Title
Gradience and locality in phonology: Case studies from Turkic vowel harmony

Author
McCollum, Adam

Publication Date
2019

UNIVERSITY OF CALIFORNIA SAN DIEGO

Gradience and locality in phonology: Case studies from Turkic vowel harmony

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

in

Linguistics

by

Adam McCollum

Committee in Charge:

Professor Sharon Rose, Chair
Professor Eric Baković
Professor Sarah Creel
Professor Marc Garellek
Professor Rachel Walker

2019

The dissertation of Adam McCollum is approved, and it is acceptable in quality and form for publication on microfilm and electronically.

Chair

University of California San Diego

2019

TABLE OF CONTENTS

Signature Page iii

Table of Contents iv

List of Abbreviations viii

List of Figures ix

List of Tables xiv

Acknowledgments xvii

Curriculum Vita xx

Abstract of the Dissertation xxii

Chapter 1: Introduction 1
- 1.1 Topics and Scope 1
- 1.2 Vowel harmony in Turkic 2
  - 1.2.1 Segmental inventory 2
    - 1.2.1.1 Consonantal inventory 2
    - 1.2.1.1 Vowel inventory 2
  - 1.2.2 Basics of harmony 3
    - 1.2.2.1 Kyrgyz 3
    - 1.2.2.2 Kazakh 6
    - 1.2.2.3 Uyghur 9
    - 1.2.2.4 Uzbek 16
- 1.3 Speakers 18
  - 1.3.1 Kyrgyz 18
  - 1.3.2 Kazakh 19
  - 1.3.3 Uyghur 20
  - 1.3.4 Uzbek 21
- 1.4 Structure of the dissertation 22
  - 1.4.1 Locality 22
  - 1.4.2 Gradience 24

Chapter 2: Locality: Locality in Uyghur backness harmony 27
- 2.1 Introduction 27
- 2.2 Background 28
- 2.3 Methods 33
  - 2.3.1 Stimuli 33
  - 2.3.2 Task 35
  - 2.3.3 Participants 38
  - 2.3.4 Segmentation 39
  - 2.3.5 Statistical analysis 39
- 2.4 Results 39
  - 2.4.1 Root-internal /i/ 39

2.4.2 Suffixal and raised /i/ ………………………………………………………. 46
2.4.2.1 Short-distance condition ……………………………………… 46
2.4.2.2 Long-distance condition ……………………………………… 49
2.5 Discussion ……………………………………………………………………….. 52

Chapter 3: Locality: Strict locality in harmony …………………………………………………. 59
3.1 Introduction ……………………………………………………………………… 59
3.2 Locality and the phonology-phonetics interface …………………………………… 60
3.2.1 Locality ………………………………………………………………….. 60
3.2.1.1 Trans-consonantal locality ……………………………………. 60
3.2.1.2 Strict locality ………………………………………………….. 61
3.2.2 The phonology-phonetics interface ……………………………………… 64
3.3 Methods ………………………………………………………………………….. 68
3.3.1 Research questions ………………………………………………………. 68
3.3.2 Stimuli …………………………………………………………………… 70
3.3.3 Participants ………………………………………………………………. 72
3.3.4 Measurement and statistical analysis ……………………………………. 72
3.4 Results …………………………………………………………………………… 76
3.4.1 Sibilants ………………………………………………………………….. 76
3.4.1.1 Kyrgyz …………………………………………………….... 76
3.4.1.2 Kazakh ……………………………………………………… 78
3.4.1.3 Uyghur ……………………………………………………… 79
3.4.1.4 Uzbek ……………………………………………………….. 80
3.4.2 Lateral ……………………………………………………………………. 81
3.4.2.1 Kyrgyz ……………………………………………………… 81
3.4.2.2 Kazakh ……………………………………………………… 83
3.4.2.3 Uyghur ……………………………………………………… 85
3.4.2.4 Uzbek ………………………………………………………. 87
3.4.3 Summary ………………………………………………………………… 89
3.5 Discussion and analysis ………………………………………………………….. 91
3.6 Conclusion ……………………………………………………………………….. 98

Chapter 4: Gradience in backness harmony: A production study ……………….……………… 100
4.1 Background ……………………………………………………………………… 100
4.2 Methods …………………………………………………………………………. 103
4.2.1 Participants ………………………………………………………………. 103
4.2.2 Stimuli …………………………………………………………………… 104
4.2.3 Task ……………………………………………………………………… 108
4.2.4 Measurement and statistical analysis ……………………………………. 109
4.2.5 Predictions ……………………………………………………………….. 111
4.3 Results …………………………………………………………………………… 114
4.3.1 Kyrgyz …………………………………………………………………… 114
4.3.2 Uzbek …………………………………………………………………….. 120
4.3.3 Kazakh …………………………………………………………………… 127
4.3.4 Uyghur …………………………………………………………………… 131
4.3.5 Summary ………………………………………………………………… 138
4.4 Analysis ………………………………………………………………………….. 138
4.4.1 Possible analyses ………………………………………………………… 139
4.4.1.1 Underspecification …………………………………………..... 140
4.4.1.2 Adaptive hypoarticulation …………………………………….. 141
4.4.1.3 Gradient phonology …………………………………………… 142

4.4.2 Predictions and independent evidence …………………………………… 143
4.5 Discussion ………………………………………………………………………….. 158
4.6 Summary ………………………………………………………………………… 159

Chapter 5: The perceptibility of backness shifts in Kyrgyz, Kazakh, and Uyghur ……………... 160
5.1 Just noticeable differences …………………………………………………… 160
5.1.1 Predictions ………………………………………………………………. 161
5.1.2 Predicted JND with production results ………………………………….. 163
5.1.2.1 Uyghur ………………………………………………………... 163
5.1.2.2 Kazakh ………………………………………………………... 166
5.1.2.3 Kyrgyz ………………………………………………………… 167
5.1.3 Summary ………………………………………………………………… 170
5.2 Perceptual thresholds ………………………………………………………… 172
5.2.1 Predictions ………………………………………………………………. 172
5.2.2 Methods …………………………………………………………………. 176
5.2.2.1 Stimuli ………………………………………………………… 176
5.2.2.2 Task …………………………………………………………... 179
5.2.2.3 Participants …………………………………………………… 180
5.2.3 Results …………………………………………………………………… 181
5.2.3.1 Kyrgyz ………………………………………………………… 181
5.2.3.2 Kazakh ………………………………………………………… 190
5.2.3.3 Uyghur ………………………………………………………... 196
5.2.4 Summary ………………………………………………………………… 200
5.3 Discussion ………………………………………………………………………….. 201

Chapter 6: Incorporating gradience into the formal analysis …………………………………... 208
6.1 Previous work …………………………………………………………………. 208
6.1.1 Phonetic gradience in phonology ………………………………………... 209
6.1.2 Gradience and exceptionality ……………………………………………. 215
6.2 Normalization …………………………………………………………………. 219
6.2.1 Uyghur …………………………………………………………………… 221
6.2.2 Kazakh …………………………………………………………………… 223
6.2.3 Kyrgyz …………………………………………………………………… 224
6.3 Analysis ……………………………………………………………………………... 225
6.3.1 Uyghur …………………………………………………………………… 225
6.3.1.1 Categorical model …………………………………………....... 225
6.3.1.2 Gradient representations ……………………………………..... 230
6.3.1.3 Constraint set ………………………………………………….... 232
6.3.1.4 Gradient model ……………………………………………….... 235
6.3.1.5 Finding the weights …………………………………………… 248
6.3.2 Kazakh ……………………………………………………………………. 253
6.3.3 Kyrgyz ……………………………………………………………………. 257
6.4 Summary ……………………………………………………………………………….. 267

Chapter 7: Discussion …………………………………………………………………….…..... 269
7.1 Reports of gradient harmony ……………………………………………………... 269
7.1.1 Vowel harmony …………………………………………………………... 270
7.1.2 Consonant harmony and vowel-consonant harmony …………………….. 275
7.1.3 Summary ………………………………………………………………….. 278
7.2 Diagnostics for phonological and phonetic effects ………………………………. 281
7.2.1 Previous proposals ……………………………………………………….. 282

7.2.2 Naturalness …………………………………………………………………….. 286
7.2.3 Iterativity and unboundedness ………………………………………….... 290
7.2.4 Predictions of a unified phonology and phonetics ……………………….. 294
7.3 Future directions …………………………………………………………………. 297
7.3.1 Empirical directions …………………………………………………….... 297
7.3.2 Formal directions ……………………………………………………….... 298
7.4 Conclusion ……………………………………………………………………….. 303

Chapter 8: Conclusion ……………………………………………………………………... 305

References ………………………………………………………………………………. 307

LIST OF ABBREVIATIONS

1	first person	IMP	imperative
2	second person	LOC	locative
3	third person	MASC	masculine
ABL	ablative	NEG	negation
ACC	accusative	NMZLR	nominalizer
ADS	adessive	OT	Optimality Theory
CAUS	causative	PFV	perfective
DAT	dative	PASS	passive
F1	first formant	PL	plural
F2	second formant	POSS	possessive
F3	third formant	PST	past tense
FEM	feminine	REL	relativizer
GEN	genitive	S	singular
GER	gerund	VP	verb phrase
HG	Harmonic Grammar	VRB	verbalizer

LIST OF FIGURES

Figure 2.1: F2(z) of root internal /i/ in Uyghur by lexeme and syllable number. Within each distribution, the middle line represents the median and the two other lines indicate the first and third quartiles ………………………………………………………………………… 40

Figure 2.2: F2(z) of root internal /i/ in Uyghur by lexeme and syllable number along with percent [+back] suffixes during elicitation. Within each distribution, the middle line represents the median and the two other lines indicate the first and third quartiles ………….….. 42

Figure 2.3: Relationship between F2(z) of root-final /i/ and suffix backness along with a histogram of counts of back (light grey) and front (dark grey) low vowel suffixes …………... 44

Figure 2.4: Correlations between F2(z) of root-final /i/ and suffixal /ɑ/~/æ/ by lexical item …. 45

Figure 2.5: F2(z) of suffix /i/ by root backness and rounding in short-distance condition ……. 47

Figure 2.6: F2(z) of suffix low vowels by root backness and rounding in short-distance condition ……………………………………………………………………………………….. 49

Figure 2.7: F2(z) of /i/ by root backness and rounding in long-distance condition ……………. 51

Figure 2.8: F2(z) of suffix low vowels by root backness and rounding in long-distance condition ……………………………………………………………………………………….. 52

Figure 3.1: Waveforms and spectrograms for /tɪs/ 'tooth' (left) and /tʏs/ 'dream' from a female Kazakh speaker ………………………………………………………………………… 73

Figure 3.2: Spectral slices from 15 msec. windows at the midpoint of the sibilants in /tɪs/ 'tooth' (left) and /tʏs/ 'dream' from a female Kazakh speaker ………………………………… 74

Figure 3.3: Waveforms and spectrograms for /ʒɯl/ 'year' (left) and /pɪl/ 'elephant' from a female Kazakh speaker ………………………………………………………………………… 75

Figure 3.4: Center of gravity in Hz by backness and rounding for the Kyrgyz sibilants (frication portions of the affricates /ʤ/ are grouped with /ʒ/) …………………………………. 77

Figure 3.5: Center of gravity in Hz by backness and rounding for the Kazakh sibilants………. 78

Figure 3.6: Center of gravity in Hz by backness and rounding for the Uyghur sibilants (frication portions of the affricates /ʧ/ are grouped with /ʃ/) …………………………………... 80

Figure 3.7: Center of gravity in Hz by backness and rounding for the Uzbek sibilants (frication portions of the affricates /ʧ/ are grouped with /ʃ/) …………………………………... 81

Figure 3.8: F2(z) at midpoint of the Kyrgyz lateral (root-internal) based on backness and rounding context ……………………………………………………………………………….. 82

