forms indicated by same-except asterisks. In that latter context, the root subcategorises for the phonology of the element following it, specifying that it must begin in a [–CONT] segment. If the morphological element following it is one that has stop-initial and continuant-initial allomorphs, the former will be required here. That would include not only those morphs subject to Hardening (the alternating class), but also those subject to Stop Insertion (the vowel class).

(321) /k/-final nouns

a.

Semantics	$\ldots_o$
Morphosyntax	N_o
Phonology	$[\ldots{*}k{*}_o]$

b. Semantics $[X(\ldots_o)_x]_y$
Morphosyntax $\{ N_o, \ldots_x \}_y$
Phonology $[\ldots_o \ [-\text{CONT}] \ldots_x]_y$

Morphs ending in an archiphoneme can be treated as a form that simply specifies a following non-continuant. A lexical entry for the MULT prefix /ŋaɾaK-/ is given in (322). Since we have only a single level of phonological specification here, with no underlying form, the prefix is represented with the form /ŋaɾa/, without a final stop of any kind, while subcategorising for an initial [−CONT] segment in the following morphological element (a stop or a nasal). This phonological subcategorisation does all the work required of the archiphoneme in the rule-based analysis.

(322) Derivational prefix with archiphoneme: /ŋaɾaK-/ 'MULT'

a. Semantics $[\text{PRED}(\text{ALL OF}(X_x)_1, Y)_z]_a$
Morphosyntax $\{ X_z, \text{Pro}_x, \text{MOOD}_y, \text{MULT}_1 \}$
Phonology $[\ldots_{x,y} [\text{ŋaɾa}_1 \ [\ [-\text{CONT}] \ldots]_z \ ]_a \ ]_b$

Suppletive incorporated nouns ending in the archiphoneme, like /paK-/ 'eye' and /ŋuɻaK-/ 'penis', are in alternation with free forms which do not share their phonology. This is a definitional property of suppletion. In RM, we can formalise suppletion with sister schemas that are linked via relational coindices in the morphosyntax and semantics, but not in the phonology. Consider (323a), which shows a schema for the suppletive incorporated noun /ŋuɻaK/ 'penis'. There is an interface link joining the string /ŋuɻa/ in the phonology to the N head in the morphosyntax and the meaning 'penis' in the semantics. In (323b), there is an interface link joining the phonological string /ŋuji/ to a N head in the morphosyntax and the same meaning in the semantics. These morphosyntactic and semantic elements of (323a) and (323b) are relationally linked to each other with an additional coindex which is not present on the phonological level. What this ends up meaning is that a noun with the meaning 'penis' is associated with one phonological form when incorporated into an adjectival or verbal stem, but is associated with a completely different phonological form when appearing as the head of a nominal word. This follows the approach taken to English *go/went* suppletion in Jackendoff & Audring (2020, p. 144).

(323) a. Semantics $[\text{PRED}(\text{Agent}_x)_z; \text{PENIS}_{2,3}]_a$
Morphosyntax $\{ [_x \ N_{2,3} \ X_z]_a, \text{Pro}_x, \text{MOOD}_y \}$
Phonology $[\ldots_{x,y} [\text{ŋuɻa}_3 \ [-\text{CONT}] \ldots_z \ ]_a \ ]_b$

b.

Semantics	$\text{PENIS}_{2,4}$
Morphosyntax	$\text{N}_{2,4}$
Phonology	$[\text{ŋuji}]_4$

Outside of the two forms just mentioned, the suppletive incorporated nouns (or compounding elements) that only appear in one or two constructions, and therefore do not appear to be productive, also tend to have a non-obvious semantic contribution, so their status as independently meaningful units is dubious. If they cannot be linked to some kind of predictable meaning, it is difficult to see how they could be generalised. Like other semantically non-compositional or 'lexicalised' compounds, they can be represented in RM (or any other constructionist framework) with a lexical entry that associates the meaning with the construction as a whole, rather than its constitutive elements. I do not provide entries for such forms here.

5.2.4 Variability

So far, I have treated Velar Stop Deletion as a categorical and inescapable alternation, since this is how Heath (1984) presents it. And, if that were all there were to deletion and its opaque interactions with Hardening and Velar Stop Insertion, the RM analysis just presented in §5.2.3 would be adequate. However, outside of Heath (1984), the literature in and on Wubuy contains various indications that deletion is in fact applied variably, and that this variation is structured: native /k/-final morphs undergo deletion only some of the time, borrowed morphs never undergo deletion, while the archiphoneme /K/ is always deleted. I present that preliminary evidence here; it serves as the impetus for the more extensive investigation of variability in /k/-deletion pursued in this chapter.

What first alerted me to the potential that /k/-deletion is not uniformly applied comes from the writing of native speakers themselves. While writing by Wubuy speakers is in general fairly scarce, the Wubuy Bible (which was developed by native speakers) contains hundreds of examples of /k/ realised orthographically (as <*g*>) in preconsonantal position, suggesting that speakers consider /k/ part of the representation of those morphs. I provide just a handful of these in (324); the first line is in Wubuy orthography as written in the Bible, the second line shows the equivalent IPA transcription, following the morph segmentation used in this thesis. The same Wubuy speakers who were most involved in Bible translation were also those who produced many of the early Wubuy bilingual education materials, and those materials also normally transcribe coda /k/ (that is, <*g*>).

(324) a. *nubagila-wala warruburru-yung warra-mijgalgur-yung wunu-lhamarrangi,*
nu-pakila-wala waɾupuɾu-juŋ waɾa-mickalkur-juŋ wunu-l̪amaɾ-aŋi
from that one those Pl-relatives-Sg 3pl>3sgM.R-come.later-Past2

nguynju yaga arrawindi yi-miyiri-yii a-wumala-rruj marri
ŋuɲcu jaka aɾawinti ji-mijiɻi-ji: a-wumala-ɾuc maɾi
same like many Fem.Obl-stars-Simil Neut.Obl-sky-Loc and

*nguynju ama-**mangarag**-jii ama-**madhaḻag**-baj.*
ŋuɲcu ama-**maŋaɻak**-ci: ama-maṯaḷak-pac
same Veg.Obl-sand-Simil Veg.Obl-beach-Perg

'And so from this one man, and he as good as dead, came descendants as numerous as the stars in the sky and as countless as the sand on the seashore'

(WB Hebrews 11:12)

b. *Mana-lhagayag ngamaambamadhiyn-bindiyung, nga mana-munggulu*
mana-l̪akajak ŋamaa-'mpamaṯ-iɲ-pintijuŋ ŋa mana-muŋulu
Veg.Top-saltwater 3Veg>1sg.R-cover-Past1-very and Veg.Top-seaweed

*ngama-**yinag**-gabaabaa*
ŋama-**jinak**-kapa:p-a:
3Veg>1sg.R-head-wrap-Past2

'The engulfing waters threatened me, the deep surrounded me; seaweed was wrapped around my head' (WB Jonah 2:5)

c. *Anubani-yung-maynji ana-rangag yagi waṉagi ana-marrya,*
anupani-juŋ-maɲci ana-ɻaŋak jaki waɳak-i ana-maɾja
that.Neut.Anaph-Sg-if Neut.Top-tree Neg hold-Npst3 Neut.Top-food

*anubani ambii-**rangag**-guldhii marri ambirri-nangana*
anupani ampi:-**ɻaŋak**-kulṯi: maɾi ampiɾi-naŋ-ana
that.Neut.Anaph 3pl>3Neut.Irr-tree-cut-Npst2 and 3pl>3Neut.Irr-burn-Npst2

anu-ḏangag a-ngura-wuy.
anu-ʈaŋak a-ŋuɻa-wuj
Neut.Top-wood Neut.Obl-fire-All

'Every tree that doesn't grow good fruit is cut down, and thrown into the fire'

(WB Matthew 7:19)

One may be tempted to attribute this absence of Velar Stop Deletion in orthographic practice to some kind of preference for writing underlying forms, but that clearly cannot explain it.

Wubuy orthography hews close to the surface phonology, so that underlying continuants that surface as stops due to Hardening are written as stops, long vowels resulting from coalescence are written as long vowels, assimilated segments are orthographically assimilated, and so on. And further, /k/-deletion is occasionally recognised in the Bible translations, as can be seen by comparing the realisation of *jinak* 'head' in (325a) and (325b).

(325) a. *Nugurru* *nurru-**yina**-baḍa-waḍaḍ!*
nukuɾu nuɾu-**jina**-paʈawaʈaʈ
/nukuɾu/ /nuɾu-**jinak**-w$_2$aʈawaʈaʈ/
2pl 2pl.R-head-strong
'You unbelieving people' (WB Luke 9:41)

b. *Aadani-yung* *ni-warrarrangi,* *aadani-yung* *yamba* *nugurru*
aːtani-juŋ ni-waɾaɾaŋi aːtani-juŋ jampa nukuɾu
this-Sg 3sgM.R-write-Past2 this-Sg because 2pl
*nurru-**yinag**-baḍa-waḍaḍ.*
nuɾu-**jinak**-paʈawaʈaʈ
2pl.R-head-strong
'Because of your hardness of heart he wrote you this' (WB Mark 10:5)

Morphs analysed in Heath (1984) as ending in an archiphoneme /K/ are treated differently by the Bible translators. Neither the Bible nor any other materials produced by native speakers that I have encountered write this segment at all, casting some degree of doubt over whether it really 'exists' in lexical representations — or at least whether it exists in the same way as ordinary /k/. Examples of *ŋaɾaK-* 'Mult' and *aK-* 'Appl' are given below, where the former is written <ngarra> (ŋaɾa) and the latter is written simply as the long vowel that results from coalescence with the final vowel of the preceding prefix.

(326) a. *Warra-wurru-wurruj marri wu-**ngarra**-jalgarrwina-yinyung marri*
waɾa-wuɾu-wuɾuc maɾi wu-**ŋaɾa**-calkaɾw-ina-jiɲuŋ maɾi
/waɾa-wuɾu-wuɾuc/ /maɾi/ /wu-**ŋaɾaK**-jalkaɾw-ina-jiɲuŋ/ /maɾi/
Pl-Red-person and 3Neut.R-Mult-crawl-Npst2-Rel and

*wu-**ngarra**-guṉmana-yinyung, nganggarra-ngu-jadugang.*
wu-**ŋaɾa**-kuɳm-ana-jiɲuŋ ŋaŋkaː-ʾɾa-ŋu-catuk-aŋ
/wu-**ŋaɾaK**-w$_1$uɳm-ana-jiɲuŋ/ /ŋaŋku-aɾa-Ø-catuk-aŋ/
3Neut.R-Mult-fly-Npst2-Rel 1sg>3Neut.Irr-Mult-Ø-finish-Npst1

‘I will wipe from the face of the earth the human race [I have created] — and with them the animals, the birds and the creatures that move along the ground’ (WB Genesis 6:7)

b. *Nga na-Yawei na-God **naa**-gaḏangi*
ŋa na-Yawei na-God **na**ː**-ʾ**-kaʈ-aŋi
/ŋa/ /na-Yawei/ /na-God/ /**nu-aK**-aʈ-aŋi/
then M-Yawei M-God 3sgM>3sgM.R-Appl-shout-Past2

na-walyinyung, “Ajiga-rruj?”
na-waʃjiɲuŋ acika-ɾuc
/na-waʃjiɲuŋ/ /acika-ɾuc/
M-man.Sg where.Masc-Loc

‘But the Lord God called to the man, “Where are you?”’ (WB Genesis 3:9)

This difference in deletion behaviour between archiphoneme /K/ and true /k/, and the potential variation in the latter, has also been noted in academic work. Among native roots, Bundgaard-Nielsen & Baker (2020) note variation in deletion of incorporated noun-final /k/ in their data, and include an annotated waveform for one of these non-deleting tokens, reproduced here in Figure 5.1. Their study is focused on speech pauses inside verbal grammatical words, and the retained /k/ in Figure 5.1 occurs before one such pause; they do not mention whether coda /k/ survives deletion when not preceding a pause, nor do they mention the behaviour of /K/.

In my own elicitations, as well as in casually observed speech, it is not at all uncommon for /k/ to surface before other consonants, though it appears to occur more in nouns than in verbs. The spectrogram in Figure 5.2 comes from a token of *ama-maṯaʃak-tuc* ‘at the beach’, obtained under elicitation (where the deletion or non-deletion of /k/ was not the focus). According to Heath (1984), this form should be *ama-maṯaʃa-tuc*, with no root-final /k/ — and indeed, this same speaker also produced tokens of this kind under elicitation. When I brought this variability to their attention, my consultants were all willing to accept and produce words containing an

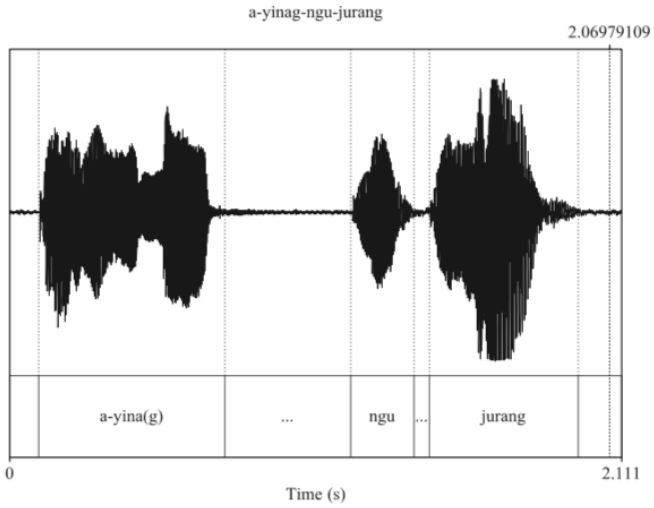

Figure 5.1: Labelled waveform of a token of *a-jinak-ŋu-cuɻaŋ* 'push my head' or 'push me with your head', from Bundgaard-Nielsen & Baker (2020). They note that the /k/ of *jinak* 'head' is retained preceding a word-internal pause in this token, even though the next segment is a consonant, which should predict deletion.

underlying heteromorphemic /k-C/ cluster either with or without deletion of the velar stop, and insisted there was no meaningful difference between them.

On the other hand, I have never encountered underlying /K/ being realised as [k] on the surface. Moreover, my attempts to realise underlying /K/ as [k] on the surface were roundly rejected as ungrammatical by my consultants. I have also recorded a token, in (327), featuring an intra-word pause at the boundary where /K/ occurs in the underlying form. No [k] can be detected there. However, the continuant following the pause undergoes Hardening all the same, thus surfacing as a stop — which, by the way, presents an interesting kind of evidence for phonological subcategorisation, since the segments that trigger the alternation are not phonetically adjacent, meaning it is not really a repair.

(327) wuɾa:[…]cuɾuɾukaɲcini
wuɾu-aK-[…]juɾuɾuka-aɲci-ni
3pl.R-Appl-tell-Recip-Pst2
'They explained [it] to each other' (JN; PN2020-10-27.1; 00:18:37)

Finally, evidence for the failure of /k/ to be deleted in loanwords comes from Baker (2009), who notes that applying deletion in such roots is considered ungrammatical among the speakers he consulted:

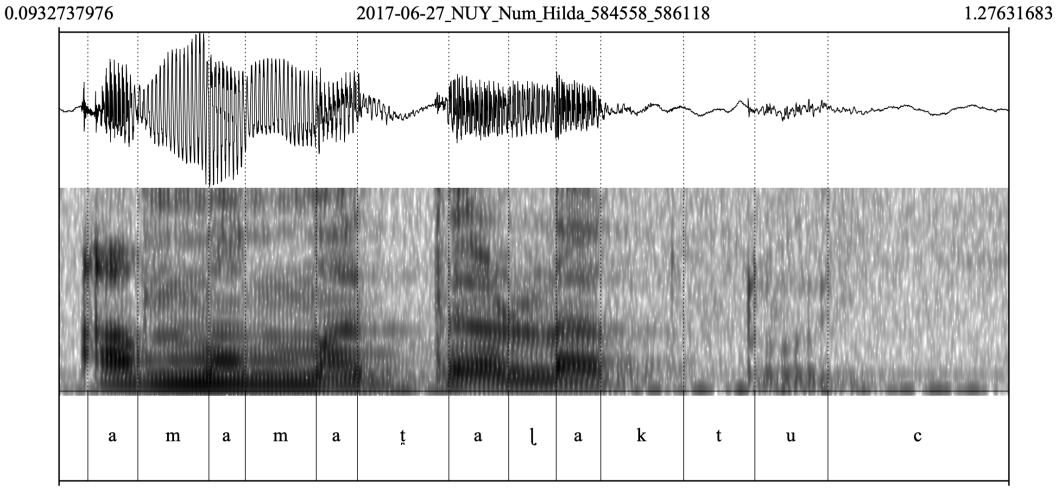

Figure 5.2: Spectrogram of token by speaker HN during elicitation: *ama-maṯaɭak-tuc* 'at the beach'. A faint release burst of the preconsonantal /k/ can be seen in the mid-upper spectra, indicating that the /k/ is not deleted here.

(328) a. ama-cɻak-tuc (*ama-cɻa-tuc)
/ama-cɻak-ɾuc/
'in the car [truck]'

b. ama-swek-tuc (*ama-swe-tuc)
/ama-swek-ɾuc/
'in the swag' (Baker 2009, p. 6))

We find, then, three main types of velar stop-final morphs, based on their susceptibility to velar deletion when preceding another consonant. While the realisation of the final /k/ varies, all three types invariably trigger whatever phonological alternations may apply in the following segment, such as Hardening and Stop Insertion:

1. **Recent loans:** No deletion possible.
2. **Stop archiphonemes:** Deletion obligatory.
3. **Native nouns (and older loans):** Deletion variable.[1]

This variation represents a different set of problems to those so far considered in this thesis. The alternations examined in Chapters 3 and 4 involve discrete, categorical behaviour. For that

[1]There are also some lexical items with two variants, one of which has a final /k/ and the other doesn't, and lexical items which have a variably present final /k/ only in their incorporated form, but these will not be discussed here since the variation is quite different in nature.

reason, it was possible to capture them in a formal model expressed in discrete symbolic terms, such as RM. The variability of Velar Stop Deletion, however, suggests a still-active process that has not yet been completely phonologised, and forces us to grapple with the nature of sound change, the relation between phonetics and phonology, and the push and pull of articulatory, acoustic and perceptual pressures on the shape of the lexicon. This will also necessitate a more direct engagement with issues of language use, which have been in the background to the analyses in previous chapters but have not yet taken centre stage. Whatever the consequences might be for our formal model, the focus here is on getting a handle on the character and distribution of these variants, and coming up with a compelling account of what is driving it. In the next section, I consider various pressures, both phonetic and morpho-phonological, that may either encourage or inhibit deletion.

5.2.5 Parallels with velar nasal deletion and assimilation

Before we move on, it is worth noting that there are also productive phonological processes targeting velar nasals, suggesting that the susceptibility to phonetic reduction or lenition is not just a feature of velar stops but of velar consonants in general. As has been discussed already in Chapters 2 and 3, the velar nasal /ŋ/ is deleted preceding other nasals, and assimilates in place to following stops. And, parallel to the argument for positing a stop archiphoneme /K/, there are morph-final nasals that show this behaviour but cannot appear at the right word edge because they occur in bound morphs, suggesting a nasal archiphoneme /N/.

But, here too, we find a pattern of variation mirroring the deletion of /k/ and /K/. In my data, it is not uncommon to find /ŋ/ surfacing as [ŋ] before heterorganic stops (failing to assimilate) and before other nasals (failing to delete). The example in (329b), was provided by my consultant after I had asked her about the word in (329a), which I read out as it was written in Heath (1980), where the transcription deletes the final /ŋ/ of *mamaɭaŋ*. My consultant therefore produced the /ŋ/ without having been prompted with it (and also coalesced the vowels differently, which is another matter). An example of /ŋ/ surfacing as [ŋ] before a heterorganic stop is given in (329c). The examples in (330) form a near-minimal pair, with one assimilating the /ŋ/ and the other not doing so.

(329) a. Deletion

aŋku:-’ɲ-mamaɭa-maka:
/aŋku-iɲ-mamaɭaŋ-maka:/
3NEUT.IRR-words-good-maka:
‘(not) good words’ (NMET 69.3.5)

b. No deletion

wa:ɻi aŋki:-'c-mamaɭa**ŋ**-maka:
/wa:ɻi/ /aŋku-ic-mamaɭaŋ-maka:/
Neg 3Neut.Irr-words-good-maka:

'They're not good words (HN; PN2020-10-28.2 00:00:57)

c. No assimilation

mana-kulmu**ŋ**-tuc kulmuŋ mana-ma-ŋu-kulmuŋ mana-raŋak
/mana-kulmuŋ-ɾuc/ /kulmuŋ/ /mana-maK-Ø-kulmuŋ/ /mana-raŋak/
Veg.Top-belly-Loc belly Veg.Top-Veg-Ø-belly Veg.Top-canoe

'In the hull [lit: belly], the hull of the canoe' (YN; PN2017-07-07.1 00:37:04)

d. No deletion

ŋu-l̪akuɭa-ŋu-ʈ̪it-aŋi ma.ɻa**ŋ**-miɾi
/ŋu-l̪akuɭa-Ø-ʈ̪it-aŋi/ /ma.ɻaŋ-miɾi/
3sg>3Fem.R-lip-Ø-close-Past2 hand-Instr

'He/she covered her mouth with his/her hands (to stop her talking)'
(HN; PN2017-06-27 00:23:54)

(330) a. Assimilation

nu-ŋu-pakaɭa**ɳ**-ʈaji:
/nuN-pakaɭaŋ-ɻaji:/
2sg>3Resid.R-eye-spear.Npst2

'You're spearing it in the eye' (LM, PN2017-10-03.2 00:05:04)

b. No assimilation

nu-ŋu-pakaɭa**ŋ**-ʈaŋ
/nuN-pakaɭaŋ-ɻaŋ/
2sg>3Resid.R-eye-spear.Past1

'You speared it in the eye' (LM, PN2017-10-03.1 00:33:55)

This behaviour does not extend to the nasal archiphoneme /N/, however, which uniformly assimilates before stops and deletes before sonorants. The three /N/-final pronominal agreement prefixes — /ŋaN-/ '1sg.Irr', /nuN-/ '2sg.R', /naN-/ '12du.Irr' — provide the most frequent context demonstrating this, as can be seen in the many examples in Chapter 3, for instance.

Velar nasals will not be examined in any great detail in this chapter, but these cursory observations are enough to suggest that whatever explanation applies to /k/-deletion should be capable of explaining the variation found in velars in general.

5.3 Dorsal (co-)articulation in Australia and elsewhere

Whereas Hardening is categorical, and therefore is the frozen reflection of a once-active process, the variation we find in Velar Stop Deletion would appear to suggest that it remains an active process, at least for some parts of the lexicon. A thorough account of this process will therefore require an investigation into the phonetic pressures that give rise to it, especially if those pressures are still exerting an influence on speaker behaviour today, as the available evidence appears to suggest.

In this section, I examine the literature on the deletion of codas vis-à-vis onsets, giving particular attention to velars as opposed to other kinds of consonant, and to some of the more cross-linguistically unique properties of dorsal articulation in Australian languages. I seek to understand here two things: why velar stop codas would be deleted at all, and why they would be deleted at different rates in different parts of the lexicon.

5.3.1 The weakness of coda consonants

On the face of it, a process deleting coda stops preceding other stops does not seem especially remarkable. A commonplace view in the phonetics and phonology literature is that coda consonants (especially stops) are less perceptually salient, and therefore more prone to reduction and deletion, than onset consonants. Onset consonants, moreover, typically license a greater number of place and manner contrasts than codas. As we will see, however, Australian languages have often been held up as notable exceptions to these cross-linguistic tendencies, since there is a large body of evidence that speakers of Australian languages exhibit tight articulatory control over VC transitions, and are finely attuned to subtle distinctions in these transitions in perception.

It has long been observed in a wide range of languages that CV transitions provide stronger perceptual cues to consonantal place of articulation than VC transitions (Fujimura et al. 1978; Malecot 1956). Perhaps surprisingly, this bias appears to hold across languages with very different phonotactics. For instance, Fujimura et al (1978) compared English and Japanese listeners in their ability to discriminate consonantal place of articulation in VCV sequences. In Japanese, a coda consonant is only permitted where is has the same place values as the onset consonant in the following syllable in the word, while in English all place of articulation contrasts are available in codas, including word-finally. Since coda consonant place is lexically

contrastive in English, one might expect English listeners to be much better than Japanese listeners at discriminating consonantal place of articulation on the basis of VC transitions. In their study, the authors presented participants with artificially spliced VCV stimuli in which the VC and CV transitions were taken from productions of different consonants. For example, the first portion of [aba] was spliced with the second portion of [aga], resulting in a stimulus in which the transition into the consonant suggested a [b], while the transition out of the consonant suggested [g]. English and Japanese listeners both overwhelmingly identified the intervocalic C on the basis of the CV transition, not the VC transition, and the two groups did not differ significantly in this regard, despite the different phonotactics in the two languages. That is, for the example just mentioned, participants more frequently identified [aga] not [aba], since the CV transition splice was taken from a production of [a**ga**], and participants apparently prefer to 'tune in' to CV transitions rather than VC transitions.

Notably, when the stimuli were played backwards, the subjects still preferred to identify place cues on the basis of the C-to-V transition of the string they heard — i.e., the [aba] interpretation of the example just mentioned — even though the portion they hear as CV [ba] came from a recording of a VC [ab] transition, and therefore lacked the release burst. The authors conclude that this asymmetry is due to an inherent bias in speech perception towards CV as opposed to VC transitions, and that this bias is robust in the face of cross-linguistic phonotactic differences. A potential limitation of the study is that participants were only able to respond by identifying a single C; they could not respond by identifying a cluster. Were such a response available, one would expect that English listeners would be more likely to suggest clusters when presented with different cues to place of articulation in the VC and CV portions of the stimuli than Japanese listeners.

Ohala (1990) builds on Fujimura et al's (1978) study by examining the effect of duration on cluster perception. In this study, stimuli were again constructed by cross-splicing CC clusters from the VC and CV portions of VCV nonsense words, but the duration of the intervocalic cluster was also manipulated in the splicing, decrementing the interval in steps of ten milliseconds. Participants are presented with these cross-spliced VC_1C_2V clusters and asked to identify the consonant, but they also have the option to respond with a cluster, unlike in Fujimura et al's (1978) study. Listeners frequently report only hearing C_2 when presented with a C_1C_2 cluster, and this likelihood increases as the duration of the intervocalic interval in the stimuli decreases. Because the duration is manipulated artificially, the source of misperception is attributed to a bias in the perceptual system itself, and not to changes in articulation that arise when increasing speech rate. Ohala (1990) uses this to argue for listener misperception as a source of the typologically common sound change from heterorganic clusters to homorganic ones (including geminates).

Even so, articulatory considerations also bear on the question of why codas and not onsets are preferentially deleted. Articulatory Phonology (Browman & Goldstein 1989, 1992) represents the phonological form of lexical entries as a two-dimensional *gestural score*, with articulatory tiers on one dimension and time on the other. These tiers represent the sets of articulators (lips, tongue tip, tongue body, velum, glottis, etc.) employed by the gestures. A percept of 'deletion' can be the result of *gestural overlap*, where a gesture on one articulatory tier (i.e. active articulator) masks the auditory cues of a temporally adjacent gesture on another tier. So for example, in faster, more casual productions of the English phrase *perfect memory*, the /t/ is often absent from the acoustic record and therefore tends to be regarded as 'deleted'. But in fact the tongue tip gesture associated with the /t/ can often still be observed using instrumental articulatory phonetic techniques. Because this tongue tip gesture temporally overlaps with the labial gesture in the /m/ of *memory*, it is 'perceptually masked': the closure of the lips associated with the /m/ conceals the percepts associated with /t/ (such as formant movements or release burst).

Articulatory Phonology makes the prediction that temporal compression (as in faster speech) should produce greater temporal overlap between gestures. Where the gestures concerned occupy different articulatory tiers (i.e. active articulators), the perceptual result will be masking of the posterior gesture by the anterior one; where the gestures occupy the same tier, the result will be a blending of the two gestures. Since /k/ is the most posterior stop gesture in the Wubuy phonemic inventory, this would predict that /k/ should be masked by any other subsequent closure under conditions of temporal compression. This would suggest gestural scores for, say, /kt/ clusters like those in (331). Each score illustrates the temporal coordination between the velar constriction gesture on the Tongue Body articulatory tier and the alveolar gesture on the Tongue Tip tier with either, (a) little overlap, or (b) complete overlap. The percept is different in each case: [kt] in the former case, but only [t] in the latter, due to the alveolar gesture masking the percept of the velar. These different phasing relationships between gestures are not discrete but continuously variable, so it is possible that where there is a degree of overlap larger than (331a) but still shy of (331b), subtle acoustic cues to the velar constriction may still be present.

(331) Gestural scores for /kt/ clusters

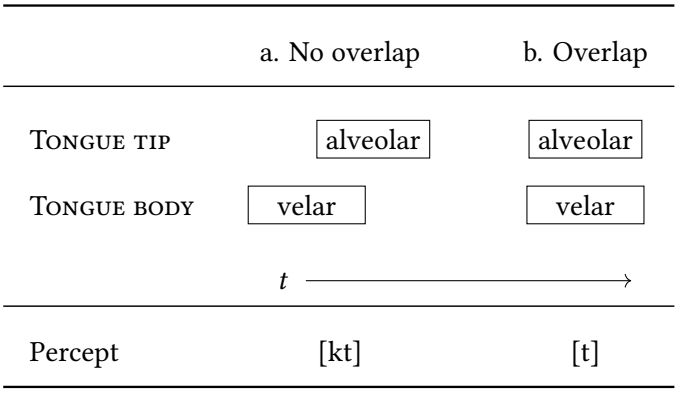

Another potentially relevant articulatory factor here is the size of the articulator. Because the tongue dorsum is a larger and bulkier articulator, dorsal gestures are expected to be slower in reaching the target of occlusion than gestures involving the tongue tip (Bybee & Easterday 2019; Recasens & Espinosa 2009). Again, this predicts that dorsal gestures are more likely to overlap in time with a subsequent consonantal gesture than tongue tip gestures in the same position, all else being equal. Bilabial gestures have been suggested to also be slower than tongue tip gestures (Bybee & Easterday 2019), but since bilabials only appear as onsets in Wubuy (except in a handful of loanwords), they are not very useful as a basis for comparison with dorsals.

The foregoing has served to suggest that the coda position is one of inherent perceptual and articulatory weakness, compared to onset position. Synchronically, listeners often fail to perceive coda consonants when they are present (due to their weaker cues), and often articulate them in such a way that transitional cues are masked by a subsequent onset (due to gestural overlap). According to theories of sound change incorporating a production-perception feedback loop, asymmetries such as these can provide the seeds to phonological alternations such as assimilation and deletion (e.g. Blevins 2004; Ohala 1981; Steriade 2001; Wedel 2006). This suggests that, if the /k/ in a /kC/ cluster is already perceptually weak, a listener may fail to hear it altogether, and then they may go on to produce the same word without any velar stop gesture at all — i.e. as simply the 'second' C. Subsequent generations of listeners then have no reason to posit a velar stop as part of the representation at all. Moreover, both the perceptual and articulatory considerations predict that deletion should be more frequent where the speech rate is faster.

5.3.2 VC transitions in Australian languages

However, it has long been argued that Australian languages buck this trend, as there are many lines of evidence suggesting that VC transitions are at least as important in speakers' perception

and articulation of consonantal contrasts as CV transitions. Summarising a large body of phonetic research, Butcher (2006, 2012) posits an overarching 'Place of Articulation Imperative' at work in Australian languages. As he points out, Australian phoneme inventories are typically 'long and thin': rich in place distinctions but poor in manner. They often contrast as many as six places of articulation in the oral and nasal plosives, at least word-medially, with no contrastive voicing, and no fricatives. Wubuy, it should be noted, is somewhat extreme in this regard in having all four coronal stops distinguished in initial, medial and final contexts in the word (Bundgaard-Nielsen, Baker, Harvey et al. 2012; Bundgaard-Nielsen, Baker, Kroos et al. 2015; Bundgaard-Nielsen, Kroos, Baker et al. 2016). With lexicons built on such fine-grained contrasts in place cues, this creates a pressure for tight articulatory control over cues to place of articulation, and adept perception of 'systematic differences in formant transition patterns at vowel-consonant boundaries' (Butcher 1999). This pressure to maintain robust spectral cues to place of articulation has been argued to exert effects on patterns of nasal place assimilation, anticipatory nasal airflow, as well as on the shape of phoneme inventories and productive phonological alternations in Australian languages.

Nasals preceding heterorganic stops have been shown to be remarkably more resistant to place assimilation in Australian languages than in English, where nasals tend to have the same place of articulation as a following consonant (Fletcher & Butcher 2014). That is, Australian languages often allow heterorganic nasal-stop clusters like [nk] or [np], without assimilating them in the direction of [ŋk] and [mp]. Speakers of Australian languages have also been shown to delay velic opening in vowels preceding nasal plosives until very close to — sometimes even later than — the onset of oral constriction (Butcher 1999; Carignan et al. 2023; Stoakes et al. 2020). This contrasts notably with the much more widespread crosslinguistic tendency to begin nasal airflow during the preceding oral vowel. The delay in nasal airflow mitigates against nasal antiresonances being produced during the vowel transition into the nasal consonant, which interfere with the formant cues to place of articulation that come from resonances in the oral cavity. That is to say, delayed onset of nasal airflow keeps place of articulation cues more perceptually distinct. In a number of Australian languages, this unique pattern of gestural coordination results in phonetic pre-stopping of post-vocalic nasal plosives; in a handful of these languages, such articulations have become phonologised (i.e. contrastive with ordinary nasals and oral stops), as in Arrernte, Arabana, Adnyamathanha, and Kaytetye (Carignan et al. 2023; Harvey, Lin et al. 2015; Harvey, San et al. 2019). Likewise, lateral airflow is often relatively late in the articulation of vowel-lateral sequences, and has been reported in some languages to result in pre-stopped lateral consonant phonemes (Butcher & Loakes 2008; Round 2014). Both the avoidance of anticipatory place assimilation and the avoidance of anticipatory nasalisation can be seen as reflections of the Place of Articulation Imperative, and thus can be seen as

maintaining the robustness of VC transitions to a degree not typically found elsewhere in the world.

Retroflex consonants occupy a special place in this discussion. The cues to retroflexion (those which distinguish apico-retroflex and apico-alveolar) are more perceptually distinct in the transitions into the consonant than out of it (Stevens & Blumstein 1975). A widespread feature among Australian languages is that the apicals are often neutralised word-initially, but not medially. This is held to be a consequence of the directional asymmetry in their perception (Steriade 2001). Therefore, we should expect that speakers of a language in which retroflexion is contrastive must pay particular attention to VC transitions.

Recognising the unusually tight coordination of VC sequences and the effect of this on the development of Australian language phonological systems, some researchers have even gone so far as to posit VC as the basic syllable type in certain Australian languages (e.g., Breen & Pensalfini 1999; Sommer 1970). According to this view, CVC sequences are preferentially syllabified as C.VC, rather than the putatively universal CV.C. Further support for this may come from the fact that many languages across Australia display evidence of historical processes of lenition and deletion of word-initial onset consonants, against the general cross-linguistic trend for onsets to resist such processes (Blevins 2001; van Egmond & Baker 2021). There has nonetheless been considerable criticism of the VC syllable proposal, on both phonological (Round 2013) and phonetic grounds (Tabain, Breen et al. 2004).

5.3.3 The special weakness of dorsals

Although the above discussion suggests that VC transitions, including in syllable codas, are more resistant to reduction in Australian languages than is the norm cross-linguistically, there are various lines of evidence suggesting that dorsals are something of an exception. The first is typological. Hamilton (1996) notes in his landmark survey of Australian phonotactics that dorsals are marked in coda position. In a heterorganic intervocalic C_1C_2 sequence, he provides the following implicational hierarchy for place features in C_1: Labial > Dorsal > Laminal > Apical (where the presence of one feature implies the feature to its right). That is to say, dorsal codas are uncommon in Australian languages, even though various other unusual kinds of heterorganic cluster, like [nk] or [np], are normally permitted.

There is also some phonetic evidence that of the six typical places of articulation in Australian languages, dorsal articulations are the least governed by the Place of Articulation Imperative — though as Butcher & Tabain (2004) note, coronals have been given the lion's share of scholarly attention among phoneticians, at the expense of dorsals and labials, and indeed Wubuy's four-way coronal distinction has received fairly extensive phonetic investigation, while dorsals have not (Bundgaard-Nielsen & Baker 2014; Bundgaard-Nielsen, Baker, Harvey et al. 2012;

Bundgaard-Nielsen, Baker, Kroos et al. 2015; Bundgaard-Nielsen, Kroos, Baker et al. 2016; Bundgaard-Nielsen, Kroos, Harvey et al. 2010; Kroos et al. 2010; Proctor et al. 2010). Butcher & Tabain (2004) present evidence from palatography and F2 locus equations (on which more below) that 'velar' stops in several geographically dispersed Australian languages vary widely in the point of contact with the roof of the mouth, with productions extending from the postalveolar region right back to practically uvular territory. They show that this variation is to a large degree determined by the place of the following vowel: front vowels condition front velar articulations, while back vowels condition more uvular articulations. Comparable English palatographic data shows a smaller degree of allophonic variation in contact area within each vowel context. This is a surprising finding given the Place of Articulation Imperative, as it predicts that a smaller inventory of contrastive places of articulation, such as English has, should be correlated with a greater degree of within-category variation in articulatory targets. However, the authors point out that, compared to crowding among coronals, with four different constriction locations, dorsal constrictions can vary widely in location without encroaching on another target. Another way of thinking about such data is in terms of coarticulation: dorsals show a greater degree of coarticulation with adjacent vowel targets than consonants at other places of articulation (see also Tabain & Butcher 1999). This is likely because vowels also require a dorsal gesture.

In the case of Iwaidja, (non-nasal) dorsal articulations are so variable that descriptive work has posited a phonemic contrast between a velar stop /k/ and a velar approximant /ɰ/. In one of the very few studies on variation in an Australian language, Shaw et al (2020) argue that this apparent distinction is likely a result of undersampling what is in fact better characterised as a widely varying continuous phonetic space for a single contrastive phoneme, on the basis of acoustic and ultrasound evidence from several speakers. Some productions of this phoneme are vowel-like, with only weak constriction and very little dip in acoustic intensity; others are more clearly stop-like. Although the authors are concerned with manner, not place, their findings can also be understood, like the findings of Butcher & Tabain (2004), in terms of velars showing low coarticulatory resistance to neighbouring vowels: intervocalic velars in Iwaidja often show almost no departure from the articulatory settings of their flanking vowels. It is also telling that it is only velars in Iwaidja, not constrictions at other places of articulation, that show such extreme variation.

The tendency for tight control over anticipatory place of articulation cues and late velic lowering for nasals in Australian languages discussed above also appears to be somewhat relaxed in the case of /ŋ/ with respect to other nasals. Butcher (1999) shows that velar nasals are especially prone to being realised as oral stops, oral approximants, or being deleted altogether. In an aerodynamic study of nasal coarticulation in Bininj Kun-Wok, Stoakes, Fletcher & Butcher (2020) find greater anticipatory nasalisation in vowels preceding /ŋ/ than preceding nasals at

other places of articulation.

A perceptual factor potentially impacting upon the discriminibility of dorsal nasals with respect to other places of articulation is nasal antiresonance, which is higher in more retracted tongue positions, due to the consequently larger size of the oral cavity in front of the velopharyngeal port (Ladefoged & Maddieson 1996, p. 116). This nasal antiresonance dampens the intensity of formants above F1, and can therefore interfere with the perception of oral place of articulation cues. From the perspective of listener-oriented sound change, one may expect that if listeners find it harder to distinguish place of articulation cues in dorsal nasals than nasals at other places of articulation, then the cues to the following segment's place of articulation may take on more prominence and the listener can 'mishear' assimilation.

Recall that in Wubuy, nasal place assimilation applies uniquely to /ŋ/, not to other nasals. That is, sequences of /ŋ/ followed by a coronal or labial do not occur, even though other kinds of heterorganic NC clusters do. This state of affairs may be interpreted as reflecting a process of phonologisation, such that dorsal gestures' greater susceptibility to coarticulation with the following oral gesture becomes a categorical rule of place assimilation.

5.3.4 Informativity effects

An increasingly large body of literature in the usage-based tradition has investigated the relationship between phonetic reduction and frequency of use, as well as related measures like predictability and informativity (Bybee 2001; Cohen Priva & Jaeger 2018; Shaw & Kawahara 2018). In general, there is a good deal of evidence that more predictable sequences tend to be more prone to phonetic reduction (Aylett & Turk 2004; Fosler-Lussier & Morgan 1999; Pluymaekers et al. 2005a). This is often understood to function as a local effect: for instance, a word may undergo greater reduction in those discourse contexts in which it is highly predictable given the word that preceded it (Aylett & Turk 2004; Baker & Bradlow 2009). On the other hand, informativity can have a more global effect, where the average informativity of a given segment across all contexts inversely correlates with its likelihood to undergo reduction in any context. Cohen Priva (2015) has shown, working from a corpus of American English, that segments which are usually predictable are prone to reduction even in contexts where they are locally unpredictable, and vice versa. In other words, the average informativity of a segment across *all* contexts influences its likelihood to be deleted in *any* given context. Relatedly, a phonological opposition with a higher functional load — measured by the number of minimal pairs which depend on it — is more informative in the sense that it does more work to distinguish one item from another; this has been shown to correlate with the likelihood of its retention over time (Babinski & Bowern 2018; Wedel, Jackson et al. 2013). Also, phonemes that are more informative in this sense are more resistant to deletion synchronically (van Son & Pols 2003). All of these

informativity effects have been argued to play a role in sound change (Hall et al. 2016).

Unfortunately, the methods used in this information-theoretic literature require corpora that are larger than the Wubuy corpus by an order of magnitude; it is therefore not possible to estimate the informativity of velar stops in Wubuy with any degree of precision. However, there are two ways of thinking about informativity in this context which do not require a full quantitative analysis. First of all, we know that morph-final /k/ or /K/ has a low functional load, in the sense of distinguishing minimal pairs. In the dictionary there is only one minimal pair of morphs whose members differ by whether a final velar stop is present, namely the incorporated nouns /l̪akaɾK-/ 'line, row' and /l̪akaɾ-/ 'saliva'. Neither of these appear to be productive: the former appears in two compounds (one of them with a root not attested outside this one compound), and the latter appears in just one: *l̪akaɾ-l̪ilil̪a* 'to slobber out [lit: saliva-drip]'.

Secondly, recall from the discussion above that a morph-final /k/ conditions some change in the sequence that follows if it is underlyingly continuant-initial or vowel-initial (e.g. Hardening, Velar Stop Insertion). Therefore, the loss of a coda velar stop gesture does not threaten any contrast, even in the rare case of /l̪akaɾK-/~/l̪akaɾ-/. In other words, one can very often *infer* the presence of an underlying /k/ (or /K/) even if it is not realised (cf. 'hypocorrection' in Ohala 1981). Listeners need not rely on auditing a morph-final [k] in order to identify the morph, since the /k/~/Ø/ alternation is also nearly always reflected in the choice of allomorph in the following morphological constituent. In a sense, then, the phonological subcategorisation requirement that a /k/-final morph imposes on a following morph is already informative enough, and renders the /k/ itself redundant, therefore able to be deleted at no cost to the listener. Thus, the broader morphophonological context may also be a factor driving velar stop deletion in Wubuy.

5.3.5 Summary

The literature discussed in this section has shown that deletion of coda consonants within clusters is cross-linguistically common, whereas the deletion of onsets is not. This has sometimes been attributed to an inherent bias in perception towards CV transitions over VC ones. In the case of velar occlusions, there may also be articulatory factors that render them more likely to be perceptually masked by a following constriction at another place of articulation: the posterior point of contact and the bulkiness of the tongue body. Australian languages are sometimes held out as an exception to these generalisations, since speakers are notably adept at maintaining tight articulatory control over VC sequences, and allow a range of heterorganic consonant sequences that seldom appear elsewhere. However, these claims appear to be restricted to coronals — indeed, dorsal stops in some Australian languages show less coarticulatory resistance than dorsals in English, and the typology of phonotactic patterns shows that they rarely appear as the first consonant in a heterorganic cluster. One of the points to take from this literature is

that 'deletion' may be too coarse a concept where variation is concerned. On both perceptual and articulatory grounds, we have reason to think that a coda may be heard as deleted even if it is actually produced, and traces of it can be detected with more fine-grained phonetic methods.

However, the research on the phonetic properties of consonants in Australian languages has largely focussed on their realisation within morphs, rather than at morphological boundaries. Since this is where phonological alternations tend to be found in such languages, we know very little about how alternations are realised phonetically, and even less about alternations that show variation. Dorsal deletion in Wubuy therefore represents a unique opportunity to study what appears to be a sound change in progress. The fact that it appears in the context of a rich set of alternations only heightens the interest, and the archiphoneme lends a dimension to the problem that is essentially *sui generis*.

5.4 Phonetic properties of dorsal stops in deletion contexts

5.4.1 Introduction

As discussed above, Heath's (1984) analysis of dorsal stop deletion claims that the process is inescapable, and that it applies irrespective of the type of morph involved. Other sources, including my own preliminary investigations, suggest that only the archiphoneme /K/ completely fails to escape deletion, while /k/ sometimes appears on the surface preceding other consonants, suggesting the rule does not apply uniformly. Moreover, there appear to be differences between native roots and loan roots in their susceptibility for deletion, where the latter appear resistant to it.

In order to probe the variability of dorsal stop deletion, I conducted two structured elicitation tasks, one testing noun roots in suffixing contexts, the other testing nouns incorporated into verbal words. The annotated recordings of these tasks formed the corpus for the quantitative analyses to follow.

These tasks set out to answer the following questions:

1. Is dorsal stop deletion variably applied?
2. If variable, does the variability differ:
    a) between morphological contexts (i.e., between verb and noun words)?
    b) between true /k/ and archiphoneme /K/?
    c) between native and loan roots?
3. What are the acoustic phonetic consequences of deletion on:

a) duration of intervocalic underlying /kC/ closures?

b) VC formant transitions into underlying /kC/ closures?

Heath's (1984) position can be taken as our null hypothesis, predicting that deletion applies uniformly to all preconsonantal /k/ and /K/ segments, and there should be no acoustic phonetic evidence of a dorsal gesture in those contexts. That is, a surface C arising from an underlying /kC/ sequence should not differ significantly in duration from that same C arising from an underlying single /C/. Accordingly, the formant transitions from a vowel going into a surface C arising from underlying /kC/ should not be significantly different from the vowel formant transitions preceding a surface C arising just from /C/. In order to reject the null hypothesis, we would need to find evidence that, quite simply, deletion *does not* apply all of the time. We can take this as our first alternative hypothesis.

If there is in fact evidence of variability in deletion, we can then ask whether that variability differs by context (morph type, word type). Our second null hypothesis, then, is that there is no differentiation by context; our second alternative hypothesis is that there is.

In 5.4.2, I present the task materials, the participants, the elicitation and recording procedure, the annotation process, the annotated corpus, and methods for the quantitative and statistical analyses. In 5.4.3, I present the results of these analyses in light of the research questions posed by the study.

5.4.2 Method

5.4.2.1 Task materials, participants, and procedure

For this study, I conducted two structured elicitation tasks, one testing noun roots in suffixing contexts, the other testing nouns incorporated into verbal words. For the noun task, I assembled a wordlist containing 48 noun roots; see Table 5.4. The list contained 33 /k/-final roots, 22 native and 11 non-native. The remaining roots, which were not /k/-final, contained a range of stops at different positions, and were included to facilitate comparison with (underlying) /k/ in clusters in the acoustic analysis. The native roots were drawn from Heath's (1982) dictionary; the non-native ones were drawn either from my own knowledge of common loanwords, or from the online Kriol dictionary (Lecompte 2014). This wordlist was used as the basis for a PowerPoint presentation, where each slide contained an image depicting the noun of interest. I checked the words for familiarity and the images on the slides for appropriateness with speaker JN.

In the task, I asked the participant what each image depicted, and then placed a paper cut-out image of an ant (*wara-ŋumic*) on top of the computer screen and asked, *awika wara-ŋumic?*, 'Where is the ant?'. This prompted answers with the Locative suffix *-ɾuc/-tuc*. For /k/-final

Native, /k/-final		Non-native, /k/-final		Filler	
anik	'ash'	pek	'bag'	paipuɭ	'Bible'
(w)aɳaɾak	'young tree goanna'	keik	'cake'	kalajuɾu*	'sea perch'
caːɾak	'young silver gull'	klok	'clock'	kataɾa	'coconut palm'
jaːk	'meat ant'	tilipek	'dillybag'	kuːku	'freshwater'
jimuɾk	'nose'	kɻok	'grog; alcohol'	kulmuŋ	'belly'
jinak	'head'	plok	'block (child's toy)'	maːta	'Macassan smoking pipe'
kaːkak	'snail'	cɻak	'car (truck)'	maɻpiti	'small ray'
kaṯiɻkaṯiɻk	'oystercatcher bird'	swek	'swag'	ŋuɻuṯu	'bird' (generic)
ḻaːɾak	'log coffin'	taktak	'rubber duck'	ŋutan	'entrails'
ḻakajak	'saltwater'	feisbuk	'Facebook'	paK-kalaːti(-ɲuŋ)	'blind person'
ḻapak	'feather'	cekbek	'sack'	ṯuɾapata**	'spear'
ḻawaːlik	'water well; soak'			ṯutucu	'crowbar'
ḻawumak	'long yam'			wiɳcitu	'golden trevally'
ɭampak	'cup'			wulpuŋ	'pandanus basket'
ɭatak	'milkwood tree'			wutu	'kidney'
maŋaɻak	'sand'				
ɻaːk	'young brown tree snake'				
ɻaŋak	'tree' (generic)				
ṯaːɾak	'beard'				
wiːk	'grey hair'				
wuɾk	'grass fire'				

Table 5.4: Stimulus items used in noun elicitation task.

* The dictionary listing for 'sea perch' is *kalajuɾuk*, with final /k/, but this was recognised as *kalajuɾu* by JN. The other speaker, DM, did not recognise the word.

** *ṯurapata* can also be realised, quite commonly, as *ṯutapata*. Here, DM produced the former, JN produced the latter, but I have heard both speakers use both forms on other occasions. This variability is not pertinent to the present study.

nouns, this context requires the Hardened allomorph *-tuc*, and it should also entail deletion of the final /k/. The participants were asked to repeat each response five times. Sometimes participants offered a synonym; I allowed them to use whatever word came to them first, and then asked if they knew a different word for the same thing. If they still did not suggest the word I was looking for, I suggested it to them and made sure they were familiar with it before recording them producing it.

The locative suffix was chosen as the context of interest for three reasons. Firstly, any noun whatsoever can appear with a locative suffix, and its semantic contribution is always fairly obvious. For most other case suffixes, it can be much more difficult to construct, for each noun of interest, a plausible, concrete situation for elicitation. For instance, it is difficult to imagine how to elicit a topological feature, like *maṯaɭak* 'beach', with the INSTR suffix *-miri* — indeed, of the 25 tokens of *maṯaɭa(k)* in Heath (1980), 19 are suffixed, but never with *-miri*. Secondly, most of the other common suffixes, like Allative/Dative *-wuj~-kuj* or Ablative *-wala~-kala* begin with /k/ in their hardened allomorphs, rendering them inappropriate for investigating the presence or absence of a preceding morph-final /k/ (though there may be durational differences). Finally, since I wanted to investigate the interaction between /k/-Deletion and Hardening, I needed a suffix subject to Hardening, therefore ruling out any nasal-initial suffixes.

The verb task was more similar to a standard elicitation procedure: I described a thematic context where a particular verb stem might be required (the one I hoped to elicit), and asked the participant to translate from an English prompt. I established a list of words to target for elicitation by taking items from a set of derivational elements and a set of verb roots, and combining these to form verb stems; see Table 5.5. The derivational elements included (underlyingly) /k/-final incorporated nouns, /K/-final suppletive incorporated nouns and the Applicative derivational prefix *aK-*. Note that all of the /k/-final incorporated nouns in Table 5.5 also appeared in the noun task stimuli in Table 5.4. The main verb roots of interest were the first three in Table 5.5, which are of the alternating class. For the purposes of comparing the results in this task with the results in the noun task, it would be desirable to test verb roots beginning in underlying /ɾ/, which would harden to /t/, just like the Locative suffix. Unfortunately, /ɾ/-initial verb roots are extremely rare in Wubuy, and do not alternate with a stop-initial realisation, instead patterning together with the sonorant non-continuants (Heath 1984, p. 27). Still, for the purposes of examining V-to-C F2 transitions, /ʈ/ should be similar to /t/, since the apicals are expected to be mainly differentiated by a lowering of F3 preceding the retroflex. I also included three roots whose hardened allomorph begins with /k/, on the suspicion that the intervocalic closure period should be longer for /k-k/ than for /K-k/. I first checked with HN that the targeted constructions were possible words of Wubuy, but I did not explain the specific research aims. Again, the participants were asked to repeat each response five times.