Figure 3.9: F2(z) at midpoint of the Kyrgyz lateral (plural suffix) based on backness and rounding context. No lateral-initial allomorphs of the plural suffix were produced in [-bk,+rd] contexts ………………………………………………………………………… 82

Figure 3.10: F2(z) at midpoint of the Kazakh lateral (root-internal) based on backness and rounding context ……………………………………………………………………………….. 84

Figure 3.11: F2(z) at midpoint of the Kazakh lateral (plural suffix) based on backness and rounding context. Since all plural suffixes occurred after a [-high] vowel and rounding harmony does not target [-high] vowels, no immediately preceding vowels were actually round ……………………………………………………………………………………... 85

Figure 3.12: F2(z) at midpoint of the Uyghur lateral (root-internal) based on backness and rounding context ………………………………………………………………………………. 86

Figure 3.13: F2(z) at midpoint of the Uyghur lateral (plural suffix) based on backness and rounding context ………………………………………………………………………………. 87

Figure 3.14: F2(z) at midpoint of the Uzbek lateral (root-internal) based on backness and rounding context ………………………………………………………………………………. 88

Figure 3.15: F2(z) at midpoint of the Uzbek lateral (plural suffix) based on backness and rounding context. Since backness and rounding harmony are almost non-existent, preceding vowel backness and rounding were used to determine context ……………………………….. 88

Figure 4.1: Example prompts for the target words /bɑltɑ-dɑ/ 'axe-LOC' and /bɑltɑ-lɑr-dɑ/ 'axe-PL-LOC.' The axe indicates the target root, with the number of axes indicating grammatical number. The downward arrow indicates locative case ……………………..... 109

Figure 4.2: Kyrgyz mean vowel F1-F2 (z) by position. Lines connect different mean F1-F2 for a particular vowel. The initial-syllable mean is closest to the transcription and the arrow indicates the final mean. ……………………………………………………………... 115

Figure 4.3: Kyrgyz F2 (z) by position and harmonic pairing ……………………………..... 117

Figure 4.4: F2 (z, with SD) in syllables 1-3 based on root type in Kyrgyz ………………….... 118

Figure 4.5: Uzbek mean vowel F1-F2 (z) by position. Lines connect different mean F1-F2 for a particular vowel. The initial-syllable mean is closest to the transcription and the arrow indicates the final mean ……………………………….. 120

Figure 4.6: F2 (z, with SD) in syllables 1-3 based on root type in Uzbek …………………….. 123

Figure 4.7: F2 (z) of root-internal /i/ in Uzbek by lexeme and syllable ……………………….. 123

Figure 4.8: Kazakh mean vowel F1-F2 (z) by position. Lines connect different mean F1-F2 for a particular vowel. The initial-syllable mean is closest to the transcription and the arrow indicates the final mean ……………………………………………………………... 127

Figure 4.9: Kazakh F2 (z) by position and harmonic pairing ………………………………… 129

Figure 4.10: F2 (z, with SD) in syllables 1-3 based on root type in Kazakh ………………….. 131

Figure 4.11: Uyghur mean vowel F1-F2 (z) by position. Lines connect different mean F1-F2 for a particular vowel. The initial-syllable mean is closest to the transcription and the arrow indicates the final mean ………………………………………………………………… 132

Figure 4.12: Uyghur F2 (z) by position and harmonic pairing ……………………………... 135

Figure 4.13: F2 (z, with SD) in syllables 1-3 based on root type in Uyghur ………………….. 137

Figure 4.14: F2 (z) of high vowels by position and syllable type in Uyghur …………………. 149

Figure 4.15: F2 (z) of high vowels by position and syllable type in Kazakh …………………. 150

Figure 4.16: F2 (z) of high vowels by position and syllable type in Kyrgyz …………………. 153

Figure 5.1: Stop voicing perception in terms of voice onset time (VOT) for the English /b-p/ contrast (adapted from Eimas & Corbit 1973:104) …………………………………………… 173

Figure 5.2: Percent vowel response by stimulus number (continuum varying between /ɪ/ and /æ/) ………………………………………………………………………………………... 174

Figure 5.3: Schematized perception of two vowel categories, /A/ and /B/ …………………… 175

Figure 5.4: F1-F2 plots in Hz for continua between Kazakh /ɪ-ɯ/ and /ʏ-ʊ/ …………………. 178

Figure 5.5: Sample image with instructions and response options from the Uyghur perception test. The prompt reads Қайсисини аңлиналған сөзге eң охшаш таллиңиз, which translates as “Select which of the two is most similar to word you heard.” In this example, the possible responses are түз “straight” and туз “salt” ……………………………………………………. 180

Figure 5.6: Percent [-back] responses to stimuli on Kyrgyz /ɑ/-/e/ continuum with zone of uncertainty (marked by dotted vertical lines) ………………………………………………… 182

Figure 5.7: Percent [-back] responses to stimuli on Kyrgyz /ɑ/-/e/ continuum compared to production means (zone of uncertainty is marked by dotted vertical lines) …………………... 183

Figure 5.8: Percent [-back] responses to stimuli on Kyrgyz /o/-/ø/ continuum with zone of uncertainty (marked by dotted vertical lines) ………………………………………………… 184

Figure 5.9: Percent [-back] responses to stimuli on Kyrgyz /o/-/ø/ continuum compared to production means (zone of uncertainty is marked by dotted vertical lines …………………… 185

Figure 5.10: Percent [-back] responses to stimuli on Kyrgyz /ɯ/-/i/ continuum with zone of uncertainty (marked by dotted vertical lines) ………………………………………………… 186

Figure 5.11: Percent [-back] responses to stimuli on Kyrgyz /ɯ/-/i/ continuum compared to production means (zone of uncertainty is marked by dotted vertical lines) …………………... 187

Figure 5.12: F2 (z) of /ɯ/ by position and syllable type in Kyrgyz …………………………... 188

Figure 5.13: Percent [-back] responses to stimuli on Kyrgyz /u/-/y/ continuum with zone of uncertainty (marked by dotted vertical lines) ………………………………………………. 189

Figure 5.14: Percent [-back] responses to stimuli on Kyrgyz /u/-/y/ continuum compared to production means (zone of uncertainty is marked by dotted vertical lines) …………………... 190

Figure 5.15: Percent [-back] responses to stimuli on Kazakh /ɑ/-/ie/ continuum with zone of uncertainty (marked by dotted vertical lines) ………………………………………………. 191

Figure 5.16: Percent [-back] responses to stimuli on Kazakh /ɑ/-/ie/ continuum compared to production means (zone of uncertainty is marked by dotted vertical lines) …………………... 192

Figure 5.17: Percent [-back] responses to stimuli on Kazakh /ɯ/-/ɪ/ continuum with zone of uncertainty (marked by dotted vertical lines) ………………………………………………. 193

Figure 5.18: Percent [-back] responses to stimuli on Kazakh /ɯ/-/ɪ/ continuum compared to production means (zone of uncertainty is marked by dotted vertical lines) …………………... 194

Figure 5.19: Percent [-back] responses to stimuli on Kazakh /ʊ/-/ʏ/ continuum with zone of uncertainty (marked by dotted vertical lines) ………………………………………………. 195

Figure 5.20: Percent [-back] responses to stimuli on Kazakh /ʊ/-/ʏ/ continuum compared to production means (zone of uncertainty is marked by dotted vertical lines) …………………... 196

Figure 5.21: Percent [-back] responses to stimuli on Uyghur /ɑ/-/æ/ continuum with zone of uncertainty (marked by dotted vertical lines) ………………………………………………. 197

Figure 5.22: Percent [-back] responses to stimuli on Uyghur /ɑ/-/æ/ continuum compared to production means (zone of uncertainty is marked by dotted vertical lines) …………………... 198

Figure 5.23: Percent [-back] responses to stimuli on Uyghur /u/-/y/ continuum with zone of uncertainty (marked by dotted vertical lines) ………………………………………………. 199

Figure 5.24: Percent [-back] responses to stimuli on Uyghur /u/-/y/ continuum compared to production means (zone of uncertainty is marked by dotted vertical lines) …………………... 200

Figure 6.1: Normalized backness of Uyghur words by backness, length and root type (with SE bars). Lighter grey indicates [-back], darker grey indicates [+back]; darker colors indicate disyllabic roots ……………………………………………………………………….. 222

Figure 6.2: Normalized backness of Kazakh words by backness, length and root type (with SE bars). Lighter grey indicates [-back], darker grey indicates [+back]; darker colors indicate disyllabic roots ……………………………………………………………………….. 224

Figure 6.3: Normalized backness of Kyrgyz words by backness, length and root type. (with SE bars). Lighter grey indicates [-back], darker grey indicates [+back]; darker colors indicate disyllabic roots ……………………………………………………………………….. 225

Figure 6.4: Categorical model predictions compared to attested outputs for Uyghur by backness, length and root type (with SE bars) ………………………………………………. 229

Figure 6.5: Normalized backness of Uyghur words by backness, length and root type (with SE bars). Lighter grey indicates [-back], darker grey indicates [+back]; darker colors indicate disyllabic roots …….. 249

Figure 6.6: Gradient HG model predictions compared to attested outputs for Uyghur by backness, length, and root type (with SE bars) …….. 252

Figure 6.7: Gradient HG model predictions compared to attested outputs for Kazakh by backness, length, and root type (with SE bars) …….. 255

Figure 7.1: Uyghur mean vowel F1-F2 (z) by position ……. 273

Figure 7.2: Kazakh mean vowel F1-F2 (z) by position ……... 274

Figure 7.3: Hypothetical pattern of unmarked vowel reduction in a language with [+ATR] harmony ……. 280

Figure 7.4: Potential distributions of sound patterns under a unified view of phonology and phonetics ……. 295

Figure 7.5: Potential distributions of sound patterns under a modular view of phonology and phonetics ……. 296

LIST OF TABLES

Table 1.1: General inventory of consonants in Central Asian Turkic (consonants in parentheses are typically found in Russian loans) 2

Table 1.2: Basic Turkic vowel inventory.. 3

Table 1.3: Information about Kyrgyz participants 19

Table 1.4: Information about Kazakh participants.. 20

Table 1.5: Information about Uyghur participants 21

Table 1.6: Information about Uzbek participants 21

Table 2.1: Schema of possible results 33

Table 2.2: The phonological backness of root-internal /i/ in Uyghur 43

Table 2.3: Schema of possible results compared against actual results 53

Table 2.4: A typology of contrast and distinctiveness (Kiparsky 2015) 54

Table 3.1: Comparison of different locality analyses 63

Table 3.2: Predictions concerning the four research questions in (34) 69

Table 3.3: Stimulus lexemes for all four languages 71

Table 3.4: Summary of significant consonant alternations across the four languages 89

Table 3.5: Mean significant center of gravity differences (Hz) for the sibilants based on vowel context 90

Table 3.6: Mean significant differences in F2 (z) at lateral midpoint based on vowel context.. 91

Table 4.1: Grammatical prompts for the Kyrgyz root /bɑʃ/ ‘head’ 105

Table 4.2: Stimulus lexemes for all four languages 106

Table 4.3: Elicited suffix allomorphs 107

Table 4.4: Language-specific harmonic pairings 111

Table 4.5: Schema for possible patterns and their analysis. Each pattern is paired with at least one analysis, which involves a phonological and phonetic component. Dashed lines indicate front vowel shifts, while solid lines indicate back vowel shifts......................... 112