Derivational elements		Verb roots (citation forms)	
/aK-/	Appl	/wa/~/pa/	bite
/paK-/	eye (suppl.)	/waɭpu/~/paɭpu/	paint
/ŋuɻaK-/	penis (suppl.)	/ɻapala/~/ʈapala/	pop out
/ʈaːɾak/	beard		
/jimuɾk/	nose	/wulʈa/~/kulʈa/	cut off
/jinak/	head	/aɾu/~/kaɾu/	leave (tr.)
/ɻaŋak/	stick; tree	/aɾaɾi/~/kaɾaɾi/	paint
/kulmuŋ/	belly		
/ɭaːɲ/	tongue		

Table 5.5: Derivational elements (prefixes or incorporable noun roots) and verb roots used to construct stimuli in verb elicitation task.

The suppletive incorporated nouns elicited were *paK* 'eye' and *ŋuɻaK* 'penis', since those are the only two that appear to have any appreciable productivity. Both of these require further comment.

The suppletive form *paK* appears in many constructions, and speakers are generally willing to accept and produce novel constructions with it. However, under elicitation both in this task and on other occasions, participants always first offered a form incorporating non-suppletive *pakaɭaŋ* — the free form of 'eye' — so I had to specifically suggest forms with *paK* in order to elicit them. They agreed that such forms were possible Wubuy words, and produced them for me when asked. The dataset reported on here contains both their initial offerings with *pakaɭaŋ*, as well as their productions with suppletive *paK*.

The other suppletive form, *ŋuɻaK* 'penis', has the potential to be an inappropriate topic of discussion with some speakers (depending, especially, on their classificatory kin relation to me), and I approached this one with caution — despite the fact that one is liable to hear Wubuy speakers employing compounds featuring this morph just about every day, since it is the favoured device for forming insults (used even by heritage speakers). I told each speaker that if they did not think it right to talk about words involving this, I would skip it and move on. After some shock and laughter, all the speakers indicated that they wanted to show me how to say it correctly, and proceeded to answer my prompts.

Five speakers were recorded in the verb task, and three in the noun task. However, due to the extensive time required to transcribe and process the recordings, I could only include two of each; in each case, I chose the recordings that seemed easiest to transcribe. One speaker appears in both tasks, so there are three speakers analysed here in total.[2]

[2]Additionally, I excluded the data from one of the five participants in the verb task because I was not confident

Neither task controlled for utterance position, and sometimes the participants would insist on framing the answer in a particular way. Speaker DM, for instance, would often insist on providing a verb and a demonstrative in his responses to the noun elicitation task — his response to being asked where the ant is (where the picture prompt is a tree) would not be the absolute minimum of, for example, *a-ɻaŋa(k)-tuc* 'on the tree', but rather *waɾa-ŋumic tawu wu-wiʈi-wiʈi a-ɻaŋak-tuc* 'the ant is there climbing on the tree'. I allowed the participants to provide whatever response they thought appropriate. This is a potential limitation to the corpus to be borne in mind.

Recording sessions took place in the library of Numbulwar School in November 2020 — likely the quietest room in the community at that time, but still far from noise-free. Recordings were taken with a Zoom H4n portable recorder, using the onboard XY microphone, but only the left channel (aimed at the speaker) was used for analysis, and the original 24-bit/96kHz recordings were later downsampled to 16-bit/48kHz in order for the formant tracking software to work to specification.

5.4.2.2 Transcription and speech database construction

The recordings were first transcribed in a modified Wubuy orthography at the utterance level in ELAN (Brugman & Russell 2004), then converted into an emuR database (Winkelmann, Jaensch et al. 2019) following automated grapheme-to-phoneme conversion and forced alignment at the segment level via WebMAUS (Kisler et al. 2017) using the language-independent model. Since it would not be straightforward to locate deleted segments without marking what has been deleted, I explicitly transcribed all underlying velar stops at this first stage of transcription in ELAN, whether or not they were audible. Moreover, I used different symbols for /k/ in native roots (<*g*>), /k/ in loans (<*k*>) and archiphoneme /K/ (<*G*>). Again, this was to facilitate querying in the analysis to follow.

I corrected all of the segment boundaries resulting from the forced alignment process by hand. Boundaries for vowels were determined by inspecting the spectrogram and waveform primarily for coherent formant structure — where this structure begins, the vowel segment label begins; where it breaks apart, the vowel segment ends. Faint traces of formant structure are still sometimes discernible during stop closures due to some reverb from the room, but this reverb tail had an obviously lowered amplitude and therefore did not substantially interfere with the segmentation procedure.

As previously mentioned, all tokens involving a morph-final (underlying) dorsal stop were

that the participant understood the task. And, most unfortunately, the recordings from a further participant in both tasks, not included in the count here, could not be used because the recording device failed during the session and corrupted the files.

transcribed as if no deletion had occurred. Where these /k/ or /K/ segments are followed by another stop (that is, in the context of interest), I aligned the right boundary of the velar with roughly the centre of the entire /kC/ closure. Correspondingly, in the resulting emuR database, these underlying stops have a phonetic segment associated with them, with a start and end time, as I wanted to allow the database to be agnostic with respect to whether morph-final velars were present or deleted from the surface during the querying procedure. However, this is merely a technical convenience — duration measures are only taken and analysed for the entire /kC/ (or /<k>C/) closure, as the independent /k/ portion cannot be directly observed acoustically except in the few cases where there is an audible release for the postvocalic /k/. For these tokens that included a visually detectable release for the preconsonantal /k/, the end of that release was taken to be the end of the /k/ annotation. The right boundary in such cases corresponds to something present in the signal, but never forms part of the analysis: duration data is only ever taken from the whole cluster, not the /k/ portion on its own. Examples of this labelling are shown in Figure 5.3 and Figure 5.4. In the former token, no percept of [k] can be heard, while in the latter it can; nonetheless, there is a segment labelled ”k” in each. (A particularly clear example of an audible release is shown in Figure 5.5.) By including the underlying stops in the transcription even where they are inaudible, we are able to distinguish between /kC/ sequences and /C/ singletons in the data, and compare them. That is, since the identity of the singletons is not at issue, these segments function as our control in analysing duration and formant transitional data.

I also explicitly coded each /k/ and /K/ in a deletion environment for whether it sounded deleted to my ears (using an additional attribute of the phonemic tier in the emuR database). This allows us to compare the acoustic properties of closures that appear to delete the velar stop against those that do not. Of course, the usual caveats around the perceptions of a non-native speaker apply. Also, where the segment following the underlying coda /k/ or /K/ is also a /k/, I always transcribe this /k/ as deleted (as per Heath’s predictions) since it is necessarily much less obvious whether it is deleted or not phonetically. We will see in the results below that /kk/ sequences are nonetheless very long, suggesting two timing slots.

After removing some obviously disfluent tokens, the resulting database contains 480 tokens in the nouns condition and 593 tokens in the verbs condition. Of the nouns, 297 tokens feature a /k/-final root (and therefore a Hardened suffix allomorph); the rest are filler, but provide tokens of /k/ and /t/ in (root-internal) intervocalic contexts to use as a baseline of comparison with the heteromorphemic clusters.

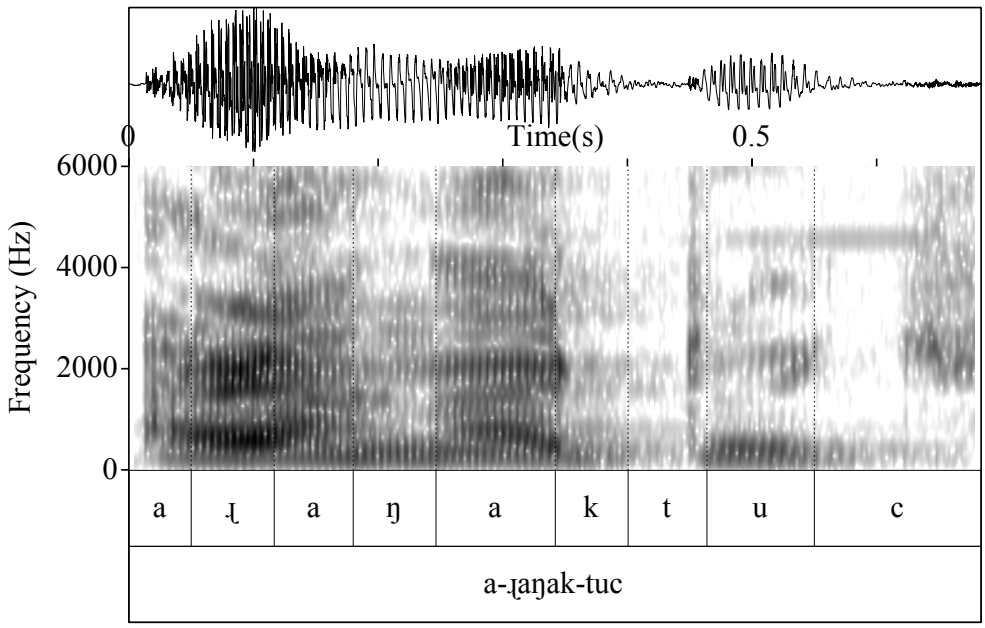

Figure 5.3: Spectrogram and waveform for *a-ɻaŋa(k)-tuc* 'NEUT.OBL-tree-LOC', where no [k] percept is apparent for the underlying /k/, by speaker DM.

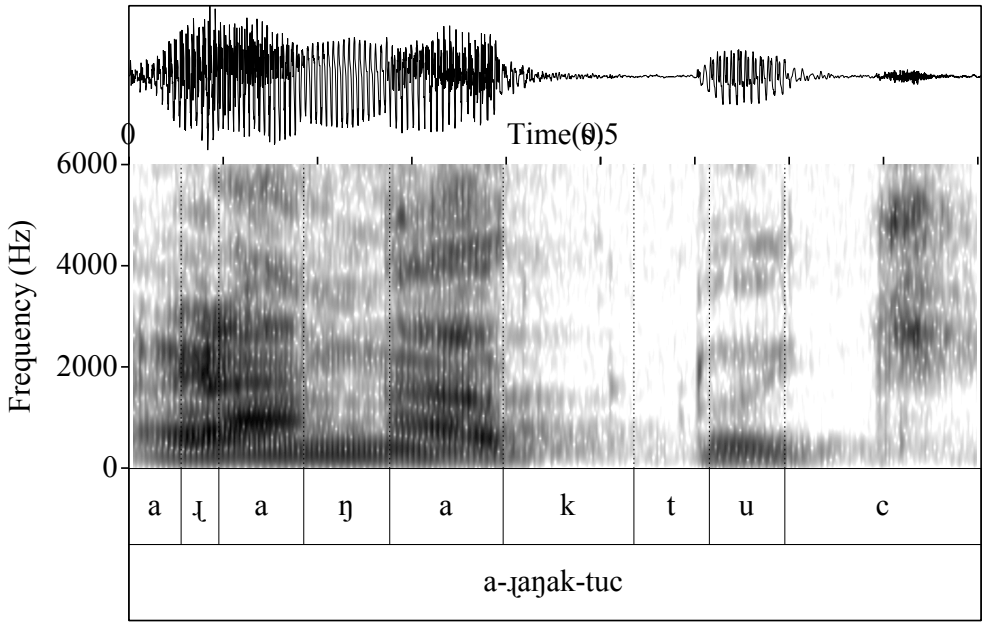

Figure 5.4: Spectrogram and waveform for *a-ɻaŋa(k)-tuc* 'NEUT.OBL-tree-LOC', where a [k] percept is apparent, by speaker DM.

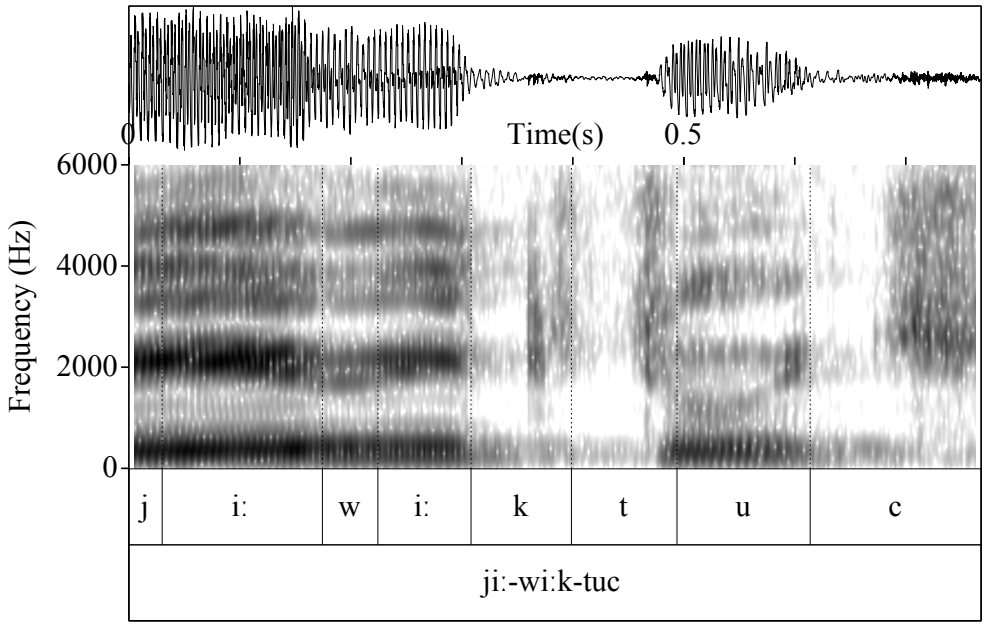

Figure 5.5: Spectrogram and waveform for *ji:-wi:k-tuc* 'Masc.Obl-grey.hair-Loc', by speaker DM.

5.4.2.3 Database querying and analysis

The emuR database was used as the basis for all of the analysis. It has a hierarchical annotation structure, where levels of the hierarchy generally correspond to linguistically-meaningful levels (Winkelmann, Jaensch et al. 2019). The levels of interest here are the Word level, which dominates a Syllable level, which dominates a Phonemic level. Only the Phonemic level contains annotations bearing time information, and all higher levels derive their time information from the Phonemic segments they dominate.

Duration and formant measures were extracted from the database using `emuR`'s Emu Querying Language (EQL2), which allows the user to extract data from segments of interest by reference both to their properties (their features, for example) and their hierarchical or linear context. This returns a dataframe — a two-dimensional array with observations in rows and variables in columns — which then serves as the basis for further quantitative analysis and visualisation in R (R Core Team 2024).

For the auditory impressions of deletion, the relevant observation here is whether the segment is explicitly marked (as an attribute of the database's Phonemic tier) as sounding deleted or not. For analysing duration, the start and end times (in milliseconds) of the /C/ or /CC/ sequence

are the observations of interest. Formant data was calculated by passing a dataframe with the segments of interest to a further querying procedure using the `wrassp` signal processing library (Bombien et al. 2021), which returns a new, larger dataframe, containing measures for F1-F4 at multiple time-points throughout the duration of the segment. In the results reported here, all of the formant tracking has been calculated 'on-the-fly', which means that no manual formant correction has been applied.

Unlike sonorants and continuants, the closure period of stops provides no spectral information as to place of articulation, since it is effectively silent. Instead, stop place of articulation can be investigated by way of such acoustic parameters as closure duration, burst duration and spectral tilt, as well as formant transitions into and out of any adjacent vowel. Since we are dealing with coda stops here, which are typically unreleased, formant transitions are the only correlates that can be observed.

Formant transitions can be quantified using F2 locus equations (Butcher & Tabain 2004; Graetzer et al. 2015; Krull 1989; Lindblom & Sussman 2012; Sussman et al. 1991). This method characterises the acoustic cues of a given consonant according to its effect on adjacent vowels. This is achieved by plotting the value of F2 at the steady state (midpoint) of the vowel against the value of F2 at the point of transition, for each vowel token. In examining C-to-V transitions, this point of transition will be at the onset of the vowel; for V-to-C transitions, it will be vowel offset. We can then fit a linear model to these plotted points for each consonant; the slope and intercept of each then constitutes a characteristic acoustic 'locus' for each place of articulation. F2 is chosen as it is well established as a robust acoustic correlate of oral constriction location (Delattre et al. 1955; Krull 1989; Modarresi et al. 2004). Since I am interested here in coda consonants, I only examine F2 in V-to-C transitions.

5.4.3 Results and discussion

5.4.3.1 Auditory impressions of deletion

As mentioned above, I coded each token of /k/ (or /K/) in a deleting environment with my impression of whether it was audible or not (as a binary variable). The frequency of deleting and non-deleting tokens, for native and non-native nouns, is shown in the bar chart in Figure 5.6. Deletion of /k/ and /K/ in verbal words is similarly plotted in Figure 5.7.

The first, rather striking, result from the noun data shown in Figure 5.6 is that /k/ is deleted only about 25% of the time in native roots, while not being deleted at all in loanwords. Such results obviously disconfirm the account in Heath (1984), according to which deletion is a categorical rule, and should therefore apply 100% of the time, native or otherwise. But the result is also surprising in the context of my own preliminary observations discussed in the previous

section of this chapter. While I expected to find some variability of /k/-Deletion among the native roots but not loaned ones — which is compatible with the findings here — I nonetheless still expected the tokens showing deletion to be the more common variant within the native category.

Recall that in the second production experiment, participants were asked to repeat verbs that contained *(i)* a derivational prefix ending in archiphoneme /K/ (always the Appl prefix /aK-/ here), *(ii)* a suppletive incorporated noun ending in archiphoneme /K/, or *(iii)* a regular incorporated noun ending in true /k/ — that is, an incorporated noun that also appears in ordinary nouns with the same meaning and form. The incorporated noun roots in this last category are a subset of the native roots that were used as stimuli in the first experiment (only some nouns are available for incorporation, and non-native roots cannot be incorporated).

Around one fifth of the /k/-final roots showed retention of /k/ in verbal words, as shown in Figure 5.7. This is almost precisely the opposite tendency from the native /k/-final roots in nouns, where about one fifth showed deletion. For its part, the archiphoneme showed uniform deletion, precisely the opposite from the loaned /k/-final roots in nouns, to which deletion never applied.

So far, the results reported here lend support to the tripartite division of velar stops established at the outset of this chapter, between archiphoneme /K/, native root /k/, and loan root /k/. The first is indeed deleted in all cases, and the last is never deleted, while the remaining category shows variability. This picture cannot be reconciled with Heath's (1984) analysis, which treats deletion as inescapable under all conditions.

However, the variability in the deletion of true /k/ is clearly very different between nouns and verbs. Not only are the rates of retention very low in verbs, I must also note that the verb tokens that I coded as having a retained /k/ showed only a very weak percept compared to retained /k/ amoung nouns. For instance, I did not find any verb tokens with any auditorily or visually identifiable release burst, whereas these were quite common among the nouns. Therefore, whatever acoustic correlates are contributing to a coda [k] percept in verbs must be at a finer level of granularity than in the nouns.

What we can see from the foregoing is that non-native roots appear to always show retention of /k/ on the surface, while morphs with archiphoneme /K/ never do. On theoretical grounds, we have reason to regard both categories as rather different in kind from ordinary native /k/-final noun roots. Today, all Wubuy speakers spend much of their daily lives speaking Kriol and English, the sources of these loans. Their use of these roots in Wubuy is therefore subject to the phonological conventions of the source languages. On the other hand, and at the risk of slipping into tautological explanation, morphs ending in archiphoneme /K/ are unusual precisely because of their distributional restrictions. They never appear in a context favouring retention,

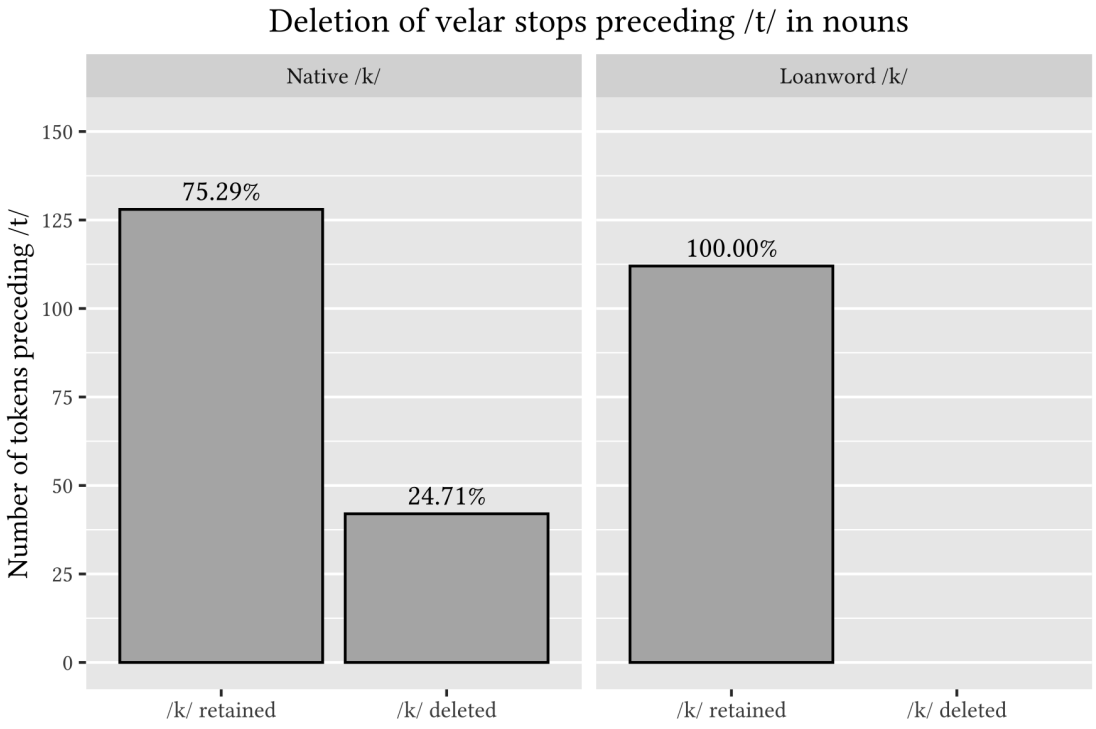

Figure 5.6: Rates of deletion for /k/ in native roots and loanwords. Percentages reflect within-category ratios: all tokens of loan roots show retention of /k/.

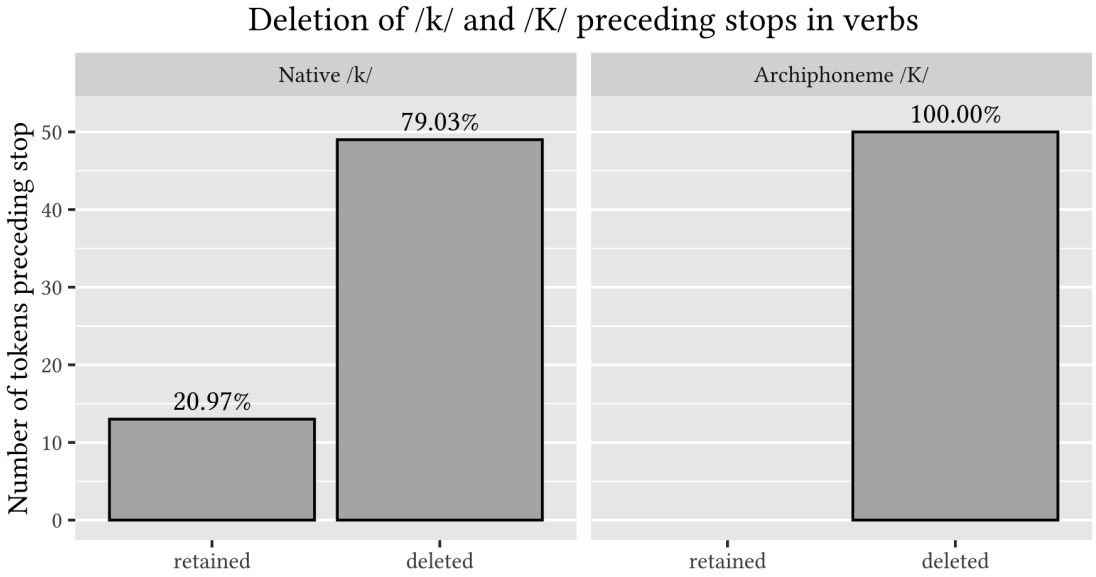

Figure 5.7: Rates of deletion for /k/ and /K/ preceding stops (other than /k/) in verbs. Percentages reflect within-category ratios.

i.e. word-finally, and it seems reasonable to exclude a morph-final stop from the representation of these forms entirely.

The variance that needs to be explained, then, is that found in the native roots. Therefore, in the statistical analysis to follow, I consider only the native roots. In any case, it is difficult to estimate confidence intervals and p values for binomial distributions where the observed proportion in the dependent variable is close to 1 or 0, so including the archiphonemes and loan roots in the model — whose proportions of deletion are 1 and 0, respectively — is unlikely to be informative.

We can fit a generalised linear mixed effects model to the native /k/-final root data (in noun and verb word contexts) using the `glmmTMB` package (Brooks et al. 2017) in R (R Core Team 2024), where dorsal stop deletion is the dependent variable of interest (binary), the fixed effects are word class (noun or verb) and word length (measured in number of syllables), and the random effects are speaker and item (i.e. noun root). Random effects are necessary in data of this kind since the observations are not independent — there are multiple observations per speaker and per item — therefore violating a core assumption of the standard linear model. Moreover, I specify both random intercepts and random slopes for speaker and item, following the advice of Barr et al (2013), who caution against the inclusion of intercepts-only random effects where there is no strong theoretical justification for doing so — as they show, intercepts-only random effect structures have a much higher propensity for Type 1 errors, while offering negligible improvement in explanatory power. By including random slopes, we allow for the possibility that individual speakers have grammars with different sensitivities to word class and length, not merely different baseline levels (which are captured by the random intercepts). As the results in Table 5.6 show, the relationship between word class and deletion was the only statistically significant one.

Parameter	Coefficient	SE	CI	CI_low	CI_high	z	df_error	p	Effects	Group	Component
(Intercept)	0.018	0.0215	0.95	0.0017	0.188	-3.355	Inf	0.0008	fixed		conditional
word_classv	27.673	35.0840	0.95	2.3061	332.063	2.619	Inf	0.0088	fixed		conditional
n_sylls	1.423	0.3233	0.95	0.9118	2.221	1.554	Inf	0.1203	fixed		conditional
SD (Intercept)	1.867	NA	0.95	NA	NA	NA	NA	NA	random	root	conditional
SD (Intercept)	0.413	NA	0.95	NA	NA	NA	NA	NA	random	speaker	conditional
SD (word_classv)	1.675	NA	0.95	NA	NA	NA	NA	NA	random	root	conditional
SD (word_classv)	1.187	NA	0.95	NA	NA	NA	NA	NA	random	speaker	conditional
Cor (Intercept~word_classv)	-1.000	NA	0.95	NA	NA	NA	NA	NA	random	root	conditional
Cor (Intercept~word_classv)	1.000	NA	0.95	NA	NA	NA	NA	NA	random	speaker	conditional

Table 5.6: Summary of results of fitting a generalised linear mixed effects model with crossed random effects in *glmmTMB*, using the formula `deletion ~ word_class + n_sylls + (1+word_class|root) + (1+word_class|speaker)`

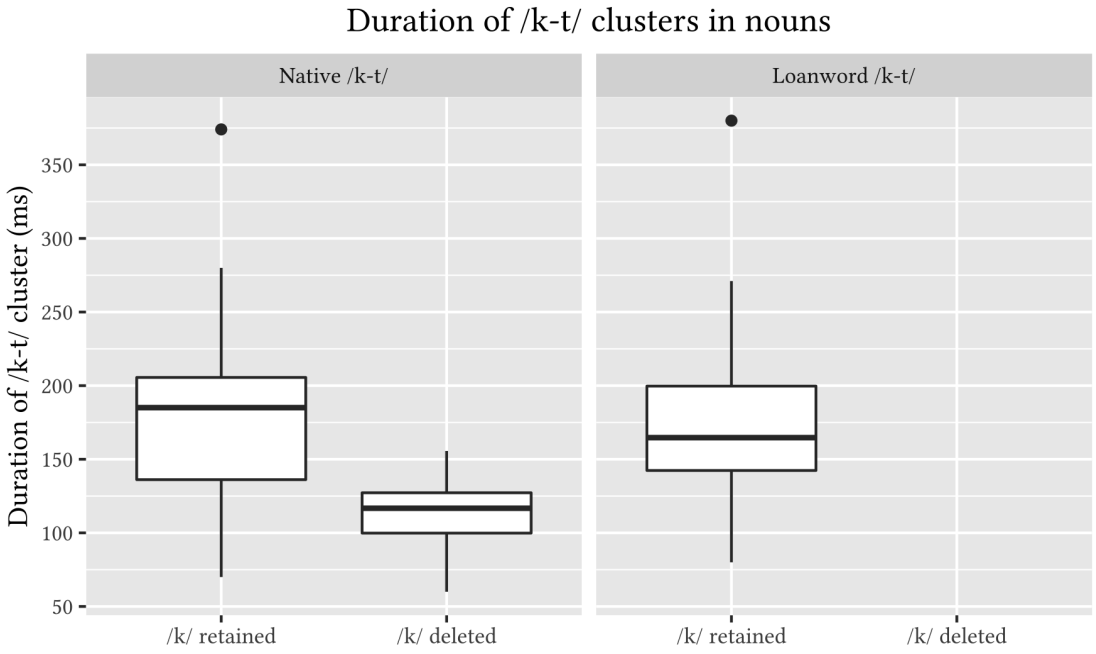

Figure 5.8: Duration of (underlyingly) /k-t/ clusters in native roots and loanwords, by whether /k/ is deleted or retained.

5.4.3.2 Acoustic analysis: duration

I turn now to acoustic analysis, here measuring duration of intervocalic periods that are underlyingly /kC/. Durations for the nouns are summarised in the boxplot in Figure 5.8, grouped by origin (native or loaned) and deletion (deleted or retained). As we would expect, deletion noticeably reduces the duration of the closure. Based on visual inspection of the plot, there is practically no difference in duration between the native and loan roots where /k/ is retained on the surface. The native roots where /k/ is deleted, however, do appear noticeably briefer than the rest. An Analysis of Variance with root type (native vs loan) and deletion as the independent variables and duration as the dependent variable shows that the effect of deletion on duration is statistically significant ($F = 71.42, p < .001$), but the effect of root type was not ($F = .789, p = .375$).

In order to situate these results amongst stop durations generally, Figure 5.9 plots these same measures of underlying /k-t/ sequences, with deleted and retained /k/, against durations for all intervocalic stops and stop sequences. (Deletion is indicated with angled brackets around the deleted segment.) Note that most of these consonants and consonant sequences are in intramorphemic contexts, whereas /kC/ or /KC/ necessarily are not, because /kC/ sequences do not appear within morphs in the Wubuy lexicon.

The first important point to observe here is that invervocalic /k/ and /t/ each have medians around 80ms, which is less than half the duration of /k-t/ where /k/ is retained, with a median of 185ms. This suggests that two consonantal timing slots are involved where deletion has not occurred — that is, that the /k-t/ interval consists of a /k/ period followed by a /t/ period.

However, the median duration of /k-t/ where deletion has applied is around 114ms, placing this category between the two poles represented by non-deleting /k-t/ on the one hand, and intervocalic /k/ or /t/ on the other. This is quite surprising: one would expect that if /k/ were deleted altogether, it would not contribute any extra duration to the period of closure for the following segment. A range of factors may bear on the interpretation of such a finding. It is possible that the deleted /k/ still leaves a timing slot that must be filled by the following segment — yielding, in this case, something like a geminate or long [t:]. Or, the longer duration of /t/ at the beginning of the locative suffix may be an effect of being at the edge of a morphological constituent: most of the intervocalic /k/ and /t/ durations are taken from tokens in which they are instead *within* a morphological constituent.

Another finding of note here is that intervocalic singleton /k/ in loanwords is especially long, with a median of 154ms — this can be seen in the cluster of data points hovering around that median in the plot. These measurements come exclusively from tokens of *a-keik-tuc* 'on the cake', where the first /k/ sounded notably more like an aspirated English [k^{h}], rather than the typical morph-initial Wubuy /k/ which typically has much shorter VOT.

Finally, we also have some tokens of the archiphoneme /K/ preceding /k/, and it has more or less the same durational distribution as ordinary intervocalic (native) /k/. This result should be interpreted with some care, however. Firstly, because all data points come from tokens of *pa(K)-kala:ti-ɲuŋ* 'blind person' (lit. eye-bad-HUM.SG). And secondly, because surface [k] could arguably be analysed as reflecting either the /K/ of *paK* or the /k/ of the *kala:ti* contextual allomorph (which is underlyingly /ala:ti/ in the traditional analysis). Note that the only orthographic representations I've seen of this stem are spelled <*ba-galadi*>.

Durations for intervocalic stops and stop clusters in verbal words are plotted in Figure 5.10. Comparing nouns and verbs, the relative weakness of the velar percept in the latter is reflected in cluster durations, which are shown at different places of articulation in Figure 5.10. Where /k-t/ with a retained /k/ in nouns tended to be nearly double the duration of the same cluster with deleted /k/, the retention of /k/ in verbs contributes very little duration to the cluster.

In Figure 5.10, the deleted /k-p/ and deleted /K-p/ durations show a very similar distribution to the duration of intervocalic /p/, while the median duration of retained /k-p/ is about 30ms higher, which is however a very small difference. In the /k-t/ sequences in nouns (Figure 5.9), the tokens with retention of /k/ had a median duration about 70ms higher than those without.

Although we have just one token with retained /k/ preceding a retroflex, the retroflexes

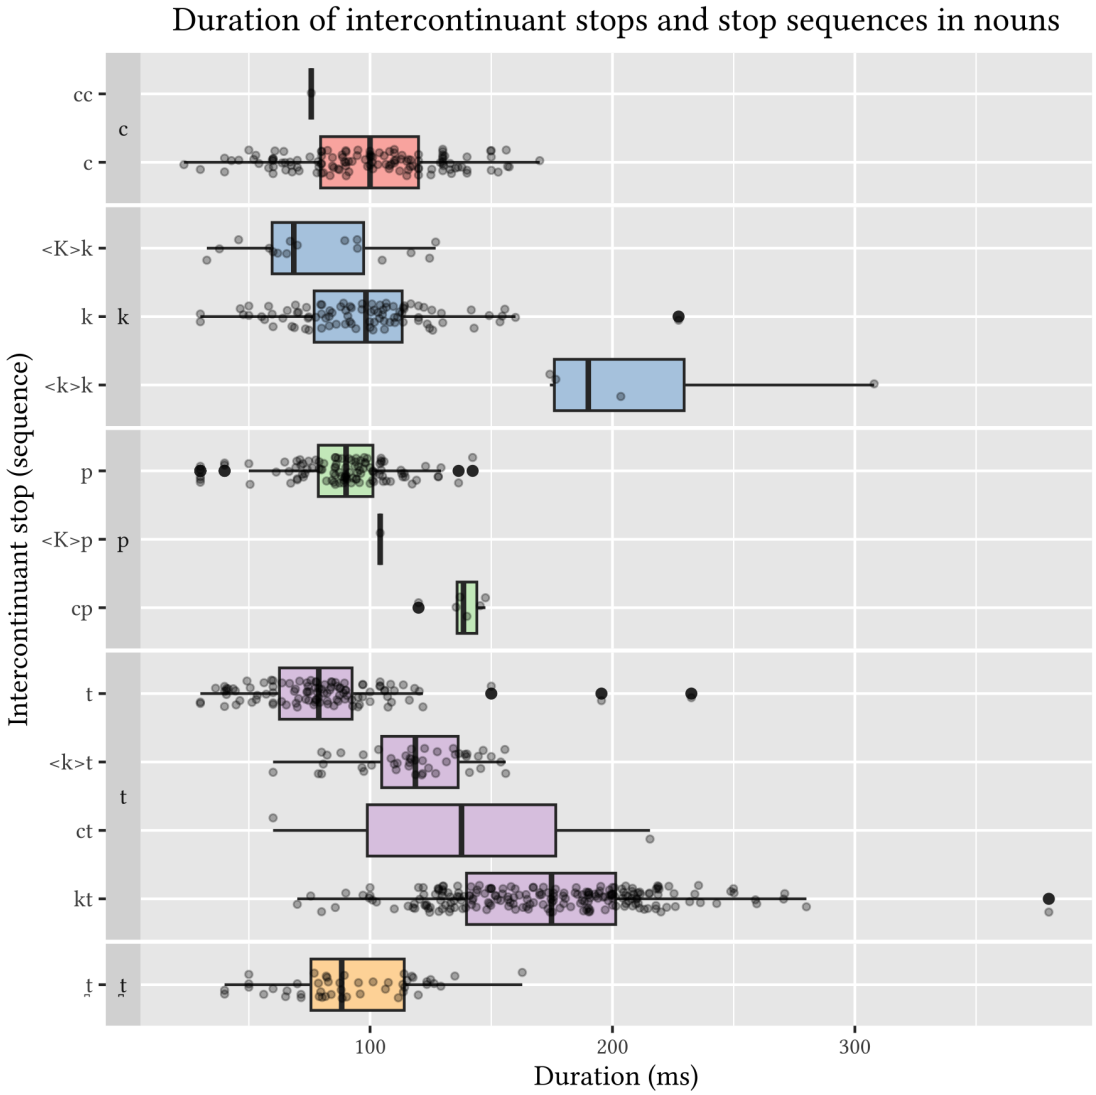

Figure 5.9: Durations of intercontinuant stops and stop sequences.

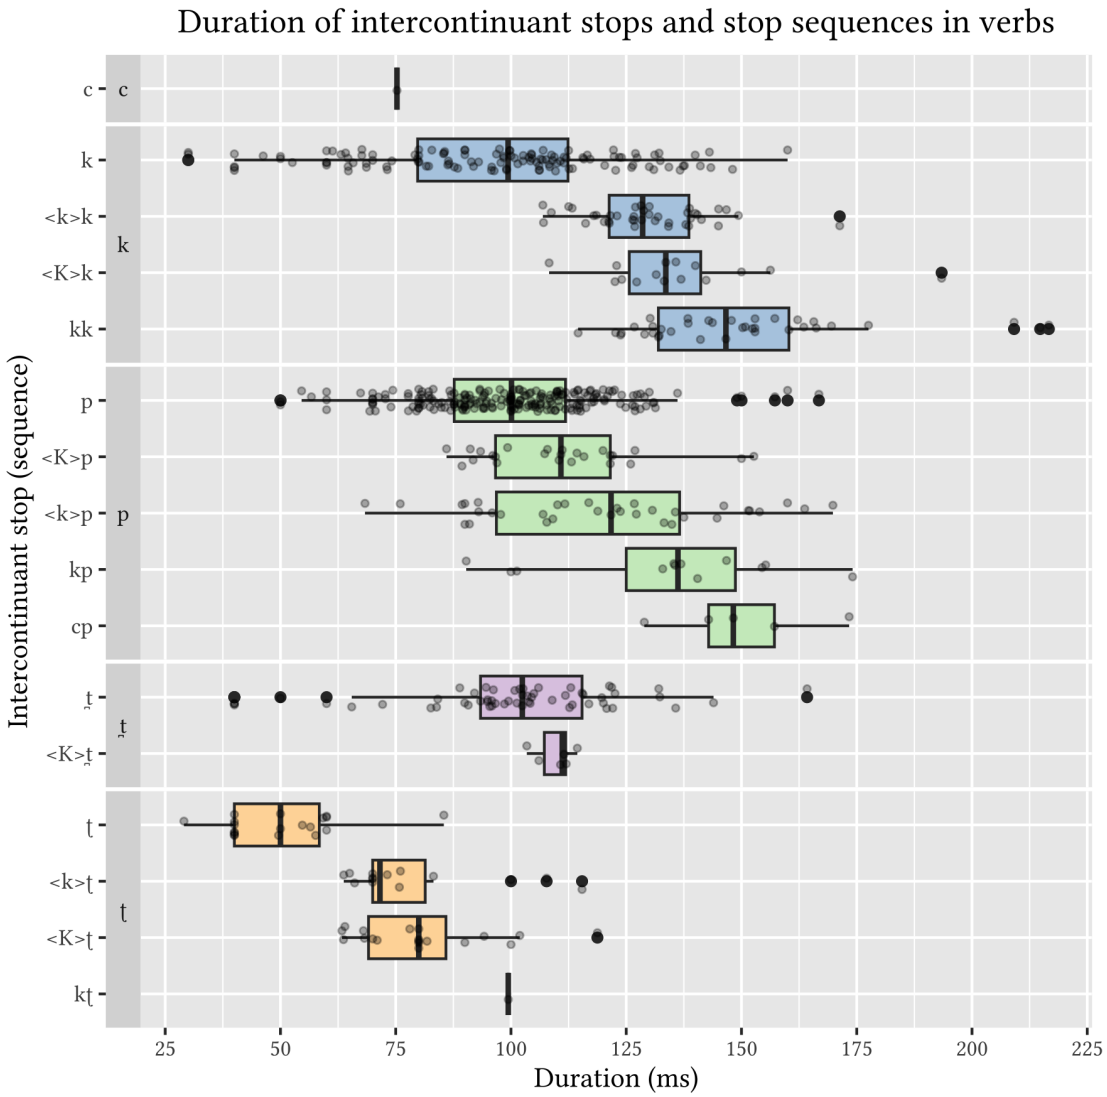

Figure 5.10: Durations for stops and stop sequences between continuants in verbs, measured in milliseconds.

following deleted velars have a surprising distribution. Upon first glance, even a deleted velar appears to contribute considerable duration, since both /k-ţ/ and /K-ţ/ are longer than single /ţ/, even though deletion should render them all identical. More surprisingly still, the deleted archiphoneme appears to contribute more length than the deleted true /k/.

However, more detailed examination shows that there is an apparent confound that we may attribute to the preceding vowel. As Figure 5.11 shows, the tokens of single /ţ/ all follow /i/, and the tokens with clusters follow either the long or short low vowel (/a/ or /a:/). But the vowel type may also be a proxy for morphological position and/or pitch accent location. All the tokens of single /ţ/ — equivalently, all tokens of /ţ/ in the environment of /i_/ — occur within root morphs, while the tokens of /ţ/ in clusters — equivalently, clusters following the low vowels — occur at the left edge of a root. Moreover, all of the tokens of /K/ after the long /a:/ vowel come from tokens of the Appl prefix, which lengthens the final vowel of the pronominal prefix before it. This long vowel comes at the left edge of the verb stem, and is virtually always associated with a high pitch accent. While it is undeniable that the tokens of single /ţ/ in these data are shorter than the clusters (even though they are not clusters on the surface), and that the longest clusters are the /K-ţ/ ones, we cannot reasonably draw any conclusions from this due to the many confounds just enumerated.

5.4.3.3 Acoustic analysis: V-to-C formant transitions

If deletion of /k/ is absolute, we should expect an absence of any acoustic correlates of /k/ whenever it precedes a consonant underlyingly; the same should apply to /K/. We have seen that the duration of a /kt/ sequence in which /k/ is deleted is similar to the duration of intervocalic /t/, as we would expect of deletion. We may also expect that, in a /kt/ sequence in which /k/ has been deleted, the formant transition from the preceding vowel into the (underlying) cluster should be more similar to the transition to /t/ than to /k/. Alternatively, it may be that some trace of a /k/ gesture remains even where its associated acoustic percept is difficult to discern. If so, we may find some trace of this in the formant transitions from the preceding vowel. In order to determine which of these scenarios is true, we first need to get a sense of the formant characteristics of the vowels themselves, and of the transitions to /k/ and /t/ where deletion is not implicated.

Figure 5.12 shows smoothed trajectories of F1-F3 for the vowels /a, a:, i, u/ (the only ones with sufficient tokens to fit the Loess-smoothing function), from the midpoint of the vowel to its end. The top row of plots shows vowels preceding intervocalic /k/, while the bottom row shows those preceding intervocalic /t/. From visual inspection of these plots, F1 and F3 do not appear to provide good transitional cues to place of articulation, as they tend to be both fairly unchanging throughout the vowel, and basically the same regardless of following segment. F2,

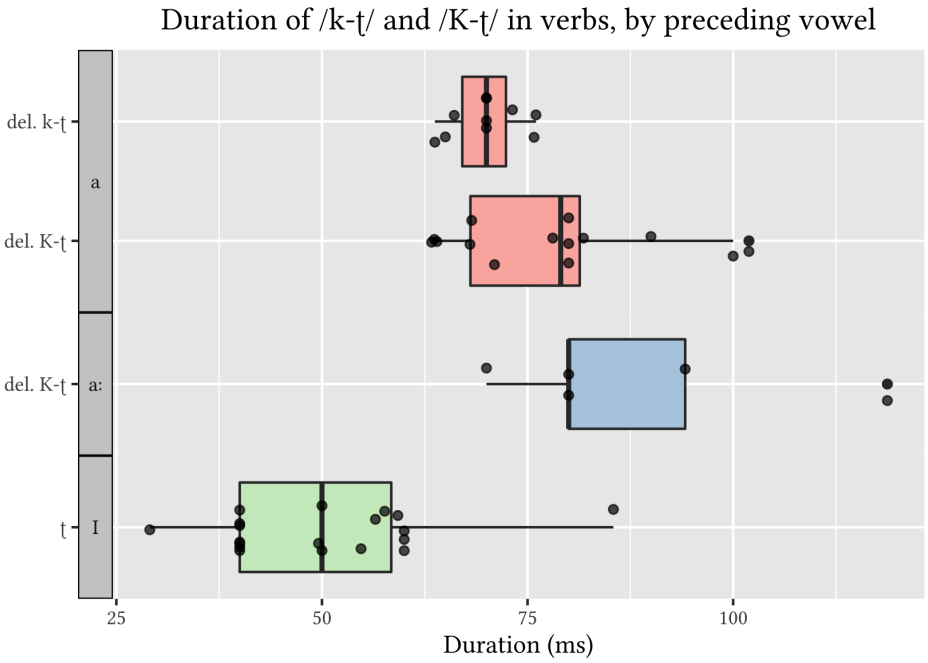

Figure 5.11: Durations for k~K-ţ, following just the low vowels /a/ or /a:/. Tokens of archiphoneme /K/ following long vowels are all from the Applicative prefix, analysed as /aK-/ in Heath (1984).

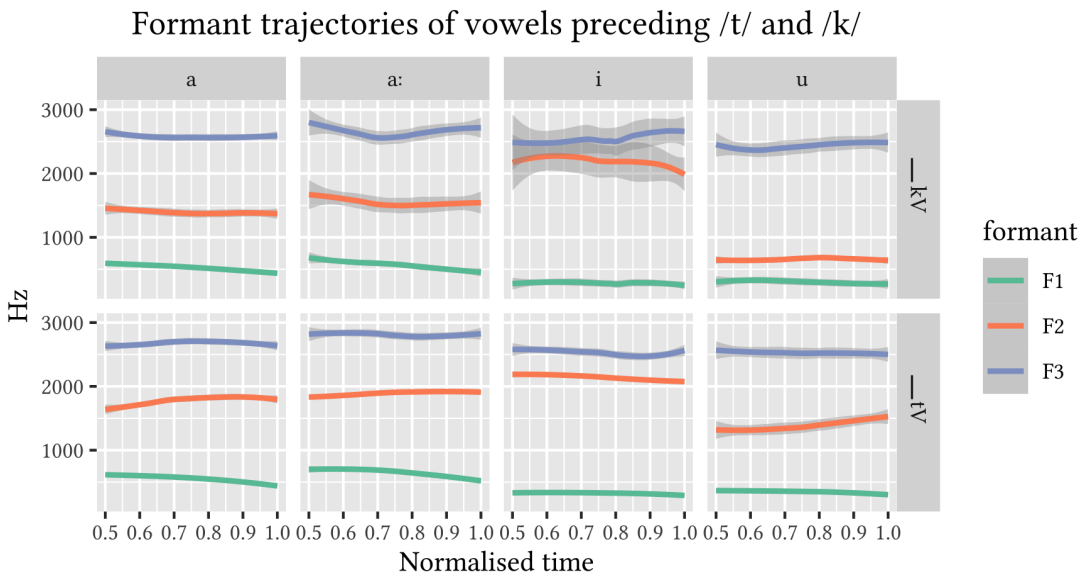

Figure 5.12: Averaged, time-normalised trajectories of F1, F2 and F3 in vowels in the contexts /_tV/ (bottom row) and /_kV/ (top row) in nouns, from the midpoint to the end of the vowel. That is, the context of interest is the following consonant. Diphthongs (in loanwords) have not been included.

however, is more distinct in both respects. All of the vowels preceding /t/ show a slight rise in F2 from steady state to offset, consistent with coarticulatory fronting. The high front vowel, /i/, does not show this — unsurprisingly, given it is already fronted. We might suppose then that rising F2 is a perceptual cue to a following /t/. This is consistent with work on place of articulation cues in Australian languages (Busby 1979) and elsewhere (Sussman et al. 1991), which finds F2 to be the more robust transitional cue distinguishing stop place of articulation.

On the other hand, the vowels before /k/ do not show any particularly tell-tale characteristic. Instead, each of the formants remains relatively stable throughout the vowel. Although velar constrictions in English are associated with a narrowing of the distance between F2 and F3 towards offset ('pinching'), the trajectories here appear to show that the velar constriction is made at whatever position the tongue is in for the steady state of the preceding vowel. The acoustics of a velar gesture may therefore be distinguished here from an alveolar gesture simply by an *absence* of rising F2 towards vowel offset.

Figure 5.4.3.3 plots F2 measurements at 80% of the vowel. In all vowels except /i/, F2 is higher preceding /t/ than preceding /k/, suggesting greater fronting. For /i/, this is reversed: F2 is lower before /t/ than /k/. However, in absolute terms, F2 preceding /t/ is even higher in /i/ (with the median at around 2200 Hz) than in the other vowels. It is unclear why the highest F2 of all

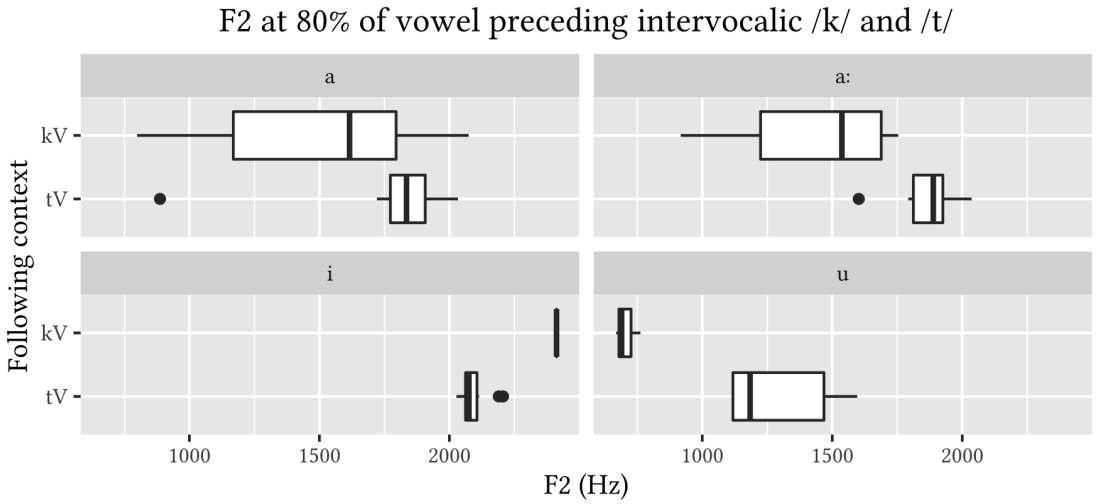

Figure 5.13: F2 frequency in Hertz at 80% of vowel duration, preceding intervocalic /t/ and /k/.

is in /i/ preceding /k/.

Of course, what is more crucial in cueing coda stop place of articulation is not the absolute F2 value at any given point but the *change* in F2 from the vowel's steady state to its offset. This can be examined by the method of locus equations (Sussman et al. 1991), where the relationship between F2 at these two positions is modelled as a linear regression, and each place of articulation can — irrespective of the contextual vowel — be described by the slope and intercept of the line. Locus equations for /k/ and /t/ are presented in Figure 5.14. That the slope for /k/ is so close to 1 (at $0.927x$), and the intercept so close to origin (at 78.5 Hz), suggests that there is little change in F2 from midpoint to offset. This aligns with what we observed in the smoothed formant traces in Figure 5.12, where /k/ did not appear to condition any formant excursion away from the steady state. The regression line for /t/ shows more of an anticipatory coarticulation effect on F2, with a much higher intercept (at 780 Hz) and a slope that is much further from 1 ($0.6x$). Linear regression was also used to compare the slopes, with the following stop as the independent variable; the difference in slopes between each stop category was statistically significant (t ratio $(753) = 6.8, p < 0.001$).

Next we need to consider the behaviour of F2 in vowels preceding stop clusters, again considering the null hypothesis that the true difference in means of F2 (measured in Hz) before retained /k/ and deleted /k/ is zero. Again, smoothed, time-normalised formant trajectories are shown in Figure 5.15, this time showing vowels preceding underlying /kt/, grouped according to whether the /k/ is deleted or not.

The F2 trajectories in Figure 5.15, for vowels preceding underlying /kt/, appear to show some

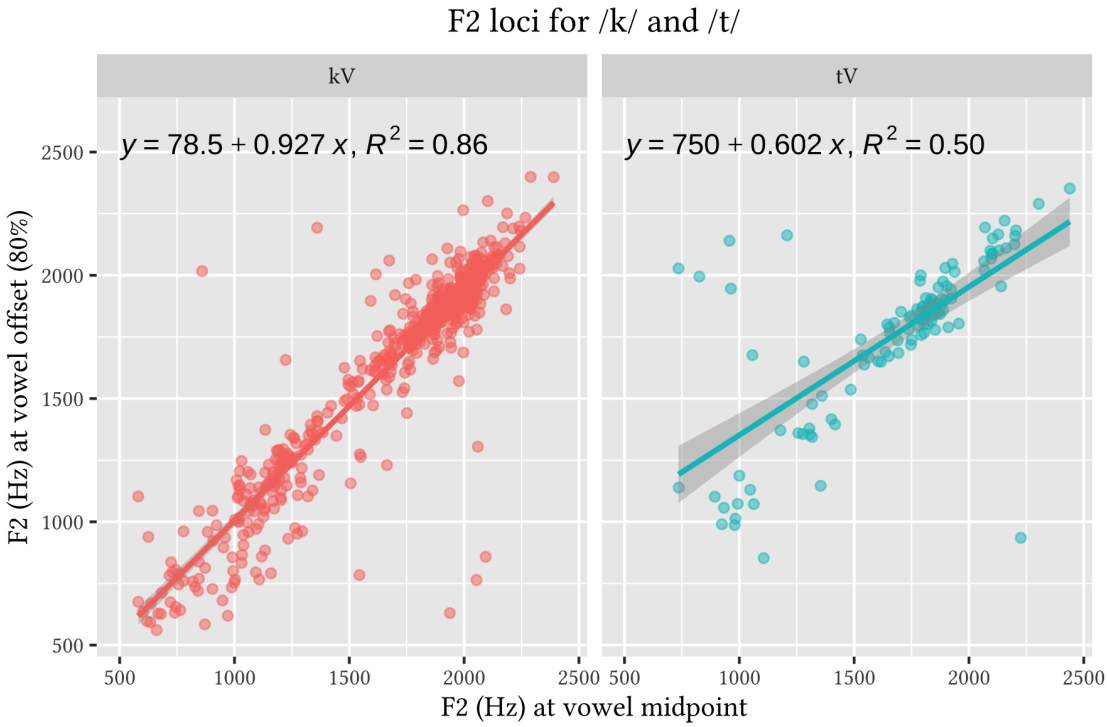

Figure 5.14: Locus equations (Sussman et al. 1991) for /t/ and /k/, measured by comparing F2 at midpoint and 80% of the preceding vowel. Tokens with an F2 value of 0 Hz were excluded.

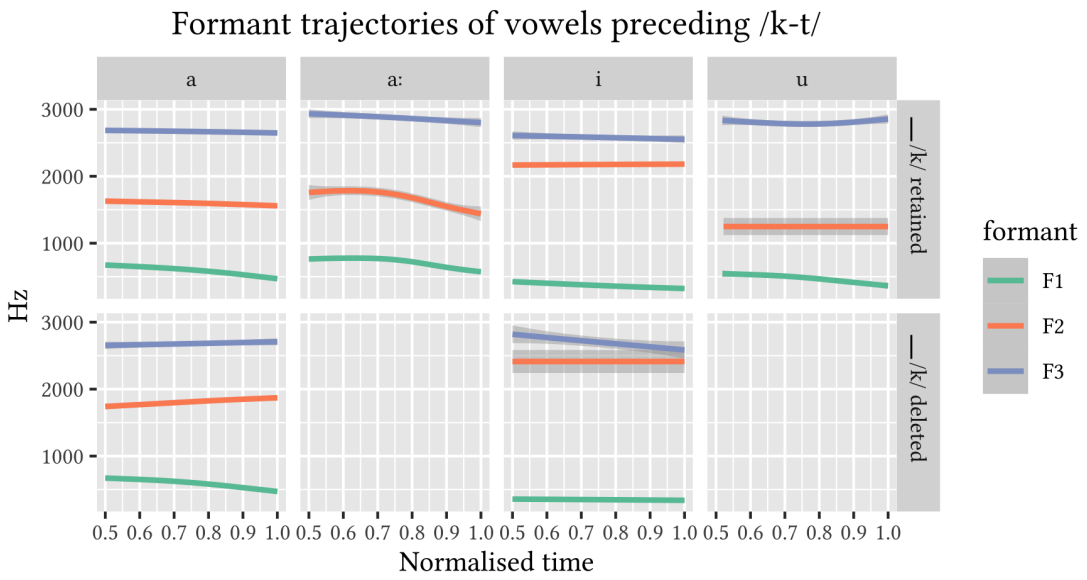

Figure 5.15: Trajectories of F1, F2 and F3 in vowels preceding /k-t/ in nouns, according to whether the /k/ is deleted (bottom row) or not (top row).

similarities to the data for the vowels preceding singleton consonants. Note that we have no cases in the data where /a:/ or /u/ precedes a deleted /k/. The cases of deletion following /i/ are all in tokens of /a-l̥awa:lik-tuc/ 'in the soak' and /a:nik-tuc/ 'in the ashes'. Preceding /kt/ in the deleting tokens, the /a/ vowel, in particular, shows an upward trajectory in F2 in similar to the transition into intervocalic /t/ that we saw earlier, in Figure 5.12.

Again, we can plot locus equations for vowels preceding /k-t/ clusters, as in Figure 5.16. However, in this instance, although the linear regressions for each condition (/k/ retained or deleted) are obviously different, the model fit is extremely poor in the case of the tokens where /k/ is deleted. We can get a somewhat better model fit by examining only the /a/ tokens, as in Figure 5.17.

The locus equation plot in Figure 5.17 certainly shows some evidence of anticipatory coarticulation towards the /t/ target when the /k/ is deleted. At the very least, it gives some credibility to my auditory impressions. However, given the practical absence of coarticulatory resistance shown by the velars in general across the formant transition data presented in this section, it does not seem that it would be possible to detect any residual trace of a dorsal constriction acoustically, as suspected at the outset. If there is an overlapping dorsal gesture that is perceptually masked in the tokens that were coded as showing deletion, this could only be investigated by instrumental methods like electromagnetic articulometry or ultrasound.

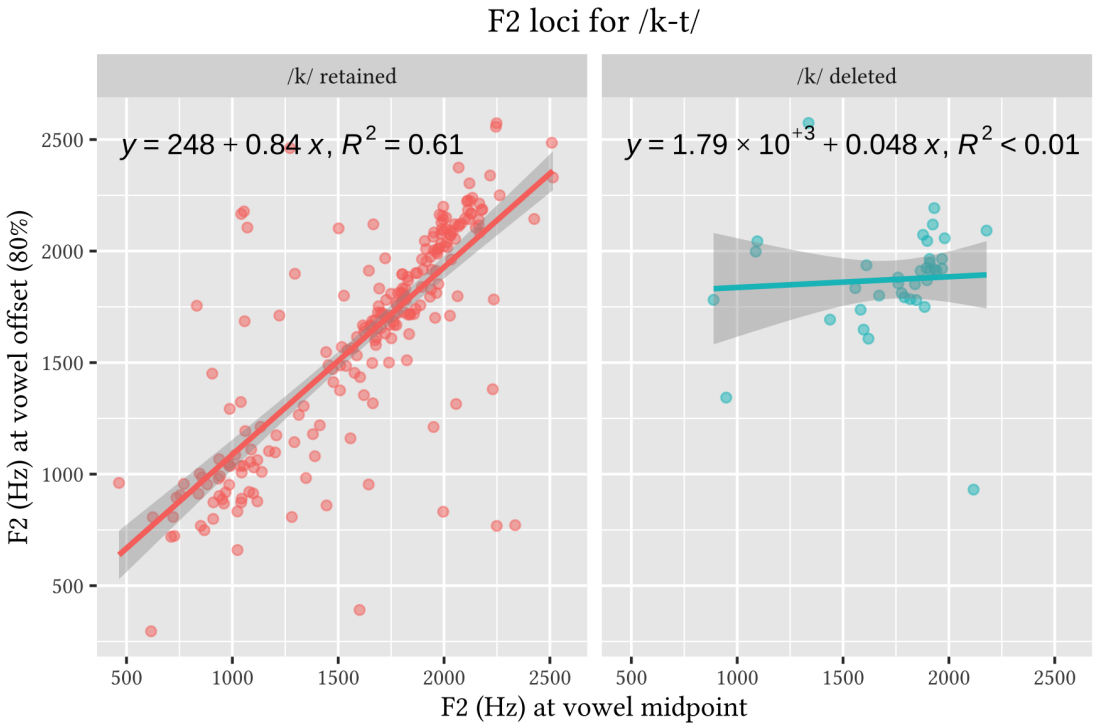

Figure 5.16: Locus equations for VC transition into /k-t/ cluster, where deletion has and has not applied, measured by comparing F2 at midpoint and 80% of the vowel. Tokens with an F2 value of 0 Hz were excluded.

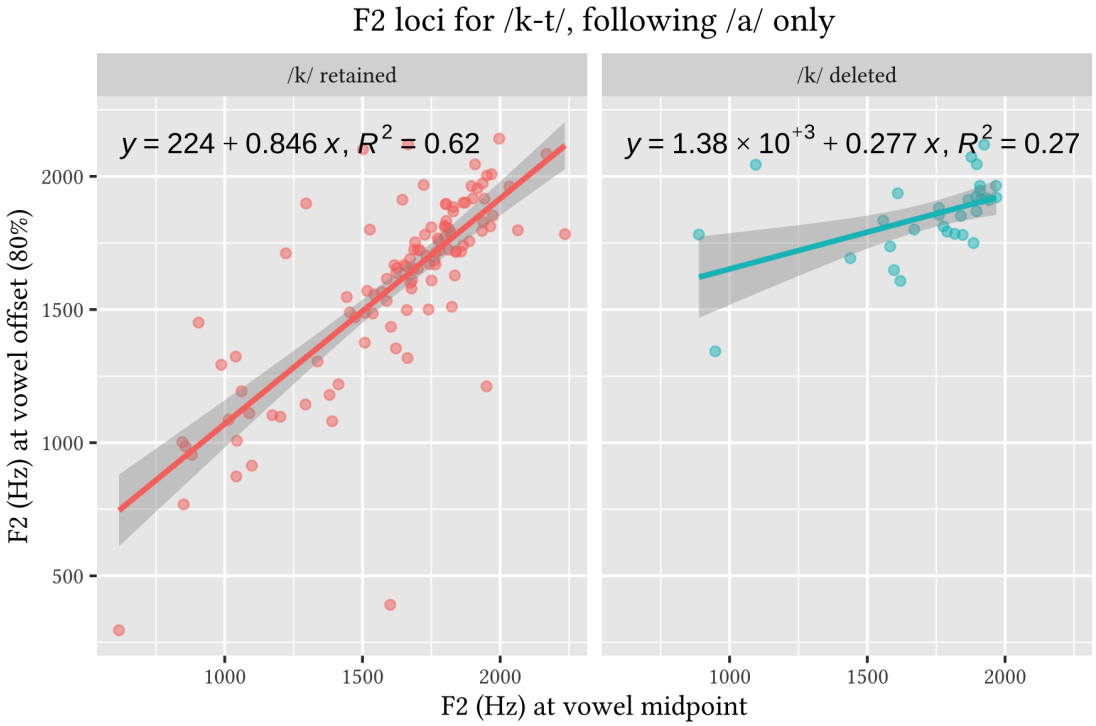

Figure 5.17: The /a/ subset of the data in Fig. 5.16.

5.5 Discussion

5.5.1 Explaining the results

The results reported here are the first to show that, contra previous description, preconsonantal velar stops in Wubuy are not deleted uniformly, but show variation depending on what kind of morph they appear in, and what kind of word that morph appears in. Loanwords show no deletion, morphs ending in archiphoneme /K/ always undergo deletion, while native noun roots undergo deletion some of the time. The probability of deletion in native roots depends in large part on whether it appears as the head of a case-suffixed noun (where deletion is considerably less frequent than non-deletion) or incorporated into a verb (where deletion is overwhelmingly the norm).

The phonetic effects of deletion are registered very clearly in duration, where underlying /kC/ sequences are on the whole shorter when the /k/ sounds deleted, and in F2 transitions, which show greater anticipatory fronting preceding underlying /kt/ sequences when they undergo deletion than when they do not. However, even where an underlying /k/ was deleted, it still appeared to contribute some additional duration, since even deleting /kt/ sequences are longer than singleton /k/ or /t/. In the formant transition data, on the other hand, no difference could be detected between deleted /<k>t/ and singleton /t/, meaning that no especially subtle cues for a dorsal constriction could be detected below the threshold of perception. This is not to say that there is no dorsal constriction, however. The VC transitions for intervocalic /k/ show almost no departure from the steady state of the vowel anyway, so it is unlikely that even an undeleted /k/ of such kind could exert any pressure on the formant structure of the upcoming /t/, which shows much more coarticulatory aggressiveness generally. An ultrasound study, on the other hand, may be more revealing of the gestural coordination of /kC/ sequences where both constrictions are articulated, but with temporal overlap.

How can we explain the distribution of deleting and non-deleting variants that we find? I would argue that what we are witnessing, far from a monolithic phonological rule, is better characterised as a sound change in progress, diffusing at different rates in different parts of the lexicon (where the lexicon is understood here to accommodate whole-word storage in the same way as the previous two chapters). The crucial factor that distinguishes these different types of velar-final morph, and which can be seen as the driver of their different propensity for deletion, is the relative frequency with which they appear at the right word edge, as opposed to word-nonfinally. The pressures that drive deletion in the first place are predicted to have a greater effect on production, and consequently a greater effect on lexical representation, where this frequency is higher.

5.5.2 Frequency, lenition, and the automatisation of articulatory routines

The claim that frequency tends to correlate with leniting sound change, of which the deletion process under examination here is a type, is not especially controversial. The debate instead centres around what the precise mechanisms responsible for producing leniting variants are, how new variants become entrenched in memory, and whether change affects the whole lexicon simultaneously or spreads from one item to another over time (Hay, Pierrehumbert et al. 2015). On this last point, it is difficult to imagine how the facts reported here can be interpreted as neogrammarian sound change, since deletion obviously applies at different rates in different parts of the lexicon. The problem lies in explaining why the pressures driving deletion have advanced so much further in some parts of the lexicon than in others.

It is by now well understood that frequency of use correlates with phonetic lenition or reduction, including consonant deletion (Baker & Bradlow 2009; Bell et al. 2009; Bybee 2001, 2002b; Pluymaekers et al. 2005b; van Son, Bolotova et al. 2004). One way of understanding the relation between frequency of use and phonetic reduction is that increased practice of a sequence through repetition leads to greater automatisation of neuromotor routines (Bybee 2001, 2006; Bybee & Easterday 2019; Davidson 2006; Davis & Redford 2019; Kapatsinski 2018; Tiede et al. 2011; Tomaschek, Tucker et al. 2018). With practice through repetition — in the early stages of acquisition and throughout one's lifetime — the articulation of speech gestures becomes on the whole faster, more precise, less error-prone, more difficult to interrupt once initiated, and less effortful in execution and planning (reviewed in Kapatsinski 2018, ch. 9). For example, an electromagnetic articulometry (EMA) study by Tiede et al (2011) examined the effect of practising phonotactically-legal nonsense words in a group of subjects over several days. With repetition, subjects' articulatory trajectories became faster, travelled a shorter distance, showed lower variability, fewer acceleration peaks, and greater overlap between closure gestures. These effects of practice appear to hold for real words also. Using a word's corpus frequency as an approximate index of practice in articulation, an EMA study by Tomaschek et al (2018) found smoother articulatory trajectories and higher velocities in tight curves for productions of higher-frequency words in German. Insofar as lenition can be treated as a reduction of articulatory effort, it should increase with practice.

The automatisation acquired through practice is, in general, achieved on an item-by-item basis — in other words, repeated experience articulating the word *misled* will make one more adept at articulating *misled*, but it is not expected to deliver an observable benefit to one's ability to articulate *deceived*. Gahl (2008) found significant differences in the acoustic realisation of 223 word pairs in English that are normally assumed to be homophonous, such as *time* and *thyme*, but whose members differ substantially in their frequency. She found a significant effect of word frequency on duration, such that the more frequent member of the pair, e.g. *time*, was very

likely to be shorter. Consistent sub-phonemic acoustic differences have been also found between *like* according to its function, whether used as an adverbial, discourse particle, quotative, and so on (Drager 2011; for related findings see Jurafsky et al. 2002). Some of these fine-grained phonetic differences also hold for units below the level of the word. Morphological function has been shown to correlate with differences in the duration of English /s~z/ (Plag, Homann et al. 2017; Tomaschek, Plag et al. 2019; Zimmermann 2016) as with the 'darkness' of intervocalic /l/ (Lee-Kim et al. 2013). These findings, taken together, provide compelling evidence that phonetic implementation is item-specific, and cannot be predicted from the phonemic string alone.

These studies give support to the idea that articulatory sequences become faster, less variable and more coarticulated with practice. However, what counts as an 'articulatory sequence' is an open question (see e.g. Bradley 2002; Cho 2001; Hall et al. 2016; Kapatsinski 2018). In particular, different accounts have been given of how to treat morphologically complex words: if a word consists of a root and a suffix, does it contain one articulatory sequence or two? One influential answer to this question comes from the work of Hay (2003). In an acoustic phonetic experiment on the production of morphologically complex English words, Hay shows that morphological complexity is gradient, not categorical, and depends on the *relative* frequency of an affixed word with respect to the base it contains. For example, *swiftly* is more frequent than *swift*, whereas *softly* is less frequent than *soft*. In her study, derived words which are more frequent than the bases they contain (such as *swiftly*) display significantly more evidence of coarticulation, meaning greater overlap between articulatory gestures, which often manifests as consonant deletion. Mere frequency of the derived word alone was not a sufficient predictor of her results. She uses these findings to argue that derived words that are more frequent than the bases they contain have a more holistic representation in the mental lexicon, whereas derived words that are less frequent than their bases are parsed into two units in processing, and constructed on the fly in production. In other words, 'items that are used together fuse together' (Bybee 2002a; cf. Mansfield's (2021) notion of 'internal predictability' as a wordhood criterion).

To be sure, showing that the frequency profile of an item correlates with lenition does not in itself demonstrate sound change, since lenition may simply be a stable (synchronic) effect of frequency (Dinkin 2008). But some recent studies using large diachronic corpora show that indeed frequency-based lenition does directly affect sound change (Hay & Foulkes 2016; Hay, Pierrehumbert et al. 2015), as predicted by exemplar-based models. In exemplar-based models, the units of a speaker's phonological knowledge are memorised phonetic tokens ('exemplars') of lexical items that they have experienced (Bybee 2001; Pierrehumbert 2001; Schweitzer 2019; Todd et al. 2019). These memories are constantly updated through further experience, and serve as the basis for speaker's articulatory plans during production, as well as for perceiving new tokens of the same category (Goldinger 1998). In Pierrehumbert's (2001) model, speakers

sample from their store of exemplars to create a production target, and the token they produce is heard by both the speaker and the addressee, and is thereby added to both listeners' stores of exemplars, creating a feedback loop. The importance of frequency in this context stems from the simple notion that, if production plans are updated on the basis of experienced exemplars, any effect of frequency on production will affect frequent lexical items more, because they go more times through the production-perception loop. The effect of lenition may then become entrenched in lexical representations, such that non-leniting variants are no longer produced.

5.5.3 Relative frequency of word-final vs. word-nonfinal velar stops

In order to apply Hay's (2003) notion of base/derived relative frequency to Wubuy, we will need to translate the terms of the ratio into ones more appropriate for a language of such morphological complexity. Instead of treating the base/derived frequency ratio as one between monomorphemic and polymorphemic morphological relatives, I will take it as describing a ratio between the frequency with which a given morph appears at the right word edge ('word-final') versus elsewhere in the word ('word-nonfinal'). In other words, I am ignoring any morphological material that may happen to be attached to the left of the morph in question, and ignore whether it is followed by a suffix or another root (as in compounding).

As noted earlier, we cannot obtain reliable raw frequency estimates for Wubuy morphs, but relative frequencies can be approximated more or less *a priori* according to known distributional restrictions of each type of velar stop. Archiphonemes are trivially more frequent word-medially than word-finally: if they could ever appear at the word edge, we would have no reason to represent them with an archiphoneme. For loans, the ratio is practically inverted. Loans occur more often at the right word edge than with following morphological material in Wubuy, but in addition, all the Wubuy consultants I work with also speak Kriol and English (either one of which may be considered the source of the loans) on a regular basis, and frequently code-switch when speaking amongst themselves. These loans are still used in the source languages, and, particularly in Kriol, only seldom appear with following morphological material. These morphs show no deletion of /k/, and speakers do not accept forms with /k/ deleted. Native morphs vary considerably in how frequently they appear word-finally or not. Some token counts of native roots are provided in Table 5.18, and proportions of deletion in the noun task is given in Figure 5.18. With such a small corpus as we have for Wubuy, the text frequency counts have limited validity as an index of distributions in actual speech, but they do offer a rough rule of thumb. The crucial point is that they occupy an intermediate position on a continuum of variability between loanwords and archiphonemes, and an intermediate position between loanwords and archiphonemes in their final/nonfinal relative frequency also.

It must further be borne in mind that, according to van Egmond & Baker's (2021) recon-

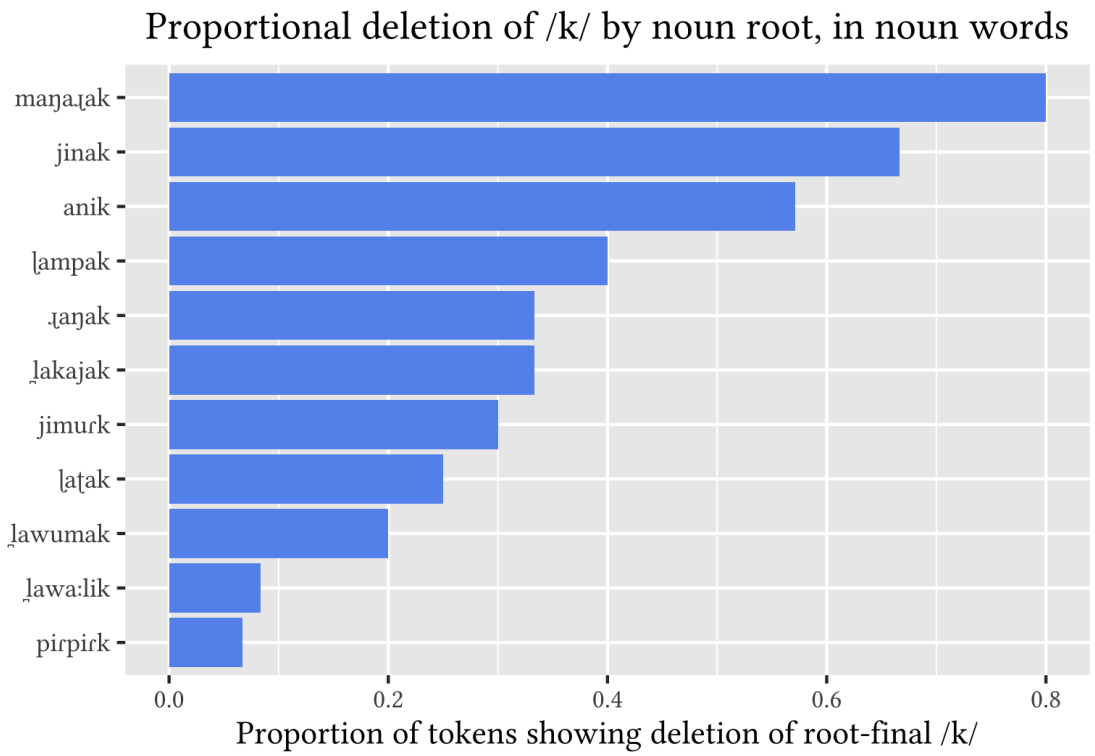

Figure 5.18: Bar chart showing proportion of tokens showing deletion relative to tokens of the same noun root that do not show deletion, within noun words. Roots with zero deleting tokens are not shown.

		word-final	word-nonfinal	percent nonfinal
anik	'coals; ashes'	0	17	100%
jinak	'head'	3	11	78.6%
maʈaɭak	'beach'	6	19	76%
ɻaŋak	'tree; stick; canoe'	22	24	52%
l̪akajak	'saltwater'	13	13	50%
l̪awumak	'long yam'	11	0	0%

Table 5.7: Token frequencies in NMET (Heath 1980), separated by whether the root appears word-finally or not. Types with 1 token or fewer in the corpus are not shown.

struction, /k/-deletion has already applied to all roots that historically contained a medial pre-consonantal /k/, changing them so thoroughly that these stops are never produced in the modern language, and are not synchronically recoverable because, not being at the root's edge, they do not interact with any productive processes. We may take these historical root-medial velar stops as a fourth category, what I will call 'heritage /*k/'. Like the archiphoneme /K/, heritage /*k/ is — or rather, was — unable to appear at the right word edge, by definition. We may presume they were deleted under the same pressures now acting upon root-final /k/. This leaves us with the inventory of potentially or actually deleting velar stops in (332).

(332) Types of velar stop

type	appears at right word edge	deleted	example	
root-medial heritage /*k/	never	always	/*makpulu/	'pelican'
archiphoneme /K/	never	always	/ŋaɾaK-/	'MULT'
native root-final /k/	sometimes	sometimes	/jinak/	'head'
loan root-final /k/	usually	never	/cɻak/	'truck'

If we assume that contemporary Wubuy's direct ancestor had no velar stop deletion, so that all of the velar stops in the table in (332) were retained on the surface, it would seem reasonable to regard deletion as having advanced through the lexicon from the top of the table to the bottom. The reason that deletion has applied in that order can be attributed to the relative frequency with which each morph appears word-nonfinally, where pressures toward gestural overlap should be stronger, creating new coarticulating variants that serve as the seeds for sound change. Thus, if /k/ in a given morph occurs more frequently in contexts where there is a pressure for the dorsal constriction gesture to overlap with the following gesture, then this will gradually shift the production target for that morph further towards overlap. That is, the stronger the expectation of other material following the /k/, the more anticipatory coarticulation there will be.

Heritage /*k/ is the most extreme case. The segment following it is always the same because it is part of the same morph; we therefore expect the gesture to be tightly coarticulated with surrounding segments (Bradley 2002; Cho 2001; Hay 2003). This anticipatory coarticulation would, according to AP-inspired approaches, lead to a high degree of gestural overlap, which can then lead to the velar constriction not being perceived by the listener at all. This listener may then produce such morphs without any articulation of the velar constriction: that is, they will delete the /k/. If subsequent generations of speakers have no exposure to non-deleting variants, and the deleting variants do not sound substantially different from them anyway

(Olejarczuk et al. 2018), then they have no reason to posit a /k/ as part of the representation of these morphs. That is, sound change, for this part of the lexicon, is complete. And indeed, this is the case in contemporary Wubuy — deletion has applied so thoroughly within morphs that these velar stops are neither observable nor synchronically recoverable.

The morph-final archiphoneme /K/ is next in line: it also carries a strong expectation of a following segment. While the exact identity of that following segment is not determined absolutely, one at least knows that it must be [–CONT]. Deletion may have been completely phonologised here at the same time or after heritage /*k/, but it is clear that the final stop is no longer part of speakers' representations, as they neither produce nor accept tokens in which /K/ is realised as [k].

The native roots and loans are where deletion has not yet been completely entrenched in lexical representations. For both types of root, there are at least some contexts where they are not followed by another consonant: when appearing word-finally, they are followed either by a pause or some other segment (any segment that can appear word-initially, including vowels). If speakers assemble production plans on the basis of their store of word exemplars containing that morph, then the likelihood of /k/-deletion in the produced token should reflect the final/nonfinal relative frequency of that morph. Moreover, even if the speaker selects a non-deleting variant from the exemplar cloud, the token they actually produce may still be perceived as deleting, because the same pressures toward articulatory overlap that result in the loss of a [k] percept historically are still acting upon productions in the present. In that way, I expect that these native and loan roots are not categorically distinct; the difference in their susceptibility to deletion is a function of current usage, which is always subject to change. Moreover, some of the native roots — the ones not shown in Figure 5.18 (cf. Table 5.4) — underwent no deletion at all, just like the loan roots.

5.5.4 Deletion in native noun roots in nouns and verbs

I have noted that there is a major difference in the likelihood of /k/-deletion in nouns versus verbs: deletion applies often in verbs (79% of tokens), but fairly infrequently in nouns (25% of tokens). This fact is not immediately explainable in terms of frequency-based lenition, as the same roots may have different likelihoods of deletion according to the class of word they appear in. Even so, the higher propensity for deletion may nonetheless be regarded as reflecting the entrenchment of patterns of coarticulation — as in Hay's (2003) derived words that are more frequent than their bases, and that do not have phonetically salient morphological boundaries. That there is a difference at all between the noun context and the verb context, for the same roots, suggests that these roots have at least partially divergent representations for constructions of different types — noun constructions versus incorporation constructions — where the latter

type of construction is more prone to becoming 'lexicalised' than the former. Some support for this comes from the fact that incorporable nouns in Wubuy form a closed class, and that some of these are suppletive with respect to their free forms. If the phonological form of incorporated nouns can diverge so completely as to be suppletive, it is perhaps unsurprising that they can also diverge in more fine-grained aspects of intergestural coordination. Certainly, this aspect of the results would benefit from further research.

5.5.5 Limitations

Relational Morphology, like any model, is a simplification of reality. Insofar as it rests on a formal algebra, it will always be too simplistic to grasp continuous phenomena, like phonetic variation. And, although proponents of Construction Grammar typically claim that schematic generalisations should be understood in probabilistic terms (Bybee 2006; Kapatsinski 2021), implementing probability into the formal representational devices is not particularly straightforward, and I do not pretend to have an answer for it here. Something akin to Albright & Hayes' (2003) model, which assigns weights to schematic generalisations based on a ratio of 'hits' (successful matches) to 'scope' (the domain of the generalisation), may prove fruitful. It is not obvious, however, how such a model would translate to RM, where relational links are bidirectional: would probabilistic weighting apply to individual schemas? Or schema pairs? Or to the relational links joining those pairs? Any earnest attempt to incorporate probability into RM would need to grapple with these questions.

On the other hand, one of the key planks of RM is the idea that words are stored and shaped by use, and patterns emerge from those words. Certainly, there's a crucial role for the word context in explaining the results here, without which the probabilistic deletion is difficult to account for. At the same time, it shows that there are coherent units smaller than the word, namely the morph, even though it cannot be analytically separated from the word context (as the context in which predictability needs to be computed). In order to coherently quantify the contribution of these two different contexts, morph and word, it may prove fruitful to draw on some of the recent literature analysing the gradient emergence of morphology from syntax in terms of mutual information, since such work also has to deal with the problem of sequences of recurrent elements that can be more or less fused to adjacent elements (as in Mansfield 2021).

5.6 Conclusion

This chapter has examined Velar Stop Deletion in Wubuy, a process whose opaque interactions with Hardening and Velar Stop Insertion has been used to motivate a highly abstract segment, the stop archiphoneme /K/. Although it has occasionally been noted in the literature that the

deletion process does not apply uniformly throughout the lexicon, as Heath (1984) originally described, this is the first time that variability has been quantified and analysed in any detail. I have shown that the variability is structured in a way that reflects the entrenchment of patterns of anticipatory coarticulation: a morph-final /k/ will be deleted according to that morph's likelihood of being followed by further material within the word. In that respect, it contributes to the still very small number of studies of phonological variation in Australian languages, and provides evidence for sound change by lexical diffusion. It is also unique in putting a particularly curious abstraction under the microscope, namely the archiphoneme — to the best of my knowledge, an abstraction whose phonetic expression has never been investigated, despite sometimes appearing in grammatical descriptions in Australia and elsewhere.[3]

Unfortunately, some of the most interesting questions raised by these phenomena cannot be answered on the basis of the present corpus. The study reported on here is fairly exploratory in nature, and future research may benefit from more controlled laboratory methods. It would also benefit from an investigation of the velar nasal /ŋ/ and its archiphoneme equivalent /N/ — which, as I mentioned briefly above, shows some indications of being subject to the same kind of structured variation as the velar stops in deletion as well as assimilation.

But more importantly, without obtaining reliable frequency estimates using large corpora, we are barred from applying more sophisticated information-theoretic methods probing the relationship between an item's probability or predictability and fine-grained details of its phonetic realisation (Foulkes et al. 2018), nor can we develop computationally implemented models to generate predictions (Daland & Zuraw 2018; Pierrehumbert 2003), which in any case can be difficult to then test experimentally given the very small number of Wubuy speakers. It is my hope nonetheless that the data and analysis presented here raises questions for languages with larger speaker populations and more extensive documentation, and that answering them will shed light on the phonologisation of phonetic processes and the diffusion of sound change through the lexicon, particularly in contexts of morphological complexity.

[3] It should be pointed out, however, that most invocations of an archiphoneme in the phonological literature, going back to the Prague School, treat it as a device for underspecification, not subcategorisation of an adjacent element, as Heath (1984) does with /K/ (see Cole & Hualde 2011).

Appendix: Forms listed in Heath 1982 with final /K/

Form	*Gloss*	*Productivity*
/(w)uK-/	Deriv. classifier (NEUT/RESID/PL/COLL)	No restrictions apparent
/niK-/	Deriv. classifier (M/MASC)	No restrictions apparent
/ŋaɾiK-/	Deriv. classifier (F/FEM)	No restrictions apparent
/maK-/	Deriv. classifier (VEG)	No restrictions apparent
/aK-/	APPLicative	No restrictions apparent
/ŋaɾaK-/	MULTiple	No restrictions apparent
/waɾaK-/	MULTiple	Rare variant of /waɾa-/; attested with 3 roots
/waːK-/	APPLicative	Rare alternative to /aK-/, attested for a number of verbs in ND, but always disfavoured amongst speakers I consulted.
/paK-/	'eye' (suppletive); independent noun form /pakaɭaŋ/	Used in many established compounds, but apparently no longer productive; transparent form /pakaɭaŋ/ used and preferred for novel compounds
/puɾK-/	—	1 verbal compound /puɾ-puta/ 'to bulge' (apparent root /$w_2$uta/ 'to be swollen' occurs in a few other compounds but not independently)
/ʈiɾK-/	'fear/bravery'	2 verbal compounds /ʈiɾ-ŋawi/ 'to be afraid', with root /ŋawi/ 'to die', and /ʈiɾ-pa/ 'to eat shit', where gloss is given in ND by analogy with equivalent term /ki-wa/, transparently derived from /gi-/ 'shit' and /wa/ 'to bite' — both compounds are used as generic swear words to indicate failure or mishap. Two other possible adjectival compounds /ʈiɾ-paːpaɻa/ 'fearless, courageous' (root does not appear elsewhere) and /ʈiɾ-kalak/ 'timid, harmless' (root /$w_1$alag/ appears in one other compound; semantics unclear)

Form	*Gloss*	*Productivity*
/l̪aK-/	‘mouth’ (suppletive); free form /ɻamaṯan/	1 verbal compound /l̪a-pa/ ‘to bite (something), to hold it in one’s teeth’; possibly implicated in some historical nominal compounds, now frozen — e.g. /l̪akuɭa/ ‘lips’
/l̪akaɾK-/	‘line, row’	2 verbal compounds /l̪akaɾ-malmalgi/ ‘(several individuals) to form a line or row’, with root unattested elsewhere; /l̪akaɾ=kaɭaṯa-/ ‘to line up’, with root /w$_2$aɭaṯa/ ‘to stand broadside (to something)’
/l̪iK-/ or /liK-/	‘paddle; sound of paddle’	1 verbal compound /l̪i-caŋa/~/li-caŋa/ ‘to hear (approaching canoe)’, with root /jaŋa/ ‘to hear’
/malaŋaK-/	‘somewhat; quite a bit’	Appears in several compounds, most relating to size or direction; /malaŋa-ɲanac/~/malaŋa-ɲanaj/ ‘far away’ is the most frequent. Some compounds are fairly transparent, as in /malaŋa-ʈuŋkal/ ‘quite big, rather big’ (← /ɻuŋkal/ ‘big’), /malaŋa-kaṯaːṯaɾi/ ‘quite a bit or somewhat behind’ (← /aṯaːṯaɾi/ ‘behind’); semantic contribution can be less obvious, as in /malaŋa-ʈamali-jac/ ‘to the north-east’ (← /ɻamali-jac/ ‘to the east’), /malaŋa-capi/ ‘(moon) to be about to set’ (← /japi/ ‘to enter’, lit. ‘put in.Refl’), or /malaŋa-caɭ̪ta/ ‘to go past, slip away’ (← /jaɭ̪ta/ ‘to slip past, escape’)
/malK-/	‘time(s)’	Appears in several compounds, usually headed by a numeral, to form a temporal adverbial: /mal-kaɲcaːpukic/ ‘once’ (with /aɲcaːpukic/ ‘one’), /mal-pula-waː/ ‘twice’ (with /w$_2$ula-waː), /mal-kaɾawinti/ ‘many times’ (with /aɾawinti/ ‘many’), /mal-kaɭakaɭa/ ‘somewhat later’ (/aɭakaɭa/ ‘halfway (between two points, in a journey); along the route’)
/maɾK-/	‘spear’	Appears in 2 compounds, with /pilwila/ ‘to shake’ and /paɾa/ ‘to be crooked’, both of which have special /w$_2$/-initial forms in compounds.
/ŋuɻaK-/	‘penis’ (suppletive); independent form /ŋuji/	Attested in a range of compounds, and is productive. Often used as the basis for novel compounds intended as insults.

Form	*Gloss*	*Productivity*
/ɻakaɾK-/	‘spear; violence; threat of violence’	Attested in a few compounds, often with non-obvious semantic contribution: /ɻakaɾ-t̪uɭuwi/ ‘to push or propel (spear) along’ (← /ɭuɭuwi/ ‘to make a heap or pile of something’), /ɻakaɾ-na/ ‘to see spear of (someone)’ (← /na/ ‘to see’), /ɻakaɾ-paʈci-waɲci/ ‘to set up (sails, rigging of ship)’ (← /w$_2$aʈci-wu/ ‘to strike, hurt, throw spears at’), /ɻakaɾ-pajama/ ‘to drop or fall down violently or dangerously (e.g. object falling from tree)’ (← /w$_2$ajama/ ‘to keep going, continue along’), /ɻakaɾ-ŋaʈpa/ ‘to be out of sight (over the horizon)’ (listed with same gloss as its head /ŋaʈpa/), /ɻakaɾ-paʈaʈt̪i/ ‘to hold spear ready to throw, to brandish spear’ (head verb root unattested outside this compound). DM also used /ɻakaɾ-kutaː/ ‘to be perched in tree’.
/aK-/~/a-/	‘ground, earth, dirt’	Only two instances of this morph in ND with clear ‘ground’ meaning: /a-paʈawaʈaʈ/ ‘hard ground, solid earth’ (← /w$_2$aʈawaʈaʈ/ ‘strong, firm, solid, hard’) and /a-paʈpiɾ-ma/ ‘(ground) to dry up’ (← /w$_2$aʈpiɾ-ma/ ‘dry-Inch’) is the only instance of /aK-/ in ND with clear ‘ground’ meaning. Other entries have more diffuse semantics: /a-ɭaɭma/ ‘to dawn’ (/ɭaɭma/ also ‘to dawn’), /a-puta/ ‘to bulge’ (apparent root does not appear outside of compounds). Not productive. Related independent form is /apan/, which can also be incorporated.
/t̪aK-/~/t̪akaːK-/~/akaːK-/	‘?’	All three appear in compounds with verb root /w$_1$uɻama/ ‘to move around (person or thing), especially in circular motion’, with practically the same meaning: /t̪a-kuɻama/ ‘go around in circles’, /t̪akaː-kuɻama/, which is transitive (i.e. ‘to surround’), /akaː-kuɻami/, which is grammatically reflexive but interpreted as intransitive. The /t̪a-/ alternant may also be implicated in some other roots that are listed as simplex in ND: /t̪apaɭajita/ ‘slip, slide’ (cf. /paɭajita/ ‘slip, slide’, /w$_2$aɭajita/ ‘to be smooth or handsome’).

Form	*Gloss*	*Productivity*
/aciɾK-/	Possibly ‘failure, haphazardly’	Possibly implicated historically in several roots given as simplex in ND: /w$_1$aciɾpatca/ ‘to miss (with thrown object)’ (cf. /w$_2$atca/ ‘to hit’), /aciɾpaɾka/ ‘to throw (something) wildly’ (cf. /w$_2$aɾka/ ‘to throw’), /akaciɾkaɻata/ ‘to be unable to find (someone, something)’ (cf. /w$_1$aɻata/ ‘to go to (something)’).
/liliK-/	‘water’ (?)	1 compound: /lili-kal̪aɾa/ ‘(water, etc.) to make (person, direct object) feel queasy or bloated’ (cf. /w$_1$al̪aɾa/ ‘to stab, jab, poke’)
/ampuluK-/	‘(top of) head’	Attested in 2 compounds in ND: /ampulu-kuɻja/ ‘to pour liquid on or over the head of (someone)’ (cf. /w$_1$uɻja/~/w$_2$uɻja/ ‘to pour liquid on (person or thing)’) and /ampulu-kuluta/ ‘to hold one’s head low, have one’s head bowed’ (verb root not attested independenty). Contemporary speakers accept novel compounds with /ampuluK-/, but the more transparent /jinak-/ is overwhelmingly more common.
/w$_1$uɾK-/	‘?’	1 compound: /wuɾ-puɭuɭu(ɭu)l̪a/ ‘to be stained’ (cf. /w$_2$uɭuɭu(ɭu)l̪a/ ‘to be bright-coloured and variegated’). Contribution of /wuɾ-/ here unclear.
/jamaɾK-/	‘side (of an object)’	Suppletive: independent form is however the almost identical /jamaɾ/. Implicated in 3 compounds, including /jamaɾ-kaɭaṭ̪a/ ‘to have one’s side facing (something)’ (cf. /w$_1$aɭaṭ̪a/ to stand broadside (to something)’). The other two compounds feature heads beginning with /ɭ/, and therefore support an interpretation of the UR as either /jamaɾ-/ or /jamaɾK-/, since either would lead to the same surface form.

Table 5.8: Morphs analysed as /K/-final in Heath (1980)

6 Prosody and morphological complexity

6.1 Introduction

In this thesis, particularly in Chapters 3 and 4, I have argued for the word as the fundamental unit in which linguistic knowledge is stored in memory, rather than units smaller than the word, such as morphemes. I showed that only by treating phonological alternations as relations between whole words, rather than deriving surface words from abstract underlying morphemes, can we make the many complex interactions between phonological processes in Wubuy tractable.

Such a characterisation is susceptible to the criticism that it presupposes that we know how to distinguish words from syntactic phrases on the one hand, and from smaller elements on the other. The difficulty of arriving at a suitably concrete definition of the 'word' is an especially pertinent issue in the study of polysynthetic languages, since it is a hallmark of these languages that their putative words are of a complexity normally shown by syntactic phrases in languages with simpler morphology (Baker 2018; Bickel & Zúñiga 2017; Bundgaard-Nielsen & Baker 2020; Mansfield 2023). In the case of Wubuy, the problem of defining the word is less pressing than in many other polysynthetic languages, including the neighbouring languages of the Gunwinyguan family, since grammatical criteria for wordhood tend to align with phonological ones. The domain that we call the 'word' is the only one in which phonological alternations apply, and it is also the domain wherein grammatical elements must follow an obligatory order (a morphological template) (Dixon & Aikhenvald 2002).

Even so, verbal words in Wubuy may be very long and dense with information, often having a phrasal interpretation. Given the close association between phrasal syntax, information structure and prosody (Bennett & Elfner 2019; Selkirk 1984), this raises the question of whether complex verbs in Wubuy prosodify in a way more characteristic of words or phrases in more analytic languages (Mansfield 2023). Polysynthetic verbs in related languages like Dalabon, Bininj Kunwok and Ngalakgan have, to varying degrees, been shown to admit multiple phrasal prosodic constituents (such as a Phonological Phrase or Intonational Phrase) within the bounds of the grammatical word, as determined by the alignment of pitch accent and phrase-final boundary tones to word-internal constituents (Bishop & Fletcher 2005; Fletcher 2014; Ross 2011). Most strikingly, these polysynthetic verbs often show speech pauses between such prosodic

constituents, further problematising our assumptions about wordhood (Dixon & Aikhenvald 2002; Evans et al. 2008; Fletcher, Evans & Ross 2004). For Wubuy in particular, it has been shown that words with pauses artificially inserted at certain morphological junctures within words are sometimes rated by listeners as more acceptable than words without pause in a two-alternant forced-choice judgement task (Bundgaard-Nielsen & Baker 2020). Furthermore, pauses at certain boundaries — namely, those demarcating major morphological boundaries — were judged more acceptable than others. Bundgaard-Nielsen & Baker interpret this pausing phenomena as evidence that grammatical words can consist of more than one planning unit.

Given such a possibility, an important question that arises is whether intonational phonology — as reflected in the distribution of pitch accents and phrasal boundary tones — also reflects the existence of prosodic units inside the Wubuy grammatical word, and whether these units align with major morphological constituents.

To answer this, I conducted an analysis of intonational pitch accents and boundary tones in a corpus of unscripted texts following approaches in Autosegmental-Metrical Phonology (Ladd 2008), in order to examine the alignment of these accents to syllables at different positions in the morphological structure of the verb. My analysis finds that pitch accents are most frequently attracted to the left edge of major morphological constituents. These are, in ascending order of size: the inflected root, the base to reduplication, the stem, and the word. However, as in data reported for related languages, pitch accents at all of these locations are optional.

Pitch accents can also appear close to the right edge of the verbal word — the penultimate or final syllable — in a way indicative of quantity sensitivity. Pitch accents on the final syllable are largely limited to words whose final syllable is heavy, even though not all words with a final heavy syllable will necessarily have a final accent. Similarly, penultimate accent is far more frequent in words with a final light syllable than in those that have a final heavy syllable, even though in such light-final words it only occurs in about half the tokens. While quantity-sensitivity has often been advanced in support of the construct of metrical feet, I instead argue here, following Gordon (2014), that this kind of evidence can be analysed at least as plausibly without feet, and that parsimony demands that we avoid postulating more levels of structure than required.

A further interesting result of this analysis is the finding that prosodic structure does not align exactly with grammatical word structure in the way that previous analyses have led us to expect. Specifically, the morph within the pronominal prefix complex marking the Veg object class — which in all cases takes the form /-ma-/ — is frequently accented, even though it always appears at the right edge of the prefix complex, and other morphological formatives in this position do not generally receive an accent. I attribute this unusual behaviour to the fact that the Veg marker is unique in being the only transparent formative in the entire prefix paradigm,

a quality which enables it to be prosodically separated from the prefix complex and to serve as the left edge of the verbal stem instead.

In the next section, §6.2, I provide an overview of debates surrounding prosodic prominence in Australian languages and elsewhere, and how this has been analysed in previous work on Wubuy, and then formulate a set of research questions. I present my methods for answering those questions in §6.3. There, I introduce the corpus, the inventory of pitch accents and boundary tones used for the study, and the querying procedures. In §6.4, I discuss the common types of tunes found in the corpus. I then examine tune-text alignment in §6.5, which forms the most substantial part of the analysis; this is organised into sections relating to each of the relevant morphological constituents of the verb. I discuss the findings in §6.6, and offer conclusions and directions for further research in §6.7.

6.2 Background

This section provides crucial background on the main empirical and theoretical contributions of the chapter. I begin, in §6.2.1, with an overview of prosodic phenomena in general. Descriptive work on Australian languages often includes commentary on word-level prosody, some of which has been featured prominently in some key theoretical debates (§6.2.2). Such commentary has generally proceeded on the basis of the (typically English-speaking) analyst's auditory impressions, a practice that has been criticised in more recent phonetic research (§6.2.3). I then turn to intonational analyses of Australian languages (§6.2), and give an extended summary of what is probably the most detailed account we have of the intonational phonology of an Australian language, namely Bishop's (2002) analysis of Bininj Kunwok, which will prove to shed much light on the Wubuy findings (§6.2.5). In §6.2.6, I summarise previous work on Wubuy prosody.

6.2.1 Prosody, prominence, and rhythm

In describing the sound patterns of a language, it is not enough only to describe the distribution and behaviour of segments. As Fletcher (2010, p. 521) states, "Speech is an activity that unfolds in time, and listeners are sensitive to fluctuations in the temporal flow of an utterance". Thus, suprasegmental properties — pitch, duration, intensity, rhythm — are an intrinsic property of continuous speech, and may be systematically exploited by languages to support lexical, grammatical or pragmatic distinctions (Clark et al. 2007). In particular, according to some influential models of the syntax-prosody interface, the organisation of these suprasegmentals is closely bound up with the chunking of fluent speech into smaller units for processing, which

itself often tends to reflect syntactic phrasing (Cutler 2012; Levelt 1993; Selkirk 1984; Wagner & Watson 2010).

The phonology of English, as well as other Germanic languages, makes an abstract distinction between 'stressed' and 'unstressed' syllables, which is cued phonetically through these suprasegmentals. Each prosodic word has a single syllable with primary stress and potentially many syllables with secondary stress or no stress. Stressed syllables are generally longer and louder, though the precise realisation depends on various contextual factors, and stress may also be registered along other phonetic dimensions (Cho & Keating 2009; Fry 1955; Lieberman 1960; Plag, Kunter et al. 2011).

But aside from these phonetic effects, stress in English has many detectable effects of a more phonological nature. Vowels in unstressed syllables are typically 'reduced' to one of a very small number of vowel categories, whereas stressed syllables admit the full range of vocalic contrasts. English stress is also subject to synchronic phonological processes such as the well-known Rhythm Rule (Liberman & Prince 1977), which moves the location of stress to avoid 'clash' — the presence of a stressed syllable immediately adjacent to another stressed syllable — as in *fiftéen* versus *fifteen mén* (cf. **fiftéen mén*).

And, importantly, stressed syllables serve as the 'docking' site for intonational pitch accents. That is, in an English sentence like *Joanne witnessed a murder*, comprising a single Intonational Phrase, any of the content words may receive the nuclear pitch accent, such as in order to suggest contrast. In each case, the syllable bearing word stress in isolation is the only one that can bear an intonational pitch accent (e.g. *Joánne, wítnessed, múrder*).

Differences in prominence may therefore reside at different levels of the grammar. If at the word level, this would typically be analysed as 'stress'. If at a phrasal level (larger than the word), this would be analysed as post-lexical or intonational prominence.

How do we disentangle the two? As Gordon (2014) argues, many descriptions of stress, particularly in less well-studied languages, make their generalisations on the basis of words uttered in isolation. Since the grammatical word is typically coextensive with the intonational phrase in such citation forms, these descriptions are likely to conflate lexical prominence and intonational prominence. That being so, it seems likely that "a large potion, if not the majority, of the prosodic typology that is currently understood to refer to stress may actually reflect phrasal prominence associated with tonal events occurring at or near phrase edges" (Gordon 2014, pp. 111–112). This risk is especially acute where no effects of stress beyond suprasegmental changes are reported which might give phonological clues to prosodic organisation — such as vowel quality alternations, phonotactic restrictions, allomorphy, and so on.

Thus, a major hurdle in the analysis of stress in an under-documented language is that of disentangling lexical and post-lexical uses of pitch (Jun & Fletcher 2014). As we will see in the

next section, there is currently little substantial agreement on the proper analysis of prosody in Australian languages in general, nor in polysynthetic languages of the Top End in particular. Specifically, the disagreement hinges on whether pitch movements are characteristic of lexical stress or post-lexical (intonational) pitch accent, and, accordingly, what prosodic constituents are required in order to explain the distribution of those pitch prominences.