Table 4.6: Mean F1 (in z, with SD) for Kyrgyz 116

Table 4.7: Mean F2 (in z, with SD) for Kyrgyz ……………………………………………… 114

Table 4.8: Model output for Kyrgyz ……………………………………………………….. 119

Table 4.9: Mean F1 (in z, with SD) for Uzbek …………………………………………….. 121

Table 4.10: Mean F2 (in z, with SD) for Uzbek …………………………………………… 121

Table 4.11: Model output for Uzbek ………………………………………………………. 124

Table 4.12: Model output for Uzbek with no root-internal /i/ ……………………………... 126

Table 4.13: Mean F1 (in z, with SD) for Kazakh ………………………………………….. 128

Table 4.14: Mean F2 (in z, with SD) for Kazakh ………………………………………….. 128

Table 4.15: Model output for Kazakh ……………………………………………………… 130

Table 4.16: Mean F1 (in z, with SD) for Uyghur. Raised vowels are indicated by a subscript. 133

Table 4.17: Mean F2 (in z, with SD) for Uyghur. Raised vowels are indicated by a subscript. 134

Table 4.18: Model output for Uyghur ……………………………………………………... 136

Table 4.19: Schema for possible patterns and their analysis. Each pattern is paired with at least one analysis, which involves a phonological and phonetic component. Dashed lines indicate front vowel shifts, while solid lines indicate back vowel shifts…………………...… 139

Table 4.20: Loan word adaptations in Kazakh …………………………………………….. 146

Table 4.21: Loan word adaptations in Uyghur (Yakup 2005:187-189) ………………………. 147

Table 4.22: Dialectical variation in Uyghur loanword adaptation (Yakup 2005:187-189) …… 148

Table 4.23: Loanword adaptations in Kyrgyz ……………………………………………… 152

Table 5.1: By-syllable F2 in Hz of alternating vowels in Uyghur. Parentheses show standard deviations ……………………………………………………………………………. 164

Table 5.2: By-syllable shifts in F2 in Hz compared with estimated just noticeable differences (JND; 5% threshold) in Uyghur. Shifts which exceed JNDs but in the wrong direction are marked with W.D. ……………………………………………………………. 165

Table 5.3: By-syllable F2 in Hz of alternating vowels in Kazakh. Parentheses show standard deviations ……………………………………………………………………………. 166

Table 5.4: By-syllable shifts in F2 in Hz compared with estimated just noticeable differences (JND; 5% threshold) in Kazakh. Shifts which exceed JNDs but in the wrong direction are marked with W.D.……………………………………………………………….. 167

Table 5.5: By-syllable F2 in Hz of alternating vowels in Kyrgyz. Parentheses show standard deviations ………………………………………………………………………………. 168

Table 5.6: By-syllable shifts in F2 in Hz compared with estimated just noticeable differences (JND; 5% threshold) in Kyrgyz …………………………………………………... 169

Table 5.7: Summary of positional shifts compared to just noticeable differences (JNDs; using a 5% threshold) ……………………………………………………………………. 171

Table 5.8: Sample list of target words for perception study ………………………………….. 178

Table 5.9: Summary of experimental results (* indicates that result derives from word-final fronting, not harmony) ……………………………………………………………………… 201

Table 5.10: Summary of production means compared to just noticeable differences (JNDs; using a 5% threshold) ……………………………………………………………………. 203

Table 5.11: Summary of experimental results (* indicates that result derives from word-final fronting, not harmony) ……………………………………………………………………… 204

Table 6.1: Summary of previous work on gradience in phonology …………………………... 215

Table 6.2: Comparing model fit of gradient and categorical analyses ………………………... 267

ACKNOWLEDGMENTS

It is a weighty thing to consider all the minor choices, chance encounters, and kind deeds that have led to the writing of this dissertation. I am certain I don't even know all the people I should be thanking. My own shortsightedness being noted, there are some people whose help has been so significant that I must say thank you one more time.

First and foremost, innumerable thanks go out to Sharon Rose, whose steady guidance helped me navigate issues ranging from challenging fieldwork sites to what constitutes good evidence to how to write. Sharon has read so many drafts of so many papers, many of which probably made her question taking me on as a student. I have benefited greatly from her mentorship and guidance over the last four years, and I cannot really imagine where I or my research would be without our weekly meetings.

In tandem with Sharon, Eric Baković has been a phenomenal help along this journey. I have appreciated his enthusiasm whenever I bring some new problem to discuss with him, and his ability to connect the dots between seemingly disparate things. Eric had to deal with at least one frantic message from overseas while trying to find my way home amidst coups and other occupational hazards, and in both distress and exuberance, Eric has been a sage guide and generous friend.

My other three committee members have all contributed in very different, yet profound ways. Marc Garellek may read drafts faster than any person in the known world. Marc has bumped my papers up to the top of his to-do list so often, and has also made time to meet with me on so many occasions to discuss the details of this dissertation, particularly elicitations prompts and materials for fieldwork. Sarah Creel has been an open-minded ally from a different department. In my first quarter at UCSD, Sarah's seminar on auditory plasticity helped me think about sound patterns from a different point of view, and her help at various points in this process has been so meaningful. She has made sure that I've interpreted statistical results correctly, and has offered a bigger view of phonology as it connects with human cognition. Last but not least, Rachel Walker kindly agreed to spend hours commuting back and forth from L.A. to San Diego to be on my committee. Rachel's analytical insight has been immensely helpful as I've attempted to navigate some uncharted waters. To all my committee members, I hope I can emulate your selflessness, patience, and insight as I move forward.

I'm also so grateful for the conversations I've had with Anna Mai, Andres Aguilar, and Kati Hout over the past few years. When I have wanted to retreat, talking to you all has reminded me that my world feels relatively tame compared to the wilds of Nahuatl laryngeals, Bantu hiatus resolution, and sinister-sounding mathematical objects called greedoids. And whenever I've found myself feeling like I understand how the world works, Eric Meinhardt has been so kind as to disabuse me of my pride, and remind me that at the end the day, Category Theory is the only real balm in Gilead. I will miss conversations with Matt Zaslansky, Michael Obiri-Yeboah, Maxine Van Doren, Nina Kaldhol, and the other fantastic people I've had the privilege to study with.

Outside of school, there are so many people who have helped when we've had kids, been deathly ill, or been deathly ill with a newborn in the house. Thanks to Mike and Kacie Forrest for doing our laundry, and for being awesome. Other thanks go out to BreeAnn Kirby, Peter and Kara Peckarsky, and to all the folks who've been so loving to me and my family over the last four years.

Speaking of family, I wouldn't be writing this if my wife weren't self-giving beyond compare. She has put her career on hold for years, has been patient while I've travelled overseas, and has been the best friend I could've wished for throughout this tumultuous journey. I love you, and I can't wait to move forward into this next phase of life with you. Aiyan has constantly reminded me that work isn't everything, and that no day is complete without some wrestling. Ari has been faithful to listen and serve around the house. And Maks has offered so many smiles to help mommy and daddy push through the final months of graduate school. Without the four of you, I would've missed out on a lot of joy, a lot of coughs and colds, and innumerable sleepless nights. Thank you all for being willing to walk with daddy through the last four years.

Throughout these years trying to manage school, children, and some semblance of sanity, my sister and parents have been so encouraging and generous while remaining absolutely mystified by the details of what I actually do. Even more than the past few years, though, they have been my earliest supporters. Thank you for making sure I graduated high school. Thank you for all the times you've helped me move, and for your willingness to put everything down to help when I've needed it most. Ashley has been a phenomenal sister, even moving out to the west coast to be closer to my boys. I'm so grateful for her part in my life.

Last but certainly not least, I am so grateful for my friends in Central Asia who made my fieldwork possible. Petya and Aisulu took care of me while in Almaty, connecting me with some friends from long ago. Almas and Mika have been like family to me for so many years now, and I am so grateful for the chance to see their family grow. During this trip Bakyt and Shinar offered fantastic food and great connections to make research on Uyghur happen. Apa and Ata have been loving adopted parents to look after me and make sure I don't eat the rabbits with the spotted livers. During this trip I made a lot of new friends in Kyrgyzstan, too, and they were both gracious as hosts and truly helpful during research. Thanks, in particular, go out to Ainaz for all her assistance in Bishkek and Osh.

CURRICULUM VITA

EDUCATION

2015-2019	Ph.D. student in Linguistics
	University of California, San Diego
	Dissertation: *Gradience and locality in phonology: Case studies from Turkic vowel harmony*
	Dissertation committee: Sharon Rose (chair), Eric Baković, Marc Garellek, Rachel Walker, Sarah Creel
2015	Master of Arts in Linguistics
	University of Florida
	Thesis: Labial harmony in Kazakh: Descriptive and theoretical issues
	Advisors: Caroline Wiltshire (chair), Brent Henderson, Ratree Wayland
2005	Bachelors of Arts (with honors) in Religion, Summa Cum Laude
	Presbyterian College

PUBLICATIONS

accepted	McCollum, Adam G. and Si Chen. Kazakh. *Journal of the International Phonetic Association*
accepted	McCollum, Adam G. and James Essegbey. Initial prominence and progressive vowel harmony in Tutrugbu. *Phonological Data & Analysis*
to appear	McCollum, Adam G. The empirical consequences of data collection methods: A case study from Kazakh vowel harmony. *Linguistic Discovery 16.2:73-110*
2018	McCollum, Adam G. Vowel dispersion and Kazakh labial harmony. *Phonology* 35.2: 287-326. doi:10.1017/S0952675718000052
2017	Chen, Si; Zhang Caicai, Adam G. McCollum, and Ratree Wayland. Statistical modelling of phonetic and phonologised perturbation effects in tonal and non-tonal languages. *Speech Communication* 88: 17-38. doi: 10.1016/j.specom.2017.01.006

CONFERENCE PROCEEDINGS

to appear	McCollum, Adam G. Gradient morphophonology: Evidence from Uyghur vowel harmony. *Proceedings of the Annual Meeting on Phonology,* UC San Diego.
2018	McCollum, Adam G. and Darya Kavitskaya. Non-iterative harmony in Crimean Tatar. *Proceedings of WCCFL 35,* University of Calgary, pp. 259-268.

2018 McCollum, Adam G. and James Essegbey. Unbounded harmony is not always myopic: Evidence from Tutrugbu. *Proceedings of WCCFL 35*. University of Calgary, pp. 251-258.

2017 McCollum, Adam G. Mayak and the typology of labial harmony. *The Supplemental Proceedings of the Annual Meeting on Phonology,* University of Southern California.

2016 McCollum, Adam G. Modeling the Gradient Evolution and Decay of Harmony Systems. *The Supplemental Proceedings of the Annual Meeting on Phonology*, University of British Columbia/Simon Fraser University.

2015 McCollum, Adam G. Labial Harmonic Shift in Kazakh: Mapping the Pathways and Motivations for Decay. *The Proceedings of the 41st Annual Meeting of the Berkeley Linguistics Society*, University of California, Berkeley, pp. 329-352.

ABSTRACT OF THE DISSERTATION

Gradience and locality in phonology: Case studies from Turkic vowel harmony

by

Adam McCollum

Doctor of Philosophy in Linguistics

University of California San Diego, 2019

Professor Sharon Rose, Chair

In very general terms, phonology is the study of both the representational and computational properties of human sound patterns. These issues have been the focus of descriptive, formal, typological, and experimental work. This dissertation draws on experimental and fieldwork data from vowel harmony in four Central Asian Turkic languages, Kyrgyz, Kazakh, Uyghur, and Uzbek, to examine the computational and representational nature of vowel harmony patterns.

One perennial computational question relates to the nature of phonological dependencies – how local must they be? In the dissertation I examine reported transparency in Uyghur backness harmony to evaluate previous analyses of transparent /i/ in the language. Results indicate that putatively transparent vowels actually undergo harmony, which in turn suggests that the analysis of Uyghur is computationally far simpler than previously thought. The dissertation also investigates the strictness with which locality is evaluated, comparing various proposals concerning the participation of consonants in vowel harmony, developing a more nuanced understanding of the interplay between phonetics and phonology that accounts for segment-intrinsic resistance to coarticulation in harmony.