6.2.2 Phonological accounts of prosody in Australian languages

Languages of Australia are frequently reported to be 'stress' or 'stress-accent' languages in the descriptive literature, meaning that they make use of word-level stress but not word-level tone (Evans 1995). Most such reports claim both primary and secondary stresses. The following summary from Dixon's (1980) classic survey is representative:

> In most Australian languages word stress is predictable. Most commonly, primary stress goes on the first syllable of a word; there is then often secondary stress on the third, fifth etc. syllables, excepting that the final syllable of a word does not usually bear stress … Most Australian languages have a very even and rhythmic sound. There is a definite preference for alternation of stressed and unstressed syllables throughout a word (and, to a lesser extent, through a sentence).
>
> (Dixon 1980, p. 128)

Dixon also notes that long vowels tend to attract stress even if they are found in non-initial position.

For multimorphemic words, it is widely claimed that the distribution of stressed syllables is determined by morphological boundaries (Baker 2014b). This notion is attributed to an unpublished manuscript by Ken Hale from 1977, in which he drew attention to the now-famous Warlpiri minimal pair:

(333) a. jápaɭa-ŋùɭu 'father's mother-elative'
b. jápa-ɭàŋu-ɭu 'person-also/even-ergative' (Warlpiri: Hale 1977 cited in Baker 2014b)

Both of these words contain the same segmental string, grouped into the same five syllables. However, the morphological segmentation differs, and the location of secondary stress reflects this difference. The secondary stress in (333b) is where it would be if the word were monomorphemic; in (333a) it is shifted one syllable to the right, such that the disyllabic suffix is stressed on its initial syllable. This behaviour is typically explained in terms of a bimoraic trochaic foot, whose left edge must be aligned to a morphological boundary (cf. Baker 2008b). More will be said on metrical feet below.

In the Warlpiri examples above, the location of stress disambiguates two segmentally identical words, and can therefore be said to play a crucial role in interpretation. To the best of my knowledge, the only contribution to semantic interpretation attributed to stress in Australian languages is of this nature: it demarcates the boundaries of morphological constituents (both words and word-internal affixes). Stress is thus not used to contrast word pairs like in English *cónvert* (noun) vs *convért* (verb). Ndjébbana (McKay 2000) and Gaagudju (Harvey 2002), spoken in North-Central Arnhem Land, and Wik-Mungkan (Sayers 1976) and Kunjen (Sommer 1969), of Western Cape York, are perhaps the only exceptions to this. All are reported to have minimal or near-minimal pairs distinguished solely on the basis of stress, where the difference in stress accent location is not derived from a difference in morphological structure as in the Warlpiri example above.

6.2.3 Phonetic correlates of stress in Australian languages

Despite the ubiquity of references to stress in Australian language descriptions, these accounts are rarely accompanied by phonetic analysis (Baker 2014b; Jepson & Ennever 2023). As Jepson & Ennever (2023, p. 155) state in a recent survey, "There exist very few experimental acoustic phonetic studies of stress or correlates of metrical prominence in Australian languages which thoroughly explore if and how [...] acoustic measures distinguish levels of stress".[1] Moreover, where such studies do exist, they tend to provide at most fairly weak evidence for earlier impressionistic accounts, and support for secondary stress is practically non-existent.

The most comprehensive account of the phonetic correlates of stress in an Australian language is that of Tabain et al (2014) on Pitjantjatjara, a language previously claimed to have a stress system much like Warlpiri's: left-aligned binary feet, with metrical reset at morphological boundaries, and primary stress on the word-initial syllable. The authors include annotation of intonational pitch accents and boundary tones as part of their analysis, and they report measures for F0, duration, RMS amplitude, spectral tilt, and the first three formants for vowels. While they find higher fundamental frequency, somewhat higher RMS amplitude, and longer duration in word-initial (primary stressed) syllables compared to syllables elsewhere in the word, they found no significant differences along any dimension between secondary stressed syllables and other word-medial syllables. Since the word-initial, primary stressed syllables also carry a post-lexical pitch accent,[2] the authors conclude that prominence there is not lexical stress at all, but should instead be attributed to intonational phrasing; further, they conclude that secondary stress does not exist at all in Pitjantjatjara. Likewise, in Djambarrpuyŋu (Jepson

[1] See, however, the recent work by Babinski (2022) on the phonetics and typology of stress across sixteen Australian languages, based on archival recordings, which came to my attention while in the final stages of revising this work.

[2] The authors use an intonational labelling scheme that is agnostic as to whether a salient pitch movement is

2019) and Bininj Kunwok (Bishop 2002) it has been shown that fundamental frequency is the only reliable correlate of stress, and syllables that are not pitch accented (i.e. secondary stressed and unstressed syllables) show no consistent phonetic differences between one another.

Outside of effects on vowels, which is cross-linguistically the more typical location to look for phonetic correlates of stress (Gordon & Roettger 2017), it has been claimed that in Australian languages stress may also be registered in consonants (e.g. Butcher 2006). Moreover, the claim is specifically that stress induces lengthening and strengthening of consonants in post-tonic position, rather than pre-tonic as has been argued to apply elsewhere (Cho & Keating 2009; Keating et al. 2004). This claim is originally due to an acoustic and articulatory study of Warlpiri by Butcher & Harrington (2003), which also looked at correlates of stress derivative of morphological structure. The experimental design used simple question-answer dialogues that systematically varied focus, as well as morphological boundary strength — as in *kuju#puŋu* 'killed the animal' (V with object) versus *kuju+puŋu* 'game-killer' (compound N) — and morphological boundary location — using the minimal pair cited in (333) above. Unfortunately, they do not report results from the latter condition. For the boundary strength condition, they found evidence of increased duration of the /uj/ in *kuju#puŋu* relative to that in *kuju+puŋu*, but no significant differences in hyperarticulation or extent of F2 excursions. They did, however, find significant increases along these dimensions in focussed contexts. They claim that most of the differences between focussed and unfocussed items are found in the post-tonic consonant in their data. However, the consonants in question include only the glide /j/, and only in that one word pair; it is questionable how far a result like this can be generalised to other consonants even in Warlpiri, let alone other Australian languages. The studies that have been undertaken to test post-tonic consonant lengthening have generally found mixed results (for an overview see Jepson & Ennever 2023). For instance, Jepson et al (2021) finds that consonants in Djambarrpuyŋu are longer in post-tonic position in disyllabic words than in the same position in longer words and other word-medial positions. They attribute this to the much more general principle of polysyllabic shortening: the more segments in a word, the shorter the duration of those segments will be.

Stress in Australian languages is rarely claimed to have any consequences outside of lending greater acoustic prominence to accented syllables (or lengthening the following consonant). The presence of a lexical stress accent on a syllable does not, in and of itself, have any effect on phonotactic possibilities, nor does the absence of stress entail reduction of vowels or lenition of consonants. Long vowels are restricted to word-initial position in many Australian languages, which is also the typical location of primary stress. In languages that allow long vowels in other

an accent or a boundary tone, but from their discussion of these word-initial pitch events, it does not seem inaccurate to call them pitch accents (see Tabain, Fletcher et al. 2014, p. 54).

positions (like Wubuy and Kayardild), those long vowels are said to attract secondary stress (though cf. Round 2009 on Kayardild). But it cannot necessarily be said in such cases that stress assignment is the cause of vowel lengthening.

There are, however, several reasons to treat statements about the nature of stress in Australian languages with caution — or indeed, in any language whose description has been undertaken by researchers who are not themselves native speakers. Particularly where the native language of the researcher is English or another Germanic language, they may be liable to hear alternations in prominence throughout the word or rhythmic organisation of syllables even in the absence of any physical correlates, as detectable by phonetic investigation. Germanic phonological systems categorically distinguish stressed and unstressed syllables, and this distinction has a wide variety of quite robust phonetic consequences: differences can be found in pitch, amplitude, intensity, spectral tilt, vowel dispersion, and jaw aperture. Just as experience with the phonemic distinctions of one's native language can lead one to miscategorise phonemes in an unfamiliar language (Best & Tyler 2007), experience with the prosodic organisation of one's native language can cause misperception of prosody elsewhere. Tabain, Fletcher & Butcher (2014) term this phonomenon 'stress ghosting'. It is the inverse of so-called 'stress deafness' (Dupoux et al. 2010; Peperkamp et al. 2010), whereby speakers of languages that do not make use of word-level stress, such as French, have difficulty perceiving stress distinctions in languages that do. English speakers are therefore likely to hear the 'ghost' of (particularly trochaic) stress even when adjacent syllables are identical along the phonetic dimensions that distinguish stress in English. Note, however, that although the stress ghosting hypothesis is grounded in established findings of second language perception, it is yet to be investigated empirically (unlike stress deafness).

The problem is compounded by the fact that different individuals may detect stress, whether ghostly or otherwise, in different locations in the word. As Mansfield (2019) shows, different descriptions of Murrinhpatha have yielded wildly varying accounts of stress accent placement, particularly when it comes to secondary stress, and since their claims are incompatible, they cannot all be true. Indeed, he argues on the basis of phonetic evidence that secondary stress does not exist in Murrinhpatha. Further, primary stress, since it occurs on the penultimate syllable of the rightmost prosodic word in a phonological phrase, should really be considered a means of demarcating the right boundary of that phrase. High F0 in Murrinhpatha is therefore not indicative of word stress at all but rather (post-lexical) intonational constituency (Mansfield 2023, p. 433).

To the best of my knowledge, the only Australian language in which metrical foot structure can be said to have any effect outside of increasing the phonetic prominence of certain syllables is Pitjantjatjara/Yankunytjatjara, two basically mutually intelligible varieties of the Western Desert Language dialect chain in Central Australia. Following initial remarks by Goddard (1985),

Wilmoth & Mansfield (2021) show that the organisation of verb stems into quantity-sensitive metrical feet conditions the appearance of an augment suffix in certain tense-aspect-mood inflections. These augment suffixes appear when the stem ends in a defective foot. They can therefore be seen as a repair for a dispreferred metrical structure, although the choice of which augment to use is determined by lexically-conditioned inflectional class. Moreover, they show that this state of affairs reflects a historical process deleting unfooted syllables — including in the language name, which is now universally pronounced [pi.can.ca.ɾa].

6.2.4 Intonational phonology of Australian languages

While 'stress' has been shown to be an increasingly problematic notion to uphold in studies of Australian language phonologies, at least in its earlier form, a growing body of work has examined prosody at the level of intonation. Detailed investigations of intonation in Australian languages exist for Wik Mungkan (Sayers 1976), Dyirbal (King 1994), Alawa (Sharpe 1972), Kayardild (Fletcher, Evans & Round 2002; Ross 2011; Round 2009), Mawng (Fletcher, Stoakes et al. 2016; Hellmuth et al. 2007), Jaminjung (Simard 2010), Djambarrpuyŋu (Jepson 2019), Dalabon (Fletcher 2014; Fletcher, Evans & Ross 2004; Ross 2011), Mayali (Fletcher & Evans 2000) and several other dialects of Bininj Kunwok such as Kuninjku, Kunwinjku, Kune, Gun-djeihmi, Kundedjnjenghmi and Manyallaluk Mayali (Bishop 2002; Bishop & Fletcher 2005). Some of this work has enabled a reconsideration of earlier impressionistic work on word-level stress, by relocating some of the prominence found there to larger intonational prosodic constituents.

Most of these works are framed within the Autosegmental-Metrical (AM) theory of intonational phonology (Ladd 2008).[3] By 'intonational phonology' is meant "the use of suprasegmental phonetic features to convey 'postlexical' or sentence-level pragmatic meanings in a linguistically structured way" (Ladd 2008, p. 4). The main suprasegmental phonetic feature relevant to intonation is pitch/f0, but duration and intensity may also be relevant, as well as pause placement, phonatory phenomena such as glottalisation or breathiness, and manipulations of pitch range.

In AM, intonation is represented as a sequence of phonological tone targets, which may be either High (H) or Low (L), or some combination thereof. The 'autosegmental' part comes from the fact that tones occupy a tier distinct from the segmental string, much as in the autosegmental representations used to analyse tone languages (e.g. Goldsmith 1976). Tones may be of two kinds: pitch accents and boundary tones. Pitch accents are aligned with respect to metrically strong syllables; boundary tones mark the edges of intonational constituents, such as the Intonational Phrase (IP). A given sequence of tones produces a 'tune' or 'melody', which is associated in the grammar with a particular pragmatic or discursive function. The tonal inventory and the

[3] Simard's (2010) analysis of Jaminjung is an exception in adopting the PENTA framework (Xu 2005).

association of tunes to meanings differs from language to language, just as with any other elements of phonology. In addition, most AM research on individual languages tends to assume that there is no single prosodic hierarchy that is instantiated in every language's intonational grammar, but rather that any putative levels of prosodic constituency must be demonstrated empirically within that language (Jun & Fletcher 2014, p. 502; also argued by Schiering et al. 2010).

In most of the Australian languages given AM treatments so far, the proposed tonal inventory is very small compared to the better-studied languages of Europe, such as English (Pierrehumbert 1980) or Dutch (Gussenhoven 2005). Often, the inventory consists of one basic pitch accent type (though it may appear upstepped or downstepped). This is usually a simple H*, as in Yukulta (Bonnin 2014); in Mawng the only pitch accent is L+H* (Hellmuth et al. 2007), though it has a similar phonetic realisation. Sometimes there is an additional pitch accent type that appears less frequently, such as the L+H* accent in Dalabon (Fletcher 2014) and Bininj Kunwok (Bishop 2002). Unlike in many languages of the world, in which unfocused material is often not associated with any pitch accent, content words in these languages typically have at least one pitch accent irrespective of discourse status, and complex verbs in polysynthetic languages often have several (Fletcher 2023; Mansfield 2023; though cf. Jepson 2019).

The inventory of boundary tones is usually larger — especially at the right edge — but still fairly small. In Mawng, for instance, the left boundaries of phrases are not tonally marked, while right boundaries may be just H or L, whether marking intermediate phrases or intonational phrases (Hellmuth et al. 2007, p. 10). Dalabon (Fletcher 2014) has those same boundary tones as well as !H%, ^H% and LH% at the right edge, and %L and %H at the left edge. Fletcher, Evans & Round (2002) and Round (2009) propose different sets of distinctive tonal events for Kayardild, but both are quite large compared to the other languages mentioned here.

Australian intonational systems are frequently reported to have two basic tune types. One, usually termed the 'basic declarative', shows a general declination in the f0 topline from the phrase-initial H* (or similar) accent towards the end of the utterance. The other is characterised by a sustained plateau-like tune that remains high throughout the phrase, and is accompanied by dramatic lengthening of the phrase-final syllable. We will see several examples of this tune in the Wubuy results reported below. This tune has been variously termed the 'stylised high level intonation contour' (Bishop & Fletcher 2005, on BKW), 'linear lengthening intonation' (Mailhammer & Caudal 2019, on Iwaidja), 'narrative high monotone' (Sharpe 1972, on Alawa), and 'iconic lengthening' (Simard 2010, on Jaminjung).

6.2.5 Intonation and morphological complexity in Arnhem Land

I turn now to an issue closer to the central concerns of this chapter, namely the prosodic organisation of complex verbs in polysynthetic languages. As Mansfield (2023) points out, 'it is not obvious how we should expect polysynthetic verbs to prosodify, since these verb structures typically have mixed characteristics of grammatical words and syntactic phrases'.

The most detailed description we have of intonational phonology in a Gunwinyguan language, and perhaps even in an Australian language, is Bishop's (2002) account of several dialects of Bininj Kunwok. Since the results to be reported in this chapter reveal a strong similarity with Bishop's description, particularly of the Manyallaluk Mayali dialect, it will be worth summarising that work here in some detail. Following this, I summarise a recent paper by Mansfield (2023), which reanalyses her account in a way that will prove fruitful for understanding the Wubuy data also.

Verb structures in Bininj Kunwok are very similar to those in Wubuy. Verbal words contain a prefix realising one or two pronominal arguments, and a stem. The verb stem minimally consists of an inflected verb root, but it is often much more complex, with several morphological elements intervening between the pronominal prefix and the root, such as incorporated noun roots, valence-changing affixes and adverbials. As in Wubuy, nouns can also be complex, as when they take noun class prefixes or undergo reduplication. This morphological structure almost entirely determines the location of prominence, in that pitch accents generally occur at the left edges of these morphological domains. However, a crucial finding is that the actual realisation of a pitch accent at any one of these privileged locations is optional. As she puts it, "not all metrically strong syllables in a word or phrase in Bininj Gun-wok bear pitch accents, though all have the logical potential to do so. The same phonological word may bear one accent in one utterance, and two in another" (Bishop 2002, p. 130). This optionality of pitch accent is demonstrated by the minimal pair in (334); both tokens are segmentally and semantically identical but have different syllables realising a pitch accent. Indeed, it is only by comparing multiple tokens that one can discover what the accentable locations in a complex word are. Similar findings have been reported for Ngalakgan (Baker 2008b) and Dalabon (Evans et al. 2008).

(334) a. ŋál-daʔ-dálʊg
(ŋál-(daʔ-(dálʊg

b. ŋal-dáʔ-dalʊg
FEM-RED-woman

'the women' (Bininj Kunwok: Bishop 2002, pp. 153–154)

Bishop accounts for the variability across tokens by positing an underlying metrical structure which makes crucial reference to feet. In the BKW prosodic hierarchy that she proposes — shown in (335) — syllables are grouped into feet, and a phonological word (generally isomorphic with the grammatical word) may contain multiple feet.

(335) Bininj Kunwok prosodic hierarchy (Bishop 2002, 100ff)

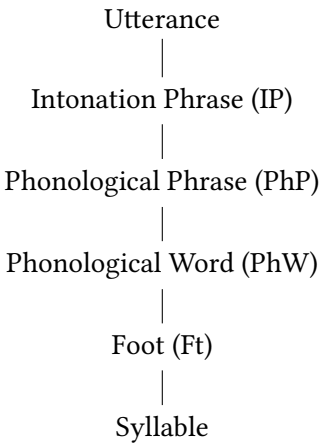

Somewhat unusually, feet in this analysis are 'unbounded' — that is, they have a left edge but not a right one (Bishop 2002, p. 120; also argued in Evans 2003). This left edge of the foot must align with the left edge of a morpheme. Feet are also left-headed, which is to say the initial syllable of the foot is a possible location for realising pitch accent. Each morpheme therefore corresponds to a single foot, though there are certain exceptional morphemes which may 'cohere' with a preceding or following morpheme — i.e., form a single foot with it. These cohering morphemes include some prefixes and suffixes that each consist of a single consonant, and the verbal TAM inflectional suffixes. In other words, the inflected root, which her analysis otherwise treats as two morphemes (root + TAM suffix), is treated as a single morphological unit for the purposes of metrical structure (Bishop 2002, 119ff). In a similar vein, Baker's (2008b) analysis of Ngalakgan treats TAM inflectional suffixes as Root-level affixation, which likewise results in those suffixes forming a single metrical domain with the root they attach to.

Bishop (2002) proposes an analysis in terms of a bracketed metrical grid (Halle & Vergnaud 1987; Hayes 1995). An example is shown in (336). The morphological segmentation, in level 0, determines the foot structure, represented by the brackets on level 1 (note the unity of the verb root and tense inflection). The initial syllable of each foot is the metrical head, and is assigned a mark, also on level 1. These marked syllables are the only ones available to receive a pitch accent on level 2, though what drives one strong syllable rather than another to be accented is still not well understood, as she acknowledges. Bishop (2002, p. 132) claims she cannot discern

any secondary stress, nor any difference between metrically strong unaccented syllables and metrically weak (and therefore unaccented) syllables. (On that basis, she suggests a rule that removes marks on level 1 that do not have a corresponding mark on level 2.)

(336) Bininj Kuwok accent assignment as metrical grid

2 (* (*
1 (* (* (*
0 ngarri-yauh- maknan
1aug- again-take.a.look.NPST

'We'll try looking at one more place' (Gun-djeihmi: Bishop 2002, p. 147)

That feet are unbounded (only have a left edge) is shown in examples of tri- and quadrisyllabic bare nouns, which show only an initial accent, as in *ŋánapːaɻu* 'water buffalo' or *gɔ́ʃɔməmɔ* 'freshwater crocodile'. If feet were binary, and foot heads were prominent, as argued for Ngalakgan (Baker 2008b), one would expect accents on alternating syllables — that is, on the first and third syllables of a quadrisyllabic root.

In the Manyallaluk Mayali dialect, there is a further accentable location outside of the left edge of a morphological constituent, and that is the penultimate syllable of the phonological word; some examples are reproduced in (337). According to Bishop, penultimate accents demonstrate an additional foot in Manyallaluk Mayali which is aligned to the penultimate syllable of the phonological word irrespective of morphological structure.

(337) a. bári-bepé-joj
'They lay separately'

b. bá-jergáni
'He/she/it was sitting down'

c. bári-geb-dukːáni
'They used to tie up the jaws [of the freshwater crocodile]'

d. bári-miɳɖé-di
'There were many of them there'

e. dárgíd-no
'alive'

f. bari-jángájʔmi
'They were cheering'

g. yi-mim-dápgéji
'You would cover their eyes' (Bishop 2002, pp. 163–168)

There are two other factors that can, in all dialects, result in a pitch accent appearing outside the leftmost syllable of a foot — one phonological, and one phonetic in nature — and in both cases it is the peninitial that receives the accent.

Firstly, there appears to be evidence that a peninitial closed syllable can attract pitch accent, as in (338). Bishop claims that this quantity-sensitivity is 'limited', since she has examples of closed peninitial syllables that fail to attract pitch accent. The claim depends on what is meant by 'closed' here. Baker (2008b, pp. 181–184) argues that nearly all of the examples of heavy syllables attracting stress in Bishop (2002) and Evans (2003) show precisely the same limitation on quantity as he describes for Ngalakgan: geminates and monogestural clusters do not contribute weight to the preceding syllable, only bigestural clusters do. There are a couple of potential counter-examples, however, such as (338c). A further difference from Ngalakgan, though, is that an accent on a heavy peninitial syllable can co-occur with an initial accent in the same foot, as in (338d).

(338) a. ŋál-jabókware
'your sister' (triangular kin term)

b. bíri-wuɭébmeŋ
'They swam'

c. ŋut:áŋge-kaʔ
'to yours'

d. móɟárk:i
'freshwater crocodile' (Kuninjku: Bishop 2002, p. 161)

Secondly, a pitch accent held to be phonologically anchored to the word-initial syllable may have its peak phonetically delayed, and appear instead on the peninitial. Bishop (2002, p. 289) demonstrates a positive correlation in Kuninjku between initial syllable duration and peak delay, measured from the onset of the metrically strong syllable to the F0 peak: that is, peak delay increases as strong syllable duration increases, even to the point of appearing on the following syllable. No significant correlation is found between peak delay and second syllable duration.

These late-peaking accents also appear 'almost exclusively' in utterance-initial position (p.290). One may therefore whether edge-marking tonal events, such as the Low IP-initial boundary tone that would be expected in utterance-initial position, may drive the realisation of the H* peak rightwards in order to avoid tonal crowding. Bishop rules this out, since she does not find delayed peaks in IP-initial, utterance-medial position. However, for closely related Dalabon, Fletcher (2014) reports a 'late F0 elbow' that frequently occurs following an IP-initial Low tone.

To summarise, in all varieties of BKW, pitch accents may be realised on the leftmost syllable of a morphological constituent. In the Manyallaluk Mayali dialect, there is the further possibility of an optional pitch accent on the penultimate syllable of the word. Pitch accents may appear on the second-from-left syllable of a morphological constituent if it is either the first of an utterance, or contains a heavy peninitial syllable.

There is some controversy, however, about whether facts such as these constitute sufficient evidence for metrical feet. As Mansfield (2023) argues, the level of prosodic constituency termed the 'foot' in Bishop (2002) is "quite unlike conventional metrical feet, in that it can contain any number of syllables, and it encompasses simple words of any length"; for that reason, he instead analyses this unit as a prosodic word (Mansfield 2023, p. 423). As a corollary, he folds the unit Bishop calls a phonological word into the phonological phrase. Note in any case that Bishop & Fletcher (2005) omit the phonological word from the Bininj Kunwok prosodic hierarchy.

The word-final foot in Manyallaluk Mayali, on the other hand, is fairly conventional in being aligned to the word-penultimate syllable, and therefore binary and trochaic. In other contexts, this has been taken as evidence for metrical feet. Ngalakgan, for example, also shows both stem-initial and word-penultimate accent in complex words, which Baker (2008b, p. 91) interprets as the outcome of constraints that place accents on the first foot of a prosodic word and the final foot of a phonological phrase.

Mansfield (2023, p. 420), however, questions whether feet are necessary for the Ngalakgan data, instead suggesting that the intonational system directly targets prosodic word-initial and phonological phrase-penultimate syllables for accentuation. Following Gordon (2014), he argues that "accent at a distance of one syllable from the edge can be independently explained by tone-crowding phenomena, e.g. the need for the penultimate H* accent to give space for a [PhP]-final Low tone" (Mansfield 2023, p. 420).

For Manyallaluk Mayali, Bishop rejects the possibility that penultimate H* accents are involved in marking the right edge of a phrasal domain such as the IP, rather than the left edge of a phonological-word-final foot, by pointing to the fact that penultimate accent appears on IP-final and IP-initial tokens of the same word (p. 168). Nonetheless, all of the tokens she illustrates with an annotated F0 trace are still at least PhP-final, which is marked by a low boundary tone (Lp). It is still possible, therefore, that penultimate accent in Manyallaluk Mayali is directly

associated with a phrasal domain larger than the phonological word, namely the phonological phrase, just as Mansfield (2023) suggests for Ngalakgan.

The notion that pitch accents may associate directly to syllables at a specified distance from the edge of a phrase is illustrated especially starkly in languages displaying 'metrical incoherence'. As Gordon (2016) shows, even languages that have independent evidence for binary metrical feet — i.e., from segmental alternations and/or lexical stress — may have intonational pitch accents that are not anchored to the metrical parse. Nganasan, for instance, has binary feet parsed from left to right; foot-initial syllables bear lexical stress, and foot-initial consonants undergo pervasive lenition processes that do not apply within feet. Postlexical pitch accent, however, reliably appears on the rightmost non-final full vowel, even if this syllable is weak according to the metrical phonology. Gordon (2014, 2016) argues that pitch accents such as these must refer directly to the position of a tone bearing unit in a larger prosodic unit, without the mediation of foot structure (with which it is demonstrably in conflict). Similar claims have been advanced for West Greenlandic, whose pitch movements Arnhold (Arnhold 2014) analyses as entirely attributable to phrase-edge boundary tones. By this reasoning, a phrase-penultimate pitch accent alone is not enough grounds to justify the existence of feet in Manyallaluk Mayali.

In the next section, I summarise existing accounts of prosodic prominence in Wubuy. Those accounts make even stronger claims for the role of the foot (or its precursor, the 'stress group') than Bishop does for Bininj Kunwok. However, as we shall see, these claims warrant even greater scepticism, and I will argue that a 'footless' analysis such as Mansfield (2023) suggests for Bininj Kunwok and Ngalakgan is also appropriate for clarifying the nature of Wubuy prosody as presented in those works.

6.2.6 Previous work on Wubuy prosody

The main existing knowledge we have on Wubuy prosody comes from two sources: some brief comments in Heath (1984) and a chapter by Hore (1981). These works, both four decades old now, predate many of the most crucial developments in theories of linguistic prosody, such as Metrical Stress Theory (e.g. Halle & Vergnaud 1987; Hayes 1995) and Autosegmental-Metrical approaches to intonation (Ladd 2008; Pierrehumbert 1980). They are thus limited by the theoretical and descriptive devices that were available at the time; nonetheless, Hore's work in particular continues to feature in phonological typologies of stress (e.g. Kager 2007).

As we have seen in previous chapters, Heath's (1984) grammar devotes a great deal of attention to phonological alternations. He is notably circumspect about word-level prominence, however, suggesting that any apparent stress phenomena would be better handled as intonation — that is, as post-lexical prominence: "It is difficult to analyse the pitch contours at word level, because word-level patterns are commonly distorted or overridden by higher-level intonation patterns"

(p. 32). He does offer some fairly tentative and informal comments on where stress is likely to be found, although it is not obvious from the description what 'stress' means here, as he identifies certain syllables as receiving 'high intonation and stress' or 'high pitch and a little stress' (p. 32). Those general observations — presumably based on auditory impressions of words in citation forms — include the tendency for long vowels to be accented, including word-finally, and the tendency for penultimate accent when 'the last few syllables are short'. In the latter case, 'a pattern of alternating high-pitched, stressed vowels in even-numbered syllables (right-to-left) may result'.

This claim that stress alternates throughout the word is where Heath's account most closely resembles the kind of stress analysis more typical of the Australianist descriptive work of the period. However, none of the examples he presents could be taken as unambiguous evidence for it. The stress accents he annotates are consistent with an analysis that takes pitch accents to mark the edges of domains — they are all aligned either to the right edge of the grammatical word or the left edge of some word-internal morphological constituent. The stress pattern in *ána-ɭáku* 'the honey', for instance, is presented as exemplifying penultimate, alternating stress, but it could just as well be taken as an example of a pitch accent aligned to the left edge of the word (and/or IP, if it was uttered in isolation) plus an accent aligned to the left edge of the stem. Heath diverges in this respect from contemporaneous descriptive grammars, which, as discussed above, typically see metrical structure as dependent on morphological structure.

The only effects of morphological structure on stress assignment Heath explicitly recognises are the following. Firstly, he notes that reduplicative words (especially nouns) tend to have stress accents on the initial syllables of both 'reduplicative segments' (i.e., reduplicant and base), and makes reference to feet in order to explain them. Though here again, 'feet' are only vaguely defined as 'units of two or three syllables' and he invokes them exclusively for reduplicative prosody. Note also that this description explicitly does not apply to productively reduplicated stems, and the only examples offered contain disyllabic and trisyllabic reduplicants — it is unclear whether he regards productively reduplicated stems, or those with monosyllabic reduplicants, in the same terms.

Secondly, he notes that monosyllabic suffixes tend to be prosodically incorporated into the stem they attach to, causing the word to receive penultimate accent, as in *á-ɭakú-wuy* 'to honey' (cf. *ána-ɭáku* above). Disyllabic suffixes, on the other hand, have 'little effect' on stem prosody, and take stress on their initial syllable, as in *a-ɭáku-jíɲuŋ* 'of honey'. We can perhaps interpret this to mean that these suffixes, having sufficient material to form a binary foot, can be treated as their own prosodic word. However, we could also interpret the stress on the suffix as, again, an instance of prominence on the penultimate syllable of the word.

A more extensive phonological treatment of Wubuy word-level prosody is that of Hore (1981),

whose period of fieldwork overlapped with Heath's. Hore posits a series of 19 transformational rules in the style of *SPE*, some of which apply cyclically, and all of which apply after the 50 segmental phonological rules in Heath (1984). Again, 'stress' here is determined on the basis of auditory impressions, presumably of words uttered in isolation.

Although Hore's analysis apparently comes with a greater degree of formal explicitness than Heath's, it is difficult to assess whether it thereby brings greater insight. His Rule 12, included as (339) here, exemplifies several of the difficulties.

(339) S S S [trisyllabic] Stress Group Medial Stress (Hore 1981, p. 44)[4]

$$\begin{array}{ccccccccccc} [-\text{seg}] & \begin{bmatrix} \text{S} \\ \alpha \text{ stress} \end{bmatrix} & \text{S} & \text{S} & [-\text{seg}] & \longrightarrow & 1 & 2 & \begin{bmatrix} 3 \\ \alpha \text{ stress} \end{bmatrix} & 4 & 5 \\ 1 & 2 & 3 & 4 & 5 & & & & & & \end{array}$$

where α = 1 or 2 [primary or secondary, respectively].

This rule does *not* apply if the stress group has the form

$C_1\ V_1\ C_2\ V_2\ C_3\ V_3$ (C) (C)

and one of the following conditions holds:

a. V_2 = i and C_3 = y.

b. $V_1 = V_2 = V_3$ and C_2 = rr.

c. $V_1 = V_2 = V_3$ = a or u, and C_2 = g, ngg, or ng, unless C_2 = g and C_3 = y, w, rr, or r.

As the rule in (339) alludes to, Hore's analysis — published before metrical phonology became a standard part of phonological theory — assumes a hierarchy of prosodic constituency that gathers syllables into 'stress groups', like some other work published in the same period, particularly on Yolŋu languages of Arnhem Land (e.g. Waters 1979; Wood 1978; also Wilkinson 2012a). A fairly nebulous theoretical construct, the stress group can consist of anywhere from one to four syllables (Hore 1981, p. 21). Only one syllable in the stress group can receive a stress accent; the location of this stressed syllable is either initial or peninitial and depends on, but is not completely determined by, the group's size. Many of his stress rules govern the inclusion or exclusion of a given syllable in a stress group, and a stress group's boundaries may be shifted many times throughout the course of a derivation. On p. 31 he describes, for instance, the rules required to take a five-syllable word with an initial S.SS.SS stress group segmentation (where S = syllable), and transform it into an SSS.SS one, and then finally SS.SSS.

[4]Note that 'S' = 'syllable' and '[−seg]' = 'major morphological boundary'

Besides having involvement in stress assignment, Hore suggests that stress groups may be phonetically delimited by 'barely detectable' pause, but this is only his 'feeling' and it 'should probably not carry too much weight as evidence' (p. 13) — this observation is perhaps supported by the findings of Bundgaard-Nielsen & Baker (2020), discussed below.

The rule in (339) is motivated by the need to explain the observation that trisyllabic stress groups 'are normally stressed on the middle syllable, but sometimes the stress is on the first syllable' (Hore 1981, p. 26). The 'sometimes' in this formulation reveals an ambiguity found throughout the analysis. While many examples of individual words, transcribed for stress, are provided throughout, we are given no information about how the data were gathered. We may presume they were uttered in isolation in a field elicitation setting, but we do not know if they were uttered once or several times, which crucially affects whether we interpret any variability as applying to types or tokens — except in a few cases where variant forms are explicitly offered, such as *ŋàla:líki~ŋalá:liki* 'turtle' (p.38). So, when a given token of one trisyllabic monomorphemic word has initial prominence and a token of another trisyllabic word has peninitial prominence, we cannot determine whether the difference represents arbitrary lexical conditioning, predictable segmental phonological conditioning (as the formulation of the rule appears to entail), or if both initial and peninitial prominence are options available to any trisyllabic monomorphemic word, such that both would be expected to arise in a corpus of sufficient size (see de Lacy 2014). Indeed, several of the trisyllabic forms to which Hore assigns medial stress are extremely common in my elicitation data, such as *ŋucica* 'fish' and *wuɻuku* 'billabong', and (based on inspection of the pitch trace in Praat) both usually appear with a pitch peak aligned to the initial syllable. In fact, the difference between lexical conditioning and phonological conditioning appears to be conflated in rules like (339), whose exceptions (and the exceptions to the exceptions) make reference to an essentially random assortment of phonological segments without any features in common; the net result is simply to exclude a small handful of lexical items from the rule. And, further, without knowing the broader discourse context, we cannot know whether the prominences reported reflect the word's position within larger intonational prosodic constituents, particularly if there are pitch events at the edges of phrase boundaries (Gordon 2014; Jun & Fletcher 2014). As we saw for Bininj Kunwok words in utterance-initial position, and Dalabon words in IP-initial position, these edges may be associated with a delayed F0 peak (Bishop 2002; Fletcher 2014). Of intonational phenomena, Hore claims that 'we can ignore these as in general they simply produce an intensified stress on a syllable that was stressed in any case' (p. 8), with the exception of two special intonational tunes: interjection intonation, and 'linear lengthening intonation' (Mailhammer & Caudal 2019), both of which lend high pitch and extended duration to the word-final syllable.

The number and complexity of the rules in Hore (1981) can thus be seen as a consequence

of the lack of precision in the atoms that rules can act upon, and the ambiguity in the nature of the data motivating the rules. This complexity tends to obscure those parts of the analysis that deserve merit, such as the acknowledgement of a crucial link between morphological structure and prosodic alignment, including the recognition that not all morphological junctures induce metrical reset, but only major ones — i.e., those between prefix and stem (noun or verb), or between incorporated noun and inflected verb root, but not between, say, verb root and tense-aspect inflectional suffix (pp. 50–54). Witness, for example, the shifts in stress group boundaries that a word such as *na-purumpuruŋa* 'the Pleiades star cluster [Masc.Top]' (one of Hore's examples) must go through in the course of a derivation. This begins with an initial segmentation into *na.purumpuruŋa*, in which the stress group boundary follows morphological structure, and is indicated by a period. But later rules (e.g rules 6 and 7) resegment this to yield *napuru.mpuruŋa*. Various rules then apply to assign and re-assign stress; ultimately, the medial syllables in both stress groups receive stress, yielding *napùru.mpurúŋa*. While the end result of all this is to assign stress to the first syllable of the root, and the penultimate syllable of the root/word, this is not the reason for doing so in Hore's analysis. Instead, it is due to a general rule of trisyllabic stress groups. The generalisation that stress aligns with morphological constituents is therefore lost in the various derivational procedures between the morphologically segmented input and the prosodified output.

However, from observing the location of stress in nearly all of Hore's examples, it seems clear that, like Heath, he perceives prominences at the left edges of major word-internal morphological constituents, and at the right edge of the word (and also, presumably the IP) — that is, on the final syllable if it contains a long vowel, and on the penultimate otherwise. In that sense, the examples appear consistent with an analysis very similar to that of Bishop (2002) for Manyallaluk Mayali, so long as we assume that what Hore describes as (word-level) stress is actually pitch accent, and what he descibes as a stress group is something like an unbounded foot (or prosodic word, as Mansfield 2023 argues). Indeed, Jepson (2019, p. 15) points out that, in descriptions of stress in Yolŋu languages, stress groups practically always serve the purpose of aligning stress to the left edge of morphological domains, and are therefore practically equivalent to the left-headed unbounded feet of Bishop (2002), which can be any length but are constrained by morphological structure.

An intonational treatment of Hore's Wubuy data would be further strengthened if we allow that, when a Wubuy word is IP-initial (which in most cases would begin with a %L boundary tone), a pitch accent associated to the initial syllable may be delayed due to tonal crowding, and instead appear on the second syllable, as also reported for Bininj Kunwok (Bishop 2002) and Dalabon (Fletcher 2014). This would go some way to explaining the variability in trisyllabic monomorphemic words between initial and medial stress realisations; although medial stress

(*qua* pitch accent) in trisyllabic IP-unique tokens may just as plausibly be treated as marking the right edge of the phrase. By my reckoning, only 4 examples out of the 123 presented in Hore (1981) present obstacles to this kind of analysis, in that they contain stressed syllables that cannot be treated as demarcating morphological domains.

There have been two more recent lines of enquiry into Wubuy prosody, though neither have attempted to describe the intonational phonology in any detail. Firstly, there have been some attempts to determine the phonetic correlates of stress, based on controlled elicitation of trisyllabic nouns in carrier phrases. Baker, Bundgaard-Nielsen, Babinski et al. (2019) test the hypothesis that stress is initial unless shifted rightward by a closed (and therefore moraically heavy) syllable. The results are fairly inconclusive: raw syllable position is the strongest predictor of duration, intensity and F0, all of which are higher in the second syllable. However, these correlates (especially duration) are greater in a closed second syllable (which should therefore be stressed) than an open second syllable, which the authors interpret as evidence that the stress analysis may not be completely 'illusory'. However, they also do not rule out the possibility that the effects they find are due to intonational boundary tones, rather than lexical stress as such. Baker, Bundgaard-Nielsen, Nyhuis et al. (2020) also examine trisyllabic nouns, but their data also includes tokens with long vowels at different positions. They test the effects on duration of a phonological analysis which assumes a quantity-sensitive lexical stress system, where a syllable may be moraically heavy either by dint of a long vowel or a coda heterorganic to the following syllable, as in Ngalakgan (Baker 2008b). They found greater duration in initial syllables if they were heavy, and in penultimate syllables otherwise; there remains a strong effect of domain-final lengthening, with vowel durations generally increasing throughout the word. This study also finds some support for lengthening of the entire VC 'interval' (Steriade 2023) of stressed vowels, otherwise known as post-tonic lengthening. Both of these studies attempted to control for intonation by embedding target items in a carrier phrase, and manipulating the position of the item within the carrier phrase, but most of the productions are uttered with a separate intonation contour for each word in the phrase anyway. Thus, without further investigation of the intonational phonology, we remain uncertain about the degree to which boundary tones may produce pitch movements that could otherwise be attributed to the presence of a stressed syllable.

The next line of enquiry is Bundgaard-Nielsen & Baker's (2020) experimental work on intra-word pausing. They present evidence that Wubuy speakers often pause at certain word-internal morphological junctures, such as the boundary between pronominal prefix and stem, and between incorporated noun and verb root, as well as either side of *ŋu-* in complex stems, similar to findings reported for the related language Dalabon (Evans et al. 2008). They term these boundaries 'licit' for pausing. Speakers do not pause within these morphs, even when they

contain a historical boundary, as between the reduplicant and base in frozen reduplicated roots, or between a historical coverb and inflecting verb; they also do not pause between parts of pronominal prefixes that may be analysed as seperable, nor between a verb root and its tense-aspect inflectional suffix. All such locations are termed 'illicit'. The authors then conducted a two-alternate forced choice word preference study, where the stimuli contained complex words either without pause, or with periods of silence artificially inserted, at either a licit or an illicit boundary. The results indicated that participants prefer both no pause and legal pause stimuli over illegal pauses, but, rather surprisingly, they show no significant preference between no pause and legal pauses. Further, certain speakers actually prefer words with legal pauses at some positions over the same words with no pause. The authors also note that in speakers' productions of complex words with internal pauses, each unit separated by a pause appears to have its own intonation contour, possibly suggesting that a grammatical word may contain two Intonational Phrases. Their results provide rather striking evidence that prosodic organisation is sensitive to the morphological structure of complex verbs, and that grammatical words may consist of multiple planning units in speech processing. This work also raises the question of whether intonational effects of morphological structure can be observed even in the absence of pausing. Indeed, the fact that sub-word morphological constituents can be distinct intonational constituents, bounded by pause, gives additional strength to the notion that the word-internal prominences that have been regarded as lexical stress can be explained by intonational organisation alone, without recourse to a lexical stress system below that.

As the preceding discussion has shown, the work so far undertaken to unravel Wubuy lexical stress, while suggestive, is limited by our lack of knowledge about the intonational phonology. Indeed, if the locations of prosodic prominence in complex words are where these works claim, whether based on researchers' impressions or instrumental analysis, we are left with a picture that is still broadly compatible with that of Bininj Kunwok as reanalysed by Mansfield (2023). That is, prosodic constituents align most of the time with morphological ones, and prominence at the edges of these constituents can be analysed by means of intonational pitch accents and boundary tones. There have been some indications that metrically heavy syllables can attract prominence away from these edges, but the true extent of weight sensitivity remains poorly understood.

6.3 Aims, methods and materials

6.3.1 Hypotheses concerning verbal prosody in Wubuy

Considering the above discussion of prosodic organisation in Wubuy and other, related languages, two questions emerge as particularly salient:

1. Is the alignment of pitch accent in Wubuy sensitive to word-internal morphological structure? If so, which positions in this structure are available for pitch accent?

2. Is there evidence for metrical feet? And, if so, is footing sensitive to syllable weight?

To answer these questions, I conducted an Autosegmental-Metrical analysis on a corpus of Wubuy texts, described in §6.3.2, which I annotated using a AM-style scheme outlined in §6.3.3. In §6.4, I present the main intonational tunes identified in the corpus. The main results bearing on the above research questions are presented in §6.5, on tune-text alignment. That section begins with an overview of the prosodic characteristics of the verbs in the corpus, including the distribution of pitch accents within the verb. The following sections then analyse each of these verb-internal landing sites for pitch accents in detail. The implications of all of these findings are discussed in §6.6.

6.3.2 Corpus materials and transcription

The materials used for this study consist of three unscripted texts, each from a different speaker. Two of these were recorded by Jeffrey Heath in the late 1970s, and their transcriptions were published in Heath (1980). They were transcribed more recently in ELAN (Brugman & Russell 2004) by Kate Horrack, with additional corrections by myself. The third text is a recording of speaker JN that I made in 2017. All of the texts are discourses on collecting food and materials from the bush.

After transcribing the texts at the utterance level in ELAN using a modified version of the community orthography, I used the BAS web services (Kisler et al. 2017) to tokenise the transcription and convert it into the X-SAMPA alphabet, which was then used for the forced-alignment and automatic syllabification procedure. The resulting hierarchical annotations (with linked segment, syllable, word and utterance levels) were converted into an emuR database (Winkelmann, Harrington et al. 2017) in the R programming environment (R Core Team 2024). All boundaries were hand-corrected using the emuR WebApp. However, the images of waveform and F0 trace presented in this chapter were made by exporting the corpus to Praat and saving pictures from there.

6.3.3 AM annotation procedure

With the corpus annotated at the segment level, I annotated tonal events using a simple AM-style (aka ToBI) schema adapted from work on the intonational phonology of the Gunwinyguan languages Dalabon (Fletcher 2014; Ross 2011) and Bininj Kunwok (Bishop 2002). This schema is shown in Table 6.1. Like many languages whose intonational phonology is still relatively

Pitch accents		Boundary tones	
H*	high target	%L	default left edge of Intonational Phrase (IP)
H::	beginning of high plateau / 'linear lengthening' on phrase-final syllable	%H	for IPs beginning with high, level pitch (ie. no observable rise)
!H	downstepped high	%	left edge of all other IPs
^H	upstepped high	L%	default right edge of IP
<	delayed peak	H%	end of high plateau and IP
		HL%	late fall IP boundary tone ('listing' or 'continuation' intonation)
		^H%	rise beginning from final H* to end of IP
		LH%	short rise at right edge of IP
		La	default right boundary of Accentual Phrase (AP)
		Ha	high boundary of AP

Table 6.1: AM labelling scheme

understudied (as can be seen in most of the papers in Jun 2014a), I transcribe only the 'Tones' part of the ToBI system, and assume that break indices can be straightforwardly (that is, redundantly) derived from the type of right boundary tone, as in many other AM analyses (Beckman et al. 2005). That being so, in parts of the analysis to follow I make use of a break index schema that assigns the IP, AP, and PrWd the break index values {4, 3, 2}, respectively.

When annotating pitch events, I tried to rely as much as possible on the visual f0 trace accompanying the waveform and transcription, taking micro-prosody or small perturbations into account. Sometimes, particularly in quieter parts, the software could not determine the pitch; if I nonetheless remained convinced that there was some pitch movement present, I transcribed based on comparing neighbouring vowels by ear. While I attempted to ignore my own intuitions to the extent possible, it should be borne in mind that my choices of where to annotate prominence will inevitably be coloured by my own native speaker biases (since I am not a native speaker of Wubuy).

I propose that Wubuy has at least two tonally marked levels of intonational constituency: the Intonational Phrase (IP) and the Accentual Phrase (AP). The IP sits at the top of the hierarchy, as diagrammed in (340). It is delimited by boundary tones at its right edge and usually also its

left, where there is also a resetting of the speaker's pitch range. There is normally a substantial pause (and intake of breath) between IPs, and where the IP is marked by a low boundary tone (%L or L%), creaky phonation. The AP is identified by a smaller degree of pitch range reset, but either no pause or, rarely, a very brief one. It has two right boundary tones: Ha or La.

(340)
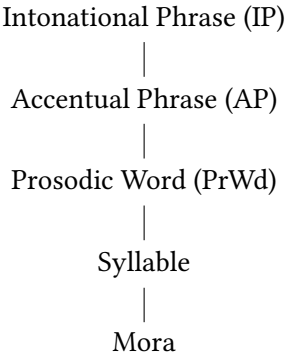

The difference between the AP and IP units can be observed in Figure 6.2, where there is a small fall from the H*-accented penultimate syllable of the first grammatical word in the utterance towards the La boundary. The pitch trace then rises back up from the La boundary to the next H* and stays fairly level (ignoring some perturbations induced by the stop release in the second syllable) until the final H*, at which point it falls to the IP-final L% boundary. As can be seen, the magnitude of the fall at the IP boundary is much greater than the fall to the AP boundary.

In this, I follow the practice of others investigating the intonational systems of Australian languages, who also note two levels of prosodic organisation based on similar junctural phenomena. For instance, Fletcher (2014) states, for Dalabon, that 'a full Intonational Phrase boundary … is realised with a higher level of juncture compared to an Accentual Phrase' (p.257), and that the 'AP-final La tone does not fall as low as the L% tone' (p.259). Fletcher also notes that what she calls an Accentual Phrase is basically equivalent to the unit labelled Phonological Phrase in AM studies of Bininj Kunwok (e.g. Bishop 2002; Bishop & Fletcher 2005), both in its position in the hierarchy and in the tonal events associated with its edges. The same equivalence applies to the Wubuy AP proposed here.

Below the IP and AP sits the prosodic word (PrWd), but this is not specifically annotated in the corpus. Indeed, understanding the nature of prosodic constituency below the AP is one of the goals of this study. In §6.5, I infer the existence of the PrWd based on the alignment of pitch accents to morphological constituents within the verb, just as Bishop (2002) does for the unbounded foot in Bininj Kunwok. As discussed above, the foot is a controversial unit of analysis in these languages, and the results of this chapter do not provide especially convincing

evidence for the utility of metrical feet, as will be discussed at length below, particularly in §6.5.5. For that reason, there is no level between PrWd and the syllable in the hierarchy in (340).

In proposing this hierarchy, I assume, as Fletcher (2014) does for Dalabon, that an IP consists of one or more APs, such that an IP boundary is necessarily also an AP boundary. In other words, no part of an utterance may be contained in an IP if it is not also contained in an AP. This decision has implications for the interpretation of unaccented material at the beginning or end of an IP. Discourse connectives like *aɻapa~apa* 'now; and then' or *ŋa* 'and' are frequently unaccented; the former usually precedes an accented content word, the latter usually follows one. IP-final unaccented material nearly always appears following a Ha boundary accompanied by extreme lengthening of the AP-final syllable (as in Figure 6.7); this constitutes a specific tune type, as discussed in §6.4. Thus in such cases, there is normally a tonal event marking the AP boundary, and I interpret the material following that boundary as also constituting its own AP, though it happens to be unaccented. However, for an unaccented word stretching from the %L boundary marking the beginning of the IP to the following word, the pitch stays low throughout, as in Figure 6.1. I have treated unaccented material like this as contained within the same AP as the following word. Both of these strategies could be given alternative treatments, for instance by treating the unaccented stretches as 'prosodic adjuncts'; I leave these questions to future work.

6.3.4 Querying and analysis

The tonally transcribed materials were then used as the basis for generating a set of tabular data containing verb words in rows, with variables of interest in the columns. The variables largely concern segmental and morphological properties of the word, alignment of the word to phrase boundaries, and the presence or absence of pitch accents in various locations of the word, which are specified in relation to morphological constituents — i.e., 'prefix-initial syllable', 'stem-initial syllable', 'stem-peninitial syllable', etc. There are also variables for speaker, text, recording, start and end times, and some other variables that relate the data points to their location in the emuR database but which are not of any research interest. This tabular data was then queried and visualised in R (R Core Team 2024), providing the basis for the quantitative analysis of tune-text alignment presented in §6.5.

6.4 Results: Intonational tune types

As we saw in Table 6.1, Wubuy has a very small inventory of pitch accents, consisting of just a monotonal H* target, which may occasionally show local downstep or upstep. I do not find any evidence in the corpus for bitonal pitch accents, such as the L+H* rising accent proposed for

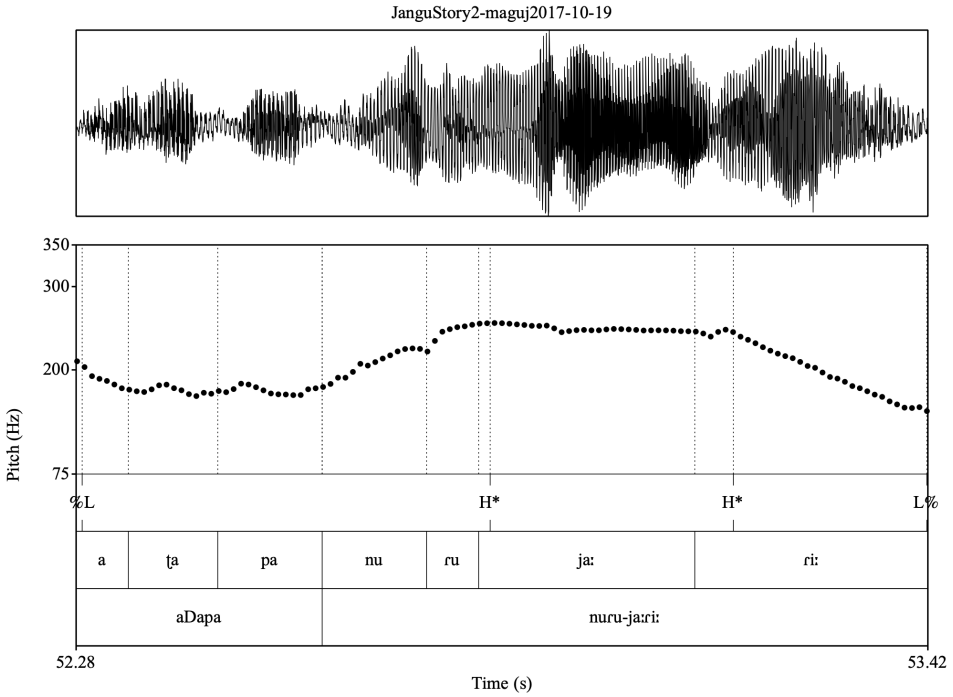

Figure 6.1: Example of unaccented material before the verb.

Dalabon (Fletcher 2014, p. 259) and Bininj Kunwok (2002), or the H*+L falling accent proposed for Djambarrpuyŋu (Jepson 2019). The tonal inventory in Table 6.1 satisfies two of Jun's (Jun 2014b) diagnostics of low 'macro-rhythmicity': the number of phrase-medial tones is small (effectively H*, Ha and La), and the most common accent type — in fact, the only one — is level (H*), rather than falling or rising.

Broadly in line with Heath's (1984, pp. 602–604) earlier comments (see above), we can distinguish the tune types in the corpus based largely on what happens to the pitch as we get towards the right edge of the Intonational Phrase. In the 'falling' tune, the pitch range compresses (i.e. the pitch topline declines) throughout the phrase. In the 'high monotone' tune, the pitch is kept high and level throughout the phrase, normally with dramatic lengthening of the phrase-final syllable. These two basic tune types account for most of the data; see Table 6.2.

The falling tune has several variants. The most common variant we find in the corpus is equivalent to the 'hat' contour reported in neighbouring languages Dalabon (Ross 2011) and Bininj Gun-Wok (Bishop 2002), and may be considered the basic declarative tune.[5] This contour begins at a low boundary (%L), rises to the first accent (H*), sustains fairly high for some time

[5]This has been reported for much more distant Australian languages as well, such as Warlpiri (Butcher & Harrington 2003).

Tune	Tokens		Right boundary	Tokens	
falling	163	69.66%	L%	134	82.21%
			HL%	18	11.04%
			LH%	10	6.13%
			^H%	1	0.61%
high monotone	41	17.52%	L%	26	63.41%
			H%	15	36.59%
complex	16	6.84%			
rise	12	5.13%			
(disfluent)	2	0.85%			

Table 6.2: Frequency table for Intonational Phrase tune types in the corpus. The 'falling' and 'sustained high' categories each have subtypes; percentages reflect their within-category proportions.

	mean	sd	min	Q1	median	Q3	max
Pitch accents in IP	2.88	2.04	0	2	2	4	17
APs in IP	1.56	0.73	1	1	1	2	5
Grammatical words in IP	2.12	1.58	1	1	2	3	13
IP-internal Ha boundaries	0.17	0.4	0	0	0	0	2
IP-internal La boundaries	0.39	0.65	0	0	0	1	3
Verbs in IP	0.85	0.75	0	0	1	1	4

Table 6.3: Summary statistics for pitch accents, grammatical words and Accentual Phrases per Intonational Phrase ($N = 233$) in the corpus.

before falling to a lower pitch (L% or La), often with creaky or irregular phonation. H* accents are annotated on the peak of the first rise and the beginning of final fall — the 'elbows' — and sometimes additional peaks are annotated in between. The pitch range compresses throughout the IP — as a result, the 'plateau' phase of the hat is not necessarily strictly flat but slopes downwards.

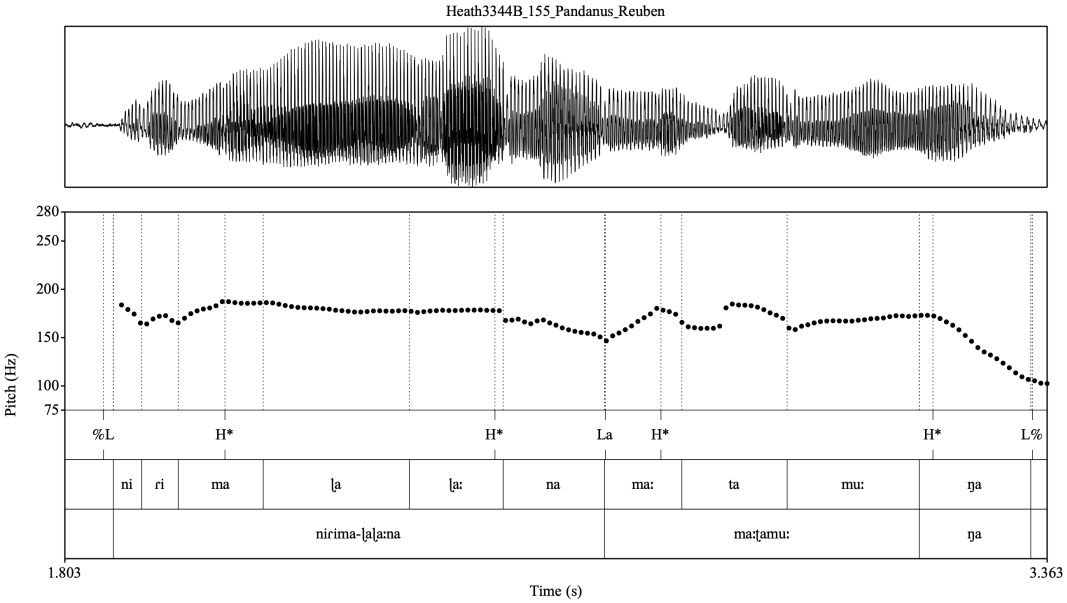

Figure 6.2: Declarative 'hat' tune in two APs. Note the shallower fall to the La boundary as opposed to the L% boundary.

In a common variation on this hat tune, there is no plateau as such, only one H* early in the phrase, followed by lowering towards the L% boundary (Figure 6.3). Another variant has an accent on the first syllable of the phrase, and therefore lacks the more typical rise from the %L boundary tone. This gives us, so far, three variations on the basic declarative tune: simple fall (HL), simple peak (LHL), and the hat (LHHL). All three are united in having a fall from a H* accent to an L tone at the boundary. These basic shapes are essentially in line with the basic declarative tunes described by Heath (1984, p. 603), although he frames the sharpness of the final fall towards the low tone of the right boundary as a dimension of continuous variability, and does not employ a notion of an accentual target.

There are some further variants which observe the same general F0 trajectory throughout the phrase as the falling tune, but have a more complex final boundary tone. One variant is that which ends with a LH% boundary, as Heath also notes. This tune is characterised by the same general fall in pitch towards the right edge of the phrase as we see in the basic declarative,

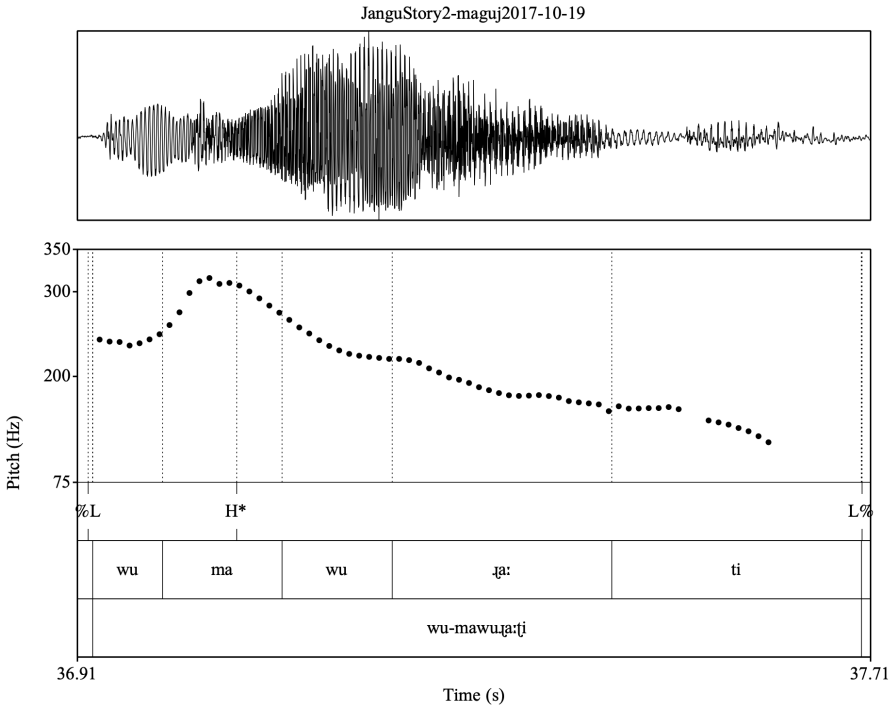

Figure 6.3: Simple rise-fall (LHL) pattern on *wu-mawuɻaːti* 'it's cold'.

except for a very late final rise during the final syllable; see Figures 6.4 and 6.5. The other is that which ends with a HL% boundary. It is commonly found on non-final items in a list, and as such signals continuation or incompleteness. This is also a kind of fall, but with a later elbow, and consequently a higher final F0 than the L% boundary — cf. Heath (1984, p. 602), who also notes this tune 'having a generally higher baseline pitch than in unmarked strings'.

There is another type of rise in the corpus, associated with a final ^H%. An example is shown in Figure 6.6. This is characterised by a longer rise beginning near the phrase-final H* and ending at the phrase boundary, reaching into the upper limit of the speaker's pitch range. Unlike the bitonal LH% boundary, the ^H% boundary does not induce a fall before the rise, since it does not involve an L target. Both of these types of rising tone are attested in Dalabon (Fletcher 2014, 260ff) and Bininj Kunwok (Bishop 2002, p. 50).[6]

The other major tune type found in the corpus is that which has been identified as a major

[6]Bishop (2002), analyses such boundary tones as an Lp phonological phrase boundary (equivalent to the AP here) followed by a H% IP boundary, however she notes that such a configuration is in practice indistinguishable from a LH% boundary, since both the L and H targets are always aligned to the final syllable of the final phonological word in the IP in her corpus (as in Wubuy). She notes (p.50) that her Lp-H% analysis is less restrictive, since it allows the possibility of an Lp boundary aligned to the penultimate word, followed by an unaccented word, followed by a final rise to the H% boundary. This does not occur.

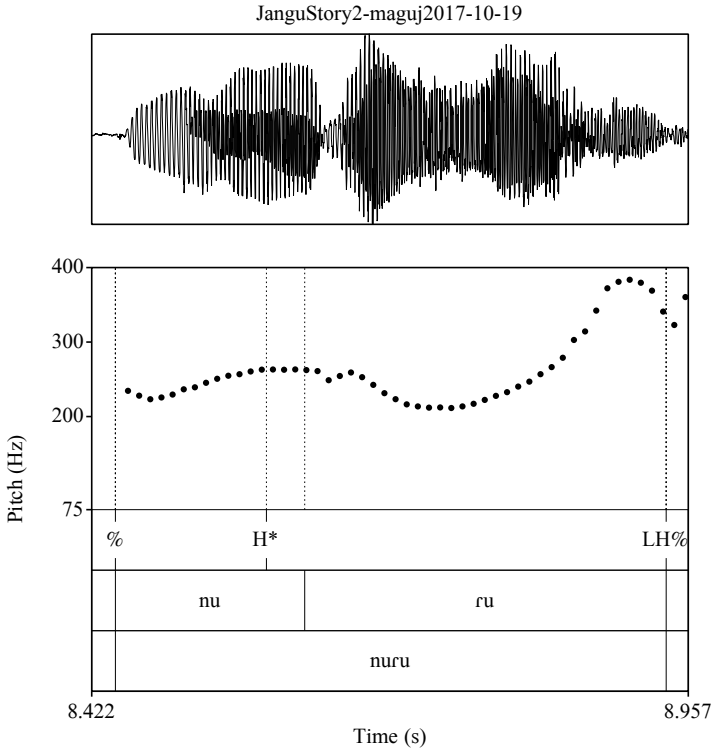

Figure 6.4: IP-unique token of *nuɾu* '1pl', with final late rise at LH% boundary.

type across a range of Australian languages, and has been variously termed the 'stylised high level intonation contour' (Bishop & Fletcher 2005, on BKW), 'linear lengthening intonation' (Mailhammer & Caudal 2019, on Iwaidja), 'narrative high monotone' (Sharpe 1972, on Alawa), and 'iconic lengthening' (Simard 2010, on Jaminjung). This tune differs from the falling tunes in maintaining a flat high or mid-level pitch throughout the phrase and dramatically lengthening the final syllable — for as long as 3 seconds — and maintaining a high, level pitch throughout. That is, we do not see the lowering of f0 or narrowing of its range as we do in other tune types, such as the hat tune. This lengthened syllable is notated as beginning with a H:: accent and ending with a H% or Ha boundary tone. Examples are shown in Figures 6.7 and 6.8.

In some cases of linear lengthening, as with the example in Figure 6.7, the high level tune really begins at the onset of the IP (with the high boundary tone %H), but one can still hear that certain syllables are relatively more accented than the others. In some tokens, this plateau is delivered in an almost song-like manner, where each syllable sounds accented, and each peak is evenly spaced, as though it were following an even musical pulse.

This linear lengthening tune is often followed immediately by a very short phrase with a rapid decline in pitch, as is also seen in Figure 6.8, where the linear lengthening on the first

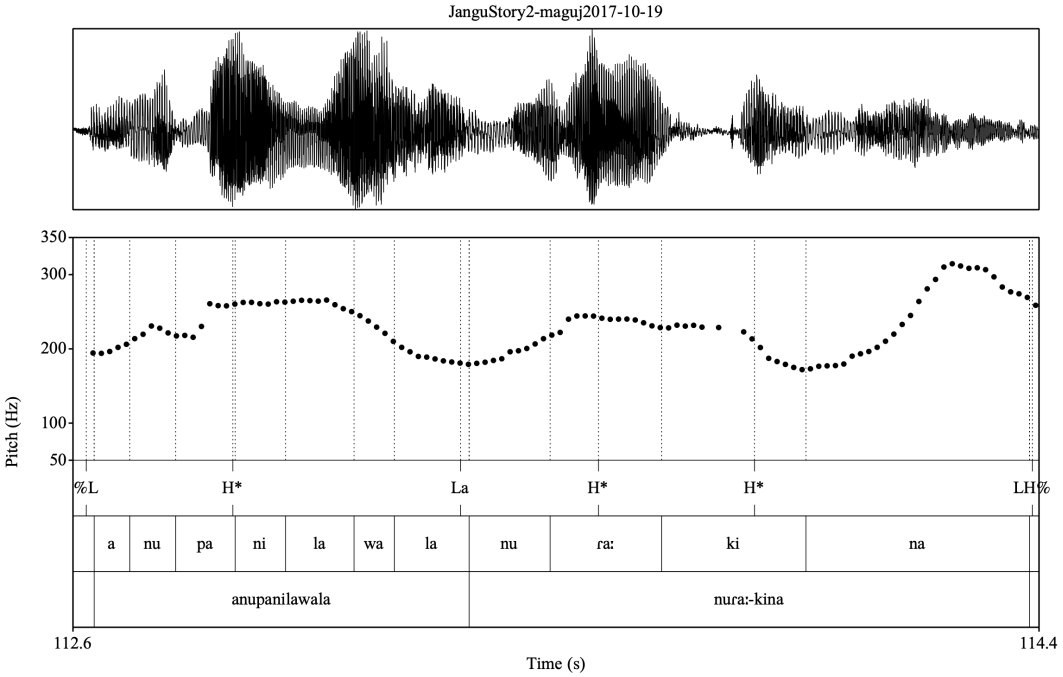

Figure 6.5: IP ending in late rise at LH% boundary: *anubanila-wala nuɾa:kina* 'from along there, we came back'

two verbs is followed by the verb *wiɾimi:-ɲkana*, spoken abruptly and with a rapid decline in pitch to the L% boundary. This second part, with the rapid fall, most often consists simply of the conjunction/connective *ŋa* 'and (then)', as also noted by Hore (1981, p. 8).[7] One may well ask questions of the nature of the relationship between the plateau phase and the fall phase — that is, do they form one coherent tune, or should they be treated as separated phrases, each with their own tune? Certainly, the high plateau and the short rapid falling tune are frequently found together, and so seem to form a tightly bound unit. Phonetically, there tends to be no pause or drawing of breath between the two. Semantically, this second event is a result or foregrounded event to the previous event, which shows the plateau tune, again suggesting a close bond. Note also that the linear lengthening verb is often reduplicated, while the following short falling IP is not. For these reasons, the combined use of reduplication and linear lengthening over a verb can be considered indicative of a kind of prolonged or repeated activity that forms the 'scene' or 'background' to the more discourse-new action expressed in the following short falling IP, as Heath (Heath 1984) points out. Formally and functionally, this is practically identical with other sustained high tunes reported in languages across the continent, such as Bininj Kunwok

[7] Heath does not specifically discuss the use of *ŋa* in this context, but several tokens appear in the annotated text which forms the basis of his discussion of discourse uses of intonation (Heath 1984, p. 592ff).

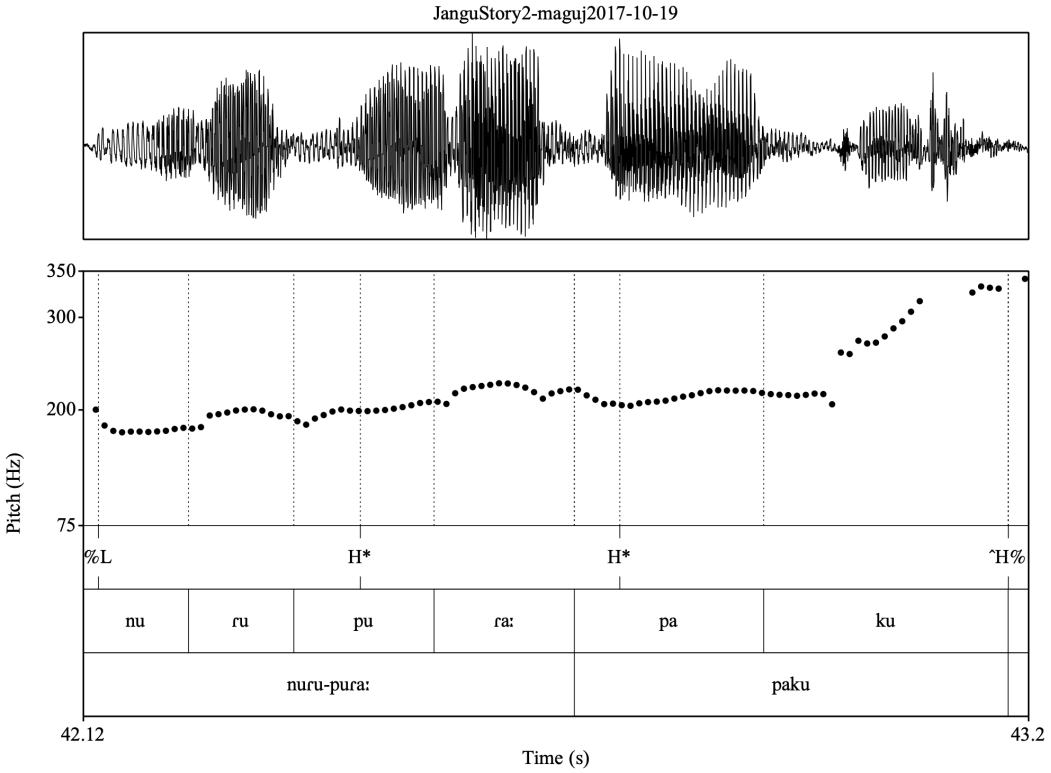

Figure 6.6: IP token showing upstepped high boundary (^H%)

(Bishop 2002; Bishop & Fletcher 2005). Given this tight bond, in the present corpus, I have treated the plateau with final fall as a variant of the basic plateau tune, and have notated the end of the plateau phase with a Ha boundary, meaning that the plateau phrase and the fall phase are contained with one IP.

A more complete account of Wubuy intonation would require a treatment of interrogatives, imperatives and quotations, as well as some investigation of the influence of information structure. Unfortunately, the texts analysed here do not provide us with the relevant examples, though quotative speech is claimed to involve a special upward shift of pitch range in Heath (1984, p. 602), which reflects my own impressions of the Wubuy speakers I have worked with. A treatment of these matters would benefit not only from other texts of different kinds, but also from more controlled experimentation (such as that of Jepson 2019).

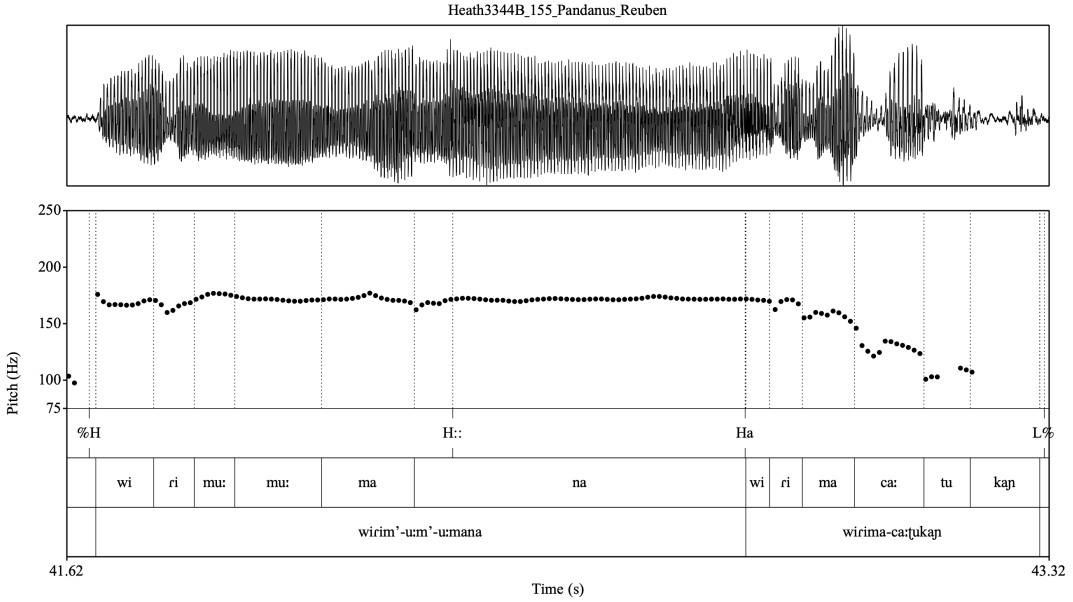

Figure 6.7: High plateau tune, followed by final AP with rapid fall to L% boundary.

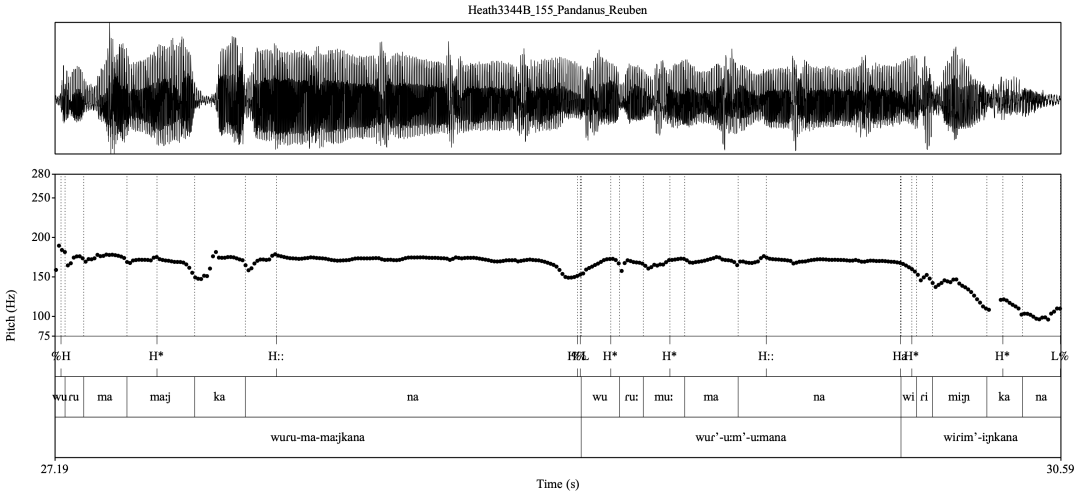

Figure 6.8: Example of IP consisting of 3 APs, which each consist of exactly one verb. The first two show the high plateau tune with final lengthening (linear lengthening intonation), while the third shows a rapid fall to the final L% boundary.

6.5 Results: Alignment of tune to text

The results in this section mostly concern the alignment of pitch accent to syllables within the word, which is our principal diagnostic of prominence, and therefore of what kinds of prosodic constituents should be posited for Wubuy.

Before presenting those results, however, we should first examine the alignment of verbal words to the major phrasal units identified in this study, namely the AP and IP. As Figure 6.9 shows, the vast majority (87%) of the verbal grammatical words in the corpus are aligned to at least an AP boundary at their right edge, as seen in the first two columns of the plot. Of these, the majority are aligned to at least an AP boundary at their left edge as well (the leftmost column). This shows that verbs are more likely to be found phrase-finally than phrase-initially, and are only rarely found phrase-medially (the 'neither' category in Figure 6.9). The heatmap contingency table in Figure 6.10 breaks this down further by distinguishing AP and IP boundaries, indicating the frequency with which those boundaries occur at either the left or right edge of the word, where lighter colours in the tiles indicate more tokens. As we see, the data is clustered around the top right corner, again reflecting the fact that the majority of the verb tokens are aligned with at least an AP, if not also an IP, at one of their edges. It should therefore be borne in mind that when we talk about verbs, we are describing a word type that is most often a major phrasal constituent unto itself.

Given that grammatical verbs are so often complete prosodic phrases of one kind or the other, it should not be altogether surprising that a given verb may have several pitch accents associated to it, as is characteristic of larger phrasal units. The histogram in Figure 6.11 plots the number of pitch accents per verb token. That is, we are not dealing with a system in which there is only one syllable in the word — a prosodic 'head', perhaps — which is available for accentuation in the intonational phonology. Rather, verbs can host multiple prominences — most frequently 2, but anywhere between 0 and 4 pitch accents occur. This number tends to increase with word length, as we can see in Figure 6.12, which plots the pitch accent count against word length in syllables. However, words with the same number of syllables show a wide range in the number of pitch accents they contain, as we see in the vertical spread of the data as we move along the x-axis in Figure 6.12.

All this is so far consistent with findings in other Gunwinyguan languages of the Top End such as Bininj Kunwok (Bishop 2002), where pitch accents can be realised at certain prominent positions within the word, but their realisation is in all cases optional. The plot in Figure 6.13 shows how often pitch accents appear at different morphotactically-defined locations in the verb. We see that nearly all pitch accents are associated with some edge. Those associated with the left edge are the 'pronominal prefix', 'stem-initial', 'stem-peninitial', and 'suffix-initial' categories,

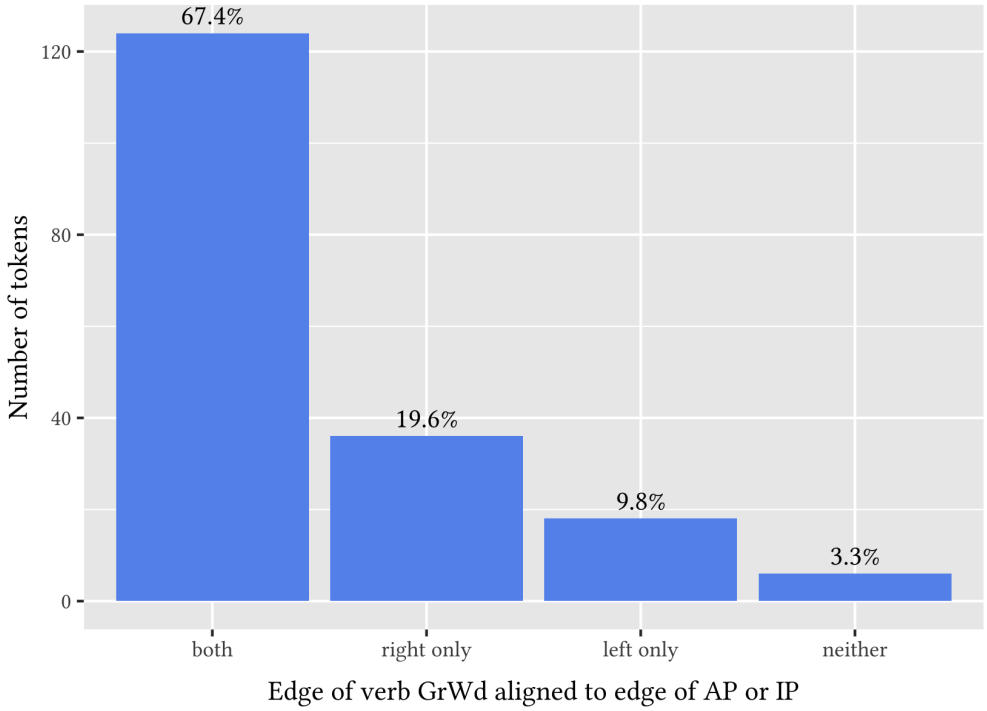

Figure 6.9: Alignment of grammatical verbs to either AP or IP.

while those at the right edge include the 'stem-penultimate' and 'stem-final' categories. A pitch accent token not at a constituent edge is assigned the 'elsewhere in stem' category — though as we shall see, the vast majority of these tokens turn out to have accents associated with the left edge of a morphological domain internal to the stem, such as the root in an incorporation construction or the base in a reduplicative construction. All of these pitch accent locations will be discussed further in the subsections to follow.

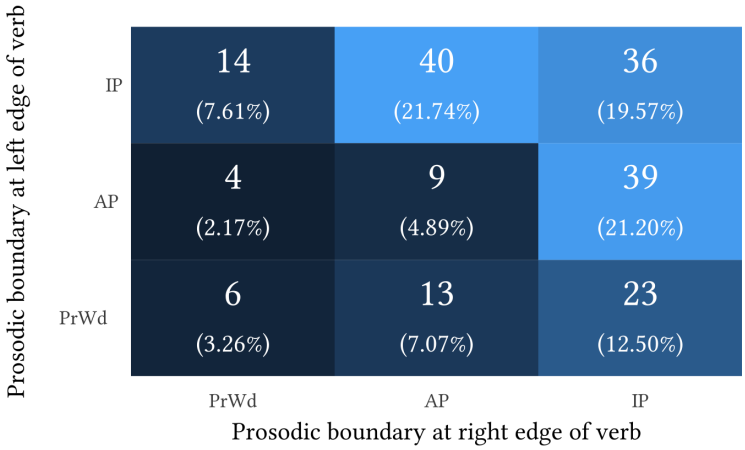

Figure 6.10: Heatmap visualisation of contingency table, showing alignment of grammatical verbs to prosodic constituents. Note that strict layering is assumed here, so that any IP boundary is necessarily also an AP boundary, and any AP boundary is necessarily also a PrWd boundary, but the reverse is not necessarily true.

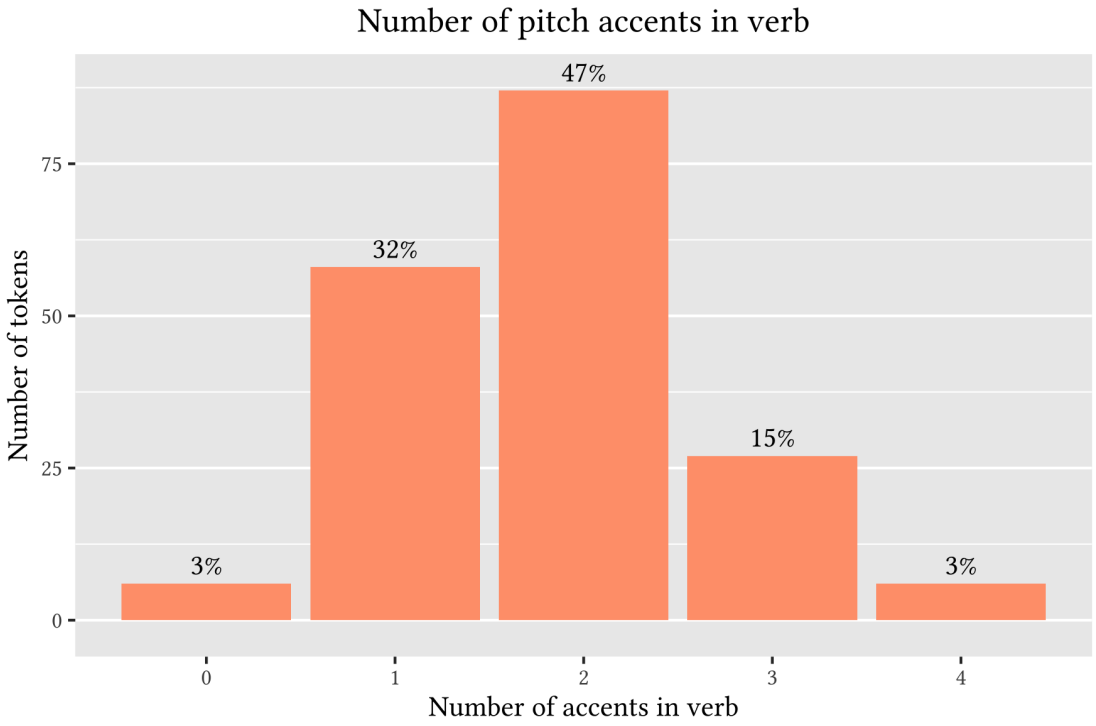

Figure 6.11: Number of pitch accents per grammatical word.

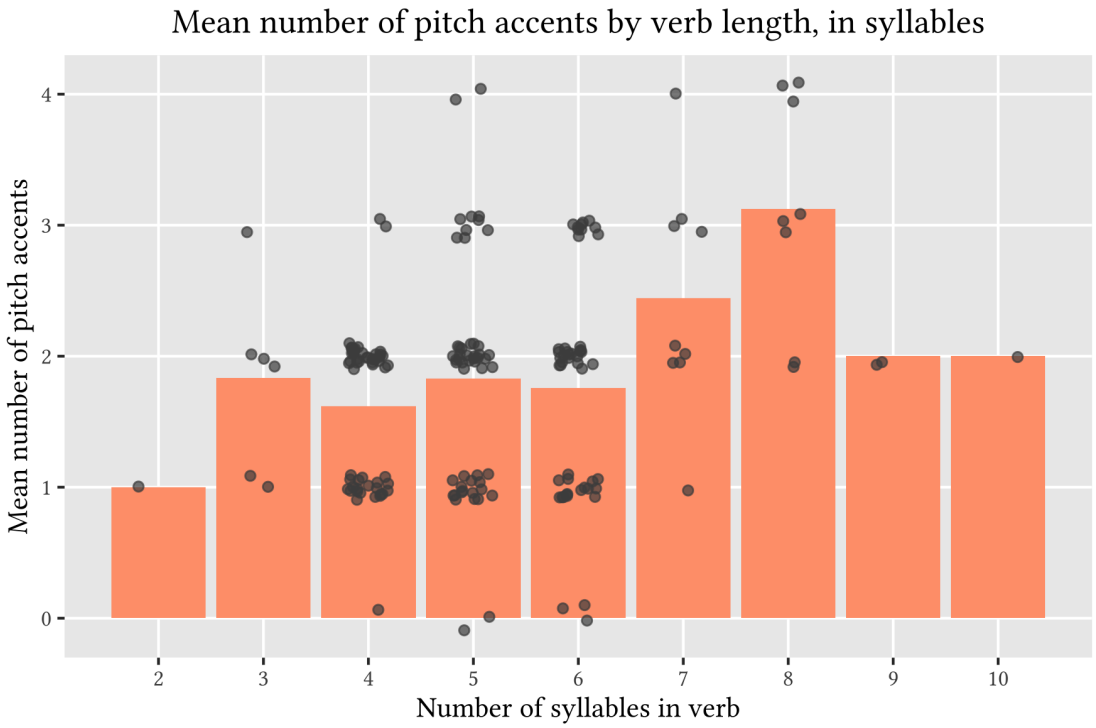

Figure 6.12: Number of pitch accents by word length, in syllables.

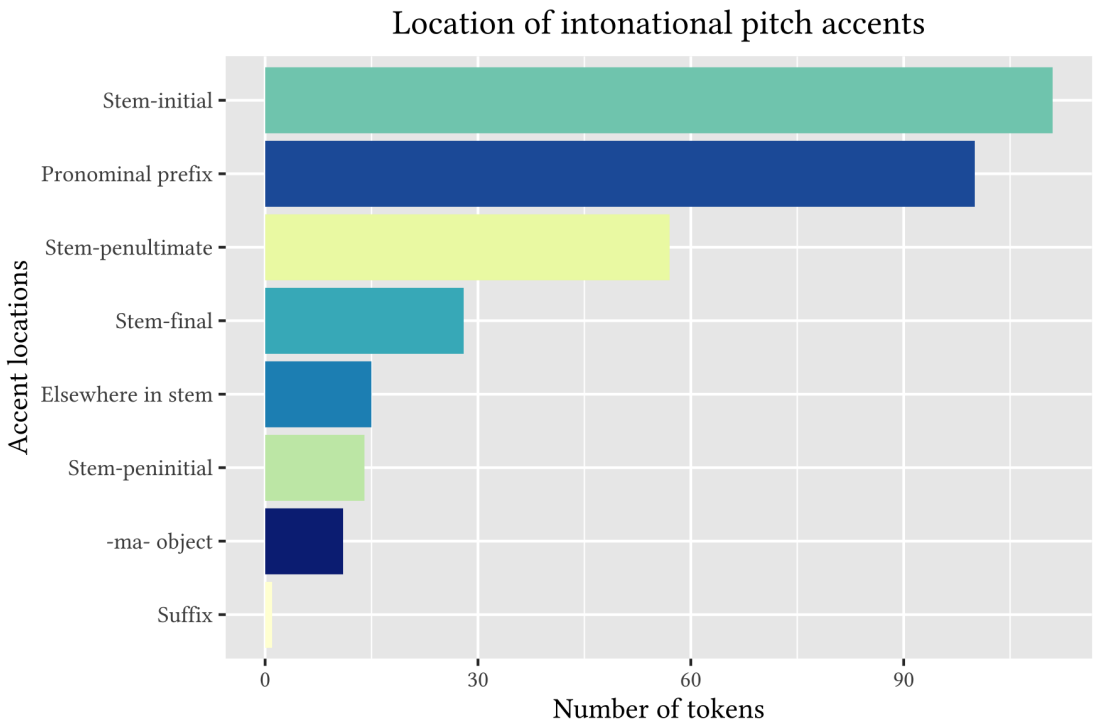

Figure 6.13: Distribution of pitch accents at different positions of the verb.

6.5.1 Pitch accents in the pronominal prefix

As the plot in Figure 6.13 shows, the second most frequent docking site for a pitch accent is in the pronominal prefix, which is the leftmost morphological element in all verb tokens in the corpus.[8] Given the pronominal prefix occurs at the beginning of the grammatical word, this is a location we would expect to be cognitively salient, and therefore a good candidate for realising intonational prominence. There is some further complexity not captured in the plot, however, since the data points contributing to the 'pronominal prefix' bar come from tokens with pitch accents on either the first or the second syllable of the pronominal prefix. This is because I regard pitch accents in both places as anchored to the same phonological target — namely, the prefix-initial (or equivalently, word-initial) syllable. This requires some justification.

As noted above, the annotation procedure employed notation for a 'delayed' pitch accent: "<". This is obviously a major analytical decision, since it requires something of a departure from treating peaks in the f0 trace as the sole source of truth. My reasoning here comes from noticing, during annotation, that one could variably find pitch accents on either the first or second syllable of the same verb prefix, and even on different tokens of the same verb. I hypothesise that the actual location of the pitch accent can shift away from the left edge of the word (i.e., to the second syllable) if this edge is aligned to a L boundary tone, in order to mitigate tonal crowding (Gordon 2014).

The data table used for analysis contains a variable for the location (syllable number) of a pitch accent in the pronominal prefix, if indeed there is one. In the text corpus, I annotated tokens with a pitch accent on the second syllable of the prefix with a H* on the first syllable — the assumed target — and a < on the location of the actual peak in the pitch trace. An example is shown in Figure 6.14. In such cases, the prefixal pitch accent location entered into the data table was determined by the actual pitch peak — i.e., the < annotation, not the H*. There are no tokens in the corpus with a pitch accent on multiple locations of the prefix, with the exception of *-ma-*, the exponent of VEG noun class, which is analysed separately below.

Although the corpus did not explicitly include a break index annotation level, it is implicit in the boundary tone labelling and the hierarchy of prosodic categories it assumes, as in other AM work (i.e. Fletcher 2014). These boundary tone labels are therefore used to derive a break index variable in the data table with the following levels: 4 = Intonational Phrase, 3 = Accentual Phrase, 2 = Prosodic Word. An additional level, 1 = Foot, may be required, but this is more controversial (see next section).

The relationship between the break index and the location of pitch accent in the pronominal prefix is plotted in Figure 6.15. Verbs with monosyllabic prefixes were excluded from the data

[8] An unprefixed verb stem is grammatical if the subject is inanimate, but there are no such tokens in this corpus.

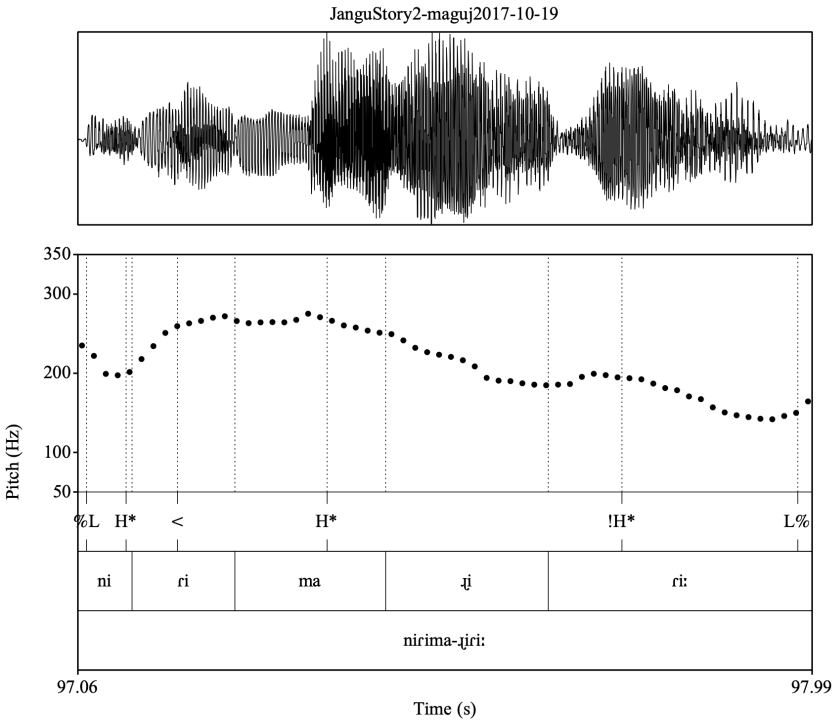

Figure 6.14: IP-unique token of *niɾima-ɻiɾi:* 'We dried it out [pandanus leaves]', with delayed word-initial H*, notated <.

for this analysis. As we can see, the relative probability of finding a pitch accent on the second syllable rather than the first increases as the degree of juncture at the left edge increases. That is, peninitial pitch accent is as frequent as initial accent when the word is at the left edge of an IP, is about half as frequent at an AP boundary, and very uncommon when the word is not at the left edge of either of these larger prosodic units.

The results of fitting a binomial regression model on the same data used in the plot is shown in Table 6.4. Accent location — which may be on either the first or second syllable — is the independent variable, and juncture type is the dependent variable. These results point in the same direction as the plot: a higher break index is associated with greater probability of finding a pitch accent on the peninitial syllable. These results can be interpreted as evidence for avoidance of tonal crowding at the left edge.

Crucially, although the prefixes in the corpus range from 1 to 4 syllables in length, we do not find accents on the third or fourth syllable unless that syllable is the *-ma-* morph, which realises a VEG argument, and is discussed separately below. Thus even where the prefix is 4 syllables long, we do not find prominence on the first and third syllables, as would be predicted by a typical metrical analysis in which syllables are organised into bimoraic trochaic feet.

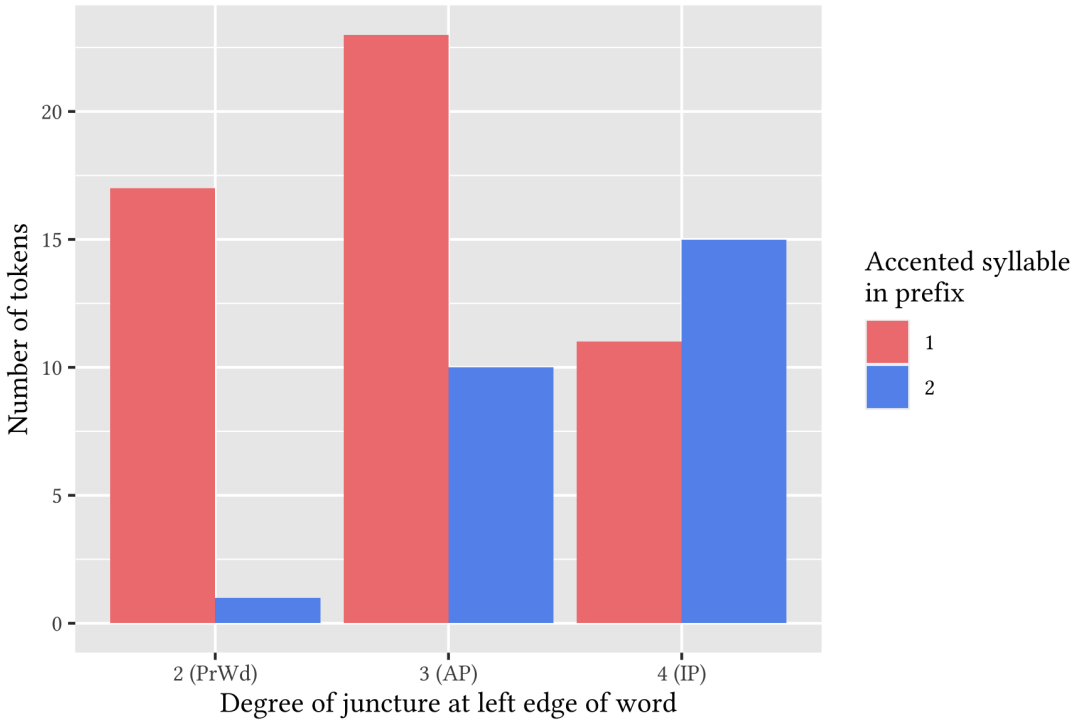

Figure 6.15: Frequency of pitch accents on either the first or second syllable of the pronominal prefix, according to the type of prosodic juncture at the left edge of the prefix/word.

	Estimate	Std. Error	z value	Pr(>|z|)
(Intercept)	-5.1451	1.3909	-3.70	0.0002
bi_left	1.3849	0.4085	3.39	0.0007

Table 6.4: Result of fitting generalised linear model (binomial) on the data presented in Figure 6.15.

6.5.2 Stem-initial accent

As the plot in Figure 6.13 indicated, the most frequent location for realisation of a pitch accent in the verbal word was the stem-initial syllable. It is also by far the most frequent location among tokens with only one pitch accent.

There are several plausible reasons why the left edge of the stem may serve as the favoured site for prominence-lending pitch accents. For one, this is the location where the semantically 'heavy' elements typically begin, namely incorporated nouns and verb roots. As such, this position is akin to word-initial position in predominantly suffixing languages, which most typically have prominence at the left word edge.

Another way of thinking about the 'heaviness' of the stem, as opposed to the pronominal prefix, is in terms of information: the prefixes form a closed class, whereas the root (or roots) that follows is from an open class and is therefore necessarily less predictable. We would therefore expect a local peak in entropy at the beginning of the stem, which would presumably make it more salient. This predictability differential is also reflected in the segmental phonology: the prefixes all feature recombinations of just a small subset of the consonant inventory, whereas roots may draw from the full set.[9]

Vowel-initial stems require special comment. As discussed in Chapter 4, the initial underlying /V-V/ sequence at the boundary between the prefix and the stem surfaces as a long vowel. The surface long vowel can be thought of as doubly associated to both of those morphological constituents. In the data under discussion here, this doubly-associated long vowel is treated as the first vowel of the stem exclusively, and a pitch accent on this vowel is consequently treated as stem-initial. In other words, I contend that the frequent presence of pitch accents on these liminal syllables should be explained with reference to their role as the left edge of the verb stem, not their role as the right edge of the pronominal prefix. If we were to instead draw the opposite conclusion, namely that there is something especially salient about the right edge of the pronominal prefix, we would need to explain why, for verbs with consonant-initial stems, pitch accents in the pronominal prefix appear exclusively on either the first or second syllable, following the conditions outlined in the preceding section (and setting aside the important exception discussed in §6.5.3). Thus, this treatment allows us to see the similarities in the patterning of vowel-initial and consonant-initial stems, since both types frequently have a pitch accent on the first vowel.

As a caveat, note that there is an inherent ambiguity in verbs that have a monosyllabic prefix and a pitch accent on the second syllable of the word. There are equal grounds for regarding

[9] The prefixes feature only the consonants {m, n, ŋ, p, k, w, ɾ} and the three vowel qualities (both long and short). There are further positional restrictions: neither of the stops appear intervocalically, while /ɾ/ only appears intervocalically. Of /p, k/, only /p/ appears prefix-initially.

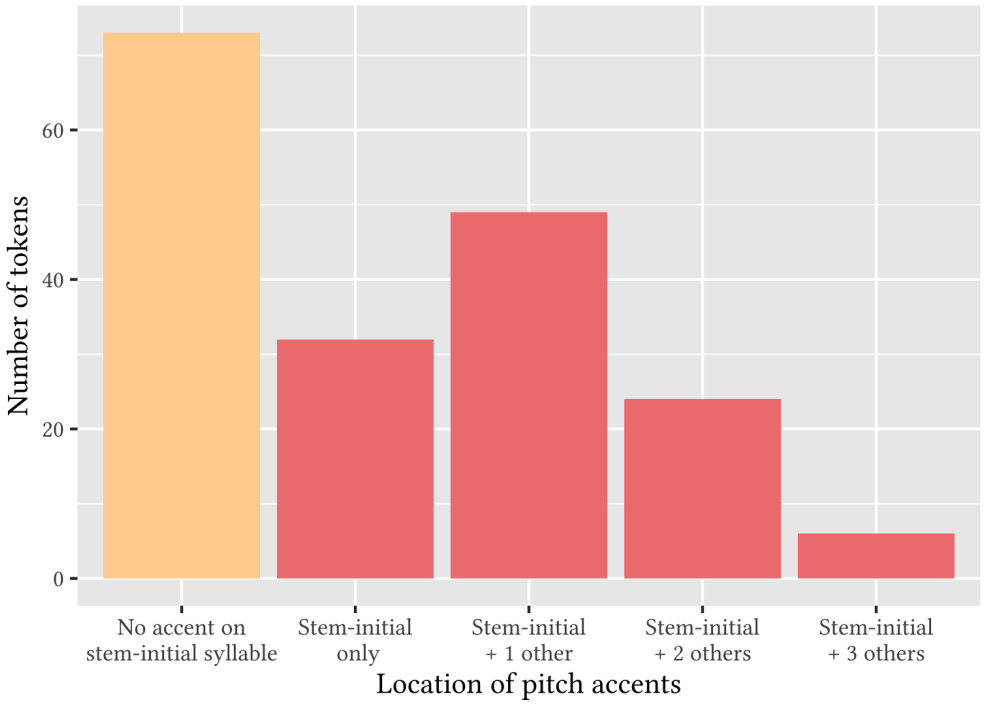

Figure 6.16: Frequency counts for verb tokens with or without a stem-initial accent, and the co-occurence of stem-initial accents with accents at other locations in the verbal word.