In addition to locality, the dissertation examines the nature of phonological representations. Structuralist and Generative research has generally assumed that phonology manipulates abstract categorical variables, in contrast to the gradient variables that pervade phonetics. As an example, Zsiga (1997) argues that vowel harmony, in contrast to gradient phonetic assimilation, produces categorical alternations between target vowels whose output forms are indistinguishable from their triggering counterparts. Results from an acoustic study suggest that backness harmony in Kazakh and Uyghur produces output sounds that systematically differ from trigger vowel qualities, with the assimilatory effect of harmony gradiently petering out across the word. After comparing findings to plausible phonetic and phonological accounts, I argue that the best account of the data involves gradient phonology. Throughout the rest of the dissertation I develop the claim that phonology may be gradient, examining gradience in harmony from perceptual, formal, and typological perspectives.

Chapter 1: Introduction

1.1 Topics and scope

This dissertation focuses on two theoretical issues: locality and gradience in phonology. These topics are addressed using empirical and experimental data from Turkic vowel harmony collected in Kazakhstan and Kyrgyzstan during the summer of 2017. The languages examined throughout the following chapters are Kyrgyz, Kazakh, Uyghur, and Uzbek. The dissertation combines quantitative methods with the formal analysis of harmony in these languages to develop new insight into the nature of phonological computation. That being said, locality and gradience are topics too broad to comprehensively understand in a lifetime, much less a dissertation. The dissertation approaches these two with rich data from experimental fieldwork in Central Asia, laying out a framework through which to further our understanding of these issues.

This introductory chapter serves to introduce the reader to the basics of vowel harmony in the four languages under study and lay out the overall structure of the dissertation. Section 1.2 discusses both backness and rounding harmony in the four languages being investigated, and should serve as a foundation from which to engage the topics discussed in each of the subsequent chapters. Section 1.3 provides some information about the speakers that graciously participated in the studies to follow. Section 1.4 then describes the empirical and theoretical foci of each chapter, pointing to the chief empirical findings and theoretical claims made throughout the dissertation.

1.2 Vowel harmony in Turkic

1.2.1 Segmental inventory

1.2.1.1 Consonantal inventory

The consonantal inventories of most Turkic languages include around twenty contrastive consonants (Menges 1995). A basic set of consonants that are common throughout the languages being investigated are presented in Table 1.1 below, which are a subset of those presented in Johanson & Csato (1998:xix). Variations on these general consonants are found regularly, particularly along the peripheries of Turkic-speaking areas. For the purposes of this dissertation, though, these are the relevant Turkic consonants. Consonants in parentheses below are typically found in common Russian loans.

Table 1.1: General inventory of consonants in Central Asian Turkic (consonants in parentheses are typically found in Russian loans)

	Bilabial	Labiodental	Dental/Alveolar	Post-alveolar/ Palatal	Velar	Uvular
Stops	p b		t d		k g	q
Fricatives		(f v)	s z	ʃ ʒ	(x)	χ ʁ
Affricates			(ʦ)	ʧ ʤ		
Nasals	m		n		ŋ	
Liquids			l r			
Glides	w			j	w	

1.2.1.2 Vowel inventory

The typical Turkic vowel inventory comprises eight vowels that contrast in height, backness, and rounding, shown in Table 1.2 (see Menges 1995 for more discussion). This idealized inventory frequently varies across languages, but provides a basic template against which to compare the languages under study.

Table 1.2: Basic Turkic vowel inventory

	-back		+back	
	-round	+round	-round	+round
+high	i	y	ɯ	u
-high	e	ø	ɑ	o

1.2.2 Basics of harmony

1.2.2.1 Kyrgyz

Kyrgyz possesses an inventory of eight short vowels, which are very consistent with the vowel qualities in Table 1.2. In addition to the short vowels, Kyrgyz possesses a set of long vowels. Below I focus only on the short vowels, as the long vowels do not differ in regard to the basic harmony patterns. Two types of vowel harmonies operate in modern Kyrgyz, backness harmony and rounding harmony. Each of these is robustly attested in the language, occurring productively throughout the lexicon, and with relatively few exceptions (Batmanov 1938; Wurm 1949; Hebert & Poppe 1963; Comrie 1981; Kara 2003; Toktonaliev 2015; Washington 2016; see Wurm 1949 for some discussion of divergent patterns). Backness harmony is illustrated in (1) below with the accusative and locative suffixes. Observe that after the front unrounded vowels, /i e/, the accusative and locative suffixes are produced with front unrounded vowels, [-di] and [-de], in (1a-d). After back unrounded vowels, however, these same two suffixes surface with back unrounded vowels, [-dɯ] and [-dɑ], in (1e-h).

(1) Backness harmony in Kyrgyz

a.	til-di	‘tongue-ACC’
b.	bel-di	‘lower back-ACC’
c.	til-de	‘tongue-LOC’
d.	bel-de	‘lower back-LOC’
e.	ʤɯl-dɯ	‘year-ACC’
f.	bɑl-dɯ	‘honey-ACC’
g.	ʤɯl-dɑ	‘year-LOC’
h.	bɑl-dɑ	‘honey-LOC’

Backness harmony iterates throughout the word, producing longer words with consistent backness, as shown below in (2). Note that backness harmony does not depend in any way on the height of the vowels, which differs from the rounding harmony pattern in the language. Also, the initial consonant of the relativizer suffix alternates between [g] and [ʁ] in (2). The velar stop occurs in front vowel contexts while the uvular fricative occurs in back vowel contexts. The velar-uvular pairs [g]-[ʁ] and [k]-[q] vary allophonically based on vowel backness.

(2) Iterative backness harmony in Kyrgyz

a. til-der-ibiz-de-gi-ler-di
‘tongue-PL-POSS.1P-LOC-REL-PL-ACC’
“the things that are on our tongues (accusative)”

b. bɑl-dɑr-ɯbɯz-dɑ-ʁɯ-lɑr-dɯ
‘honey-PL-POSS.1P-LOC-REL-PL-ACC’
“the things that are in our varieties of honey (accusative)”

In addition to backness harmony, rounding harmony operates in Kyrgyz, as exemplified by the data in (3). All the roots in (1-2) were unrounded, but the round roots in (3) trigger round vowel allomorphs of the accusative and locative suffixes. In (3a-d), the height of the front round vowel roots that trigger harmony and the suffix do not affect harmony, but in (3e-h), rounding harmony among the back vowels is dependent on vowel height. In (3g), the high vowel in /qul/ is followed by the unrounded back vowel allomorph of the locative suffix, [-dɑ] rather than [-do]. Harmony occurs in all height and backness combinations, except one. A syllable containing the vowel [u] may not be followed by a syllable containing [o] (see Kaun 1995, 2004 for an analysis).

(3) Rounding harmony in Kyrgyz

a.	gyl-dy	'flower-ACC'		
b.	køl-dy	'lake-ACC'		
c.	gyl-dø	'flower-LOC'		
d.	køl-dø	'lake-LOC'		
e.	qul-du	'slave-ACC'		
f.	ʤol-du	'road-ACC'		
g.	qul-dɑ	'slave-LOC'	(qul-do in Southern Kyrgyz, Comrie 1981)
h.	ʤol-do	'road-LOC'		

Note that the failure of rounding harmony in (3g) reportedly varies by dialect. Comrie (1981) reports that speakers of the Southern dialect produce harmony in this context, [qul-do]. However, during data collection only one token of [u]…[o] was recorded. One speaker from Osh produced the word [dos-um-do] 'friend-POSS.1S-LOC' once, but in all other instances both produced [u]…[ɑ] in this context. All other consultants produced only [u]…[ɑ]. Like backness harmony in (2), rounding harmony iterates throughout the word (4), subject only to the condition on high back vowel triggers and non-high back vowel targets illustrated in (3g, 4b). The example in (4c) illustrates that the failure of harmony in the sequence […ubuz-dɑ…] in (4b) is due only to the height of the trigger-target sequence. In (4c) the same possessive suffix successfully triggers harmony on the following accusative suffix, and in (3h) the locative suffix undergoes rounding harmony from a [+back, +round] vowel.

(4) Iterative rounding harmony in Kyrgyz

a. gyl-dør-ybyz-dø-gy-lør-dy
'flower-PL-POSS.1P-LOC-REL-PL-ACC'
"the things that are in/on our flowers (accusative)"

b. ʤol-dor-ubuz-dɑ-ʁɯ-lɑr-dɯ *ʤol-dor-ubuz-do-ʁu-lor-du
'road-PL-POSS.1P-LOC-REL-PL-ACC'
"the things that are on our roads"

c. ʤol-dor-ubuz-du
'road-PL-POSS.1P-ACC'
"our roads (accusative)"

Generally, rounding harmony is described as very consistent in Kyrgyz, although M. Kirchner (1998b) suggests that harmony is optional. Like most descriptions, rounding harmony was very consistently applied across tokens and speakers in the data collected during fieldwork.

1.2.2.2 Kazakh

Unlike Kyrgyz, Kazakh has developed at least one additional contrastive vowel, /æ/ (M. Kirchner 1998a). Additionally, the high vowels have lowered some in Kazakh, and are transcribed as /ɪ ʏ ɯ ʊ/ below. The typical Turkic front mid unrounded vowel, /e/ in Table 2, is diphthongal in Kazakh /ie/ (see McCollum & Chen 2019 for more on the inventory). As for harmony, Kazakh exhibits both backness and rounding harmony. Backness harmony is very consistent, as demonstrated in (5). The forms in (5a-d) illustrate that the accusative and locative suffixes surface as [-dɪ] and [-die] after front unrounded vowels, while those in (5e-h) illustrate that these suffixes surface as [-dɯ] and [-dɑ] after back unrounded vowels.

(5) Backness harmony in Kazakh

a.	tɪl-dɪ	'tongue-ACC'
b.	biel-dɪ	'lower back-ACC'
c.	tɪl-die	'tongue-LOC'
d.	biel-die	'lower back-LOC'
e.	ʒɯl-dɯ	'year-ACC'
f.	bɑl-dɯ	'honey-ACC'
g.	ʒɯl-dɑ	'year-LOC'
h.	bɑl-dɑ	'honey-LOC'

The range of backness harmony is evident in longer words, like those in (6). As above, backness harmony is completely independent of trigger or target vowel height. Observe also that the initial consonant of the relativizer suffix alternates between [g] and [ʁ] in (6). The velar stop occurs in front vowel contexts while the uvular fricative occurs in back vowel contexts. The velar-uvular pairs [g]-[ʁ] and [k]-[q] vary allophonically based on vowel backness non-initially but contrast in word-initial position (compare /qij/ 'manure' with /kij/ 'wear.IMP').

(6) Iterative backness harmony in Kazakh

a. tɪl-dier-ɪmɪz-die-gɪ-lier-dɪ
'tongue-PL-POSS.1P-LOC-REL-PL-ACC'
"the things that are on our tongues (accusative)"

b. bɑl-dɑr-ɯmɯz-dɑ-ʁɯ-lɑr-dɯ
'honey-PL-POSS.1P-LOC-REL-PL-ACC'
"the things that are in our varieties of honey (accusative)"

Alongside backness harmony, Kazakh exhibits rounding harmony. Compared to Kyrgyz, though, rounding harmony is far more restricted in Kazakh. In (7), only [+high] vowels may undergo harmony, while the [-high] vowels do not. In colloquial speech, particularly in more connected speech /ie/ may undergo harmony to [ø], but /ɑ/ almost never undergoes harmony (for more details, see McCollum 2018; McCollum & Chen 2019).

(7) Rounding harmony in Kazakh

a.	gʏl-dɪ ~ gʏl-dʏ	'flower-ACC'
b.	køl-dɪ ~ køl-dʏ	'lake-ACC'
c.	gʏl-die	'flower-LOC'
d.	køl-die	'lake-LOC'
e.	qʊl-dɯ ~ qʊl-dʊ	'slave-ACC'
f.	ʒol-dɯ ~ ʒol-dʊ	'road-ACC'
g.	qʊl-dɑ	'slave-LOC'
h.	ʒol-dɑ	'road-LOC'

In addition to limiting the targets of harmony to high vowels only, morphology also plays a role in the operation of rounding harmony in Kazakh. Within roots, rounding harmony is more regularly obeyed (whether viewed as active harmony or as a morpheme structure condition; Harrison & Kaun 2000). Root-internal examples are shown in (8) below. The two high round vowels [ʏ ʊ] regularly occur in second syllables of polysyllabic roots, while [ø] optionally occurs in examples like (8c,d).

(8) Root-internal rounding harmony in Kazakh

a.	ʒʏzʏm	'grape'
b.	ørʏk	'apricot'
c.	tʏliek ~ tʏløk	'chick'
d.	øziek ~ øzøk	'kernel'
e.	qʊrʊm	'soot'
f.	orʊn	'place'
g.	qʊlɑq	'ear'
h.	bolɑt	'steel'

One additional difference between rounding harmony in Kazakh and Kyrgyz is the domain in which harmony obtains. In Kyrgyz, rounding harmony iterates throughout the word. In Kazakh, though harmony usually extends only to a second- and rarely to a third-syllable. As noted above, harmony is dependent on morphological structure, so harmony is much more common in the third-syllable if the root is polysyllabic (9a-d; Balakayev 1962; Kirchner 1998a; McCollum 2018).

(9) Non-iterative rounding harmony in Kazakh

a.	ʒʏzʏm-dɪ ~ ʒʏzʏm-dʏ	'grape-ACC'
b.	ørʏg-ɪmɪz-dɪ ~ ørʏg-ʏmɪz-dɪ	'apricot-POSS.1P-ACC'
c.	qʊrʊm-dɯ ~ qʊrʊm-dʊ	'soot-ACC'
d.	orʊn-dɯʁ-ɯm ~ orʊn-dʊʁ-ɯm	'place-NMZLR-POSS.1S'
e.	køl-ʏm-dɪ ~ køl-ɪm-dɪ	'lake-POSS.1S-ACC'

1.2.2.3 Uyghur

While Kazakh has developed at least one additional contrastive vowel, Uyghur is described as having lost one, /ɯ/ (Nadzhip 1971; Hahn 1986, 1991, 1998; Lindblad 1990; Yakup 2005; Engesaeth et al. 2009; Abdurehim 2014). In most of these works, underlying /ɯ/ is preserved, but surface [ɯ] is neutralized with [i]. Underlying /ɯ/ is maintained to explain harmony, although it is unclear in previous work how satisfactory the inclusion of an eighth vowel is for the analyses presented (see Vaux 2000 for a critique). Unlike Kyrgyz and Kazakh, the non-high front unrounded vowel in Uyghur is actually low, /æ/, although there is some allophonic variation between [æ] and [ɛ].

Backness harmony in Uyghur triggers alternations on the low vowels and the high rounded vowels. In (10a-d), the locative suffix alternates between [-dæ] and [-dɑ]. In (10e-h), the backness of the initial vowel conditions the backness of the gerundial suffix, as well as the place of articulation of the dorsal consonant, [ɡ]~[ʁ].

(10) Backness harmony in Uyghur

a.	køl-dæ	'lake-LOC'
b.	bæl-dæ	'waist-LOC'
c.	jol-dɑ	'road-LOC'
d.	bɑl-dɑ	'honey-LOC'
e.	kæl-ɡy	'come-GER'
f.	bær-ɡy	'give-GER'
g.	qɑl-ʁu	'remain-GER'
h.	bɑr-ʁu	'go-GER'

Previous work describes /i/ as transparent to harmony in Uyghur. Within roots, /i/ idiosyncratically conditions the backness of following vowels, as in (11a-d; Mayer et al. 2019). In polysyllabic roots, /i/ is transparent.[1] In disyllabic roots with second-syllable /i/, the backness of the initial-syllable vowel controls the realization of the suffix (11e,f).

[1] Thank you to Xiayimaierdan Abudushalamu for providing the forms in (11e,f).

(11) /i/ in Uyghur roots

a.	it-dɑ	'dog-LOC'
b.	til-dɑ	'tongue-LOC'
c.	biz-dæ	'1P-LOC'
d.	siz-dæ	'2S.FORM-LOC'
e.	qædir-gæ	'regard-DAT'
f.	gɑzir-ʁɑ	'sunflower-DAT'

Like non-initial /i/ in roots, suffixal /i/ is also described as transparent. The examples in (12) show two /i/ vowels allowing root backness to determine the realization of a following suffix.

(12) Transparent /i/ in Uyghur suffixes

a.	gyl-imiz-dæ	'flower-POSS.1P-LOC'
b.	bæl-imiz-dæ	'waist-POSS.1P-LOC'
c.	jol-imiz-dɑ	'road-POSS.1P-LOC'
d.	bɑl-imiz-dɑ	'honey-POSS.1P-LOC'

In addition to underlying high vowels, high vowels may arise via vowel raising and epenthesis in the language. As for vowel raising, low vowels in medial open syllables raise to high vowels. In (13a), the unaffixed root for 'child' shows that the word-final vowel is [+low]. The low feature of the root-final vowel is also evident in (13b,c), where this vowel occurs in a closed syllable. Yet, when this vowel occurs in a medial open syllable, as in (13d), it raises to [+high]. The same generalizations hold for other underlying low vowels, like /æ/ in (13e-h). In (13) below, raised vowels are indicated by a subscript.

(13) Transparent raised vowels in Uyghur

a.	bɑlɑ	‘child’
b.	bɑlɑ-m	‘child-POSS.1S’
c.	bɑlɑ-m-dɑ	‘child-POSS.1S-LOC’
d.	bɑli$_R$-dɑ	‘child-PL’
e.	sællæ	‘turban’
f.	sællæ-m	‘turban-POSS.1S’
g.	sællæ-m-dæ	‘turban-POSS.1S-LOC’
h.	sælli$_R$-dæ	‘turban-LOC’

Previous literature generally describes raised vowels as surface [i], although some of Yakup’s (2005) transcriptions suggest that these vowels may alternate for backness. Epenthetic vowels will be discussed in connection with rounding harmony below.

Chapter 2 focuses on transparency in Uyghur, examining root-internal /i/, as well as suffixal /i/ to determine whether historical *ɯ and *i have merged to /i/ in the language, and whether suffixal /i/ alternates for vowel backness. Production data for underlying, epenthetic, and raised vowels are compared, demonstrating that suffixal and raised /i/ does, in fact, alternate for backness harmony.[2]

In addition to backness harmony, rounding harmony also occurs in Uyghur. Rounding harmony exhibits more restrictions than in Kyrgyz and Kazakh. Hahn (1991, 1998) and Lindblad (1990) argue that rounding harmony targets epenthetic vowels only. However, Hahn (1991:50-51) also asserts:

> In some dialects, all or most phonemically high vowels and in certain variants also raised vowels (in the Lopnur “dialect” even non-high vowels) are subject to this roundness assimilation, all of which is attributable to dialectal variation in rule sequence. In this regard, *some dialectally influenced variation is tolerated in the spoken forms of what is considered the standard language. This applies particularly to the labialization of underlying vowels in closed syllables*; e.g. jyzymdin ~ jyzymdyn ‘from my face’, ~ qolumdun ‘from my hand’, qollan

[2] Xiayimaierdan Abudushalamu (personal communication) also indicates that in his productions of disyllabic roots with second-syllable /i/, as in (11e,f), that these vowels exhibit phonetic alternations based on initial-syllable backness.

> 'apply' → qoliniʃ [qollunuʃ] 'application', jyræk 'heart' → jyrigim [jyrygym] (Sadvaskov 198:64). However, such variation is not normally expressed orthographically, certainly not in print [emphasis mine; I have substituted my transcriptions for his].

Hahn suggests the rounding of underlying high vowels, even in the standard variety of the language, may be conditioned based on syllable type. He reports that rounding harmony may occur on underlying high vowels, even in standard language, if they are in closed syllables. However, the examples cited are also compatible with the following interpretation: rounding harmony targets high vowels so long as they are not word-final. This reinterpretation of Hahn's claim is consistent with my data, in which underlying high vowels, epenthetic, and raised vowels all undergo rounding harmony so long as they are not word-final.

Rounding harmony targeting epenthetic vowels is exemplified by the first-person singular possessive suffix in (14). Since epenthetic vowels are never word final they undergo harmony. In (14a, repeated from 13f), a single consonant, /m/, attaches to a vowel-final root to mark this particular possessive suffix. Vowel-final roots like in (14a) provide evidence that this possessive suffix is composed of a consonant only. In (14b-f) though, a vowel is inserted between consonant-final roots and the nasal to repair an illicit coda cluster. The epenthetic vowel in (14b) is /i/, which is unrounded like the root vowel. In (14c-f), however, we see that the epenthetic vowel agrees in both rounding and backness with the root vowel.

(14) Rounding harmony in Uyghur: Epenthetic vowels

a.	sællæ-m	‘turban-POSS.1S’
b.	bæl-im	‘waist-POSS.1S’
c.	gyl-ym	‘flower-POSS.1S’
d.	køl-ym	‘lake-POSS.1S’
e.	qul-um	‘slave-POSS.1S’
f.	jol-um	‘road-POSS.1S’

As indicated earlier, rounding harmony in Uyghur may also target underlying vowels. Root-internal rounding harmony (again, as either active harmony or a static morpheme structure condition) is illustrated in (15). In the examples below, round initial-syllable vowels may be followed by round vowels if the second-syllable vowel is underlyingly high. If the second-syllable vowel is non-high (e.g. 15c,d,g,h), only backness harmony obtains.

(15) Rounding harmony in Uyghur: Root-internal vowels

a.	jyzym	‘grape’
b.	tøʃyk	‘hole’
c.	kyræk	‘spade’
d.	tøʃæk	‘hill’
e.	qurum	‘soot’
f.	orun	‘place’
g.	qulɑq	‘ear’
h.	polɑt	‘steel’

We have seen that epenthetic vowels undergo harmony, and harmony is obeyed within roots with non-initial high vowels. Harmony may also affect suffixal high vowels, shown below with the ablative suffix in (16). In all allomorphs of the ablative suffix in (16), the vowel optionally agrees in both backness and rounding with the root vowel. Chapter 2 focuses on /i/’s participation in harmony, but for the present I transcribe this vowel as /i/ for consistency with extant descriptions of the language.

(16) Rounding harmony in Uyghur: Non-word-final suffix vowels

a.	bæl-din	'waist-ABL'
b.	gyl-din ~ gyl-dyn	'flower-ABL'
c.	køl-din ~ køl-dyn	'lake-ABL'
d.	qul-din ~ qul-dun	'slave-ABL'
e.	jol-din ~ jol-dun	'road-ABL'

Compare the alternation of the ablative suffix in (16) with the behavior of the third-person singular possessive suffix in (17). In (17a-e), this possessive suffix does not undergo either rounding or backness harmony because the target vowel is word-final. In (17f-i) though, the possessive suffix is word-medial, and in this position optionally undergoes harmony. The forms in (17) show most clearly how position constrains harmony in Uyghur: high vowels may undergo harmony only in non-final positions. Word-finally, all high vowels resist rounding and backness harmony.