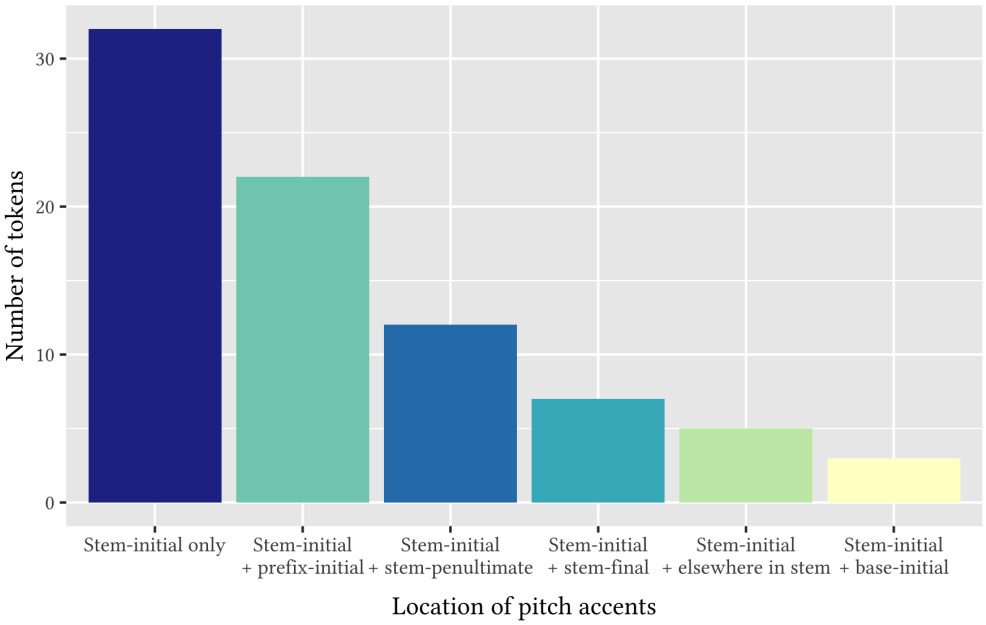

Figure 6.17: Locations of pitch accents among verb tokens with a stem-initial accent and either zero or one other accents in the verb.

such tokens as examples of a pitch accent associated to the stem-initial syllable (where it does in fact appear), or a pitch accent on the word-initial prefixal syllable but with a delayed peak. I have interpreted all such cases as representing the first option.

Note that we cannot reliably explain the predominance of stem-initial accents by its distance from the edge, as in traditional models of alternating stress. Recall that Heath's account predicts accent on the stem-initial syllable only if it is an odd-numbered syllable from the right word edge, and 'the last few syllables are short' (Heath 1984, p. 32). However, the corpus contains tokens of stem-initial accents in stems consisting of 2, 3, 4 or 5 short syllables. We find the same widespread distribution when we count the distance from the edge by moras or syllables (of any weight).

Notably, there are only 2 tokens of stem-initial accents appearing immediately following a monosyllabic prefix that is also accented; both tokens contain a long vowel in the prefix, as in (341).[10] Otherwise, verbs with a monosyllabic prefix may appear with either a prefixal accent or a stem-initial accent, but not both. This is reminiscent of Bishop's (2002) findings for Bininj Kunwok, which she regards as evidence for prosodic incorporation of the monosyllabic prefix into the stem (see also Mansfield 2023).

(341) a. wáː-l̪áɳʈa-l̪áɳʈa-ŋi
3pl>3pl.R-Red-wait_for-Past2
'They waited around a long time for them'
(Heath3345A_125_RoundYam_Maadi1 56368)

b. wúː-níni-bíː-ni danu
3Coll.R-soft-Fact.Refl-Past2 that
'It [the yam] became soft' (Heath3345A_125_RoundYam_Maadi1 67558)

[10] There is a further possible token, *wáː-máwuɻaːtíː* 'it's cold', but the prefix here is, according to the traditional analysis, derived from /wu-/ '3Resid'/'3Neut' plus /ara/ 'Mult'. This combination yields either the somewhat more transparent [waːɾa-] or the heavily contracted [waː]. However, the contraction is so conventionalised for verbs describing the weather that it may perhaps also be considered a distinct kind of impersonal prefix.

6.5.3 The separability of *ma* within the pronominal prefix

As a consequence of the fact that two of the texts in the corpus make frequent reference to *makuc* 'pandanus', which is of the VEG noun class, there are many tokens of verbs whose prefixal inflection realises an argument of that class. It is striking that of the 20 such tokens, 11 have a pitch accent associated with the final syllable of the prefix, which in every case has the form /ma/ (see, for example, Figure 6.14 above, or Figure 6.22 below).

In the verbal pronominal prefix paradigm, all prefixes that realise a VEG argument — whether subject or object — end in /ma/, and the form /ma/ never appears in any prefix that does not realise a VEG argument. A small selection of these prefixes is shown in (342). No other grammatical feature value has a completely reliable associated form, nor does any other have a completely predictable linear position within the prefix. The realisation of NEUT, for instance, is sometimes associated with a final /wu/ (343a,b), a final /ku/ (343c), or a vowel change (343d). This /ma/ element therefore constitutes the only biunique mapping between form and function in the entire pronominal prefix paradigm. That it is a zone of exemption from the otherwise highly fusional character of the pronominal prefix paradigm may explain why it is also prosodically exceptional, in allowing a prefix-final syllable to receive a pitch accent. That is, the /ma/ element shows a level of prosodic independence reflecting its morphotactic separability.

(342) VEG prefixes:

a. ma-
   3VEG.R-

b. ŋama-
   1sg>3VEG.R-

c. niɾima-
   1pl>3VEG.R-

d. nimpiɾima-
   2pl>3VEG.IRR-

(343) NEUT prefixes:

a. wu-
   3NEUT.R-

b. ŋawu-
   1sg>3NEUT.R-

c. ŋa:ŋku-
   1pl>3NEUT.IRR-

d. niɾi- (*cf.* nuɾu-)
   1pl>3NEUT.R- 1pl.R-

This behaviour is not unheard of in the region. Bishop (2002, p. 127), for instance, also notes that, in Bininj Kunwok, trisyllabic pronominal prefixes that admit of a transparent segmentation show variability in pitch accent also. While Bininj Kunwok is generally far more agglutinative than Wubuy, having very few grammatical word-internal segmental alternations, its pronominal prefixes also tend to display a fair degree of fusion.

It is also true that monosyllabic prefixes can be incorporated into the same prosodic word as the stem, as suggested by the fact that we have no tokens in which there is an accent on both the

monosyllabic prefix and the stem-initial syllable. Likewise, we have no tokens of an accented /ma/ followed by an accented stem-initial syllable. That being so, the availability of /ma/ for pitch accent should be seen as an instance of the same strategy of prosodic incorporation.

6.5.4 Accents within the stem

The only category here not obviously associated with a phrase edge is the 'elsewhere in stem' category. Closer inspection of this category reveals that all these items involve reduplication, and the 'elsewhere' syllable is the first syllable of the base. This is of course a morphological stem otherwise, so an accent on its initial syllable may also be considered stem-initial. We must also note that, in this dataset, there is only an accent on both reduplicant and base if the reduplicant is disyllabic (not if it is monosyllabic).

6.5.5 Right edge accents

We have seen in the preceding sections that pitch accents demarcate the left boundary of the major morphological constituents of the verb — the pronominal prefix and the stem. There is also a large number of verbs in the dataset bearing a pitch accent proximal to the right edge of the word. There are two questions we must ask of the distribution, in light of previous accounts of prosody in Wubuy and other nearby languages. Firstly, are pitch accents attracted to heavy syllables at the right edge? Second, what prosodic domain or level of constituency in the prosodic hierarchy should these right-edge pitch accents be associated to? Both of these questions bear crucially on the controversy over metrical feet.

If the location of pitch accent were sensitive to syllable weight, we should expect to find pitch accents appearing on the final syllable when that syllable is heavy, and on the penultimate syllable when the final is light. Recall that in general, pitch accents are optional, and therefore the mere absence of a pitch accent on a metrically prominent syllable does not invalidate the weight-sensitivity hypothesis — instead we would need to find positive evidence of pitch accents straying from those metrically prominent syllables.

Figures 6.25 and 6.26 plot the number of tokens with a stem-final and stem-penultimate pitch accent, respectively, in each case separated according to the weight of the final syllable. The plots reveal that, indeed, pitch accents at the right edge are distributed according to final syllable weight. When the final syllable is heavy, there is a final accent about half the time, as we see in the left panel of Figure 6.25. But pitch accent on a final light syllable, as plotted in the right panel, is only attested in one token (discussed below). In an impressive reversal, tokens with a final light syllable receive penultimate accent about half the time, as seen in the left panel of Figure 6.26, whereas we have only two tokens (7%) of penultimate accent where the final

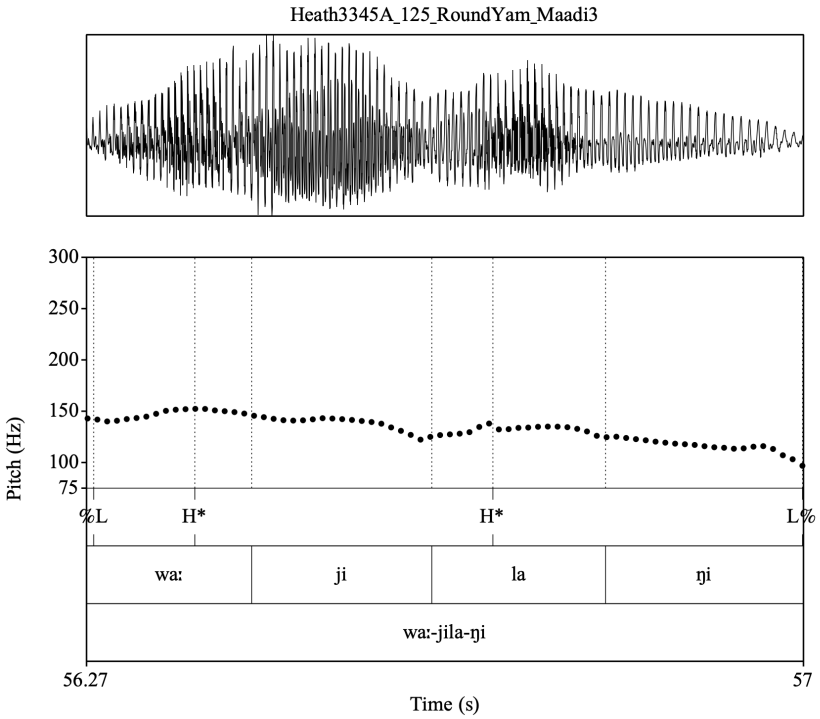

Figure 6.18: Token of *wa:-jila-ŋi* 'they sliced it (Coll)', with penultimate pitch accent.

syllable is heavy, as seen in the right panel. Examining the correlation between right accent location (penultimate, final, or neither) and final syllable weight using a Chi-squared goodness of fit test reveals a highly statistically significant relationship (Pearson's $\chi^2 = 83.73$, df = 2, $p < .001$). It is hard to avoid the conclusion that pitch accent is weight-sensitive.

The exceptions to this overall pattern are also interesting. The only example of stem-final accent on a light syllable is a token with a monosyllabic (word-level) suffix, *nuɾ-á:ɾa-ŋu-puɾí-ɾuc* 'where we were all sitting around'. In this case the accented syllable is the penultimate syllable of the word, even if not the stem. This suggests, as was reported also by Heath (1984) and Hore (1981), that monosyllabic suffixes may be prosodically incorporated into the constituent to their left. Note that the stop coda in the *-ɾuc* 'Loc' suffix would not count as moraic in Ngalakgan (Baker 2008b), since word-final consonants are extra-prosodic. Final syllables closed with a stop thus do not attract pitch accent in that language, and perhaps do not in Wubuy either, though we have little data to draw on.

Evidence of weight-sensitivity may point in favour of bimoraic metrical feet, and therefore word-level prominence — that is, a pitch accent may associate with the strong syllable of the rightmost foot. However, this pitch accent at the right edge may still be largely demarcative

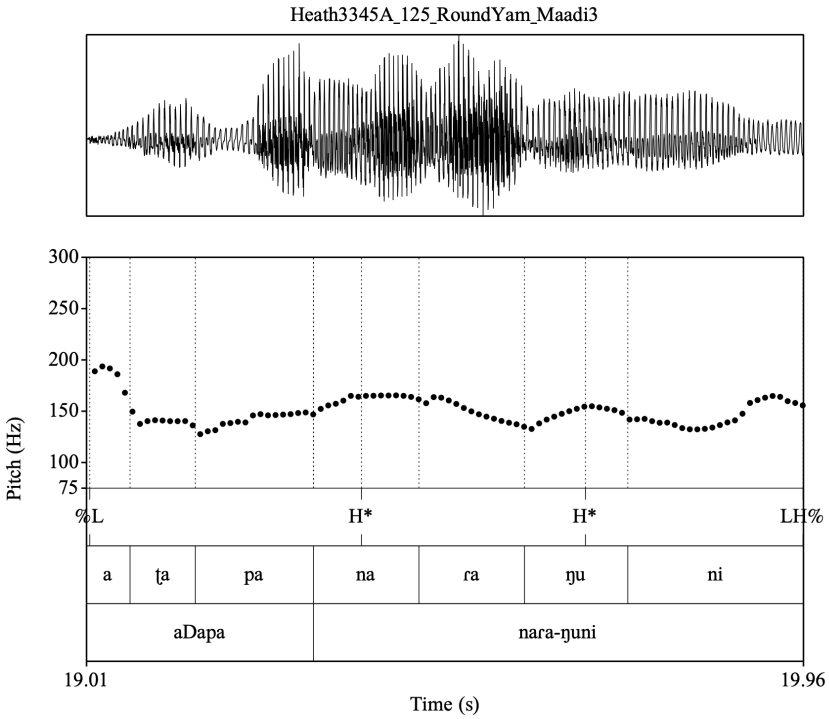

Figure 6.19: Token of *naɾa-ŋuni* 'we ate it (COLL)', with penultimate pitch accent.

if it turns out that it preferentially appears at the edge of major phrasal constituents. Indeed, recall from Figures 6.9 and Figure 6.10 that the vast majority (87%) of the verbs in the database are aligned to at least an AP boundary at their right edge.

Figure 6.27 plots the distribution of pitch accents at the right edge according to the degree of prosodic juncture at that edge. For a pitch accent to count as aligned to the right edge, therefore having a value of TRUE, it must have penultimate accent (if the final syllable is light) or final accent (if the final syllable is heavy). However, if that right edge prominent syllable is also the *leftmost* syllable of some other morphological domain, it is excluded from consideration, and receives a value of NA. For example, a token with a heavy monosyllabic stem, with or without a pitch accent on its only syllable, will be NA, likewise the initial syllable of a disyllabic stem ending in a light syllable. The same applies for the base to reduplication, or the root in an incorporation construction. In other words, the plot only assigns a value of TRUE or FALSE to unambiguous cases where the metrically strong syllable at the right edge is not also metrically strong at the left edge.

The plot suggests a substantial difference in the distribution of right-edge pitch accent at IP boundaries, where it appears more often than not, compared to APs and PrWds, where

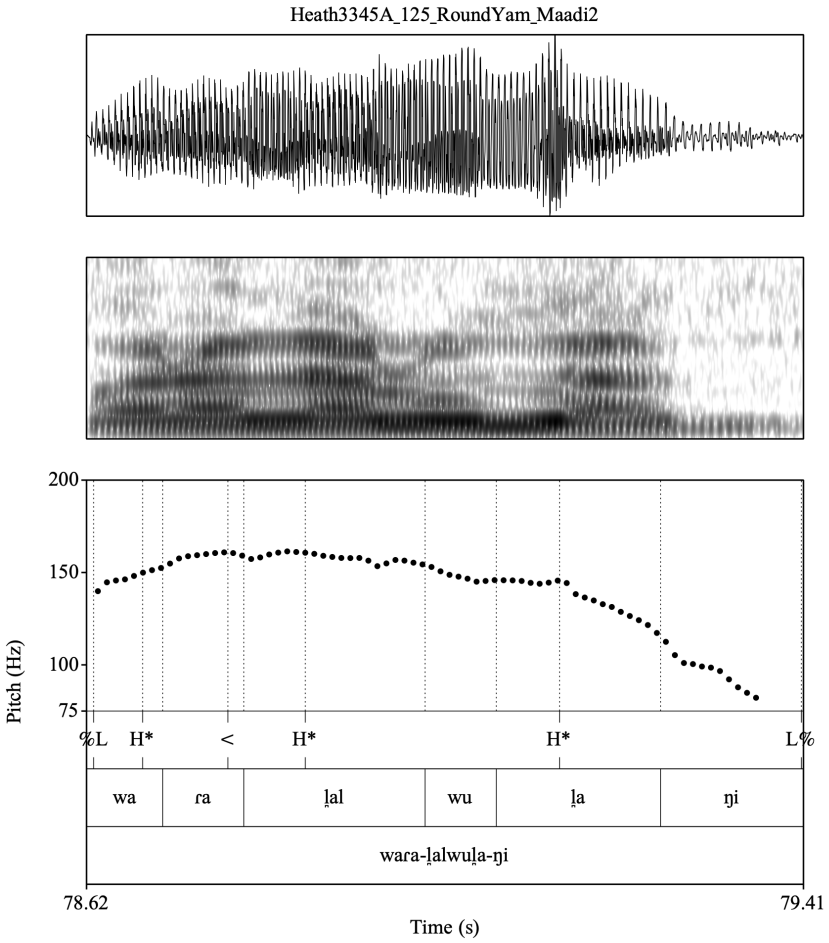

Figure 6.20: Token of *waɾa-l̪alwul̪a-ŋi* 'they soaked it (Coll)', with penultimate pitch accent.

absence of pitch accent is more frequent. This association between a high degree of juncture and right-edge pitch accent is partially supported by a binomial logistic regression, shown in Table 6.5, which finds a statistically significant contrast between the IP and AP boundaries ($z = 3.11, p < .01$).

There are still five tokens of verbs with a right edge accent but no right edge AP boundary. Of these, three are tokens of *waŋi-ɻarcaŋi* 'she unloaded it [Coll]', uttered very rapidly in succession within the same stretch. The high speech rate here makes junctures much harder to discern, and AP boundaries were therefore not annotated; it may be that these should be considered as APs nonetheless. The other two tokens were of (different) verbs with final long vowels, immediately followed by unaccented *ŋa*. As mentioned earlier, this discourse connective may be considered prosodically incorporated with the verb, in which case the verb-final long

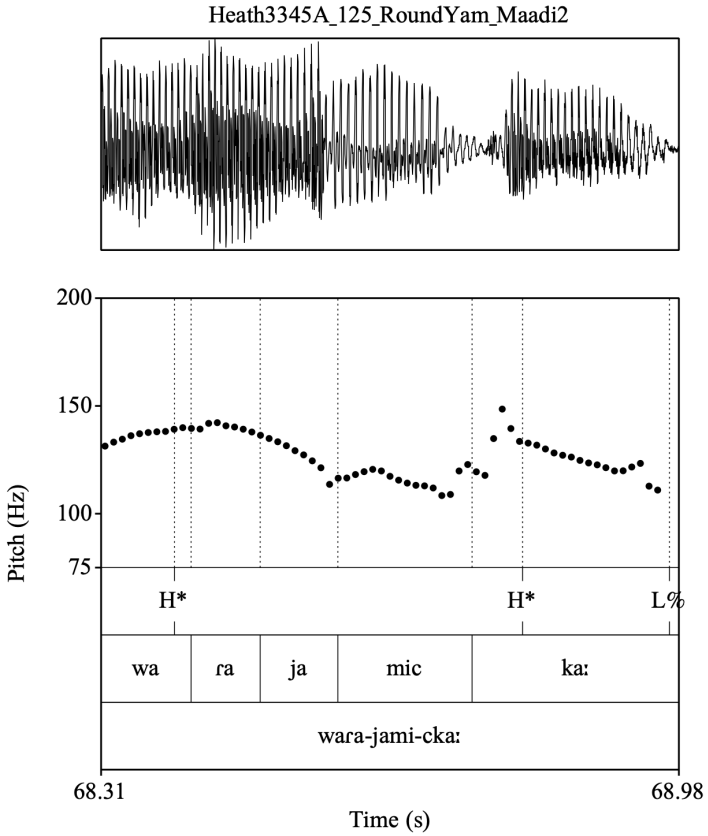

Figure 6.21: Token of *wara-jami-cka:* 'they told them', with final pitch accent.

vowel is in fact the penultimate vowel in the phrase-final prosodic word, analogously to the situation with *-ruc* discussed above.

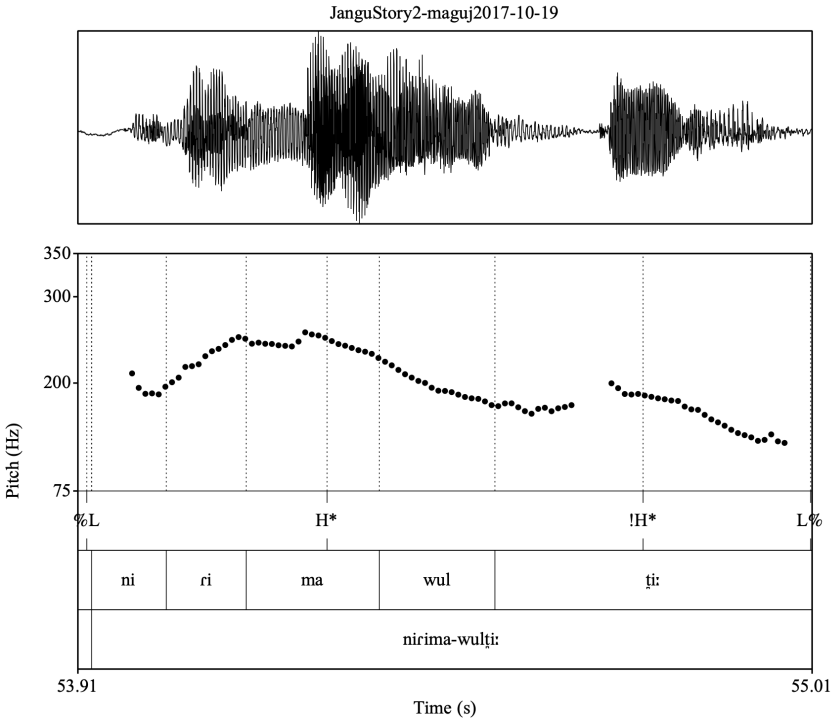

Figure 6.22: Token of *niɾima-wulʈ-iː* 'we cut it (VEG)', with final pitch accent, and accent on *-ma-* in prefix.

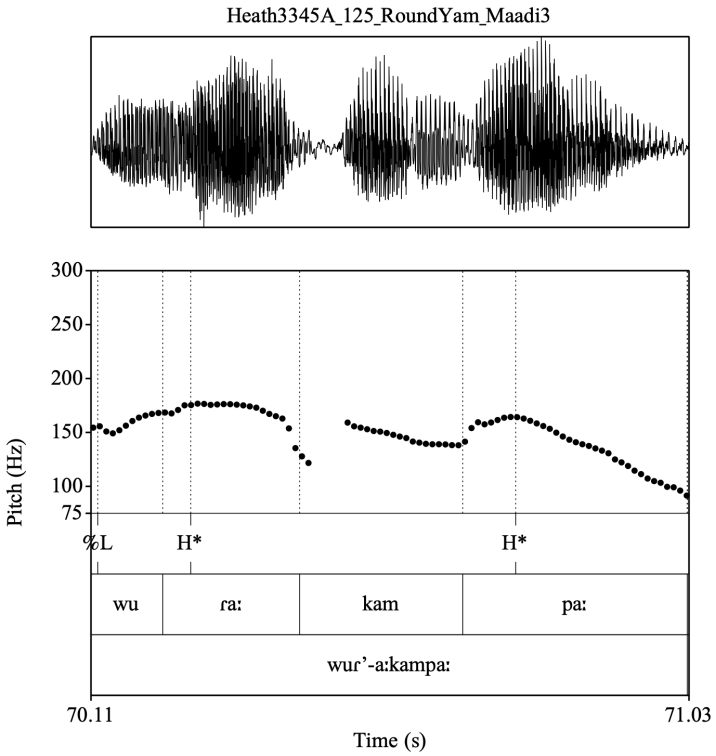

Figure 6.23: Token of *wuɾ-a:kamp-a:* 'they roasted it', with final pitch accent.

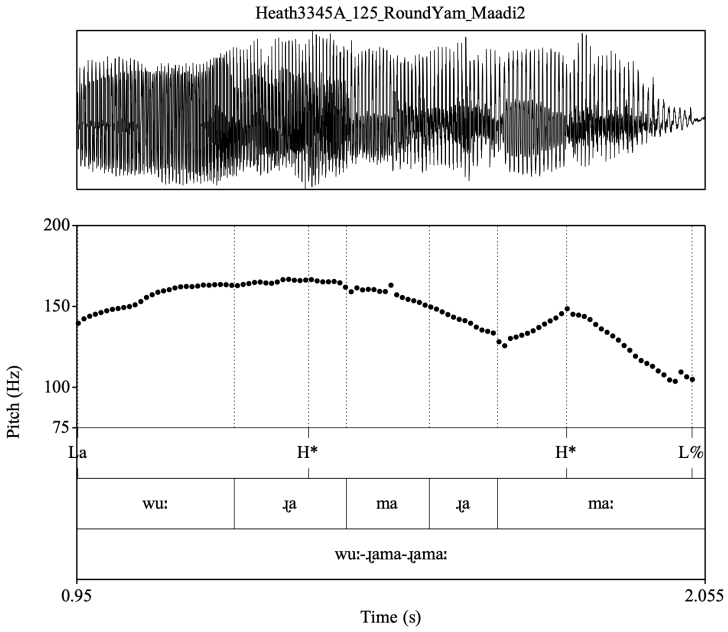

Figure 6.24: Token of *wu:-ɻama-ɻam-a:* ‘they were abundant, everywhere’, with final pitch accent.

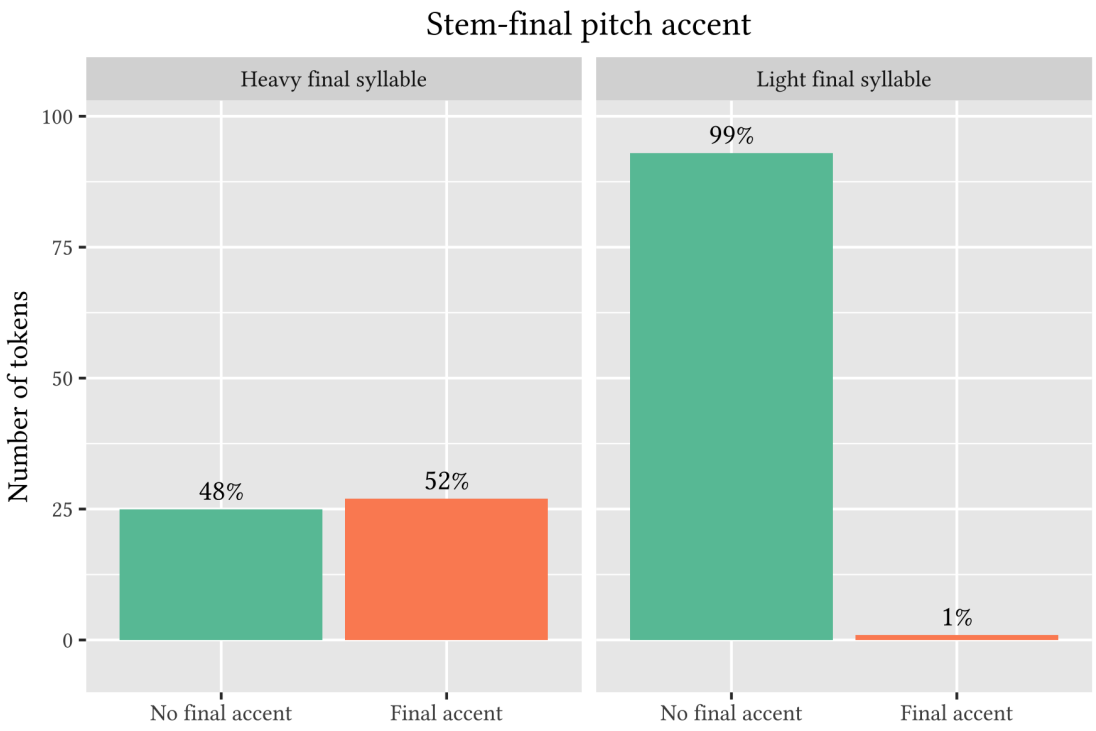

Figure 6.25: Distribution of stem-final pitch accent, by final syllable weight.

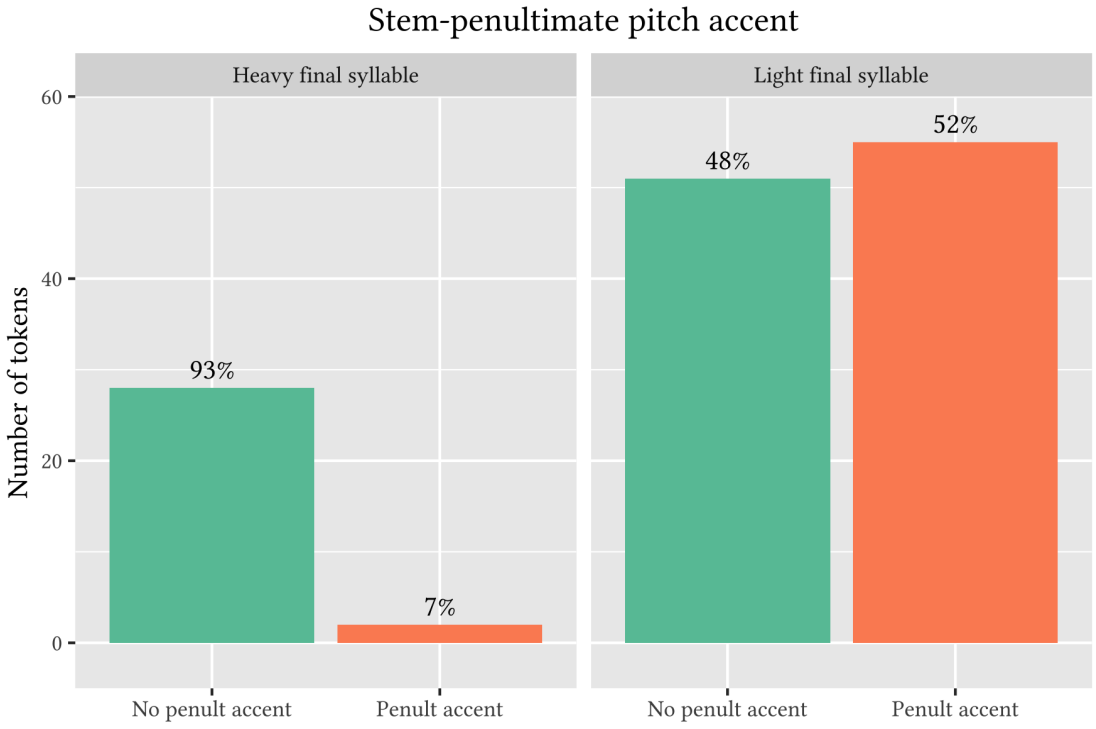

Figure 6.26: Distribution of stem-penultimate pitch accent, by final syllable weight.

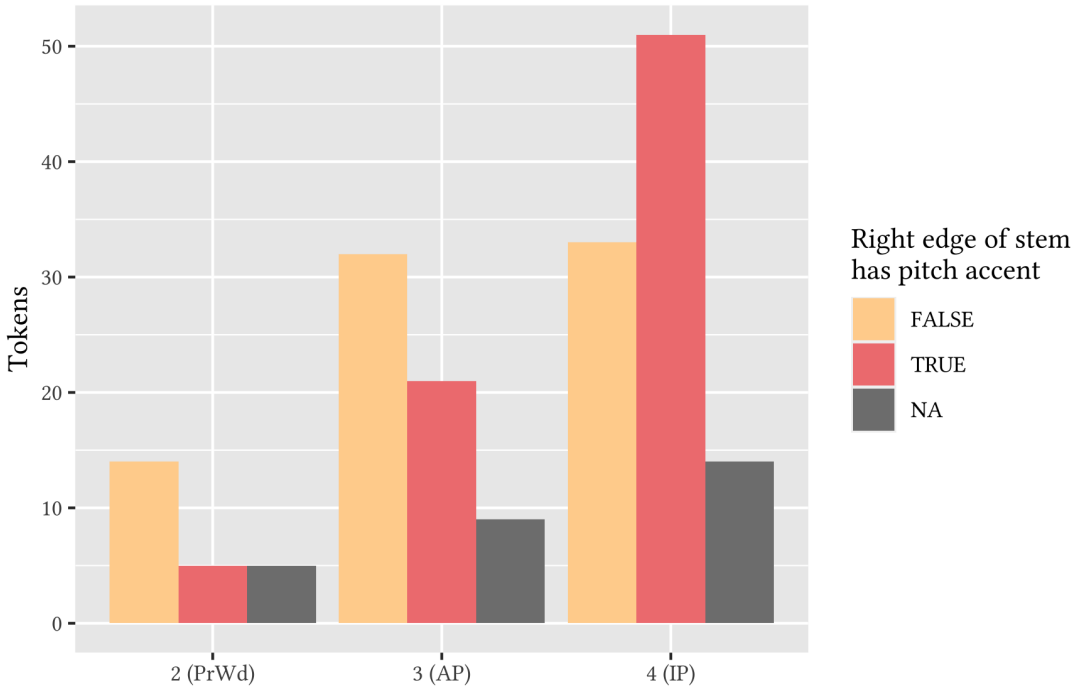

Figure 6.27: Frequency of pitch accents at the right edge of the stem (penultimate or final) by prosodic juncture. Syllables that are also at the left edge of a morphological domain (i.e., root, stem, base to reduplication) are treated as NA, not TRUE or FALSE, since their edge-association is ambiguous.

	Estimate	Std. Error	z value	Pr(>|z|)
(Intercept)	-2.6907	0.8791	-3.06	0.0022
bi_right	0.7767	0.2498	3.11	0.0019

Table 6.5: Results of logistic regression fitted to the data in Figure 6.27. The reference level is break index 4 (the IP boundary).

6.6 Discussion

The picture that emerges from the results reported here is generally in line with what has been described for other languages of the region. The inventory of tonal events and their organisation into characteristic intonational tunes is typical of Australian languages: there is a very small inventory of pitch accentual targets (here, only a monotonal H* and its upstepped and downstepped variants), two levels of prosodic phrasing (labelled AP and IP here), and two major tune types (an unmarked falling tune and a high monotone with extreme lengthening of the phrase-final syllable), with some variations within each.

The alignment of pitch accents to the text is strongly reminiscent of Bininj Kunwok (Bishop 2002) — perhaps unsurprisingly given the close structural similarity and likely shared history between the two languages. That is, pitch accents in Wubuy are always optional, and are of two kinds: those that are aligned to the left edge of prosodic words, and those that are aligned to the right edge of the AP. Prosodic words generally align to morphological constituents, but their boundaries can be shifted to include monosyllabic constituents either within the morphosyntactic word (as with monosyllabic pronominal prefixes) or outside it (as with *ŋa* 'and/then'). Accordingly, prosodic words are typically smaller than the morphosyntactic word, and each may host exactly one pitch accent except if AP-final, where an additional pitch accent is possible. Indeed, as the combined boxplot and scatterplot in Figure 6.28 shows, the number of pitch accents per verb tends to increase as a function of morphological complexity — which is our diagnostic for the number of prosodic words. Notably, no tokens have more pitch accents than the number of prosodic words plus one (for the AP-final accent).

This pitch accent aligned to the right edge of the AP is, as I have shown, sensitive to syllable weight: it is reliably realised on the syllable containing the penultimate mora. This finding forms an interesting complement to the Manyallaluk Mayali dialect of Bininj Kunwok (Bishop 2002), which (according to my interpretation of the data) has an optional accent on the penultimate syllable of the AP. We could in fact describe this accent in exactly the same terms as I have for Wubuy: there are no contrastively long vowels in Bininj Kunwok, so the penultimate syllable would also be the syllable containing the penultimate mora.[11] Taking this even further, we may speculate that the Wubuy pattern developed historically from one like that of Manyallaluk Mayali. So far as we are able to understand from the current state of reconstruction in the Gunwinyguan family (or subgroup, depending on one's inclination), Wubuy's long vowels, at least within roots, developed from what were once short vowels separated by a glide — in fact, this remains partly productive as a synchronic process at morphological boundaries as well. If this pathway is true for the long vowels realising verbal TAM inflections also, and if the

[11] So far as I am aware, weight-bearing codas do not appear in this position of the verb either.

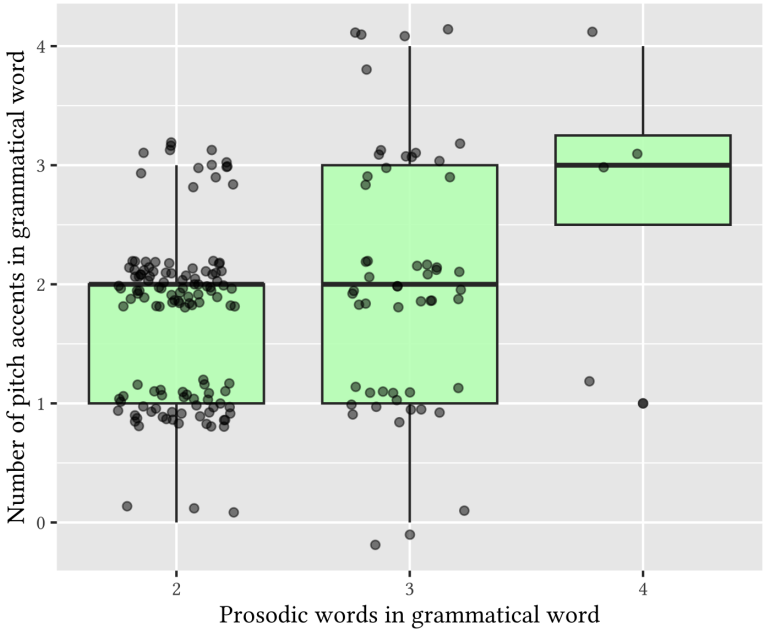

Figure 6.28: Number of pitch accents by number of prosodic words within verb token.

immediate ancestor of Wubuy had an optional pitch accent on the phrase-penultimate syllable just as in Manyallaluk Mayali, then the current system simply retains those pitch accents on the historical reflex of those accented syllables.

Like Mansfield (2023) and Gordon (2014), I do not regard such facts as sufficient evidence for feet. Certainly, we do not find evidence of the kind of alternating prominence that one finds in languages for which foot structure is well-motivated, nor do we know of any segmental or morphological effects that could be attributed to such a structure. As those authors argue for other languages, phrase-penultimate pitch accents may be independently explained by the need to avoid tonal crowding in the vicinity of phrasal boundary tones; postulating feet in the prosodic hierarchy solely to account for phrase-edge accent alignment adds considerable unnecessary complexity to the system. In fact, we already have good reason to regard tonal crowding as a relevant factor in repelling pitch accent from the absolute phrase edge, because that is precisely what we observe at the left edge of the IP.

Without metrical feet, the analysis affords a greater role for intonational levels of prominence, and accordingly de-emphasises the role of lexical stress. Indeed, if the pitch events we find inside the morphosyntactic word can be explained as demarcating the edges of prosodic constituents, it seems that an account of lexical stress is largely unnecessary. This casts doubt on the assumption of Hore (1981), one that was extremely widespread in earlier descriptive work on the languages of Australia, that Wubuy has a lexical stress system that is meaningfully different from the

intonation system. Instead, it affirms the suspicions of Heath (1984) that no attempt to excavate a level of lexical stress from beneath intonational prominence is likely to bear fruit.

6.7 Conclusion

The present chapter has investigated the prosodic characteristics of verbs in Wubuy, providing the missing ingredient to the model of the phonology and morphology of complex words developed in the previous chapters, which are concerned exclusively with segmental phenomena. Once again, however, morphological constituency is crucial to predicting phonological behaviour. The prosodic behaviour, in turn, provides further justification for the notion that word-internal morphological constituency is accessible to speakers (without thereby implying that the units so accessed are abstract in the way assumed by morphemic theory). Taking this further, the findings here accord with the conclusion of Bundgaard-Nielsen & Baker (2020) that Wubuy's polysynthetic words are prosodically like sentences: they typically consist of a single Intonational Phrase containing one or more smaller constituents, perhaps corresponding to multiple production planning or processing units.

The results presented in this chapter also contribute to a growing literature urging caution when interpreting 'stress' in underdescribed languages (de Lacy 2014). While I do not regard the question of the levels of prosodic constituency — particularly the metrical foot — settled even in the present case of Wubuy verbs, I share Bishop's (2002) and Mansfield's (2023) doubts that languages like these are best analysed with iterative binary feet, even at a very abstract level. This is in line with the unease found in much recent work around earlier claims of a universal prosodic hierarchy (Culhane 2022; Schiering et al. 2010).

While the present study only scratches the surface of Wubuy prosody, even just within the domain of verbs, it opens up a number of new avenues to explore. A fuller picture would require data from the full range of Wubuy's morphological affordances. For instance, the present study lacks tokens of noun incorporation almost entirely, nor are we able to observe the prosodic effect of derivational suffixes on the verb complex. It would also of course be necessary to examine other word classes, about which we know little; such work could begin with the present corpus. One major limitation on the present corpus, however, is that it contains only monologic texts. Dialogic texts would be expected to contain a far wider variety of tunes, and this may end up revealing further possibilities for tune-text alignment not considered here. A further consequence of having only narratives in this corpus is that we do not have any tokens at all of verb-final nasals, which, while common enough, tend to occur in past punctual and future punctual (including imperative) tense-aspect inflections (PAST1 and NPST1 suffixes) — these are much rarer in narratives than in conversation. Informally, I have very often observed final

accent on these verbs as well, suggesting these nasals are moraic. Finally, the AM phonological analysis offered here may fruitfully serve as the basis for an acoustic phonetic analysis, rather than remaining an unaccounted-for confound, as it is in the existing phonetic studies on Wubuy prosody. Ideally, the use of text corpora would be complemented by targeted experimental work.

Thus, the research presented in this chapter is, although a substantial addition to the relatively small number of intonational phonological descriptions of Australian languages, just a glimpse into what is a fairly complex and unique prosodic system. Further work is needed to begin to unravel its rich interplay of syllable weight, tone assignment and morphological constituency.

7 Conclusion

7.1 What does the Wubuy lexicon sound like?

This thesis has sought to understand the phonology of complex words in Wubuy, and much has been made of the value of understanding words as the basic units of storage in the mental lexicon, rather than smaller, more abstract elements, like the structuralist morpheme. Words are what language users encounter, the smallest an utterance can be. These stored words are the basis on which language users discover patterns, as they compare words to each other, noticing the parts that overlap and the parts that differ. Patterns help hearers understand the words they hear and speakers produce the words that their hearer will understand — even if neither speaker nor hearer have ever heard that word before. In Wubuy, as we have seen, words can be extremely long, with many parts that overlap with other parts of other words in the lexicon.

As words are produced over and over again, generation after generation, their form can change. Practiced tongues move through articulatory targets with greater speed, sometimes not reaching one target before moving to the next, knowing that practiced ears will still understand them (Sóskuthy & Hay 2017; Todd et al. 2019; Tomaschek, Tucker et al. 2018). As this process unfolds, articulatory gestures may become weaker — shorter in duration and smaller in magnitude (Bybee & Easterday 2019). A fortis stop can become a lenis stop; a lenis stop can become a glide or a rhotic or a lateral.

At the easternmost extent of the large Gunwinyguan language family, speakers of the immediate ancestor to Wubuy and Anindilyakwa, according to the available comparative evidence, variably produced morph-initial stops as continuants (van Egmond & Baker 2021). In some words and not others, the lenited variants, the continuants, became the norm. In Wubuy, this process of lenition nonetheless left pockets of resistance: stops following non-continuants remained stops. As a result, parts of words that once overlapped completely with shared parts of other words — roots, stems, affixes — now differed in a crucial respect according to their immediate phonological environment: following non-continuants, they retained the older stop, while elsewhere they were lenited to continuants at the same place of articulation. That is, what emerged was an alternation. This alternation arose only with those morphs to undergo lenition, and not with those that escaped it. Thus the root in *ŋa-ju̪ɻi:* ‘I’m transporting it’ alternates with

the root in *nuɲ-cuɻi:* 'you're transporting it'.

How did speakers of this nascent Wubuy manage that alternation, and how did they know when and how to apply it? As we saw in Chapter 3, identifying a stop-initial morph is not a straightforward matter when it could potentially alternate with either a stop or a continuant. Under circumstances we can only speculate about, another element must have already regularly co-occurred with the non-alternating stops: the string /ŋu/. This was most probably a meaningful morph at some stage, but came to be associated primarily with its phonological context; it was thus recruited into the emerging system of paradigmatic phonological oppositions, serving as a clue to the listener that the following stop is to be interpreted as belonging to a morph whose form is stop-initial in all environments, and not to a morph that is elsewhere continuant-initial. Were it not for this /ŋu/ element, such an inference would not be possible. Thus *nuɲ-cuɻi:* 'you're transporting it' is able to contrast with *nuŋu-cuɻi:* 'you're pushing it', despite the identical forms of the roots. The interactions between ŋu-Epenthesis and continuant-stop alternations, together with the assimilation and deletion of nasals, highlight the importance of treating word parts in context. As we saw, speakers are seemingly able to leverage the fact that where root- or stem-initial stops and continuants can contrast, prefix form is held constant; where that contrast in stem-initial segments is neutralised (by continuants becoming stops), prefix form alternates. That is, where contrast in one place disappears, contrast elsewhere in the word becomes a useful heuristic.[1]

Wubuy's *ŋu-* thus turns out to be akin to other cases of so-called 'empty morphs' in the literature, which have been argued to serve as a diagnostic of a stem's inflection class (i.e. Wilmoth & Mansfield 2021). This empty morph, *ŋu-*, however, enables the reliable identification of what I have been calling *phonological class* — essentially like inflection class, but where the classes are determined by phonological features across related word forms, rather than being lexically determined (much like Labov's (1994) notion of a 'word class'). In that sense, I have treated phonological alternations as fundamentally embedded in morphology — as patterning of an irreducibly phonological nature among sets of morphologically related word forms. The classes relevant to the alternations in Chapter 3 turned out to be the *stop class*, the *alternating class*, and the *sonorant class*.

This historical process of lenition in Wubuy (as well as in Anindilyakwa) also had the effect of reshaping the distribution of manner constrasts in the lexicon, such that, compared to other languages of the family, stops ended up relatively underrepresented in favour of sonorants. In Chapter 4, I suggested that this may have been a factor responsible for the emergence of

[1]Relatedly, Round et al's (2022) phylogenetic study of the Pama-Nyungan family of Australia presents evidence of historical 'transphonologisation', where the functional load of tonic vowel length contrasts comes to be transferred to post-tonic consonant manner contrasts.

Wubuy's highly unusual reduplicant shape alternation, where stop-initial bases condition a monosyllabic reduplicant, and sonorant-initial bases condition a disyllabic reduplicant. Being so underrepresented, initial stops are more informative than sonorants; during processing, a stem-initial CV sequence contributes more to identifying the stem when the C is a stop than when it is a sonorant. Given that productive CV reduplication is not found elsewhere in the family — again, other than Anindilyakwa, tellingly — the monosyllabic reduplicant probably emerged from reducing the disyllabic reduplicant, but reduction to a monosyllable could only become the pronunciation norm where the remainder was able to 'minimally evoke' the stem it copies from (Baker & Bundgaard-Nielsen 2018). That is to say, only (non-alternating) stop-initial reduplicants could reduce to a monosyllable. Again, a difference in the workings of lenition on stops versus segments of other kinds results in an alternation, where reduplicant size is conditioned by the phonological class of the base. The phonological classes required to describe this distribution, as I showed in Chapter 4, are exactly those that are required for the alternations at the prefix-stem boundary: a stop class (monosyllabic reduplicant), a sonorant class (disyllabic), an alternating class (disyllabic), and a vowel class (long monosyllable). Reduplicant shape thus becomes yet another diagnostic of phonological class, and therefore another clue by which speakers can infer the rest of the paradigm for that stem.

In Chapter 5, I showed that lenition, while thoroughly phonologised in the above-mentioned alternations, remains a force acting upon contemporary speech production. Morph-final velar stops are variably deleted preceding other segments in the word, which (due to continuant-stop alternations at these junctures), can only be another non-continuant: for example, in *ama-maṯaḻa(k)-tuc* 'at the beach', the /k/ is deleted some of the time. While the existence of variability in velar deletion had occasionally been mentioned in passing in more recent literature, the present thesis is the first to detail the extent of that variability and describe its phonetic manifestations. I showed that the rate of velar stop deletion depends on the type of morph it appears in, where the rate of deletion is in inverse proportion to that morph's likelihood of appearing at the right word edge. I took this as evidence of a sound change in progress, proceeding by lexical diffusion, as gestural sequences across morphs become more coordinated with practice. Part of what makes this deletion possible is, once again, the fact that morph-final /k/ tends not to be particularly informative, as it can always be inferred by the restrictions it imposes on the initial segment of the following morph: this is one of those environments where historical lenition of morph-initial stops to continuants was inhibited.

It is striking, then, that it is lenition — certainly among the most commonplace and phonetically natural forces driving sound change — that, by advancing throughout the Wubuy lexicon on multiple fronts over time, has resulted in a synchronic system that, on the face of it, seems so 'unnatural'. Most of the alternations mentioned above can be described as opaque in one

way or another, and to that extent are difficult to frame in terms of synchronic well-formedness, or markedness avoidance. Analysts faced with other phonological systems showing widespread opacity have tended to move in one of two ways: either towards serial derivations, passing fairly abstract underlying inputs through a lengthy chain of transformations (which, in some cases, simply recapitulate the historical changes), or towards more complex types of constraint interactions — sympathy constraints, conjoined constraints, candidate chains, et cetera. Either way, it is difficult to imagine realistic cognitive analogues of such procedures (Lakoff 1993).

I have instead tried to capture the contemporary system of Wubuy phonology as it is presented to speakers, using mechanisms assumed to be general to all cognitive domains: the ability to remember, to notice similarities and differences, to categorise, to generalise, and to predict, among others (Archangeli & Pulleyblank 2017; Bybee 2010). In characterising speakers' knowledge, I have avoided representations with hidden levels of structure (such as underlying and intermediate phonological forms) and procedures which crucially depend on them, instead treating synchronic phonological generalisations as analogical relations between full surface word forms (Blevins 2004; Bybee 2001). To that end, I adopted the Relational Morphology formalism (Jackendoff & Audring 2020), following through the consequences of its basic principles for phonology — an area that has remained under-theorised across all of Construction Grammar. RM does justice to the holistic properties of words while still recognising that they share recurrent features, including phonological ones. These shared recurrent elements organise words into paradigms. By expressing these paradigmatic oppositions as linked schemas abstracting over full word forms, we were able to analyse a wide range of phonological alternations and their interactions — many of them opaque — using only surface phonology. These included ŋu-Epenthesis, Hardening, Nasal Assimilation, Nasal Deletion, Velar Stop Deletion, Reduplication, and Vowel Coalescence. All of these, I showed, could be captured simultaneously via the unification of constructional schemas, which at the same time express these words' templatic morphological organisation.

In taking words as our basic units of phonological and morphological analysis, the view advanced in this thesis does not deny the existence of word-internal structure. Rather, morphological structure emerges from analogical relations between words: we can identify the edges of shared parts of word pairs, even if those parts have no independent existence in the lexicon outside of the parts that instantiate them (Ackerman & Malouf 2016). As we saw in Chapter 6, where I conducted the first detailed analysis of Wubuy's intonational phonology, these major morphological boundaries within words tend to coincide with the boundaries of prosodic constituents — constituents best characterised as prosodic words, in my analysis, tonally marked with an optional pitch accent achored to their leftmost syllable. The sensitivity of prosodic organisation to morphological structure lends credence to the view that the long,

complex words of Wubuy and other polysynthetic languages are like phrases. And indeed, the grammatical verbal word tends to align with at least an Accentual Phrase, and often constitutes a single Intonational Phrase unto itself, as characteristic of syntactic phrases in languages with more analytic morphology. I also presented evidence that this Accentual Phrase is tonally marked at its right edge by means of a pitch accent anchored to its penultimate mora, which is to be found in the final syllable when it is long, and the penultimate otherwise.

7.2 Directions for future research

Certainly, much more could be said about Wubuy, and about the kinds of analysis I have proposed to account for it. Here, I touch on some of the pathways for future investigation that this thesis has opened up.

There are a number of empirical domains of Wubuy phonology that I have not really addressed. For instance, I have analysed vowel coalesence at morphological boundaries, which is the dominant way of responding to hiatus. But there is also evidence that speakers sometimes articulate two vowels in sequence. There is also close similarity between the alternations of vowel-initial morphs and /w_1/-initial morphs, which variably delete the glide — but not quite total identity. These issues deserve further study. The fact that deletion applies to /w_1/ but not ordinarily to /w_2/ is also quite striking, since it adds further complexity to the account of continuant-stop alternations I gave in Chapter 3. It would be worth investigating to what degree speakers are aware of these morphophonological differences, and to what extent they generalise them to novel contexts (like loanwords).