(17) Rounding harmony in Uyghur: Word-final suffix vowels

a.	bæl-i	'waist-POSS.3S'
b.	gyl-i	'flower-POSS.3S'
c.	køl-i	'lake-POSS.3S'
d.	qul-i	'slave-POSS.3S'
e.	jol-i	'road-POSS.3S'
f.	gyl-i-ni ~ gyl-y-ni	'flower-POSS.3S-ACC'
g.	køl-i-ni ~ køl-y-ni	'lake-POSS.3S-ACC'
h.	qul-i-ni ~ qul-u-ni	'slave-POSS.3S-ACC'
i.	jol-i-ni ~ jol-u-ni	'road-POSS.3S-ACC'

1.2.2.4 Uzbek

The vowel inventory in Uzbek, like in Uyghur, may have lost a contrastive vowel via the merger of historical *ɯ and *i to /i/. It is generally assumed that this is the case, although I am not certain for the Osh dialect. There are many surface [ə] and [ɯ] alongside [i] and [ɪ], and I cannot confidently state that all instances of surface [ə] and [ɯ] derive from /i/ via some allophonic rule. However, I will transcribe all high unrounded vowels as /i/, in conformity with previous descriptive work (Sjoberg 1963; Ibrohimov 1967; Raun 1969; Reshetov & Shoabdurahmonov 1978; Tursunov et al. 1992; Töjchiboyev & Hasanov 2004; Hasanboj 2009). Also, in many cases *ɑ has raised and rounded to /o/ (phonetically more often [ɔ]; compare [boltɑ] with Kazakh and Kyrgyz [bɑltɑ], and Uyghur [pɑltɑ] 'axe') in the modern language, a pattern that mirrors the raising and rounding of historical *a: in Uzbek's Indo-Iranian neighbor, Tajik.

Backness harmony was relatively robust in Chagatai, the genetic predecessor of modern Uzbek (Bodrogligeti 2001), although harmony has generally been lost in the contemporary language. In particular, the standard variety of Uzbek no longer exhibits backness harmony, although some dialects reportedly preserve varying degrees of harmony (Ibrohimov 1967; Reshetov & Shoabdurahmonov 1978; Töjchiboyev & Hasanov 2004). The lack of harmony is evident in the locative suffix, which surfaces as [-dɑ] regardless of root backness.

(18) The locative suffix in Uzbek

a.	til-dɑ	'tongue-LOC'
b.	gyl-dɑ	'flower-LOC'
c.	bæl-dɑ	'lower back-LOC'
d.	køl-dɑ	'lake-LOC'
e.	qul-dɑ	'slave-LOC'
f.	χɑl-dɑ	'answer-LOC'
g.	jol-dɑ	'road-LOC'

High vowels do not alternate for harmony, either, as demonstrated with the accusative suffix in (19). Specifically, the accusative suffix is realized as [-ni] irrespective of root vowel quality.

(19) The accusative suffix in Uzbek

a.	til-ni	'tongue-ACC'
b.	gyl-ni	'flower-ACC'
c.	køl-ni	'lake-ACC'
d.	bæl-ni	'lower back-ACC'
e.	qul-ni	'slave-ACC'
f.	χɑl-ni	'answer-ACC'
g.	jol-ni	'road-ACC'

As can be seen in (19) above, rounding harmony does not regularly apply in Uzbek. In Chagatai, rounding harmony was lexically restricted (Bodrogligeti 2001), and appears to be so in the modern Osh dialect of the language. Typically, rounding harmony in Turkic obtains on high vowels, but in the examples in (19) it does not. However, during data collection rounding harmony did consistently apply to one suffix, the first-person singular possessive, shown in (20). In (20a), a root with a final vowel is followed by /m/, while in (20b-f), a vowel intervenes between the root-final consonant and the nasal consonant. This epenthetic vowel undergoes harmony, with limited backness harmony also applying. The acoustic differences between the vowels transcribed as [y] and [u] below definitely do not approximate the difference between root [y] and [u]. Yet, there is an impressionistic and acoustic difference in backness between the two that suggests not all backness harmony has decayed in the language. See Chapter 4 to compare non-initial [y] and [u].

(20) Rounding harmony in Uzbek

a.	boltɑ-m	'axe-POSS.1S'
b.	bæl-im	'waist-POSS.1S'
c.	gyl-ym	'flower-POSS.1S'
d.	køl-ym	'lake-POSS.1S'
e.	qul-um	'slave-POSS.1S'
f.	jol-um	'road-POSS.1S'

Data was collected from Uzbek to serve as a control group for the studies conducted on the other three languages. To ascertain if some patterns were phonetic or phonological byproducts of harmony, data from closely related Uzbek, which doesn't generally exhibit harmony, was included in the dissertation. Data from Uzbek factors into the discussion of strict locality in Chapter 3, as well as gradience in Chapter 4.

1.3 Speakers

Forty speakers took part in data collection. At several places in the dissertation I briefly summarize the age and gender for the participants from each language. In this section I provide more data, presenting the age, gender, and birth region of each consultant.

1.3.1 Kyrgyz

Thirteen Kyrgyz speakers living in Bishkek, Kyrgyzstan participated in data collection. Eleven of these participants were female and two were male. Speakers 9-13 were from southern Kyrgyzstan. Southern dialects of Kyrgyz exhibit some phonological patterns that diverge from the Northern dialect, which has served as the basis for standardization in Kyrgyzstan (see e.g. Batmanov 1938, 1940; Junusaliev 1971 for more on Kyrgyz dialects). After data collection, differences between speakers of each dialect were compared and no differences outside of lexical preferences emerged (e.g. Southern speakers produced /ɑsɑl/, a lexeme shared with Uzbek, instead of /bɑl/ 'honey'). Jonathan Washington (personal communication) suggests significant dialect levelling is reducing variation from the standard variety of the language. Additionally, since data was collected in the capital, Bishkek, which is situated in northern Kyrgyzstan, it is not terribly surprising that the data generally conforms to the Northern dialect. Speaker data is presented in Table 1.3.

Table 1.3: Information about Kyrgyz participants

Speaker number	Age	Gender	Birth region
1	22	Female	Bishkek
2	19	Female	Bishkek
3	18	Female	Bishkek
4	37	Female	Naryn
5	57	Female	Naryn
6	18	Male	Naryn
7	32	Female	Naryn
8	50	Female	Chu
9	23	Female	Talas
10	42	Female	Osh
11	20	Female	Osh
12	45	Female	Qara køl
13	45	Male	Osh

1.3.2 Kazakh

Nine Kazakh speakers living in Taldykorgan, Kazakhstan participated in data collection. Seven of those speakers were female and two were male. All speakers were from southeastern Kazakhstan except Speaker 9, who was born in Shymkent in Southern Kazakhstan. Dialect differences are not nearly as large in Kazakh as in Kyrgyz, and even though Speaker 9 grew up speaking the Southern dialect of the language, she has lived in Taldykorgan for a number of years, speaking the Eastern variety of the language common in the Taldykorgan and Almaty areas of southeastern Kazakhstan (see Amanzholov 1959 for background on Kazakh dialects). Table 1.4 provides data about Kazakh participants.

Table 1.4: Information about Kazakh participants

Speaker number	Age	Gender	Birth region
1	24	Male	Tekeli
2	49	Female	Taldykorgan
3	45	Male	Raimbek
4	22	Female	Tekeli
5	21	Female	Taldykorgan
6	19	Female	Taldykorgan
7	36	Female	Almaty
8	42	Female	Taldykorgan
9	43	Female	Shymkent

1.3.3 Uyghur

Nine Uyghur speakers living in Chunja, Kazakhstan participated in data collection. Five of the participants were female, and four were male. All nine speakers were born and grew up in the Uyghur district in southeastern Kazakhstan. Speaker data for the nine Uyghur participants in presented in Table 1.5. The variety of Uyghur spoken in Kazakhstan is part of the central dialect group, which includes the standard variety upon which the various orthographies (notably Perso-Arabic and Cyrillic; see Yakup 2005 and Abdurehim 2014 for discussion of Uyghur dialects).

1.3.4 Uzbek

Nine Uzbek speakers living in Osh, Kyrgyzstan participated in data collection. Five of the Uzbek participants were male and four were female. Speaker 2 was born in Tashkent, Uzbekistan but all other speakers were born and raised in the Osh region of Kyrgyzstan. Speaker data is presented in Table 1.6.

Table 1.5: Information about Uyghur participants

Speaker number	Age	Gender	Birth region
1	51	Female	Uyghur district
2	57	Female	Uyghur district
3	19	Female	Uyghur district
4	27	Male	Uyghur district
5	34	Male	Uyghur district
6	48	Male	Uyghur district
7	63	Male	Uyghur district
8	47	Female	Uyghur district
9	54	Female	Uyghur district

Table 1.6: Information about Uzbek participants

Speaker number	Age	Gender	Birth region
1	23	Male	Osh
2	36	Female	Tashkent, Uzbekistan
3	38	Male	Osh
4	45	Female	Osh
5	21	Male	Osh
6	28	Female	Osh
7	27	Female	Osh
8	2	Male	Osh
9	19	Male	Osh

1.4 Structure of the dissertation

1.4.1 Locality

Locality is a recurrent topic throughout the history of Generative phonology. Gafos (1999) frames locality as a precondition for all assimilation. For vowels, Gafos argues that patterns like vowel harmony depend on the nature of the vocalic and consonantal gestures. Building off of work on vowel-to-vowel coarticulation (e.g. Öhman 1966), he contends that vowel gestures persist across intervening consonants, making syllable-adjacent vowels truly adjacent in a more meaningful way. This local vowel-to-vowel interaction has been variously formalized as autosegmental planes, tiers, suprasegmental rules, and gestures. One might ask, why does it matter if phonological patterns are local?

This question can be answered in at least two very related ways. First, if non-local phonological interactions are allowed into a theory, how does one restrict the expressivity of that theory? In work like Archangeli & Pulleyblank (1994), enhanced representational structure (e.g. autosegmental tiers) in tandem with well-formedness conditions are used to curtail the overly expressive power of a non-local theory of phonology. In more recent work framed within formal language theory, the vast majority of phonological patterns fall within the innermost regions of the subregular hierarchy, supporting a larger claim that phonological patterns are restricted in such a way that more complex, non-local dependencies are excluded from the typology (Heinz 2011a, 2011b, 2018; Chandlee & Heinz 2018; Jardine 2016). Computational phonologists have argued that non-local patterns are absent (or strikingly rare) because they are not learnable (Lai 2015; Avcu 2018), in contrast to the relative expressivity of syntactic patterns (Heinz & Idsardi 2011, 2013). This claim relates to the second possible way to address the importance of locality in phonology, learnability. If phonological patterns may be decidedly non-local, then the hypothesis space for the learner grows to unwieldy proportions. Hayes & Wilson (2008) demonstrate that their inductive phonotactic learner cannot learn a vowel harmony pattern with a linear representational system. In such a system, syllable-adjacent vowels are not local, but instead must relate to one another across a potentially large number of consonants. In their simulations, the learner must entertain featural combinations across three and four segments to capture harmonic generalizations in Shona. The combinatorics of these long-distance relations along with the size of the feature matrix describing the inventory of the language yields 1.9 billion phonotactic constraints to be evaluated. The point they make, echoed in much of the literature on harmony, is that by representing syllable-adjacent vowels locally, one can bootstrap the learning process by reducing the hypothesis space for the learner.

In short, locality provides strong restrictions on the expressivity of phonological patterns, and as a result, restricts the size of the hypothesis space so as to make learning more feasible. From a formal perspective, locality is a great boon, but one must ask the question why, if locality is so fundamental to a theory of phonology, is it discussed again and again through the history of Generative phonology. Locality has been a topic for discussion, not because all phonological patterns are clearly local, but instead because many do not seem to be, no matter how much we would like them to be.