The question of what information in the word is leveraged by speakers would also benefit from different methods, specifically computational modelling. A serious limitation of the RM analysis is that it remains a largely verbal model, and is therefore at risk of being unclear in its predictions. A computational model of the analogical procedures RM proposes for the Wubuy data may reveal very different kinds of generalisation than the highly abstract ones I have proposed here. And indeed, I would expect that such modelling would detect various subregularities and narrow generalisations, for instance at the level of the phoneme rather than the broad features like [CONT] or [SON] (Albright 2002). Another limitation of the RM model is that it, like most discrete formal linguistic models, does not have a particularly felicitous way of representing variation, nor gradient phenomena — despite the fact that many of its assumptions about storage and processing appear to be compatible with exemplar-type models (i.e. Hall et al. 2016). This remains an area to be worked out in more detail, and would again benefit from computational implementation.

Finally, while the present thesis has made significant advances is our understanding of Wubuy

prosody, it is by no means a complete account. The most obvious limitation is the nature of the corpus, which consists of monologic narratives only. By analysing conversation, we would be able to observe a wider range of tune types, with a wider range of discoursive functions, but we would also be able to examine a wider range of segmental contexts for tonal events — especially verbs ending in a nasal coda, which are common in conversation but rare in narratives, and don't appear at all in the present corpus. The account in this thesis also opens up new possibilities for approaching Wubuy intonation with more controlled experimental methodologies, allowing us to subject some of my claims about prosodic organisation to further scrutiny.

7.3 Concluding remarks

What I hope to have conveyed over the course of this thesis, even if only in part, is the value of Wubuy in representing a certain horizon of possibility in human language, where language is understood as a complex adaptive system (Beckner et al. 2009). Although we can find fairly grounded explanations for the pathways by which Wubuy phonology came to be what it is, the resulting system is undeniably one of daunting complexity — full of patterns quite unlike any other, and certainly very different from its neighbours. As I hope to have shown, the dense interactions among all these remarkable patterns provoke new kinds of analysis. This is the genius of Wubuy, the product of many generations of speakers adapting it to their changing needs, repurposing and reanalysing older material and positing new communicative strategies as the shape of the lexicon evolves. Wubuy thus offers an especially compelling case for seeing language as the 'old-growth forest of the mind' (Davis 2019).

Bibliography

Ackerman, Farrell, James P. Blevins & Robert Malouf (2009). 'Parts and Wholes: Implicative Patterns in Inflectional Paradigms'. *Analogy in Grammar: Form and Acquisition*. Ed. by James P. Blevins & Juliette Blevins. Oxford: Oxford University Press.

Ackerman, Farrell & Robert Malouf (2013). 'Morphological Organization: The Low Conditional Entropy Conjecture'. *Language* 89.3, pp. 429–464.

Ackerman, Farrell & Robert Malouf (2016). 'Implicative Relations in Word-based Morphological Systems'. *The Cambridge Handbook of Morphology*. Ed. by Andrew Hippisley & Gregory Stump. Cambridge: Cambridge University Press, pp. 297–328. DOI: 10.1017/9781139814720.012.

Albright, Adam (2002). 'Islands of Reliability for Regular Morphology: Evidence from Italian'. *Language* 78.4, pp. 684–709. DOI: 10.1353/lan.2003.0002.

Albright, Adam (2009). 'Modeling Analogy as Probabilistic Grammar'. *Analogy in Grammar: Form and Acquisition*. Ed. by James P. Blevins & Juliette Blevins. Oxford: Oxford University Press, pp. 185–213.

Albright, Adam & Bruce Hayes (2003). 'Rules vs. Analogy in English Past Tenses: A Computational/Experimental Study'. *Cognition* 90.2, pp. 119–161. DOI: 10.1016/S0010-0277(03)00146-X.

Alderete, John (2001). 'Dominance Effects as Transderivational Anti-Faithfulness'. *Phonology* 18.2, pp. 201–253. DOI: 10.1017/S0952675701004067.

Alderete, John & Kayleigh MacMillan (2015). 'Reduplication in Hawaiian: Variations on a Theme of Minimal Word'. *Natural Language & Linguistic Theory* 33.1, pp. 1–45. DOI: 10.1007/s11049-014-9255-7.

Alpher, Barry, Nicholas Evans & Mark Harvey (2003). 'Proto Gunwinyguan Verb Suffixes'. *The Non-Pama-Nyungan Languages of Northern Australia: Comparative Studies of the Continent's Most Linguistically Complex Region*. Ed. by Nicholas Evans. Studies in Language Change. Canberra: Pacific Linguistics, pp. 305–352.

Ambridge, Ben (2020). 'Abstractions Made of Exemplars or "You're All Right, and I've Changed My Mind": Response to Commentators'. *First Language* 40.5-6, pp. 640–659. DOI: 10.1177/0142723720949723.

Amenta, Simona & Davide Crepaldi (2012). 'Morphological Processing as We Know It: An Analytical Review of Morphological Effects in Visual Word Identification'. *Frontiers in Psychology* 3.232. DOI: 10.3389/fpsyg.2012.00232.

Anderson, Stephen R. (1992). *A-Morphous Morphology*. Cambridge: Cambridge University Press.

Anderson, Stephen R. (2015). 'The Morpheme: Its Nature and Use'. *The Oxford Handbook of Inflection*. Ed. by Matthew Baerman. Oxford: Oxford University Press, pp. 10–34. DOI: 10.1093/oxfordhb/9780199591428.013.2.

Archangeli, Diana & Douglas Pulleyblank (1994). *Grounded Phonology*. MIT Press.

Archangeli, Diana & Douglas Pulleyblank (2017). 'Phonology as an Emergent System'. *The Routledge Handbook of Phonological Theory*. Ed. by S. J. Hannahs & Anna R. K. Bosch. Routledge, pp. 476–503. DOI: 10.4324/9781315675428-17.

Archangeli, Diana & Douglas Pulleyblank (2022). *Emergent Phonology*. Conceptual Foundations of Language Science 7. Berlin: Language Science Press.

Arndt-Lappe, Sabine & Mirjam Ernestus (2020). 'Morpho-Phonological Alternations: The Role of Lexical Storage'. *Word Knowledge and Word Usage: A Cross-Disciplinary Guide to the Mental Lexicon*. Ed. by Vito Pirrelli, Ingo Plag & Wolfgang U Dressler. Berlin: De Gruyter Mouton, pp. 191–227. DOI: 10.1515/9783110440577-006.

Arnhold, Anja (2014). 'Prosodic Structure and Focus Realization in West Greenlandic'. *Prosodic Typology II: The Phonology of Intonation and Phrasing*. Ed. by Sun-Ah Jun. Oxford: Oxford University Press. DOI: 10.1093/acprof:oso/9780199567300.003.0008.

Aronoff, Mark (1994). *Morphology by Itself: Stems and Inflectional Classes*. Linguistic Inquiry Monographs 22. Cambridge, Mass: MIT Press.

Aylett, Matthew & Alice Turk (2004). 'The Smooth Signal Redundancy Hypothesis: A Functional Explanation for Relationships between Redundancy, Prosodic Prominence, and Duration in Spontaneous Speech'. *Language and Speech* 47.1, pp. 31–56. DOI: 10.1177/00238309040470010201.

Baayen, R. Harald, Yu-Ying Chuang, Elnaz Shafaei-Bajestan & James P. Blevins (2019). 'The Discriminative Lexicon: A Unified Computational Model for the Lexicon and Lexical Processing in Comprehension and Production Grounded Not in (De)Composition but in Linear Discriminative Learning'. *Complexity* 2019. DOI: 10.1155/2019/4895891.

Baayen, R. Harald, Ton Dijkstra & Robert Schreuder (1997). 'Singulars and Plurals in Dutch: Evidence for a Parallel Dual-Route Model'. *Journal of Memory and Language* 37.1, pp. 94–117. DOI: 10.1006/jmla.1997.2509.

Baayen, R. Harald, James M. McQueen, Ton Dijkstra & Robert Schreuder (2003). 'Frequency Effects in Regular Inflectional Morphology: Revisiting Dutch Plurals'. *Morphological Structure*

in Language Processing. Ed. by R. Harald Baayen & Robert Schreuder. Trends in Linguistics: Studies and Monographs 151. Berlin: Mouton de Gruyter, pp. 355–390.

Baayen, R. Harald & Fermín Moscoso del Prado Martín (2005). 'Semantic Density and Past-Tense Formation in Three Germanic Languages'. *Language* 81.3, pp. 666–698.

Baayen, R. Harald, Fabian Tomaschek, Susanne Gahl & Michael Ramscar (2017). 'The Ecclesiastes Principle in Language Change'. *The Changing English Language*. Ed. by Marianne Hundt, Sandra Mollin & Simone E. Pfenninger. Cambridge: Cambridge University Press, pp. 21–48. DOI: 10.1017/9781316091746.002.

Baayen, R. Harald, Lee H. Wurm & Joanna Aycock (2007). 'Lexical Dynamics for Low-Frequency Complex Words: A Regression Study across Tasks and Modalities'. *Mental Lexicon* 2.3, pp. 419–463. DOI: 10.1075/ml.2.3.06baa.

Babinski, Sarah (2022). *Archival Phonetics & Prosodic Typology in Sixteen Australian Languages*. PhD thesis. New Haven, Connecticut: Yale University.

Babinski, Sarah & Claire Bowern (2018). 'Mergers in Bardi: Contextual Probability and Predictors of Sound Change'. *Linguistics Vanguard* 4.s2. DOI: 10.1515/lingvan-2017-0024.

Baese-Berk, Melissa, Tuuli H. Morrill & Laura Dilley (2018). 'Predictability and Perception for Native and Non-Native Listeners'. *Linguistics Vanguard* 4.s2. DOI: 10.1515/lingvan-2017-0022.

Baker, Brett J. (2008a). 'The Interpretation of Complex Nominal Expressions in Southeast Arnhem Land Languages'. *Discourse and Grammar in Australian Languages*. Ed. by Ilana Mushin & Brett Baker. Amsterdam: John Benjamins, pp. 135–166.

Baker, Brett J. (2008b). *Word Structure in Ngalakgan*. Stanford, Calif: Center for the Study of Language and Information.

Baker, Brett J. (2009). 'Two Opaque Processes in Wubuy: Issues in Learnability and the Lexicon'. OzPhon. University of Melbourne.

Baker, Brett J. (2014a). 'Incorporation in Wubuy'. *Selected Papers from the 44th Conference of the Australian Linguistic Society, 2013*. Australian Linguistic Society Annual Conference. Ed. by Lauren Gawne & Jill Vaughan. Melbourne: University of Melbourne Minerva Access, pp. 231–260.

Baker, Brett J. (2014b). 'Word Structure in Australian Languages'. *The Languages and Linguistics of Australia: A Comprehensive Guide*. Ed. by Harold Koch & Rachel Nordlinger. Vol. 3. World of Linguistics. Berlin: De Gruyter Mouton, pp. 139–214.

Baker, Brett J. (2018). 'Super-Complexity and the Status of 'Word' in Gunwinyguan Languages of Australia'. *The Construction of Words: Advances in Construction Morphology*. Ed. by

Geert Booij. Studies in Morphology. New York: Springer, pp. 255–286. DOI: 10.1007/978-3-319-74394-3_10.

Baker, Brett J. (2021). 'Stay-k-Ation: Phonological Processes as Evidence of Underlying Phonemic Identity in Wubuy'. Australian Linguistics Society Annual Conference. La Trobe University.

Baker, Brett J. & Rikke L. Bundgaard-Nielsen (2018). 'Reduplication and Lexical Access'. Australian Languages Workshop. Marysville, Victoria.

Baker, Brett J., Rikke L. Bundgaard-Nielsen, Sarah Babinski & Janet Fletcher (2019). 'Acoustic Correlates of Stress in Wubuy'. *Proceedings of ICPhS XIX*. Ed. by Sasha Calhoun, Paola Escudero, Marija Tabain & Paul Warren. Canberra, Australia: Australasian Speech Science and Technology Association Inc., pp. 2729–2733.

Baker, Brett J., Rikke L. Bundgaard-Nielsen, Peter Nyhuis, Hywel Stoakes & Janet Fletcher (2020). 'Lengthening of VC Intervals in Wubuy Stress'. LabPhon17. University of British Columbia.

Baker, Brett J. & Mark Harvey (2022). 'Initial Dropping in Wubuy'. International Conference on Historical Linguistics. University of Oxford.

Baker, Brett J. & Mark Harvey (in press). *The Phonology and Morphology of Australian Languages*. Oxford: Oxford University Press.

Baker, Brett J., Kate Horrack, Rachel Nordlinger & Louisa Sadler (2010). 'Putting It All Together: Agreement, Incorporation, Coordination and External Possession in Wubuy (Australia)'. *Proceedings of LFG10*. Ed. by Miriam Butt & Tracy Holloway King. Stanford, CA: CSLI Publications, pp. 64–84.

Baker, Brett J. & Rachel Nordlinger (2008). 'Noun-Adjective Compounds in Gunwinyguan Languages'. *Proceedings of the LFG08 Conference*. Ed. by Miriam Butt & Tracy Holloway King. Stanford, CA: CSLI Publications, pp. 109–128.

Baker, Rachel E. & Ann R. Bradlow (2009). 'Variability in Word Duration as a Function of Probability, Speech Style, and Prosody.' *Language & Speech* 52.4, pp. 391–413. DOI: 10.1177/0023830909336575.

Baković, Eric (1999). 'Deletion, Insertion and Symmetrical Identity'. *ROA–300*.

Baković, Eric (2007). 'A Revised Typology of Opaque Generalisations'. *Phonology* 24.2, pp. 217–259. DOI: 10.1017/S0952675707001194.

Baković, Eric (2011). 'Opacity and Ordering'. *The Handbook of Phonological Theory*. Ed. by John A. Goldsmith, Jason Riggle & Alan C. L. Yu. Malden, MA: Wiley-Blackwell, pp. 40–67.

Balling, Laura Winther & R. Harald Baayen (2012). 'Probability and Surprisal in Auditory Comprehension of Morphologically Complex Words'. *Cognition* 125.1, pp. 80–106. DOI: 10.1016/j.cognition.2012.06.003.

Barr, Dale J., Roger Levy, Christoph Scheepers & Harry J. Tily (2013). 'Random Effects Structure for Confirmatory Hypothesis Testing: Keep It Maximal'. *Journal of Memory and Language* 68.3, pp. 255–278. DOI: 10.1016/j.jml.2012.11.001.

Beckman, Mary E., Julia Hirschberg & Stefanie Shattuck-Hufnagel (2005). 'The Original ToBi System and the Evolution of the ToBi Framework'. *Prosodic Typology: The Phonology of Intonation and Phrasing*. Ed. by Sun-Ah Jun. Oxford University Press, pp. 9–54. DOI: 10.1093/acprof:oso/9780199249633.003.0002.

Beckner, Clay, Richard Blythe, Joan Bybee, Morten H. Christiansen, William Croft, Nick C. Ellis, John Holland, Jinyun Ke, Diane Larsen-Freeman & Tom Schoenemann (2009). 'Language Is a Complex Adaptive System: Position Paper'. *Language Learning* 59.s1, pp. 1–26. DOI: 10.1111/j.1467-9922.2009.00533.x.

Bednall, James (2019). *Temporal, Aspectual and Modal Expression in Anindilyakwa, the Language of the Groote Eylandt Archipelago, Australia*. PhD thesis. Canberra and Paris: The Australian National University and Université Paris Diderot (Paris 7).

Beguš, Gašper (2018). *Unnatural Phonology: A Synchrony-Diachrony Interface Approach*. PhD thesis. Harvard University.

Bell, Alan, Jason M. Brenier, Michelle Gregory, Cynthia Girand & Dan Jurafsky (2009). 'Predictability Effects on Durations of Content and Function Words in Conversational English'. *Journal of Memory and Language* 60.1, pp. 92–111. DOI: 10.1016/j.jml.2008.06.003.

Bennett, Ryan & Emily Elfner (2019). 'The Syntax–Prosody Interface'. *Annual Review of Linguistics* 5.1, pp. 151–171. DOI: 10.1146/annurev-linguistics-011718-012503.

Benua, Laura (1997). *Transderivational Identity*. PhD thesis. Amherst: University of Massachusetts.

Bermúdez-Otero, Ricardo (1999). *Constraint Interaction in Language Change [Opacity and Globality in Phonological Change]*. PhD thesis. Manchester / Santiago de Compostela: University of Manchester / Universidad de Santiago de Compostela.

Bermúdez-Otero, Ricardo (2003). 'The Acquisition of Phonological Opacity'. *Variation within Optimality Theory: Proceedings of the Stockholm Workshop on Variation within Optimality Theory*. Stockholm: Dept of Linguistics, Stockholm University, pp. 25–36.

Best, Catherine T. & Michael D. Tyler (2007). 'Nonnative and Second-Language Speech Perception: Commonalities and Complementarities'. *Second Language Speech Learning: The Role of Language Experience in Speech Perception and Production*. Ed. by M. J. Munro & O.-S. Bohn. Amsterdam: John Benjamins, pp. 13–34.

Bickel, Balthasar & Fernando Zúñiga (2017). 'The 'Word' in Polysynthetic Languages: Phonological and Syntactic Challenges'. *The Oxford Handbook of Polysynthesis*. Ed. by

Michael Fortescue, Marianne Mithun & Nicholas Evans. Oxford: Oxford University Press, pp. 158–185.

Bird, Steven (1990). *Constraint-Based Phonology*. PhD thesis. Edinburgh: University of Edinburgh.

Bird, Steven & T. Mark Ellison (1994). 'One-Level Phonology: Autosegmental Representations and Rules as Finite Automata'. *Computational Linguistics* 20.1, pp. 55–90.

Bishop, Judith (2002). *Aspects of Intonation and Prosody in Bininj Gun-wok: An Autosegmental-Metrical Analysis*. PhD thesis. Melbourne: University of Melbourne.

Bishop, Judith & Janet Fletcher (2005). 'Intonation in Six Dialects of Bininj Gun-wok'. *Prosodic Typology: The Phonology of Intonation and Phrasing*. Ed. by Sun-Ah Jun. Oxford University Press, pp. 331–361. DOI: 10.1093/acprof:oso/9780199249633.003.0012.

Blevins, James P. (2006). 'Word-Based Morphology'. *Journal of Linguistics* 42.3, pp. 531–573. DOI: 10.1017/S0022226706004191.

Blevins, Juliette (2001). 'Where Have All the Onsets Gone? Initial Consonant Loss in Australian Aboriginal Languages'. *Forty Years on: Ken Hale and Australian Languages*. Ed. by Jane Simpson, David Nash, Mary Laughren, Peter Austin & Barry Alpher. Canberra: Pacific Linguistics Canberra, p. 481.

Blevins, Juliette (2003). 'A Note on Reduplication in Bugotu and Cheke Holo'. *Oceanic Linguistics* 42.2, pp. 499–505. DOI: 10.2307/3623248.

Blevins, Juliette (2004). *Evolutionary Phonology: The Emergence of Sound Patterns*. Cambridge: Cambridge University Press.

Blevins, Juliette (2005). 'The Role of Phonological Predictability in Sound Change: Privileged Reduction in Oceanic Reduplicated Substrings'. *Oceanic Linguistics* 44.2, pp. 517–526. DOI: 10.1353/ol.2005.0028.

Blevins, Juliette (2008). 'Consonant Epenthesis: Natural and Unnatural Histories'. *Linguistic Universals and Language Change*. Ed. by Jeff Good. Oxford: Oxford University Press, pp. 79–107. DOI: 10.1093/acprof:oso/9780199298495.003.0004.

Blevins, Juliette & Andrew Wedel (2009). 'Inhibited Sound Change: An Evolutionary Approach to Lexical Competition'. *Diachronica* 26.2, pp. 143–183.

Bombien, Lasse, Raphael Winkelmann & Michel Scheffers (2021). *Wrassp: An R Wrapper to the ASSP Library*.

Bonami, Olivier & Sarah Beniamine (2016). 'Joint Predictiveness in Inflectional Paradigms'. *Word Structure* 9.2, pp. 156–182. DOI: 10.3366/word.2016.0092.

Bonet, Eulàlia (2004). 'Morph Insertion and Allomorphy in Optimality Theory'. *International Journal of English Studies* 4.2, pp. 73–104.

Bonnin, Cicely (2014). *Intonation and Prosody in Yukulta: A Tangkic Language of North West Queensland*. MA thesis. Brisbane: University of Queensland.

Booij, Geert (1999). 'Lexical Storage and Regular Processes'. *Behavioral and Brain Sciences* 22.6, pp. 1016–1016. DOI: 10.1017/S0140525X99242227.

Booij, Geert (2010). *Construction Morphology*. Oxford: Oxford University Press.

Booij, Geert (2012). 'Allomorphy and the Architecture of Grammar'. *Phonological Explorations*. Ed. by Bert Botma & Roland Noske. Berlin, Boston: De Gruyter. DOI: 10.1515/9783110295177.9.

Booij, Geert (2014). 'Language Use and the Architecture of Grammar: A Construction Morphology Perspective'. *Jezična uporaba i ustroj gramatike: perspektiva konstrukcijske morfologije*. 40.78, pp. 193–212.

Booij, Geert (2016). 'Construction Morphology'. *The Cambridge Handbook of Morphology*. Ed. by Andrew Hippisley & Gregory Stump. Cambridge: Cambridge University Press, pp. 424–448. DOI: 10.1017/9781139814720.016.

Booij, Geert & Jerzy Rubach (1984). 'Morphological and Prosodic Domains in Lexical Phonology'. *Phonology Yearbook* 1, pp. 1–27.

Borowsky, Toni (1993). 'On the Word Level'. *Studies in Lexical Phonology*. Ed. by Sharon Hargus & Ellen M. Kaisse. Vol. 4. Phonetics and Phonology. Boston: Academic Press, pp. 199–234. DOI: 10.1016/B978-0-12-325071-1.50013-3.

Boudelaa, Sami & William D. Marslen-Wilson (2015). 'Structure, Form, and Meaning in the Mental Lexicon: Evidence from Arabic'. *Language, Cognition and Neuroscience* 30.8, pp. 955–992. DOI: 10.1080/23273798.2015.1048258.

Bradley, Travis G (2002). 'Gestural Timing and Derived Environment Effects in Norwegian Clusters'. *WCCFL 21 Proceedings*. Ed. by Line Mikkelsen & Christopher Potts. Somerville, MA: Cascadilla Press, pp. 43–56.

Breen, Gavan & Rob Pensalfini (1999). 'Arrernte: A Language with No Syllable Onsets'. *Linguistic Inquiry* 30.1, pp. 1–25.

Brooks, Mollie E., Kasper Kristensen, Koen J. van Benthem, Arni Magnusson, Casper W. Berg, Anders Nielsen, Hans J. Skaug, Martin Maechler & Benjamin M. Bolker (2017). 'glmmTMB Balances Speed and Flexibility among Packages for Zero-Inflated Generalized Linear Mixed Modeling'. *The R Journal* 9.2, pp. 378–400. DOI: 10.32614/RJ-2017-066.

Browman, Catherine P. & Louis Goldstein (1989). 'Articulatory Gestures as Phonological Units'. *Phonology* 6.02, pp. 201–251. DOI: 10.1017/S0952675700001019.

Browman, Catherine P. & Louis Goldstein (1992). 'Articulatory Phonology: An Overview'. *Phonetica* 49.3-4, pp. 155–180.

Brugman, H & A Russell (2004). 'Annotating Multimedia/Multi-modal Resources with ELAN'. *Proceedings of LREC 2004*. Fourth International Conference on Language Resources and Evaluation.

Bundgaard-Nielsen, Rikke L. & Brett J. Baker (2014). 'Frequency in the Input Affects Perception of Phonological Contrasts for Native Speakers'. *Proceedings of the 15th Australasian International Speech Science and Technology Conference*. Parramatta, Australia. Ed. by Christopher Carignan & Michael D. Tyler. ASSTA, pp. 205–208.

Bundgaard-Nielsen, Rikke L. & Brett J. Baker (2020). 'Pause Acceptability Indicates Word-Internal Structure in Wubuy'. *Cognition* 198.104167. DOI: 10.1016/j.cognition.2019.104167.

Bundgaard-Nielsen, Rikke L., Brett J. Baker, Mark Harvey, Catherine T. Best & Christian Kroos (2012). 'The Perception of Coronal Stops in Wubuy'. *Proceedings of the 14th Australasian International Conference on Speech Science and Technology (SST)*. Ed. by Felicity Cox, Katherine Demuth, Susan Lin, Kelly Miles, Sallyanne Palethorpe, Jason Shaw & Ivan Yuen. Canberra: ASSTA, pp. 217–220.

Bundgaard-Nielsen, Rikke L., Brett J. Baker, Christian Kroos, Mark Harvey & Catherine T. Best (2012). 'Vowel Acoustics Reliably Differentiate Three Coronal Stops of Wubuy across Prosodic Contexts'. *Laboratory Phonology* 29, pp. 133–161.

Bundgaard-Nielsen, Rikke L., Brett J. Baker, Christian Kroos, Mark Harvey & Catherine T. Best (2015). 'Discrimination of Multiple Coronal Stop Contrasts in Wubuy (Australia): A Natural Referent Consonant Account'. *PLOS ONE* 10.12. Ed. by Niels O. Schiller. DOI: 10.1371/journal.pone.0142054.

Bundgaard-Nielsen, Rikke L., Christian Kroos, Brett J. Baker, Catherine T. Best & Mark Harvey (2016). 'Consonantal Timing and Release Burst Acoustics Distinguish Multiple Coronal Stop Place Distinctions in Wubuy (Australia)'. *The Journal of the Acoustical Society of America* 140.4, pp. 2794–2809. DOI: 10.1121/1.4964399.

Bundgaard-Nielsen, Rikke L., Christian Kroos, Mark Harvey, Catherine Best, Brett J. Baker & Louis Goldstein (2010). 'A Kinematic Analysis of Temporal Differentiation of the Four-Way Coronal Stop Contrast in Wubuy (Australia)'. *Proceedings of the 13th Australasian International Conference on Speech Science and Technology [SST]*.

Burzio, Luigi (2002). 'Missing Players: Phonology and the Past-Tense Debate'. *Lingua* 112.3, pp. 157–199. DOI: 10.1016/S0024-3841(01)00041-9.

Burzio, Luigi (2005). 'Lexicon and Grammar: Unequal but Inseparable'. Unpublished manuscript. Johns Hopkins University, Baltimore.

Busby, Peter Andrew (1979). *A Classificatory Study of Phonemic Systems in Australian Aboriginal Languages*. MA thesis. Canberra: Australian National University.

Butcher, Andrew (1999). 'What Speakers of Australian Aboriginal Languages Do with Their Velums and Why: The Phonetics of the Nasal/Oral Contrast'. *Proceedings of the XIVth International Congress of Phonetic Sciences Berkeley: ICPhS*, pp. 479–482.

Butcher, Andrew (2006). 'Australian Aboriginal Languages: Consonant-Salient Phonologies and the "Place-of-Articulation Imperative"'. *Speech Production*. Ed. by Jonathan Harrington & Marija Tabain. New York: Psychology Press, pp. 187–210.

Butcher, Andrew (2012). 'On the Phonetics of Long, Thin Phonologies'. *Quantitative Approaches to Problems in Linguistics*. Ed. by Cathryn Donohue, Shunichi Ishihara & William Steed. Munich: Lincom, pp. 133–154.

Butcher, Andrew & Jonathan Harrington (2003). 'An Acoustic and Articulatory Analysis of Focus and the Word/Morpheme Boundary Distinction in Warlpiri'. *Proceedings of the 6th International Seminar on Speech Production*. Sydney, pp. 19–24.

Butcher, Andrew & Debbie Loakes (2008). 'Enhancing the Left Edge: The Phonetics of Prestopped Sonorants in Australian Languages.' *The Journal of the Acoustical Society of America* 124.4, pp. 2527–2527. DOI: 10.1121/1.4782973.

Butcher, Andrew & Marija Tabain (2004). 'On the Back of the Tongue: Dorsal Sounds in Australian Languages'. *Phonetica* 61.1, pp. 22–52. DOI: 10.1159/000078661.

Bybee, Joan L. (1999). 'Usage-Based Phonology'. *Functionalism and Formalism in Linguistics*. Ed. by Michael Darnell, Edith Moravcsik, Frederick J. Newmeyer, Michael Noonan & Kathleen Wheatley. Vol. 1: General Papers. Amsterdam: John Benjamins Publishing, pp. 211–242.

Bybee, Joan L. (2001). *Phonology and Language Use*. Cambridge: Cambridge University Press.

Bybee, Joan L. (2002a). 'Sequentiality as the Basis of Constituent Structure'. *Evolution of Language out of Pre-language*. Ed. by T. Givón & Bertram F. Malle. Philadelphia: John Benjamins, pp. 109–134.

Bybee, Joan L. (2002b). 'Word Frequency and Context of Use in the Lexical Diffusion of Phonetically Conditioned Sound Change'. *Language Variation and Change* 14.3, pp. 261–290. DOI: 10.1017/S0954394502143018.

Bybee, Joan L. (2006). 'From Usage to Grammar: The Mind's Response to Repetition'. *Language* 82.4, pp. 711–733.

Bybee, Joan L. (2010). *Language, Usage and Cognition*. Cambridge: Cambridge University Press.

Bybee, Joan L. (2012). 'Patterns of Lexical Diffusion and Articulatory Motivation for Sound Change'. *The Initiation of Sound Change: Perception, Production, and Social Factors*. Ed. by Maria-Josep Solé & Daniel Recasens. Amsterdam: John Benjamins, pp. 211–234.

Bybee, Joan L. & Shelece Easterday (2019). 'Consonant Strengthening: A Crosslinguistic Survey and Articulatory Proposal'. *Linguistic Typology* 23.2, pp. 263–302. DOI: 10.1515/lingty-2019-0015.

Bybee, Joan L. & Carol Lynn Moder (1983). 'Morphological Classes as Natural Categories'. *Language* 59.2, pp. 251–270. DOI: 10.2307/413574.

Bye, Patrik (2006). 'Allomorphy: Selection, Not Optimization'. *Freedom of Analysis*. Ed. by Sylvia Blaho, Patrik Bye & Martin Kraemer. Berlin: Mouton de Gruyter, pp. 63–91.

Caballero, Gabriela (2006). '"Templatic Backcopying" in Guarijio Abbreviated Reduplication'. *Morphology* 16.2, pp. 273–289. DOI: 10.1007/s11525-007-9108-8.

Caballero, Gabriela (2008). *Choguita Rarámuri (Tarahumara) Phonology and Morphology*. PhD thesis. University of California, Berkeley.

Caballero, Gabriela & Sharon Inkelas (2013). 'Word Construction: Tracing an Optimal Path through the Lexicon'. *Morphology* 23.2, pp. 103–143. DOI: 10.1007/s11525-013-9220-x.

Carignan, Christopher, Juqiang Chen, Mark Harvey, Clara Stockigt, Jane Simpson & Sydney Strangways (2023). 'An Investigation of the Dynamics of Vowel Nasalization in Arabana Using Machine Learning of Acoustic Features'. *Laboratory Phonology* 14.1, pp. 1–31.

Carr, Therese, Melanie Wilkinson & Philippa Stansell (2017). 'Language Revitalisation in a Bilingual Program—The Case of Numbulwar School'. *History of Bilingual Education in the Northern Territory: People, Programs and Policies*. Ed. by Brian Clive Devlin, Samantha Disbray & Nancy Regine Friedman Devlin. Language Policy. Singapore: Springer, pp. 149–163. DOI: 10.1007/978-981-10-2078-0_13.

Carstairs, Andrew (1990). 'Phonologically Conditioned Suppletion'. *Contemporary Morphology*. Ed. by Wolfgang U. Dressler, Hans C. Luschützky, Oskar E. Pfeiffer, John R. Rennison & Hans C. Luschützky. Berlin: De Gruyter, pp. 17–23.

Carstairs-McCarthy, Andrew (2001). 'Grammatically Conditioned Allomorphy, Paradigmatic Structure, and the Ancestry Constraint'. *Transactions of the Philological Society* 99.2, pp. 223–245. DOI: 10.1111/1467-968X.00081.

Cho, Taehong (2001). 'Effects of Morpheme Boundaries on Intergestural Timing: Evidence from Korean'. *Phonetica* 58.3, pp. 129–162. DOI: 10.1159/000056196.

Cho, Taehong & Patricia Keating (2009). 'Effects of Initial Position versus Prominence in English'. *Journal of Phonetics* 37.4, pp. 466–485. DOI: 10.1016/j.wocn.2009.08.001.

Chomsky, Noam (2015). *Aspects of the Theory of Syntax*. 50th Anniversary Edition. Cambridge, Massachusetts: The MIT Press.

Chomsky, Noam & Morris Halle (1968). *The Sound Pattern of English*. Studies in Language. New York: Harper & Row.

Chong, Adam J. (2011). 'Lenition in Gaalpu: An Optimality Theoretic Analysis'. *Australian Journal of Linguistics* 31.4, pp. 473–490. DOI: 10.1080/07268602.2011.625601.

Clark, John, Colin Yallop & Janet Fletcher (2007). *An Introduction to Phonetics and Phonology*. 3rd ed. Malden, MA: Blackwell.

Cohen Priva, Uriel (2015). 'Informativity Affects Consonant Duration and Deletion Rates'. *Laboratory Phonology* 6.2, pp. 243–278. DOI: 10.1515/lp-2015-0008.

Cohen Priva, Uriel (2017). 'Informativity and the Actuation of Lenition'. *Language* 93.3, pp. 569–597. DOI: 10.1353/lan.2017.0037.

Cohen Priva, Uriel & Emily Gleason (2020). 'The Causal Structure of Lenition: A Case for the Causal Precedence of Durational Shortening'. *Language* 96.2, pp. 413–448. DOI: 10.1353/lan.2020.0025.

Cohen Priva, Uriel & T. Florian Jaeger (2018). 'The Interdependence of Frequency, Predictability, and Informativity in the Segmental Domain'. *Linguistics Vanguard* 4.s2. DOI: 10.1515/lingvan-2017-0028.

Cole, Jennifer & José Ignacio Hualde (2011). 'Underlying Representations'. *The Blackwell Companion to Phonology*. Ed. by Marc van Oostendorp, Colin J. Ewen, Elizabeth Hume & Keren Rice. Vol. 1. Malden, MA: Wiley-Blackwell, pp. 1–26.

Coleman, John (1990). 'Unification Phonology: Another Look at "Synthesis-by-Rule"'. *COLING 1990 Volume 3: Papers Presented to the 13th International Conference on Computational Linguistics*.

Coleman, John (1998). *Phonological Representations: Their Names, Forms and Powers*. Cambridge: Cambridge University Press.

Creemers, Ava, Amy Goodwin Davies, Robert J. Wilder, Meredith Tamminga & David Embick (2020). 'Opacity, Transparency, and Morphological Priming: A Study of Prefixed Verbs in Dutch'. *Journal of Memory and Language* 110. DOI: 10.1016/j.jml.2019.104055.

Crepaldi, Davide, Kathleen Rastle, Max Coltheart & Lyndsey Nickels (2010). ''Fell' Primes 'Fall', but Does 'Bell' Prime 'Ball'? Masked Priming with Irregularly-Inflected Primes'. *Journal of Memory and Language* 63.1, pp. 83–99. DOI: 10.1016/j.jml.2010.03.002.

Croft, William (2001). *Radical Construction Grammar: Syntactic Theory in Typological Perspective*. Oxford: Oxford University Press.

Culhane, Kirsten (2022). 'The Prosodic Foot beyond Prosodic Prominence: A Preliminary Survey'. *Linguistic Typology*. DOI: 10.1515/lingty-2022-0039.

Culicover, Peter W. & Ray Jackendoff (2005). *Simpler Syntax*. Oxford: Oxford University Press.

Culicover, Peter W. & Ray Jackendoff (2012). 'Same-except: A Domain-General Cognitive Relation and How Language Expresses It'. *Language* 88.2, pp. 305–340. DOI: 10.1353/lan.2012.0031.

Cutler, Anne (2012). *Native Listening: Language Experience and the Recognition of Spoken Words*. Cambridge, Mass.: MIT Press.

Czaplicki, Bartłomiej (2013). 'Arbitrariness in Grammar: Palatalization Effects in Polish'. *Lingua* 123, pp. 31–57. DOI: 10.1016/j.lingua.2012.10.002.

Daland, Robert & Kie Zuraw (2018). 'Loci and Locality of Informational Effects on Phonetic Implementation'. *Linguistics Vanguard* 4.s2. DOI: 10.1515/lingvan-2017-0045.

Davidson, Lisa (2006). 'Schwa Elision in Fast Speech: Segmental Deletion or Gestural Overlap?' *Phonetica* 63.2-3, pp. 79–112. DOI: 10.1159/000095304.

Davis, Maya & Melissa A. Redford (2019). 'The Emergence of Discrete Perceptual-Motor Units in a Production Model That Assumes Holistic Phonological Representations'. *Frontiers in Psychology* 10, p. 2121. DOI: 10.3389/fpsyg.2019.02121.

Davis, Wade (2019). *Why Indigenous Languages Matter*. Canadian Geographic. URL: https://canadiangeographic.ca/articles/why-indigenous-languages-matter/ (visited on 23/12/2023).

Deal, Amy Rose & Matthew Adam Wolf (2017). 'Outwards-Sensitive Phonologically-Conditioned Allomorphy in Nez Perce'. *The Morphosyntax-Phonology Connection: Locality and Directionality at the Interface*. Ed. by Vera Gribanova & Stephanie S. Shih. Oxford: Oxford University Press, pp. 29–60.

De Lacy, Paul (2014). 'Evaluating Evidence for Stress Systems'. *Word Stress: Theoretical and Typological Issues*. Ed. by Harry van der Hulst. Cambridge: Cambridge University Press, pp. 149–193. DOI: 10.1017/CBO9781139600408.008.

Delattre, Pierre C., Alvin M. Liberman & Franklin S. Cooper (1955). 'Acoustic Loci and Transitional Cues for Consonants'. *The Journal of the Acoustical Society of America* 27.4, pp. 769–773. DOI: 10.1121/1.1908024.

Di Sciullo, Anna-Maria & Edwin Williams (1987). *On the Definition of Word*. Linguistic Inquiry Monographs. Cambridge, Massachusetts: The MIT Press.

Dickson, Gregory F. (2015). *Marra and Kriol: The Loss and Maintenance of Knowledge across a Language Shift Boundary*. PhD thesis. Canberra: The Australian National University.

Dinkin, Aaron J. (2008). 'The Real Effect of Word Frequency on Phonetic Variation'. *University of Pennsylvania Working Papers in Linguistics* 14.1, p. 8.

Dixon, Robert M. W. (1977). 'Some Phonological Rules in Yidiny'. *Linguistic Inquiry* 8.1, pp. 1–34.

Dixon, Robert M. W. (1980). *The languages of Australia*. Cambridge language surveys. Cambridge: Cambridge University Press.

Dixon, Robert M. W. & Alexandra Y. Aikhenvald (2002). 'Word: A Typological Framework'. *Word: A Cross-Linguistic Typology*. Ed. by Robert M. W. Dixon & Alexandra Y. Aikhenvald. Cambridge: Cambridge University Press, pp. 1–41.

Drager, Katie K. (2011). 'Sociophonetic Variation and the Lemma'. *Journal of Phonetics*. Cross-Language Speech Perception and Variations in Linguistics Experience 39.4, pp. 694–707. DOI: 10.1016/j.wocn.2011.08.005.

Dupoux, Emmanuel, Sharon Peperkamp & Núria Sebastián-Gallés (2010). 'Limits on Bilingualism Revisited: Stress 'Deafness' in Simultaneous French–Spanish Bilinguals'. *Cognition* 114.2, pp. 266–275. DOI: 10.1016/j.cognition.2009.10.001.

Ernestus, Mirjam & R. Harald Baayen (2004). 'Analogical Effects in Regular Past Tense Production in Dutch'. *Linguistics* 42.5, pp. 873–903. DOI: 10.1515/ling.2004.031.

Evans, Nicholas (1995). 'Current Issues in the Phonology of Australian Languages'. *The Handbook of Phonological Theory*. Ed. by John A. Goldsmith. Cambridge, Mass.: Blackwell, pp. 723–761.

Evans, Nicholas (2003). *Bininj Gun-Wok: A Pan-Dialectal Grammar of Mayali, Kunwinjku and Kune*. Pacific Linguistics 541. Canberra: Pacific Linguistics.

Evans, Nicholas, Janet Fletcher & Belinda Ross (2008). 'Big Words, Small Phrases: Mismatches between Pause Units and the Polysynthetic Word in Dalabon'. *Linguistics* 46.1. DOI: 10.1515/LING.2008.004.

Fedorenko, Evelina, Idan Asher Blank, Matthew Siegelman & Zachary Mineroff (2020). 'Lack of Selectivity for Syntax Relative to Word Meanings throughout the Language Network'. *Cognition* 203, p. 104348. DOI: 10.1016/j.cognition.2020.104348.

Fillmore, Charles J., Paul Kay & Mary Catherine O'Connor (1988). 'Regularity and Idiomaticity in Grammatical Constructions: The Case of Let Alone'. *Language* 64.3, pp. 501–538. DOI: 10.2307/414531.

Finkel, Raphael & Gregory T. Stump (2009). 'Principal Parts and Degrees of Paradigmatic Transparency'. *Analogy in Grammar: Form and Acquisition*. Ed. by James P. Blevins & Juliette Blevins. Oxford: Oxford University Press, pp. 13–53.

Fletcher, Janet (2010). 'The Prosody of Speech: Timing and Rhythm'. *The Handbook of Phonetic Sciences*. Ed. by William J. Hardcastle, John Laver & Fiona E. Gibbon. 1st ed. Wiley, pp. 521–602. DOI: 10.1002/9781444317251.ch15.

Fletcher, Janet (2014). 'Intonation and Prosody in Dalabon'. *Prosodic Typology II: The Phonology of Intonation and Phrasing*. Ed. by Sun-Ah Jun. Oxford: Oxford University Press, pp. 252–272.

Fletcher, Janet (2023). 'Intonation'. *The Oxford Guide to Australian Languages*. Ed. by Claire Bowern. Oxford: Oxford University Press, pp. 159–164.

Fletcher, Janet & Andrew Butcher (2014). 'Sound Patterns of Australian Languages'. *The Languages and Linguistics of Australia: A Comprehensive Guide*. Ed. by Harold Koch & Rachel Nordlinger. Vol. 3. World of Linguistics. Berlin: De Gruyter Mouton, pp. 91–138.

Fletcher, Janet & Nicholas Evans (2000). 'Intonational Downtrends in Mayali'. *Australian Journal of Linguistics* 20.1, pp. 23–38. DOI: 10.1080/07268600050003346.

Fletcher, Janet, Nicholas Evans & Belinda Ross (2004). 'Pausing Strategies and Prosodic Boundaries in Dalabon'. *Proceedings of the Tenth Australian International Conference on Speech Science and Technology*, pp. 436–9.

Fletcher, Janet, Nicholas Evans & Erich R. Round (2002). 'Left-Edge Tonal Events in Kayardild (Australian) — a Typological Perspective'. *ISCA Archive*. Speech Prosody 2002. Aix-en-Provence.

Fletcher, Janet, Hywel Stoakes, Ruth Singer & Deborah Loakes (2016). 'Intonational Correlates of Subject and Object Realisation in Mawng (Australian)'. *Speech Prosody 2016*. Speech Prosody 2016. ISCA, pp. 188–192. DOI: 10.21437/SpeechProsody.2016-39.

Fosler-Lussier, Eric & Nelson Morgan (1999). 'Effects of Speaking Rate and Word Frequency on Pronunciations in Conversational Speech'. *Speech Communication* 29.2, pp. 137–158. DOI: 10.1016/S0167-6393(99)00035-7.

Foulkes, Paul, Gerry Docherty, Stefanie Shattuck Hufnagel & Vincent Hughes (2018). 'Three Steps Forward for Predictability. Consideration of Methodological Robustness, Indexical and Prosodic Factors, and Replication in the Laboratory'. *Linguistics Vanguard* 4.s2. DOI: 10.1515/lingvan-2017-0032.

Frisch, Stefan A. (2017). 'Exemplar Theories in Phonology'. *The Routledge Handbook of Phonological Theory*. Ed. by S. J. Hannahs & Anna R. K. Bosch. London: Routledge, pp. 553–568. DOI: 10.4324/9781315675428-20.

Fry, D. B. (1955). 'Duration and Intensity as Physical Correlates of Linguistic Stress'. *The Journal of the Acoustical Society of America* 27.4, pp. 765–768. DOI: 10.1121/1.1908022.

Fujimura, O., M.J. Macchi & L.A. Streeter (1978). 'Perception of Stop Consonants with Conflicting Transitional Cues: A Cross-Linguistic Study'. *Language and Speech* 21.4, pp. 337–346. DOI: 10.1177/002383097802100408.

Gafos, Diamandis (1998). 'A-Templatic Reduplication'. *Linguistic Inquiry* 29.3, pp. 515–527.

Gagné, Christina L. (2017). 'Psycholinguistic Approaches to Morphology: Theoretical Issues'. *Oxford Research Encyclopedia of Linguistics*. Oxford: Oxford University Press. DOI: 10.1093/acrefore/9780199384655.013.258.

Gahl, Susanne (2008). 'Time and Thyme Are Not Homophones: The Effect of Lemma Frequency on Word Durations in Spontaneous Speech'. *Language* 84.3, pp. 474–496. DOI: 10.1353/lan.0.0035.

Geertzen, Jeroen, James P. Blevins & Petar Milin (2016). 'The Informativeness of Linguistic Unit Boundaries'. *Italian Journal of Linguistics* 28.2, pp. 1–24.

Goddard, Cliff (1985). *A Grammar of Yankunytjatjara*. Alice Springs: Institute of Aboriginal Development.

Goldberg, Adele E. (1995). *Constructions: A Construction Grammar Approach to Argument Structure*. University of Chicago Press.

Goldinger, Stephen D. (1998). 'Echoes of Echoes? An Episodic Theory of Lexical Access'. *Psychological Review* 105.2, pp. 251–279. DOI: 10.1037/0033-295X.105.2.251.

Goldsmith, John Anton (1976). *Autosegmental Phonology*. PhD thesis. Massachusetts Institute of Technology.

Golston, Chris (1996). 'Direct Optimality Theory: Representation as Pure Markedness'. *Language* 72.4, p. 713. DOI: 10.2307/416100.

Gordon, Matthew (2005). 'A Perceptually-Driven Account of Onset-Sensitive Stress'. *Natural Language & Linguistic Theory* 23.3, pp. 595–653. DOI: 10.1007/s11049-004-8874-9.

Gordon, Matthew (2014). 'Disentangling Stress and Pitch-Accent: A Typology of Prominence at Different Prosodic Levels'. *Word Stress*. Ed. by Harry van der Hulst. Cambridge: Cambridge University Press, pp. 83–118. DOI: 10.1017/CBO9781139600408.005.

Gordon, Matthew (2016). 'Metrical Incoherence: Diachronic Sources and Synchronic Analysis'. *Dimensions of Phonological Stress*. Ed. by Harry van der Hulst, Jeffrey Heinz & Rob Goedemans. Cambridge: Cambridge University Press, pp. 9–48. DOI: 10.1017/9781316212745.002.

Gordon, Matthew & Timo Roettger (2017). 'Acoustic Correlates of Word Stress: A Cross-Linguistic Survey'. *Linguistics Vanguard* 3.1. DOI: 10.1515/lingvan-2017-0007.

Gow, David W., Enes Avcu, Adriana Schoenhaut, David O. Sorensen & Seppo P. Ahlfors (2023). 'Abstract Representations in Temporal Cortex Support Generative Linguistic Processing'. *Language, Cognition and Neuroscience* 38.6, pp. 765–778. DOI: 10.1080/23273798.2022.2157029.

Graetzer, Simone, Janet Fletcher & John Hajek (2015). 'Locus Equations and Coarticulation in Three Australian Languages'. *The Journal of the Acoustical Society of America* 137.2, pp. 806–821. DOI: 10.1121/1.4904488.

Gurevich, Olga (2006). *Constructional Morphology: The Georgian Version*. PhD thesis. Berkeley: University of California, Berkeley.

Gussenhoven, Carlos (2005). 'Transcription of Dutch Intonation'. *Prosodic Typology: The Phonology of Intonation and Phrasing*. Ed. by Sun-Ah Jun. Oxford University Press, pp. 118–145. DOI: 10.1093/acprof:oso/9780199249633.003.0005.

Hall, Kathleen Currie (2012). 'Phonological Relationships: A Probabilistic Model'. *McGill Working Papers in Linguistics* 22.1, pp. 1–14.

Hall, Kathleen Currie, Elizabeth Hume, T. Florian Jaeger & Andrew Wedel (2016). 'The Message Shapes Phonology'. Unpublished manuscript.

Halle, Morris & Jean-Roger Vergnaud (1987). *An Essay on Stress*. Current Studies in Linguistics Series 15. Cambridge, Mass.: MIT Press.

Hamilton, Philip (1996). *Phonetic Constraints and Markedness in the Phonotactics of Australian Aboriginal Languages*. PhD thesis. Toronto: University of Toronto.

Hammond, Michael (2000). 'There Is No Lexicon!' *Coyote Papers* 10, pp. 55–77.

Hansson, Gunnar Ólafur (2007). 'Blocking Effects in Agreement by Correspondence'. *Linguistic Inquiry* 38.2, pp. 395–409.

Harmon, Zara & Vsevolod Kapatsinski (2017). 'Putting Old Tools to Novel Uses: The Role of Form Accessibility in Semantic Extension'. *Cognitive Psychology* 98, pp. 22–44. DOI: 10.1016/j.cogpsych.2017.08.002.

Harvey, Mark (1997a). 'Head and Agreement Classes: An Areal Perspective'. *Nominal Classification in Aboriginal Australia*. Ed. by Mark Harvey & Nicholas Reid. Studies in Language Companion Series 37. Amsterdam: John Benjamins Publishing Company, pp. 147–164. DOI: 10.1075/slcs.37.09har.

Harvey, Mark (1997b). 'Nominal Classification and Gender in Aboriginal Australia'. *Nominal Classification in Aboriginal Australia*. Ed. by Mark Harvey & Nicholas Reid. Studies in Language Companion Series 37. Amsterdam: John Benjamins, pp. 17–62. DOI: 10.1075/slcs.37.06har.

Harvey, Mark (2002). *A Grammar of Gaagudju*. Berlin: Mouton de Gruyter.

Harvey, Mark & Brett J. Baker (2020). 'Epenthetic Prefixation in Alawa and Marra'. *Australian Journal of Linguistics* 40.3, pp. 273–295. DOI: 10.1080/07268602.2020.1803209.

Harvey, Mark, Susan Lin, Myfany Turpin, Ben Davies & Katherine Demuth (2015). 'Contrastive and Non-contrastive Pre-stopping in Kaytetye'. *Australian Journal of Linguistics* 35.3, pp. 232–250. DOI: 10.1080/07268602.2015.1023170.

Harvey, Mark, Nay San, Margaret Carew, Sydney Strangways, Jane Simpson & Clara Stockigt (2019). 'Pre-Stopping in Arabana'. *Australian Journal of Linguistics* 39.4, pp. 419–462. DOI: 10.1080/07268602.2019.1643290.

Haspelmath, Martin (2020). 'The Morph as a Minimal Linguistic Form'. *Morphology* 30.2, pp. 117–134. DOI: 10.1007/s11525-020-09355-5.

Hauser, Ivy & Coral Hughto (2020). 'Analyzing Opacity with Contextual Faithfulness Constraints'. *Glossa* 5.1 (1), p. 82. DOI: 10.5334/gjgl.966.

Hay, Jennifer (2003). *Causes and Consequences of Word Structure*. New York: Routledge.

Hay, Jennifer & Paul Foulkes (2016). 'The Evolution of Medial /t/ over Real and Remembered Time'. *Language* 92.2, pp. 298–330. DOI: 10.1353/lan.2016.0036.

Hay, Jennifer, Janet B. Pierrehumbert, Abby J. Walker & Patrick LaShell (2015). 'Tracking Word Frequency Effects through 130 Years of Sound Change'. *Cognition* 139, pp. 83–91. DOI: 10.1016/j.cognition.2015.02.012.

Hayes, Bruce (1995). *Metrical Stress Theory*. Chicago: University of Chicago Press.

Hayes, Bruce (1999). 'Phonetically Driven Phonology: The Role of Optimality Theory and Inductive Grounding'. *Functionalism and Formalism in Linguistics*. Ed. by Michael Darnell, Edith Moravcsik, Frederick J. Newmeyer, Michael Noonan & Kathleen Wheatley. Vol. 1: General Papers. Amsterdam: John Benjamins Publishing, pp. 243–286.

Heath, Jeffrey (1978a). *Linguistic Diffusion in Arnhem Land*. Canberra: Australian Institute of Aboriginal Studies.

Heath, Jeffrey (1978b). *Ngandi Grammar, Texts, and Dictionary*. Canberra: Australian Institute of Aboriginal Studies.

Heath, Jeffrey (1980). *Nunggubuyu Myths and Ethnographic Texts*. Canberra: Australian Institute of Aboriginal Studies.

Heath, Jeffrey (1982). *Nunggubuyu Dictionary*. Canberra: Australian Institute of Aboriginal Studies.

Heath, Jeffrey (1984). *Functional Grammar of Nunggubuyu*. Canberra: Australian Institute of Aboriginal Studies.

Heath, Jeffrey (1986). 'Syntactic and Lexical Aspects of Nonconfigurationality in Nunggubuyu (Australia)'. *Natural Language & Linguistic Theory* 4.3, pp. 375–408. DOI: 10.1007/BF00133375.

Hellmuth, Sam, Frank Kügler & Ruth Singer (2007). 'Intonational Patterns, Tonal Alignment and Focus in Mawng'. *Proceedings of the 16th ICPhS Satellite Workshop: Intonation Phonology: Understudied or Fieldwork Languages*.

Hendricks, Sean Quillan (1999). *Reduplication without Template Constraints: A Study in Bare-Consonant Reduplication*. PhD thesis. University of Arizona.

Hockett, Charles F. (1954). 'Two Models of Grammatical Description'. *Word* 10.2-3, pp. 210–234. DOI: 10.1080/00437956.1954.11659524.

Hockett, Charles F. (1967). 'The Quantification of Functional Load'. *Word* 23.1-3, pp. 300–320. DOI: 10.1080/00437956.1967.11435484.

Höder, Steffen (2014). 'Phonological Elements and Diasystematic Construction Grammar'. *Constructions & Frames* 6.2, pp. 202–231. DOI: 10.1075/cf.6.2.04.

Hoffmann, Thomas & Graeme Trousdale (2013). 'Construction Grammar: Introduction'. *The Oxford Handbook of Construction Grammar*. Ed. by Thomas Hoffmann & Graeme Trousdale.

Oxford: Oxford University Press, pp. 1–13. DOI: 10.1093/oxfordhb/9780195396683.013.0001.

Hooper, Joan B (1976). *An Introduction to Natural Generative Phonology*. New York: Academic Press.

Hore, Michael (1978). 'New versus Old Information in Nunggubuyu'. *Oceanic Linguistics* 17.1, p. 11. DOI: 10.2307/3622825.

Hore, Michael (1981). 'Syllable Length and Stress in Nunggubuyu'. *Work Papers of SIL-AAB*. Vol. 5. Australian Phonologies: Collected Papers A. Darwin: Summer Institute of Linguistics, pp. 1–62.

Horrack, Kate (2010). *Expressing Possession in Wubuy*. Honours thesis. Armidale: University of New England.

Horrack, Kate (2018). *Argument Realisation in Wubuy*. PhD thesis. Melbourne: University of Melbourne.

Hughes, Earl (1969). 'Tales of the Nunggubuyu Tribe'. Unpublished mimeograph.

Hughes, Earl (1970). 'More Tales of the Nunggubuyu Tribe'. Unpublished mimeograph.

Hughes, Earl (1971). *Nunggubuyu-English Dictionary*. Sydney: University of Sydney.

Hume, Elizabeth, Kathleen Currie Hall, Andrew Wedel, Adam Ussishkin, Martine Adda-Dekker & Cédric Gendrot (2011). 'Anti-Markedness Patterns in French Epenthesis: An Information-Theoretic Approach'. *Annual Meeting of the Berkeley Linguistics Society*. Vol. 37, pp. 104–123.

Hurch, Bernhard & Veronika Mattes (2005). 'Graz Database on Reduplication'. *Faits de Langues (to appear)*.

Idsardi, William J. (1997). 'Sympathy Creates Chaos'. Unpublished manuscript. University of Delaware.

Idsardi, William J. (1998). 'Tiberian Hebrew Spirantization and Phonological Derivations'. *Linguistic Inquiry* 29.1, pp. 37–73. DOI: 10.1162/002438998553653.

Inkelas, Sharon (2014). *The Interplay of Morphology and Phonology*. Oxford: Oxford University Press.

Inkelas, Sharon & Laura J. Downing (2015). 'What Is Reduplication? Typology and Analysis Part 1/2: The Typology of Reduplication'. *Language and Linguistics Compass* 9.12, pp. 502–515. DOI: 10.1111/lnc3.12166.

Inkelas, Sharon & Cemil Orhan Orgun (1995). 'Level Ordering and Economy in the Lexical Phonology of Turkish'. *Language* 71.4, pp. 763–793. DOI: 10.2307/415744.

Inkelas, Sharon & Stephanie S. Shih (2014). 'Unstable Surface Correspondence as the Source of Local Conspiracies'. *NELS 44: Proceedings of the 44th Meeting of the North East Linguistic Society*. Ed. by Jyoti Iyer & Leland Kusmer. Vol. 1. University of Connecticut, pp. 191–204.

Inkelas, Sharon & Cheryl Zoll (2005). *Reduplication: Doubling in Morphology*. Cambridge Studies in Linguistics 106. Cambridge: Cambridge University Press.

Ito, Junko & Armin Mester (2003). 'On the Sources of Opacity in OT: Coda Processes in German'. *The Syllable in Optimality Theory*. Ed. by Caroline Féry & Ruben van de Vijver. Cambridge: Cambridge University Press, pp. 271–303. DOI: 10.1017/CBO9780511497926.012.

Itô, Junko & Armin Mester (2001). 'Covert Generalizations in Optimality Theory: The Role of Stratal Faithfulness Constraints'. *Studies in Phonetics, Phonology, and Morphology* 7.2, pp. 273–299.

Jackendoff, Ray (1975). 'Morphological and Semantic Regularities in the Lexicon'. *Language* 51.3, p. 639. DOI: 10.2307/412891.

Jackendoff, Ray (2002). *Foundations of Language: Brain, Meaning, Grammar, Evolution*. Oxford: Oxford University Press.

Jackendoff, Ray (2008). 'Construction After Construction and Its Theoretical Challenges'. *Language* 84.1, pp. 8–28. DOI: 10.1353/lan.2008.0058.

Jackendoff, Ray (2017). 'In Defense of Theory'. *Cognitive Science* 41, pp. 185–212.

Jackendoff, Ray & Jenny Audring (2020). *The Texture of the Lexicon: Relational Morphology and the Parallel Architecture*. Oxford, New York: Oxford University Press.

Jepson, Kathleen (2019). *Prosody, Prominence & Segments in Djambarrpuyŋu*. PhD thesis. Melbourne: University of Melbourne.

Jepson, Kathleen & Thomas Ennever (2023). 'Lexical Stress'. *The Oxford Guide to Australian Languages*. Ed. by Claire Bowern. Oxford: Oxford University Press, pp. 145–158. DOI: 10.1093/oso/9780198824978.003.0014.

Jepson, Kathleen, Janet Fletcher & Hywel Stoakes (2021). 'Prosodically Conditioned Consonant Duration in Djambarrpuyŋu'. *Language and Speech* 64.2, pp. 261–290. DOI: 10.1177/0023830919826607.

Jun, Sun-Ah, ed. (2014a). *Prosodic Typology II: The Phonology of Intonation and Phrasing*. Oxford: Oxford University Press.

Jun, Sun-Ah (2014b). 'Prosodic Typology: By Prominence Type, Word Prosody, and Macro-Rhythm'. *Prosodic Typology II: The Phonology of Intonation and Phrasing*. Ed. by Sun-Ah Jun. Oxford Linguistics. Oxford: Oxford University Press, pp. 520–539.

Jun, Sun-Ah & Janet Fletcher (2014). 'Methodology of Studying Intonation'. *Prosodic Typology II: The Phonology of Intonation and Phrasing*. Ed. by Sun-Ah Jun. Oxford Linguistics. Oxford: Oxford University Press.

Jurafsky, Daniel, Alan Bell & Cynthia Girand (2002). 'The Role of the Lemma in Form Variation'. *Laboratory Phonology 7*. Ed. by Carlos Gussenhoven & Natasha Warner. Phonology and Phonetics 4-1. Berlin: De Gruyter Mouton, pp. 3–34.

Kager, René (1999). *Optimality Theory*. Cambridge: Cambridge University Press.

Kager, René (2007). 'Feet and Metrical Stress'. *The Cambridge Handbook of Phonology*. Ed. by Paul De Lacy. Cambridge: Cambridge University Press, pp. 195–227.

Kapatsinski, Vsevolod (2018). *Changing Minds Changing Tools: From Learning Theory to Language Acquisition to Language Change*. Cambridge, Mass.: MIT Press.

Kapatsinski, Vsevolod (2021). 'What Are Constructions, and What Else Is Out There? An Associationist Perspective'. *Frontiers in Communication* 5, p. 575242. DOI: 10.3389/fcomm.2020.575242.

Keating, Patricia, Taehong Cho, Cécile Fougeron & Chai-Shune Hsu (2004). 'Domain-Initial Articulatory Strengthening in Four Languages'. *Phonetic interpretation: Papers in laboratory phonology VI*, pp. 143–161.

Kenstowicz, Michael J. (1994). *Phonology in Generative Grammar*. Cambridge, MA: Blackwell.

Kim, Juhyae (2023). 'Adjectives and Adverbs'. *The Oxford Guide to Australian Languages*. Ed. by Claire Bowern. Oxford: Oxford University Press, pp. 278–290. DOI: 10.1093/oso/9780198824978.003.0025.

King, Heather B. (1994). *The Declarative Intonation of Dyirbal an Acoustic Analysis*. MA thesis. Canberra: Australian National University.

Kiparsky, Paul (1973). 'Abstractness, Opacity and Global Rules'. *Three Dimensions of Linguistic Theory*. Ed. by Osamu Fujimura. Tokyo: Taikusha, pp. 1–136.

Kiparsky, Paul (1982a). *Explanation in Phonology*. Dordrecht: Foris Publications.

Kiparsky, Paul (1982b). 'How Abstract Is Phonology?' *Explanation in Phonology*. Dordrecht: Foris Publications, pp. 119–164.

Kiparsky, Paul (1982c). 'Lexical Morphology and Phonology'. *Linguistics in the Morning Calm*. Ed. by The Linguistic Society of Korea. Seoul: Hanshin, pp. 3–91.

Kiparsky, Paul (1982d). 'Word-Formation and the Lexicon'. Mid-America Linguistics Conference.

Kiparsky, Paul (1985). 'Some Consequences of Lexical Phonology'. *Phonology Yearbook* 2, pp. 85–138.

Kiparsky, Paul (2000). 'Opacity and Cyclicity'. *Linguistic Review* 17.2-4, pp. 351–366.

Kiparsky, Paul (2009). 'Reduplication in Stratal OT'. *Reality Exploration and Discovery: Pattern Interaction in Language & Life*. Ed. by Linda Uyechi & Lian-Hee Wee. Stanford: Center for the Study of Language and Information.

Kirchner, Robert (1996). 'Synchronic Chain Shifts in Optimality Theory'. *Linguistic Inquiry* 27.2, pp. 341–350.

Kirchner, Robert (1997). 'Contrastiveness and Faithfulness'. *Phonology* 14.1, pp. 83–111. DOI: 10.1017/S0952675797003291.

Kisler, Thomas, Uwe Reichel & Florian Schiel (2017). 'Multilingual Processing of Speech via Web Services'. *Computer Speech & Language* 45, pp. 326–347. DOI: http://dx.doi.org/10.1016/j.csl.2017.01.005.

Kiyomi, Setsuko (1993). *A Typological Study of Reduplication as a Morpho-Semantic Process: Evidence from Five Language Families (Bantu, Australian, Papuan, Austroasiatic and Malayo-Polynesian)*. PhD thesis. Bloomington: Indiana University.

Kroos, Christian, Rikke L. Bundgaard-Nielsen, Louis Goldstein & Catherine Best (2010). 'Tongue Body Position Differences in the Coronal Stop Consonants of Wubuy'. *Proceedings of the 13th Australian Speech Science and Technology Conference*, pp. 86–89.

Krott, Andrea, Robert Schreuder, R. Harald Baayen & Wolfgang U. Dressler (2007). 'Analogical Effects on Linking Elements in German Compound Words'. *Language and Cognitive Processes* 22.1, pp. 25–57. DOI: 10.1080/01690960500343429.

Krull, Diana (1989). 'Second Formant Locus Patterns and Consonant-Vowel Coarticulation in Spontaneous Speech'. *Phonetic Experimental Research at the Institute of Linguistics, University of Stockholm* 10, pp. 87–108.

Labov, William (1994). *Principles of Linguistic Change*. Language in Society 20. Oxford: Blackwell.

Ladd, D. Robert (2008). *Intonational Phonology*. 2nd ed. Cambridge: Cambridge University Press. DOI: 10.1017/CBO9780511808814.

Ladefoged, Peter & Ian Maddieson (1996). *The Sounds of the World's Languages*. Oxford: Blackwell.

Lakoff, George (1993). 'Cognitive Phonology'. *The Last Phonological Rule*. Ed. by John A. Goldsmith. Chicago: University of Chicago Press, pp. 117–145.

Langacker, Ronald W. (1987). *Foundations of Cognitive Grammar: Theoretical Prerequisites*. Stanford University Press.

Lapointe, Steven G. (2001). 'Stem Selection and OT'. *Yearbook of Morphology 1999*. Ed. by Geert Booij & Jaap van Marle. Yearbook of Morphology. Dordrecht: Springer Netherlands, pp. 263–297. DOI: 10.1007/978-94-017-3722-7_10.

Lecompte, Maarten, ed. (2014). *Kriol-English Interactive Dictionary*. Red. by Charles E. Grimes. AuSIL Interactive Dictionary Series A9. Darwin: Australian Society for Indigenous Languages.

Lee-Kim, Sang-Im, Lisa Davidson & Sangjin Hwang (2013). 'Morphological Effects on the Darkness of English Intervocalic /l/'. *Laboratory Phonology* 4.2, pp. 475–511. DOI: 10.1515/lp-2013-0015.

Levelt, Willem J. M. (1993). *Speaking: From Intention to Articulation*. The MIT Press. DOI: 10.7551/mitpress/6393.001.0001.

Lévêque, Dimitri & Thomas Pellard (2023). 'The Implicative Structure of Asama Verb Paradigms: A Quantitative Study of Segmental and Suprasegmental Alternations'. *Morphology* 33, pp. 261–286. DOI: 10.1007/s11525-023-09410-x.

Liberman, Mark & Alan Prince (1977). 'On Stress and Linguistic Rhythm'. *Linguistic Inquiry* 8.2, pp. 249–336.

Lieberman, Philip (1960). 'Some Acoustic Correlates of Word Stress in American English'. *The Journal of the Acoustical Society of America* 32.4, pp. 451–454. DOI: 10.1121/1.1908095.

Lindblom, Björn & Harvey M. Sussman (2012). 'Dissecting Coarticulation: How Locus Equations Happen'. *Journal of Phonetics* 40.1, pp. 1–19. DOI: 10.1016/j.wocn.2011.09.005.

Łubowicz, Anna (2003). 'Counter-Feeding Opacity as a Chain Shift Effect'. *WCCFL 22 Proceedings*. Ed. by Gina Garding & Mimu Tsujimura. Somerville, MA: Cascadilla Press, pp. 315–327. DOI: 10.7282/T3CV4MP2.

Luce, Paul A. & David B. Pisoni (1998). 'Recognizing Spoken Words: The Neighborhood Activation Model'. *Ear and Hearing* 19.1, pp. 1–36.

MacBride, Alexander I (2004). *A Constraint-Based Approach to Morphology*. PhD thesis. Los Angeles: University of California, Los Angeles.

Mailhammer, Robert & Patrick Caudal (2019). 'Linear Lengthening Intonation in English on Croker Island: Identifying Substrate Origins'. *JournaLIPP* 6, pp. 40–56.

Malecot, Andre (1956). 'Acoustic Cues for Nasal Consonants: An Experimental Study Involving a Tape-Splicing Technique'. *Language* 32.2, p. 274. DOI: 10.2307/411004.

Mansfield, John (2019). *Murrinhpatha Morphology and Phonology*. De Gruyter. DOI: 10.1515/9781501503306.

Mansfield, John (2021). 'The Word as a Unit of Internal Predictability'. *Linguistics* 59.6, pp. 1427–1472. DOI: 10.1515/ling-2020-0118.

Mansfield, John (2023). 'The Prosodic Structure of Australian Polysynthetic Verbs: Bininj Gun-wok, Murrinhpatha, and Ngalakgan'. *Word Prominence in Languages with Complex Morphologies*. Ed. by Ksenia Bogomolets & Harry Van der Hulst. Oxford: Oxford University Press.

Mansfield, John, Rosey Billington & Hywel Stoakes (2022). 'Vowel Predictability and Omission in Anindilyakwa'. DOI: 10.31234/osf.io/fpcrj.

Mansfield, John & Ian Green (2021). 'Fricative Contrasts and Neutralization in Marri Tjevin'. *Australian Journal of Linguistics* 41.2, pp. 220–261. DOI: 10.1080/07268602.2021.1957774.

Marslen-Wilson, William D., Lorraine K. Tyler, Rachelle Waksler & Lianne Older (1994). 'Morphology and Meaning in the English Mental Lexicon'. *Psychological review* 101.1, pp. -33.

Martinet, André (1952). 'Function, Structure, and Sound Change'. *Word* 8.1, pp. 1–32. DOI: 10.1080/00437956.1952.11659416.

Mascaró, Joan (1976). *Catalan Phonology and the Phonological Cycle*. PhD thesis. Massachusetts Institute of Technology.

Mascaró, Joan (2007). 'External Allomorphy and Lexical Representation'. *Linguistic Inquiry* 38.4, pp. 715–735. DOI: 10.1162/ling.2007.38.4.715.

McCarthy, John J. (1993). 'A Case of Surface Constraint Violation'. *Canadian Journal of Linguistics/Revue canadienne de linguistique* 38.2, pp. 169–195. DOI: 10.1017/S0008413100014730.

McCarthy, John J. (1996). 'Remarks on Phonological Opacity in Optimality Theory'. *Linguistics Department Faculty Publication Series* 66.

McCarthy, John J. (1999). 'Sympathy and Phonological Opacity'. *Phonology* 16.03, pp. 331–399.

McCarthy, John J. (2003a). 'Comparative Markedness'. *Theoretical Linguistics* 29.1-2, pp. 1–51. DOI: 10.1515/thli.29.1-2.1.

McCarthy, John J. (2003b). 'Sympathy, Cumulativity, and the Duke-of-York Gambit'. *The Syllable in Optimality Theory*. Ed. by Caroline Féry & Ruben van de Vijver. Cambridge: Cambridge University Press, pp. 23–76. DOI: 10.1017/CBO9780511497926.012.

McCarthy, John J. (2004). 'Optimal Paradigms'. *Paradigms in Phonological Theory*. Ed. by Laura J. Downing, T. Alan Hall & Renate Raffelsiefen. Oxford: Oxford University Press, pp. 170–210. DOI: 10.1093/acprof:oso/9780199267712.003.0008.

McCarthy, John J. (2007). *Hidden Generalizations: Phonological Opacity in Optimality Theory*. Advances in Optimality Theory. London: Equinox.

McCarthy, John J., Wendell Kimper & Kevin Mullin (2012). 'Reduplication in Harmonic Serialism'. *Morphology* 22.2, pp. 173–232. DOI: 10.1007/s11525-012-9203-3.

McCarthy, John J. & Alan Prince (1986). 'Prosodic Morphology'. MS. University of Massachusetts and Brandeis University.

McCarthy, John J. & Alan Prince (1990). 'Prosodic Morphology and Templatic Morphology'. *Perspectives on Arabic Linguistics II: Papers from the Second Annual Symposium on Arabic Linguistics*. Ed. by Mushira Eid & John J. McCarthy. Amsterdam: John Benjamins, pp. 1–54.

McCarthy, John J. & Alan Prince (1993a). 'Generalized Alignment'. *Yearbook of Morphology 1993*. Ed. by Geert Booij & Jaap Van Marle. Dordrecht: Springer Netherlands, pp. 79–153. DOI: 10.1007/978-94-017-3712-8_4.

McCarthy, John J. & Alan Prince (1993b). 'Prosodic Morphology: Constraint Interaction and Satisfaction'. Unpublished manuscript. University of Massachusetts, Amherst and Rutgers University.

McCarthy, John J. & Alan Prince (1994). 'The Emergence of the Unmarked: Optimality in Prosodic Morphology'. Unpublished manuscript. University of Massachusetts, Amherst and Rutgers University.

McCarthy, John J. & Alan Prince (1995). 'Faithfulness and Reduplicative Identity'. *University of Massachussetts occasional papers*. Linguistics Department Faculty Publication Series 18.

McClelland, James L & Jeffrey L Elman (1986). 'The TRACE Model of Speech Perception'. *Cognitive Psychology* 18.1, pp. 1–86. DOI: 10.1016/0010-0285(86)90015-0.

McKay, Graham (2000). 'Ndjébbana'. *Handbook of Australian Languages*. Ed. by Robert M. W. Dixon & Barry Blake. Vol. 5. Oxford: Oxford University Press, pp. 155–356.

McManus, Hope Eliza (2008). *Loanword Adaptation: A Study of Some Australian Aboriginal Languages*. Honours thesis. Sydney: University of Sydney.

Meakins, Felicity (2014). 'Language Contact Varieties'. *The Languages and Linguistics of Australia: A Comprehensive Guide*. Ed. by Harold Koch & Rachel Nordlinger. Vol. 3. Berlin: De Gruyter Mouton, pp. 365–416.

Mielke, Jeff (2008). *The Emergence of Distinctive Features*. Oxford: Oxford University Press.

Mielke, Jeff, Mike Armstrong & Elizabeth Hume (2003). 'Looking through Opacity'. *Theoretical Linguistics* 29.1-2, pp. 123–139.

Milin, Petar, Victor Kuperman, Aleksandar Kostić & R. Harald Baayen (2009). 'Words and Paradigms Bit by Bit: An Information-Theoretic Approach to the Processing of Inflection and Derivation'. *Analogy in Grammar: Form and Acquisition*. Ed. by James P. Blevins & Juliette Blevins. Oxford: Oxford University Press, pp. 214–252.

Modarresi, Golnaz, Harvey M. Sussman, Bjorn Lindblom & Elizabeth Burlingame (2004). 'Stop Place Coding: An Acoustic Study of CV, VC#, and C#V Sequences'. *Phonetica* 61.1, pp. 2–21. DOI: 10.1159/000078660.

Mohanan, Karuvannur Puthanveettil (1982). *Lexical Phonology*. PhD thesis. Massachusetts Institute of Technology.

Moreton, Elliott (2004). 'Non-Computable Functions in Optimality Theory'. *Optimality Theory in Phonology: A Reader*. Ed. by John J. McCarthy. Malden: Blackwell, pp. 141–163.

Moreton, Elliott & Paul Smolensky (2002). 'Typological Consequences of Local Constraint Conjunction'. *Proceedings of the 21st West Coast Conference on Formal Linguistics*. WCCFL 21. Ed. by Line Mikkelsen & Christopher Potts. Cambridge, MA: Cascadilla Press, pp. 306–319.

Mortensen, David (2006). *Logical and Substantive Scales in Phonology*. PhD thesis. Berkeley: University of California, Berkeley.

Murray, Robert W. & Theo Vennemann (1983). 'Sound Change and Syllable Structure in Germanic Phonology'. *Language* 59.3, pp. 514–528. DOI: 10.2307/413901.

Nash, David (1986). *Topics in Warlpiri Grammar*. Outstanding Dissertations in Linguistics. New York: Garland.

Nordlinger, Rachel (2023). 'Australian Languages and Syntactic Theory'. *The Oxford Guide to Australian Languages*. Ed. by Claire Bowern. Oxford: Oxford University Press, pp. 31–36. DOI: 10.1093/oso/9780198824978.003.0004.

Norris, Dennis (1994). 'Shortlist: A Connectionist Model of Continuous Speech Recognition'. *Cognition* 52.3, pp. 189–234. DOI: 10.1016/0010-0277(94)90043-4.

Norris, Dennis & James M. McQueen (2008). 'Shortlist B: A Bayesian Model of Continuous Speech Recognition.' *Psychological Review* 115.2, pp. 357–395. DOI: 10.1037/0033-295X.115.2.357.

Nübling, Damaris & Renata Szczepaniak (2008). 'On the Way from Morphology to Phonology: German Linking Elements and the Role of the Phonological Word'. *Morphology* 18.1, pp. 1–25. DOI: 10.1007/s11525-008-9120-7.

O'Donnell, Timothy J. (2015). *Productivity and Reuse in Language: A Theory of Linguistic Computation and Storage*. Cambridge, MA: MIT Press.

Ogden, Richard (1999). 'A Declarative Account of Strong and Weak Auxiliaries in English'. *Phonology* 16.1, pp. 55–92. DOI: 10.1017/S095267579900370X.

Ohala, John J (1981). 'The Listener as a Source of Sound Change'. *Papers from the Parasession on Language and Behavior*. Ed. by Carrie S Masek, Roberta A Hendrick & Mary Frances Miller. Chicago: Chicago Linguistic Society, pp. 178–203.

Ohala, John J (1990). 'The Phonetics and Phonology of Aspects of Assimilation'. *Papers in Laboratory Phonology*. Ed. by John Kingston & Mary E. Beckman. Vol. 1. Cambridge: Cambridge University Press, pp. 258–275.

Olejarczuk, Paul, Vsevolod Kapatsinski & R. Harald Baayen (2018). 'Distributional Learning Is Error-Driven: The Role of Surprise in the Acquisition of Phonetic Categories'. *Linguistics Vanguard* 4.s2. DOI: 10.1515/lingvan-2017-0020.

Orgun, Cemil Orhan (1996). *Sign-Based Morphology and Phonology with Special Attention to Optimality Theory*. PhD thesis. University of California, Berkeley.

Orgun, Cemil Orhan (1998). 'Cyclic and Noncyclic Phonological Effects in a Declarative Grammar'. *Yearbook of Morphology 1997*. Ed. by Geert Booij & Jaap Van Marle. Dordrecht: Springer Netherlands, pp. 179–218. DOI: 10.1007/978-94-011-4998-3_7.

Osgarby, David (2018). 'Reconstructing Proto-Mirndi Verbal Morphology: From Particles and Clitics to Prefixes'. *Australian Journal of Linguistics* 38.2, pp. 223–292. DOI: 10.1080/07268602.2018.1400504.

Panther, Forrest (2021). *Topics in Kaytetye Phonology and Morpho-Syntax*. PhD thesis. Newcastle, Australia: University of Newcastle.

Parker, Jeff & Andrea D. Sims (2020). 'Irregularity, Paradigmatic Layers, and the Complexity of Inflection Class Systems: A Study of Russian Nouns'. *The Complexities of Morphology*. Ed. by Peter Arkadiev & Francesco Gardani. Oxford: Oxford University Press. DOI: 10.1093/oso/9780198861287.003.0002.

Parncutt, Amy (2015). *Towards a Phonological Typology of Reduplication in Australian Languages*. Honours thesis. Brisbane: University of Queensland.

Paster, Mary (2006). *Phonological Conditions on Affixation*. PhD thesis. Berkeley: University of California, Berkeley.

Paster, Mary (2009). 'Explaining Phonological Conditions on Affixation: Evidence from Suppletive Allomorphy and Affix Ordering'. *Word Structure* 2.1, pp. 18–37. DOI: 10.3366/E1750124509000282.

Peperkamp, Sharon, Inga Vendelin & Emmanuel Dupoux (2010). 'Perception of Predictable Stress: A Cross-Linguistic Investigation'. *Journal of Phonetics* 38.3, pp. 422–430. DOI: 10.1016/j.wocn.2010.04.001.

Picanço, Gessiane (2002). 'Tonal Polarity as Phonologically Conditioned Allomorphy in Mundurukú'. *Annual Meeting of the Berkeley Linguistics Society*. Vol. 28. 1, pp. 237–248.

Pierrehumbert, Janet B. (1980). *The Phonology and Phonetics of English Intonation*. PhD thesis. Massachusetts Institute of Technology.

Pierrehumbert, Janet B. (1999). 'Formalizing Functionalism'. *Functionalism and Formalism in Linguistics*. Ed. by Michael Darnell, Edith Moravcsik, Frederick J. Newmeyer, Michael Noonan & Kathleen Wheatley. Vol. 1: General Papers. Amsterdam: John Benjamins Publishing, pp. 287–304.

Pierrehumbert, Janet B. (2001). 'Exemplar Dynamics: Word Frequency, Lenition and Contrast'. *Frequency and the Emergence of Linguistic Structure*. Ed. by Joan L. Bybee & Paul J. Hopper. Typological Studies in Language 45. Amsterdam: John Benjamins, pp. 137–158.

Pierrehumbert, Janet B. (2003). 'Phonetic Diversity, Statistical Learning, and Acquisition of Phonology'. *Language and Speech* 46.2-3, pp. 115–154. DOI: 10.1177/00238309030460020501.

Pierrehumbert, Janet B. (2008). 'Word-Specific Phonetics'. *Laboratory Phonology 7*. Ed. by Carlos Gussenhoven & Natasha Warner. Berlin: De Gruyter Mouton, pp. 101–139. DOI: 10.1515/9783110197105.

Pierrehumbert, Janet B. (2016). 'Phonological Representation: Beyond Abstract Versus Episodic'. *Annual Review of Linguistics* 2.1, pp. 33–52. DOI: 10.1146/annurev-linguistics-030514-125050.

Plag, Ingo, Julia Homann & Gero Kunter (2017). 'Homophony and Morphology: The Acoustics of Word-Final S in English'. *Journal of Linguistics* 53.1, pp. 181–216. DOI: 10.1017/S0022226715000183.

Plag, Ingo, Gero Kunter & Mareile Schramm (2011). 'Acoustic Correlates of Primary and Secondary Stress in North American English'. *Journal of Phonetics* 39.3, pp. 362–374. DOI: 10.1016/j.wocn.2011.03.004.

Pluymaekers, Mark, Mirjam Ernestus & R. Harald Baayen (2005a). 'Articulatory Planning Is Continuous and Sensitive to Informational Redundancy'. *Phonetica* 62.2-4, pp. 146–159. DOI: 10.1159/000090095.

Pluymaekers, Mark, Mirjam Ernestus & R. Harald Baayen (2005b). 'Lexical Frequency and Acoustic Reduction in Spoken Dutch'. *The Journal of the Acoustical Society of America* 118.4, pp. 2561–2569. DOI: 10.1121/1.2011150.

Prince, Alan & Paul Smolensky (2004). *Optimality Theory: Constraint Interaction in Generative Grammar*. Malden, MA: Blackwell.

Proctor, Michael, Rikke L. Bundgaard-Nielsen, Catherine Best, Louis Goldstein, Christian Kroos, Mark Harvey et al. (2010). 'Articulatory Modelling of Coronal Stop Contrasts in Wubuy'. *Proceedings of the 13th Australian International Conference of Speech Science and Technology [SST]*, pp. 90–93.

Pruitt, Kathryn (2023). 'Serialism and Opacity in Phonological Theory'. *Annual Review of Linguistics* 9.1, pp. 497–517. DOI: 10.1146/annurev-linguistics-031220-120748.

Pullum, Geoffrey K. (1976). 'The Duke of York Gambit'. *Journal of Linguistics* 12.1, pp. 83–102.

R Core Team (2024). *R: A Language and Environment for Statistical Computing*. Version 4.3.3. Vienna, Austria: R Foundation for Statistical Computing.

Rácz, Péter, Clay Beckner, Jennifer Hay & Janet B. Pierrehumbert (2020). 'Morphological Convergence as On-Line Lexical Analogy'. *Language* 96.4, pp. 735–770.

Recasens, Daniel & Aina Espinosa (2009). 'An Articulatory Investigation of Lingual Coarticulatory Resistance and Aggressiveness for Consonants and Vowels in Catalan'. *The Journal of the Acoustical Society of America* 125.4, pp. 2288–2298. DOI: 10.1121/1.3089222.

Rose, Sharon & Rachel Walker (2004). 'A Typology of Consonant Agreement as Correspondence'. *Language* 80.3, pp. 475–531. DOI: 10.1353/lan.2004.0144.

Ross, Belinda (2011). *Prosody and Grammar in Dalabon and Kayardild*. PhD thesis. Melbourne: University of Melbourne.

Round, Erich R. (2009). *Kayardild Morphology, Phonology and Morphosyntax*. PhD thesis. New Haven, Connecticut: Yale University.

Round, Erich R. (2011). 'Word Final Phonology in Lardil: Implications of an Expanded Data Set'. *Australian Journal of Linguistics* 31.3, pp. 327–350. DOI: 10.1080/07268602.2011.598630.

Round, Erich R. (2012). *Kayardild Morphology and Syntax*. Oxford: Oxford University Press.

Round, Erich R. (2013). 'Why Reduplicate VC? Kuuk Thaayorre Answers a Lingering Question'. 44th Annual Conference of the Australian Linguistics Society. Melbourne.

Round, Erich R. (2014). 'Prestopping of Nasals and Laterals Is Only Partly Parallel'. *Language Description Informed by Theory*. Ed. by Rob Pensalfini, Myfany Turpin & Diana Guillemin. Amsterdam: John Benjamins, pp. 81–95.

Round, Erich R. (2023). 'Morphophonology: Lenition and Assimilation'. *The Oxford Guide to Australian Languages*. Ed. by Claire Bowern. Oxford: Oxford University Press, p. 0. DOI: 10.1093/oso/9780198824978.003.0012.

Round, Erich R., Rikker Dockum & Robin J. Ryder (2022). 'Evolution and Trade-Off Dynamics of Functional Load'. *Entropy* 24.4 (4), p. 507. DOI: 10.3390/e24040507.

Russell, Kevin (1995). 'Morphemes and Candidates in Optimality Theory'. Unpublished manuscript. University of Manitoba.

Ryan, Kevin & Jeffrey Heath (2021). 'Onset Prominence in Wubuy Reduplication'. *Proceedings of the Annual Meetings on Phonology*. Ed. by Peter Jurgec et al. Washington, DC: Linguistic Society of America. DOI: 10.3765/amp.v9i0.5148.

Ryan, Kevin M. (2014). 'Onsets Contribute to Syllable Weight: Statistical Evidence from Stress and Meter'. *Language* 90.2, pp. 309–341.

Saba Kirchner, Jesse (2010). *Minimal Reduplication*. PhD thesis. Santa Cruz: University of California Santa Cruz.

Saba Kirchner, Jesse (2013). 'Minimal Reduplication and Reduplicative Exponence'. *Morphology* 23.2, pp. 227–243. DOI: 10.1007/s11525-013-9225-5.

Saperstein, Andrew D. (1997). *A Word-and-Paradigm Approach to Reduplication*. PhD thesis. The Ohio State University.

Sayers, Barbara (1976). *Interpenetration of Stress and Pitch in Wik-Munkan Grammar and Phonology*. Canberra: Pacific Linguistics.

Schiering, René, Balthasar Bickel & Kristine A. Hildebrandt (2010). 'The Prosodic Word Is Not Universal, but Emergent'. *Journal of Linguistics* 46.3, pp. 657–709. DOI: 10.1017/S0022226710000216.

Schreuder, Robert & R.Harald Baayen (1997). 'How Complex Simplex Words Can Be'. *Journal of Memory and Language* 37.1, pp. 118–139. DOI: 10.1006/jmla.1997.2510.

Schultze-Berndt, Eva (2000). *Simple and Complex Verbs in Jaminjung: A Study of Event Categorisation in an Australian Language*. MPI Series in Psycholinguistics 14. Wageningen: Ponsen & Looijen.

Schweitzer, Antje (2019). 'Exemplar-Theoretic Integration of Phonetics and Phonology: Detecting Prominence Categories in Phonetic Space'. *Journal of Phonetics* 77, p. 100915. DOI: 10.1016/j.wocn.2019.100915.

Scobbie, James M. (1991). *Attribute Value Phonology*. PhD thesis. Edinburgh: University of Edinburgh.

Selkirk, Elisabeth O. (1984). *Phonology and Syntax: The Relation between Sound and Structure*. Current Studies in Linguistics Series 10. Cambridge, Mass: MIT Press.

Sereno, Joan A. & Allard Jongman (1997). 'Processing of English Inflectional Morphology'. *Memory & Cognition* 25.4, pp. 425–437. DOI: 10.3758/BF03201119.

Seyfarth, Scott (2014). 'Word Informativity Influences Acoustic Duration: Effects of Contextual Predictability on Lexical Representation'. *Cognition* 133.1, pp. 140–155. DOI: 10.1016/j.cognition.2014.06.013.

Sharpe, Margaret (1972). *Alawa Phonology and Grammar*. Australian Aboriginal Studies 37. Canberra: Australian Institute of Aboriginal Studies.

Shaw, Jason & Shigeto Kawahara (2018). 'Predictability and Phonology: Past, Present and Future'. *Linguistics Vanguard* 4.s2. DOI: 10.1515/lingvan-2018-0042.

Shaw, Jason A., Christopher Carignan, Tonya G. Agostini, Robert Mailhammer, Mark Harvey & Donald Derrick (2020). 'Phonological Contrast and Phonetic Variation: The Case of Velars in Iwaidja'. *Language* 96.3, pp. 578–617. DOI: 10.1353/lan.2020.0042.

Simard, Candide (2010). *The Prosodic Contours of Jaminjung, a Language of Northern Australia*. PhD thesis. Manchester: University of Manchester.

Simpson, J. & M. Withgott (1986). 'Pronominal Clitic Clusters and Templates'. *The Syntax of Pronominal Clitics*. Ed. by Hagit Borer. Syntax and Semantics 19. Orlando: Academic Press, pp. 147–174. DOI: 10.1163/9789004373150_008.

Skousen, Royal (2009). 'Expanding Analogical Modeling into a General Theory of Language Prediction'. *Analogy in Grammar: Form and Acquisition*. Ed. by James P. Blevins & Juliette Blevins. Oxford Linguistics. Oxford: Oxford University Press, pp. 164–184.

Smith, Caitlin (2020). 'Partial Height Harmony as Partial Transparency'. *Proceedings of the 2019 Annual Meeting on Phonology*. Ed. by Hyunah Baek, Chikako Takahashi & Alex Hong-Lun Yeung. 0. Washington, DC: Linguistic Society of America. DOI: 10.3765/amp.v8i0.4732.

Smolensky, Paul (1995). 'On the Structure of the Constraint Component CON of UG'. Handout of talk at UCLA.

Sommer, Bruce A. (1969). *Kunjen Phonology: Synchronic and Diachronic*. Canberra: Pacific Linguistics.

Sommer, Bruce A. (1970). 'An Australian Language without CV Syllables'. *International Journal of American Linguistics* 36.1, pp. 57–58.

Sóskuthy, Márton & Jennifer Hay (2017). 'Changing Word Usage Predicts Changing Word Durations in New Zealand English'. *Cognition* 166, pp. 298–313. DOI: 10.1016/j.cognition.2017.05.032.

Spaelti, Philip (1996). 'A Constraint-Based Theory of Reduplication Patterns'. *Proceedings of the Fourteenth West Coast Conference on Formal Linguistics*. Ed. by José Camacho, Lina Choueiri & Maki Watanabe. Stanford: CSLI Publications, pp. 477–492.

Spaelti, Philip (1997). *Dimensions of Variation in Multi-Pattern Reduplication*. PhD thesis. Santa Cruz: University of California Santa Cruz.

Steriade, Donca (2001). 'Directional Asymmetries in Place Assimilation: A Perceptual Account'. *The role of speech perception in phonology*, pp. 219–250.

Steriade, Donca (2023). 'Vowel-to-Vowel Intervals in the Quantitative Meters of Ancient Greek and Latin'. Chicago Linguistics Society. Chicago.

Stevens, Kenneth N. & Sheila E. Blumstein (1975). 'Quantal Aspects of Consonant Production and Perception: A Study of Retroflex Stop Consonants'. *Journal of Phonetics* 3.4, pp. 215–233.

Stoakes, Hywel M., Janet M. Fletcher & Andrew R. Butcher (2020). 'Nasal Coarticulation in Bininj Kunwok: An Aerodynamic Analysis'. *Journal of the International Phonetic Association* 50.3, pp. 305–332. DOI: 10.1017/S0025100318000282.

Stump, Gregory T. (2001). *Inflectional Morphology: A Theory of Paradigm Structure*. Cambridge: Cambridge University Press.

Surendran, Dinoj & Partha Niyogi (2006). 'Quantifying the Functional Load of Phonemic Oppositions, Distinctive Features, and Suprasegmentals'. *Current Issues in Linguistic Theory* 279. Ed. by Ole Nedergaard Thomsen, pp. 43–58. DOI: 10.1075/cilt.279.05sur.

Sussman, Harvey M., Helen A McCaffrey & Sandra A Matthews (1991). 'An Investigation of Locus Equations as a Source of Relational Invariance for Stop Place Categorization'. *Journal of the Acoustical Society of America* 90.3, pp. 1309–1325.

Svenonius, Peter (2012). 'Look Both Ways: Outward-looking Allomorphy in Icelandic Participles'. Unpublished manuscript. CASTL, University of Tromsø.

Tabain, Marija, Gavan Breen & Andrew Butcher (2004). 'VC vs. CV Syllables: A Comparison of Aboriginal Languages with English'. *Journal of the International Phonetic Association* 34.2, pp. 175–200.

Tabain, Marija & Andrew Butcher (1999). 'Stop Consonants in Yanyuwa and Yindjibarndi: Locus Equation Data'. *Journal of Phonetics* 27.4, pp. 333–357. DOI: 10.1006/jpho.1999.0099.

Tabain, Marija, Janet Fletcher & Andrew Butcher (2014). 'Lexical Stress in Pitjantjatjara'. *Journal of Phonetics* 42, pp. 52–66. DOI: 10.1016/j.wocn.2013.11.005.

Tiede, Mark, Christine Mooshammer, Louis Goldstein, Stefanie Shattuck-Hufnagel & Joseph Perkell (2011). 'Motor Learning of Articulator Trajectories in the Production of Novel Utterances'. *Proceedings of ICPhS XVII*. International Congress of the Phonetic Sciences. Hong Kong, pp. 1986–1989.

Todd, Simon, Janet B. Pierrehumbert & Jennifer Hay (2019). 'Word Frequency Effects in Sound Change as a Consequence of Perceptual Asymmetries: An Exemplar-Based Model'. *Cognition* 185, pp. 1–20. DOI: 10.1016/j.cognition.2019.01.004.

Tomaschek, Fabian, Denis Arnold, Franziska Bröker & R. Harald Baayen (2018). 'Lexical Frequency Co-Determines the Speed-Curvature Relation in Articulation'. *Journal of Phonetics* 68, pp. 103–116. DOI: 10.1016/j.wocn.2018.02.003.

Tomaschek, Fabian, Ingo Plag, Mirjam Ernestus & R. Harald Baayen (2019). 'Phonetic Effects of Morphology and Context: Modeling the Duration of Word-Final S in English with Naïve Discriminative Learning'. *Journal of Linguistics*, pp. 1–39. DOI: 10.1017/S0022226719000203.

Tomaschek, Fabian, Benjamin V. Tucker, Matteo Fasiolo & R. Harald Baayen (2018). 'Practice Makes Perfect: The Consequences of Lexical Proficiency for Articulation'. *Linguistics Vanguard* 4.s2. DOI: 10.1515/lingvan-2017-0018.

Uffmann, Christian (2007). 'Intrusive [r] and Optimal Epenthetic Consonants'. *Language Sciences*. Issues in English Phonology 29.2, pp. 451–476. DOI: 10.1016/j.langsci.2006.12.017.

Urbanczyk, Suzanne (2001). *Patterns of Reduplication in Lushootseed*. New York: Garland.

Urbanczyk, Suzanne (2006). 'Reduplicative Form and the Root-Affix Asymmetry'. *Natural Language & Linguistic Theory* 24.1, pp. 179–240. DOI: 10.1007/s11049-005-4373-x.

Ussishkin, Adam (2005). 'A Fixed Prosodic Theory of Nonconcatenative Templatic Morphology'. *Natural Language & Linguistic Theory* 23.1, pp. 169–218.

Ussishkin, Adam, Natasha Warner, Ian Clayton, Daniel Brenner, Andrew Carnie, Michael Hammond & Muriel Fisher (2017). 'Lexical representation and processing of word-initial morphological alternations: Scottish Gaelic mutation'. *Laboratory Phonology: Journal of the Association for Laboratory Phonology* 8.1, p. 8. DOI: 10.5334/labphon.22.

Välimaa-Blum, Riitta (2005). *Cognitive Phonology in Construction Grammar: Analytic Tools for Students of English*. Planet Communication. Berlin New York: Mouton de Gruyter.

Van der Hulst, Harry & Nancy A. Ritter (2000). 'The SPE-heritage of Optimality Theory'. *The Linguistic Review* 17.2-4, pp. 259–290.

Van der Spuy, Andrew (2017). 'A Construction Grammar Account of Zulu Singular/Plural Inflection'. *Nordic Journal of African Studies* 26.3 (3), pp. 24–24. DOI: 10.53228/njas.v26i3.85.

Van Egmond, Marie-Elaine (2012). *Enindhilyakwa Phonology, Morphosyntax and Genetic Position*. PhD thesis. Sydney: University of Sydney.

Van Egmond, Marie-Elaine & Brett J. Baker (2021). 'The Genetic Position of Anindilyakwa'. *Australian Journal of Linguistics*, pp. 1–36. DOI: 10.1080/07268602.2020.1848796.

Van Son, Rob, Olga Bolotova, Louis CW Pols & Mietta Lennes (2004). 'Frequency Effects on Vowel Reduction in Three Typologically Different Languages (Dutch, Finish, Russian).' *Interspeech*.

Van Son, Rob & Louis Pols (2003). 'Information Structure and Efficiency in Speech Production'. *Proceedings of Eurospeech 2003*. Geneva, Switzerland, pp. 769–772.

Vaux, Bert (1998). *The Phonology of Armenian*. The Phonology of the World's Languages. Oxford: Clarendon Press.

Vaux, Bert (2002). 'Consonant Epenthesis and the Problem of Unnatural Phonology'. *Yale University Linguistics Colloquium*.

Vaux, Bert (2003). 'Syllabification in Armenian, Universal Grammar, and the Lexicon'. *Linguistic Inquiry* 34.1, pp. 91–125. DOI: 10.1162/002438903763255931.

Vaux, Bert & Bridget D. Samuels (2017). 'Consonant Epenthesis and Markedness'. *Beyond Markedness in Formal Phonology*. Ed. by Bridget D. Samuels. Vol. 241. Linguistik Aktuell/Linguistics Today. Amsterdam: John Benjamins Publishing Company, pp. 69–100. DOI: 10.1075/la.241.04vau.

Verstraete, Jean-Christophe (2005). 'The Semantics and Pragmatics of Composite Mood Marking: The Non-Pama-Nyungan Languages of Northern Australia'. *Linguistic Typology* 9.2, pp. 223–268. DOI: 10.1515/lity.2005.9.2.223.

Vihman, Marilyn & William Croft (2007). 'Phonological Development: Toward a "Radical" Templatic Phonology'. 45.4, pp. 683–725. DOI: 10.1515/LING.2007.021.

Wagner, Michael & Duane G. Watson (2010). 'Experimental and Theoretical Advances in Prosody: A Review'. *Language and Cognitive Processes* 25.7-9, pp. 905–945. DOI: 10.1080/01690961003589492. pmid: 22096264.

Walker, Rachel (2000). 'Long-Distance Consonantal Identity Effects'. *Proceedings of WCCFL*. Vol. 19, pp. 532–545.

Waters, Bruce (1979). 'A Distinctive Features Approach to Djinang Phonology and Verb Morphology'. *Work Papers of SIL-AAB*. Series A 4.

Wedel, Andrew (2006). 'Exemplar Models, Evolution and Language Change'. *The Linguistic Review* 23.3, pp. 247–274.

Wedel, Andrew (2009). 'Resolving Pattern Conflict: Variation and Selection in Phonology and Morphology'. *Analogy in Grammar: Form and Acquisition*. Ed. by James P. Blevins &

Juliette Blevins. Oxford: Oxford University Press. DOI: 10.1093/acprof:oso/9780199547548.001.0001.

Wedel, Andrew, Scott Jackson & Abby Kaplan (2013). 'Functional Load and the Lexicon: Evidence That Syntactic Category and Frequency Relationships in Minimal Lemma Pairs Predict the Loss of Phoneme Contrasts in Language Change'. *Language and Speech* 56.3, pp. 395–417. DOI: 10.1177/0023830913489096.

Wedel, Andrew, Abby Kaplan & Scott Jackson (2013). 'High Functional Load Inhibits Phonological Contrast Loss: A Corpus Study'. *Cognition* 128.2, pp. 179–186. DOI: 10.1016/j.cognition.2013.03.002.

Wilkinson, Karina (1988). 'Prosodic Structure and Lardil Phonology'. *Linguistic Inquiry* 19.2, pp. 325–334.

Wilkinson, Melanie (2012a). *Djambarrpuyŋu: A Yolŋu Variety of Northern Australia*. Munich: LINCOM Europa.

Wilkinson, Melanie (2012b). 'Wubuy Grammar Notes'. Unpublished manuscript.

Wilmoth, Sasha & John Mansfield (2021). 'Inflectional Predictability and Prosodic Morphology in Pitjantjatjara and Yankunytjatjara'. *Morphology*, pp. 355–381. DOI: 10.1007/s11525-021-09380-y.

Winkelmann, Raphael, Jonathan Harrington & Klaus Jänsch (2017). 'EMU-SDMS: Advanced Speech Database Management and Analysis in R'. *Computer Speech & Language* 45, pp. 392–410. DOI: 10.1016/j.csl.2017.01.002.

Winkelmann, Raphael, Klaus Jaensch, Steve Cassidy & Jonathan Harrington (2019). *emuR: Main Package of the EMU Speech Database Management System*.

Winter, Bodo & Andrew Wedel (2016). 'The Co-evolution of Speech and the Lexicon: The Interaction of Functional Pressures, Redundancy, and Category Variation'. *Topics in Cognitive Science* 8.2, pp. 503–513. DOI: 10.1111/tops.12202.

Wolf, Matthew Adam (2008). *Optimal Interleaving: Serial Phonology-Morphology Interaction in a Constraint-Based Model*. PhD thesis. University of Massachusetts Amherst.

Wood, Robin (1978). 'Some Yuulngu Phonological Patterns'. *Papers in Pacific Linguistics*. Ed. by Stephen A. Wurm. Vol. 51. Series A 11. Canberra: Pacific Linguistics, pp. 53–117.

Wurm, Lee H., Mirjam Ernestus, Robert Schreuder & R. Harald Baayen (2006). 'Dynamics of the Auditory Comprehension of Prefixed Words: Cohort Entropies and Conditional Root Uniqueness Points'. *Mental Lexicon* 1.1, pp. 125–146. DOI: 10.1075/ml.1.1.08wur.

Xu, Yi (2005). 'Speech Melody as Articulatorily Implemented Communicative Functions'. *Speech Communication*. Quantitative Prosody Modelling for Natural Speech Description and Generation 46.3, pp. 220–251. DOI: 10.1016/j.specom.2005.02.014.

Xu, Zheng (2016). 'The Role of Morphology in Optimality Theory'. *The Cambridge Handbook of Morphology*. Ed. by Andrew Hippisley & Gregory Stump. Cambridge: Cambridge University Press, pp. 550–587. DOI: 10.1017/9781139814720.020.

Xu, Zheng & Mark Aronoff (2011). 'A Realization Optimality Theory Approach to Blocking and Extended Morphological Exponence1'. *Journal of Linguistics* 47.3, pp. 673–707. DOI: 10.1017/S002222671100003X.

Yip, Moira (1998). 'Identity Avoidance in Phonology and Morphology'. *Morphology and Its Relation to Phonology and Syntax*. Ed. by Steven G Lapointe, Diane Brentari & Patrick M Farrell. Stanford: CSLI Publications, pp. 216–246.

Yu, Alan C L (2007). *A Natural History of Infixation*. Oxford: Oxford University Press.

Zimmermann, Julia (2016). 'Morphological Status and Acoustic Realization: Findings from New Zealand English'. *SST 2016*. Paramatta, Australia, pp. 201–204.

Zukoff, Sam (2020). 'Reduplicant Shape Alternations in Ponapean: Evidence Against Morphological Doubling Theory'. *NELS 50*. Ed. by Mariam Asatryan, Yixiao Song & Ayana Whitmal. Amherst, MA: GLSA.