Chapter 2 investigates reported non-local (transparent) harmony in Uyghur. Previous descriptions argue that the high unrounded vowel /i/ is transparent to harmony, and as a result, that backness harmony may skip /i/ and assimilate vowels at some distance. Although transparency in a language like Hungarian is constrained by distance (Ringen & Kontra 1989; Hayes & Londe 2006), transparency across at least three high unrounded vowels is reported for Uyghur. Vaux (2000) develops a complex derivational analysis of Uyghur, and Mayer & Major (2018) show that the computational complexity of Uyghur with transparent /i/ is greater than most other harmony patterns. Previous descriptions and analyses of Uyghur harmony rely on impressionistic data (Nadzhip 1971; Lindblad 1990; Hahn 1991, 1998; Vaux 2000; Yakup 2005; Abdurehim 2014), but Chapter 2 presents evidence from a production experiment to examine the participation of /i/ in harmony, both within roots and in suffixes. The complexity of previous analyses depends largely on the non-participation of /i/, but results from Chapter 2 indicate, contrary to previous work, that the high unrounded vowels alternate for harmony [i]~[ɯ] in suffixes. Although root-internal evidence is far less clear on the contrastive status of [i] and [ɯ], it is clear that non-initial /i/ undergoes significant backness and rounding alternations. As a result, the analysis of Uyghur does not rely on transparent /i/, providing a way forward for continued work on the analysis of harmony in the language.

Chapter 3 examines a related issue in Turkic, the strictness of locality. While locality is clearly a desirable formal property of any theory, definitions of locality vary greatly throughout the literature. The tightest restriction on action-at-a-distance has been called 'strict locality' (Walker 1998; Gafos 1999; Ní Chiosáin & Padgett 2001; Benus & Gafos 2007; Smith 2018), which serves as the theoretical focus for Chapter 3. Consistent with the body of work advocating for strictly local assimilation in harmony, Kazakh linguists have argued for several decades that backness and rounding harmony in both Kazakh and Turkic more generally is strictly local (Dzhunisbekov 1980, 1991; Abuov 1994). Chapter 3 thus investigates the reported locality of harmony in Kyrgyz, Kazakh, Uyghur, and Uzbek through an analysis of intervening sibilants and laterals. Acoustic analysis indicates most sounds obey both harmonies, but that at least one sound fails to undergo alternations for at least one harmony. The chapter compares predictions from three different views of locality. Existing proposals stipulate strict locality either through representational assumptions (Gafos 1999; Smith 2018) or via restrictions on Gen (Ní Chiosáin & Padgett 2001), but I advocate instead that strict locality can be modeled in Optimality Theory (Prince & Smolensky 1993) as a violable constraint, incorporating phonetic evidence on coarticulatory resistance to develop a brief analysis of limited consonantal transparency in harmony.

1.4.2 Gradience

Throughout much of Structuralist and Generative phonology, one working assumption has been that phonological computation is categorical and discrete. Phonology involves the manipulation of an inventory of abstract symbols. For many, the categorical, symbolic nature of phonology is what distinguishes it from phonetics. Phonetics involves gradient translation of phonological categories into continuous space and time (see Cohn 1993, 1998, 2003, 2006; Chitoran & Cohn 2009 for discussion).

This distinction has held sway since the 1980's, and in conjunction with pre-existing diagnostics from work in Lexical Phonology reifies linguists' intuitions about what is phonological and what is phonetic. According to most linguist's intuitions, phonological patterns are categorical, interact with morphology and other abstract linguistic structure, describable in terms of distinctive features or other categorical representations (e.g. moras, syllables, words etc), and exceptionful. While many other characterizations are possible, these lay out much of what phonology is within linguistic theory.

Vowel harmony is one of the quintessential cases of a phonological pattern. Vowel harmony interacts with morphology, as lexical phonological patterns do (e.g. van der Hulst & van der Weijer 1995; Baković 2000). Vowel harmony is bounded by phonological domains like roots, words, and sometimes phrases. Harmony requires access to relatively abstract features, triggering alternations that are not always straightforward in phonetic terms (e.g. Kazakh /ɑ/ alternates with /ie/ for backness harmony). Lastly, vowel harmony is replete with exceptions (Finley 2010; Mahanta 2012). For all these reasons vowel harmony stands out as one of the clearer examples of what phonology is.

Yet, despite all the indicators that suggest that vowel harmony is at the core of what phonologists think phonology is, Chapter 4 demonstrates that vowel harmony may be gradient, producing sounds that show incomplete assimilation to the harmonic feature. At a definitional level, this kind of subphonemic gradience should be excluded from phonology if one holds to the assumptions above. In Chapter 4, I lay out potential phonological and phonetic explanations for gradient patterns in Central Asian Turkic, systematically demonstrating that known phonetic patterns of reduction (centralization) and interpolation across unspecified segments cannot account for the attested patterns. Findings suggest that vowel harmony may be gradient, and by extension, that phonology may be gradient. If phonology can be gradient, this immediately raises the question of how to differentiate phonology from phonetics. The final three chapters attempt to reframe discussion of these two modules of the grammar, outlining how to incorporate gradient variables into the formal analysis of vowel harmony.

Chapter 5 discusses a proposed prerequisite for phonological distinctions, perceptibility. If the gradient harmony patterns in Chapter 4 are truly phonological, their outputs should be distinguishable from underlying vowel qualities. The chapter uses previous work on just noticeable differences (e.g. Flanagan 1955; Kewley-Port 2001; Szeredi 2016), using predictions from previous literature to evaluate the perceptibility of positional variation in each language. In Chapter 5 I also report findings from a perception study that suggests patterns of phonetic centralization in Kyrgyz do not cross thresholds of discriminability while patterns of gradient harmony in Kazakh and Uyghur do. Combined with evidence concerning the perceptibility of positional variation, these findings distinguish phonetic from phonological position variation in the languages under study, further supporting the phonological status of gradient backness harmony.

Chapter 6 develops a Harmonic Grammar (Legendre et al. 1990) analysis of gradient harmony by the incorporation of gradient featural variables. While other work has incorporated gradience percolating up from phonetics (Flemming 2001; Lionnet 2017), this chapter argues that phonology is the source of gradience in Turkic backness harmony. The proposed analysis is shown to account for the vast majority of variance in Kazakh and Uyghur. I argue that categorical harmony in Kyrgyz is differentiated from the gradient patterns found in Kazakh and Uyghur, not by the representations employed, but by constraint weightings. Thus, gradient variables are not restricted to languages with gradient patterns, or to the gradient patterns themselves, but pervade phonological representations.

Chapter 7 discusses the phonology-phonetics interface, as well as a range of potential diagnostics for understanding the differences between the two. Taking a range of factors discussed in laboratory, diachronic and synchronic phonology, I suggest a set of tests to determine whether a given pattern is phonological or phonetic. Afterwards, I discuss the formal analysis developed in Chapter 6 and potential avenues for future development. More generally, Chapter 7 lays out a research program investigating gradience in phonology. Finally, Chapter 8 concludes the dissertation, summarizing findings from the preceding chapters.

Chapter 2: Locality in Uyghur backness harmony

2.1 Introduction

One of the defining characteristics of theories of vowel harmony is locality (Goldsmith 1976; Archangeli & Pulleyblank 1994; Gafos 1999; Nevins 2010). While the representational structures may differ across various theories, each theory demands that locality define harmonic interactions. Within computational phonology, locality has been leveraged to define and defend the relative simplicity of phonological patterns, including harmony (Heinz 2018; Heinz & Lai 2013). When faced with non-local (also called transparent) interactions in harmony, researchers have developed elaborate formal mechanisms in order to satisfy a formal definition of locality. Gafos (1999) contends that locality is a precondition for all assimilation, and proposes that articulatory gestures extend across intervening segments to capture action-at-a-distance. Odden (1994) develops a set of feature-geometric adjacency parameters through which to impose a stricter sense of locality on autosegmental interactions. Hayes & Wilson (2008) motivate vowel projections as a way to bootstrap the learning process, demonstrating that their phonotactic learner cannot learn the 1.9 billion constraints necessary to define vowel harmony in Shona if vowels are not represented locally (see also Hansson 2014 on projections for consonant harmony). While transparency has generated a profound amount of work developing formal theory, a foundational question still remains for vowel harmony– how non-local may harmony be?

A more recent program of research has asked this empirical question, and in many cases has found that harmony is, completely independent of any theoretical formalism, far more local than previously thought. Phonetic and phonological research on languages with reported transparency has found that, in many cases, 'transparent' segments actually alternate for harmony. Furthermore, some work argues that transparency is also constrained by distance. As distance increases, transparent vowels begin to block harmony. Both of these findings suggest the need to examine reported transparency to more adequately ascertain what degree of transparency is allowable in a theory of harmony. This chapter examines backness harmony in Uyghur to determine to what extent harmonic interactions may be non-local. Extant descriptions of Uyghur claim that harmony may skip multiple /i/ vowels, but no phonetic studies have been conducted on the language. This chapter fills that gap by analyzing acoustic data on the realization of /i/ within roots and suffixes.

2.2 Background

Reported transparency has prompted a number of experimental studies. Two important points have emerged from this literature: 'transparent' vowels often actually phonetically alternate for harmony, and transparency may be subject to distance-based constraints. First, in a number of languages with reported transparency, phonetic studies have shown that the 'transparent' vowels actually alternate for harmony (Gordon 1999; Gick et al. 2006; Benus & Gafos 2007; Ritchart & Rose 2017). While Benus & Gafos (2007) find small phonetic differences in Hungarian /i/ based on backness context, Gick et al. (2006) and Ritchart & Rose (2017) report much more salient alternations in Kinande and Moro. In Kinande, the low vowel /a/ undergoes ATR harmony, surfacing as [ə]. In Moro, 'transparent' /ə/ is shown to actually be two vowels, /ə/ and /ɘ/, which form a contrastive pairing for height harmony. However, Dye (2015) presents the only phonetic evidence to-date for true transparency in harmony. She investigates the realization of transparent high vowels in Wolof using ultrasound and acoustic data, finding that both /i/ and /u/ fail to undergo even low-level phonetic alternations based on ATR context. These studies lend significant support to a decidedly local conception of vowel harmony, although the Wolof results suggest the need for further work to more adequately assess the typicality of low-level phonetic alternations for harmony (see Walker et al. 2008 for locality in Kinyarwanda consonant harmony).

Second, phonological evidence supports an additional, distance-based restriction on transparency. For instance, in Hungarian a single /i/, /iː/, /e/, or /eː/ is transparent (with, as noted above, small phonetic effects) to backness harmony. In (21a,b), a single front unrounded vowel is followed by the [-back] variant of the dative suffix. However, in (21c,d) a single front unrounded vowel is followed by the [+back] variant of the dative suffix. In (21e,f) the transparent vowel occurs in a word-medial syllable, and in these cases the initial-syllable vowel controls the realization of the dative suffix.

(21) Transparency in Hungarian (Hayes & Londe 2006)[3]

a.	tsiːm-nek	'address-DAT'
b.	kert-nek	'garden-DAT'
c.	hiːd-nɑk	'bridge-DAT'
d.	ʃiːp-nɑk	'whistle-DAT'
e.	fyser-nek	'spice-DAT'
f.	pɑlleːr-nɑk	'foreman-DAT'

However, when a regular vowel is followed by two neutral vowels, the subsequent suffix is far more likely to surface as [-back], regardless of initial vowel backness (22f; Vago 1980; Ringen & Kontra 1989; Siptár et al. 2000; Hayes & Londe 2006). In (22d,e), harmony spans a single /i/, but in (22f), the backness of the initial syllable does not reliably control the realization of the dative suffix when two /i/ vowels intervene. In these cases, Ringen & Kontra (1989) report variation, and Gafos & Dye (2011) report a consistent [-back] allomorph of the dative suffix.

[3] /ɑ/ and /e/ are phonetically realized as [ɔ] and [ɛ], respectively. Though [ɔ] is phonetically round, it is paired with long /ɑː/ and not /oː/, suggesting its phonological status in the language.

(22) Transparency and a count effect (Ringen & Kontra 1989; Gafos & Dye 2011)

a.	mɑm	'mom'
b.	mɑm-i	'mom-DIM'
c.	mɑm-ʧi	'mom-DIM'
d.	mɑm-i-nɑk	'mom-DIM-DAT'
e.	mɑm-ʧi-nɑk	'mom-DIM-DAT'
f.	mɑm-i-ʧi-nɑk ~ mɑm-i-ʧi-nek	'mom-DIM-DIM-DAT'

This effect is distinct from the low-level phonetic alternations reported in Benus & Gafos (2007). In (22f), the final syllable is output as [ɑ] or [e] (phonetically [ɔ] or [ɛ]), not slight variants of a single sound, which suggests a decidedly phonological effect. If these vowels were entirely transparent, it should not matter how many intervene between a trigger and target. However, the fact that their number does matter, suggests that they are not entirely transparent, a point made by Ringen & Kontra (1989; see also Hayes & Londe 2006).

Previous work thus suggests that reported transparent vowels often undergo alternations, and their putative invisibility to harmony is constrained by distance. However, descriptions of Uyghur seem to flout both generalizations. First, some reports claim that Uyghur /i/ does not even phonetically vary according to backness harmony. Second, words with multiple transparent /i/ are reported, with no apparent effect of count on the behavior of the transparent vowel. The described behavior of Uyghur /i/ is summarized in the following examples.

In (23a), the high front vowel triggers the front vowel allomorph of the dative suffix. In (23b,c), the same root, /iʃ/ 'work' triggers a [+back] variant of the plural suffix, but a [-back] variant of the verbalizer suffix. Lindblad (1990:Ch. 4) shows that some roots trigger [+back] allomorphs of some suffixes but [-back] allomorphs of others, like (23b,c). Finally, some high front vowels trigger [+back] suffixes, as in (23d).

(23) Initial /i/ in Uyghur

a.	biz-gæ	'1P-DAT'
b.	iʃ-lɑr	'work-PL'
c.	iʃ-læ	'work-VRB'
d.	jil-ʁɑ	'year-DAT'

Thus, Uyghur /i/ parallels Hungarian /eː i iː/ in (21), triggering both front and back allomorphs of suffixes. In addition to this idiosyncratic behavior of /i/ in roots, previous descriptions argue that /i/ is transparent in non-initial positions (Hahn 1986, 1991; Lindblad 1990; Vaux 2000; Abdurehim 2014; cf. Yakup 2005). In much Turkological work on Uyghur, researchers assume an underlying distinction between /i/ and /ɯ/ that is neutralized on the surface (e.g. Binnick 1991; Hahn 1991, 1998; Johanson 1998). Note though, that such an analysis still caanot easily account for forms like (23b,c), without stipulating some lexical difference been nominal and verbal roots for 'work.'

Examples of non-initial /i/ are presented below in (24). In (24a,b), the third-person singular possessive suffix /-i/ intervenes between two front vowels, while in (24c,d) the same vowel intervenes between two back vowels.

(24) Non-initial /i/ and transparency in Uyghur

a.	bæl-i-dæ	'waist-POSS.3S-LOC'
b.	køl-i-dæ	'lake-POSS.3S-LOC'
c.	bɑl-i-dɑ	'honey-POSS.3S-LOC'
d.	qol-i-dɑ	'hand-POSS.3S-LOC'

Specifically, Lindblad (1990:13) and Vaux (2000) argue that transparent /i/ is not even phonetically affected by preceding vowel backness, insisting that all allophonic variation of /i/ is conditioned by consonantal context (see also Hahn 1986:46). Thus, some descriptions of Uyghur claim that the language exhibits a truly transparent /i/, without any low-level phonetic alternations based on vowel backness.

Second, existing literature claims that harmony may span a number of transparent vowels (25; see Vaux 2000). In (25a,b), backness spreads across two intervening transparent vowels while in (25c,d), harmony spans three intervening transparent vowels. In these forms, the locative suffix agrees in backness with the initial vowel regardless of how many /i/ vowels intervene.

(25) Harmony across multiple /i/ vowels in Uyghur

a.	køl-imiz-dæ	'lake-POSS.1P-LOC'
b.	jol-imiz-dɑ	'road-POSS.1P-LOC'
c.	sælli-lir-i-dæ	'turban-PL-POSS.3S-LOC'
d.	pɑqi-lir-i-dɑ	'frog-PL-POSS.3S-LOC'

Given that existing work on transparency in Uyghur depends entirely on textual and impressionistic data, this chapter investigates the pattern with acoustic data from a production study conducted in Kazakhstan. Four questions are addressed. First, what is the distribution of [i] and [ɯ] within roots? To answer this question, it is necessary to know whether both [i] and [ɯ] even surface in the language. If both do, then the question becomes: are these two sounds contrastive or allophonic?

Second, do non-initial [i] and [ɯ] alternate for harmony? A subsidiary question relates to the magnitude of their alternations. If they do, in fact, alternate for harmony, is this a small, presumably imperceptible alternation, like Hungarian /i/, or is this larger, as in Kinande? If there is no alternation between [i] and [ɯ], does harmony skip the high unrounded vowels, as in (23-25)? Fourth and finally, if the high unrounded vowels are transparent, does count affect the realization of following vowels? The possibilities are schematized in Table 2.1. If high unrounded vowels are contrastive in the initial syllable and alternate for harmony non-initially, Uyghur would resemble Moro. If, however, initial-syllable high unrounded vowels are not contrastive, but exhibit phonetic alternations non-initially, the language would be more like Kinande. Alternatively, if the high unrounded vowels don't contrast for [back] within roots or alternate (not even phonetically) in suffixes, the language would resemble Dye's (2015) description of Wolof. Another possibility is summarized in the rightmost column. If [i] and [ɯ] display low-level phonetic alternations under certain distance-based restrictions, the language would parallel Hungarian.

Table 2.1: Schema of possible results

	Phonological alternations	Phonetic alternations	Transparency	Distance-delimited pseudo-transparency
Initial [i]-[ɯ]	/i/ and /ɯ/ are contrastive	[i] and [ɯ] are allophonic	[i] and [ɯ] are allophonic	[i] and [ɯ] exhibit imperceptible allophonic differences
Non-initial [i]-[ɯ]	alternate for [back]	phonetically alternate for [back]	do not alternate for [back]	exhibit imperceptible allophonic differences
Following vowels	alternate for [back]	alternate for [back]	alternate for [back]	alternate for [back] if number of /i/ < *x*
Comparison language	Moro (Ritchart & Rose 2017)	Kinande (Gick et al. 2006)	Wolof (Dye 2015)	Hungarian (Benus & Gafos 2007; Hayes & Londe 2006)

2.3 Methods

2.3.1 Stimuli

During the recording phase, participants were presented a set of pictures corresponding to the Uyghur nouns containing the six uncontroversial harmonic pairings in the language, /ɑ-æ, o-ø, u-y/. Pictorial prompts were used to avoid an orthographic confound, since Uyghur orthographies represent only /i/. Target words were derived from monosyllabic and disyllabic roots, as shown in (26).

Monosyllabic roots ended either in a sibilant or a liquid (e.g. 26a,b), and disyllabic roots contained two vowels that agreed for the feature [high] (26h-m).

(26) Example stimuli

Monosyllabic roots

a.	bɑʃ	'head'
b.	bɑl	'honey'
c.	bæl	'lower back/waist'
d.	jol	'road'
e.	køl	'lake'
f.	qul	'slave'
g.	gyl	'flower'

Disyllabic roots

h.	pɑltɑ	'axe'
i.	sællæ	'turban'
j.	χormɑ	'persimmon'
k.	tøpæ	'hill'
l.	qurum	'soot'
m.	jyʒym ~ jyzym	'grape'

Additionally, stimuli with putative /i/ were drawn from lexical items that either, one, are reported to exhibit variation in Hahn (1991) or Lindblad (1990), or two, are cognates with /ɯ/ in closely related languages that maintain contrastive /ɯ/, Kyrgyz (kr) and Kazakh (kz). Target stimuli with putative /i/ are shown in (27). Two monosyllabic stimuli were selected that are reported to trigger [+back] harmony and have [+back] cognates in Kyrgyz and Kazakh (27a,c). Two monosyllabic stimuli were selected that are reported to trigger [+back] harmony but have [-back] cognates in Kyrgyz and Kazakh, (27b,d). Finally, in (27e,f) two disyllabic target words were selected that reportedly trigger [-back] harmony and correspond to [-back] cognates in Kyrgyz and Kazakh.

(27) Stimuli with putative /i/

	Stimulus	Gloss	Status in Lindblad	Cognates
a.	qiʃ	‘winter’	[+back]	qɯʃ (kr), qɯs (kz)
b.	ʧiʃ	‘tooth’	[+back]	tiʃ (kr), tɪs (kz)
c.	jil	‘year’	[+back]	ʤɯl (kr), ʒɯl (kz)
d.	pil	‘elephant’	[+back]	pil (kr), pɪl (kz)
e.	ilim	‘science’	[-back]	ilim (kr); ɪlɪm, ʁɯlɯm (kz)
f.	ʃilim	‘paste’	[-back]	ʤelim (kr), ʒielɪm (kz)

2.3.2 Task

Each session was divided into training and recording phases. During the training phase, participants were taught a small set of pictorial-grammatical correspondences involving number, case, and possession. This phase typically lasted less than 5 minutes. After participants completed the training, the recording phase began. Throughout each session, participants were presented images on a laptop computer screen that showed both a picture representing a lexical item with a pictorial prompt from the training phase. When speakers were unable to guess the target word from the prompt, they were given either the equivalent Russian word or a paraphrase in the target language.

Roots were elicited in four cases, singular and plural numbers, and in first- and third-person possessive forms. The cases elicited were nominative, accusative, locative, and ablative. Example inflected forms from the roots /køl/ ‘lake’ and /jol/ ‘road’ are shown below in (28).

(28) Example elicited forms

		/køl/	/jol/
a.	NOM	køl	jol
b.	ACC	køl-ni	jol-ni
c.	LOC	køl-dæ	jol-dɑ
d.	ABL	køl-din	jol-din
e.	PL	køl-lær	jol-lɑr
f.	PL-LOC	køl-lær-dæ	jol-lɑr-dɑ
g.	POSS.1S	køl-ym	jol-um
h.	POSS.3S	køl-i	jol-i

I examined three types of high vowels, underlying, epenthetic, and raised. Underlying high vowels are present in three of the suffixes elicited, the ablative, accusative, and third-person possessive (28b,d,h). Epenthetic high vowels occur with the first-person singular possessive suffix (28g). Note that Lindblad (1990), Hahn (1991), and Vaux (2000) describe this suffix as completely harmonic, in contrast to their descriptions of the other possessive suffixes. Thus, if their descriptions are correct, we should expect to see a morpheme-based difference in harmony, with POSS.1S exhibiting alternations while the other suffixes do not. When the POSS.1S suffix attaches directly to a vowel-final stem, it is realized as /-m/, but when it attaches to a consonant-final stem, a high vowel is inserted before the bilabial nasal. Finally, high vowels may arise due to low vowel raising, as shown below in (29). In (29a), the word-final vowels of the nominative forms of /pɑltɑ/ 'axe' and /sællæ/ 'turban' are low. In (29b), these root-final vowels are also low, since the syllable is closed. However, in (29c,d), these vowels occur in medial open syllables, and are raised to [+high]. Raised vowels are indicated by a subscript below.

(29) Raised vowels

		/pɑltɑ/	/sællæ/
a.	NOM	pɑltɑ	sællæ
b.	POSS.1S	pɑltɑ-m	sællæ-m
c.	PL	pɑlti$_R$-lɑr	sælli$_R$-lær
d.	PL-POSS.3-LOC	pɑlti$_R$-li$_R$r-i-dɑ	sælli$_R$-li$_R$r-i-dæ

The four research questions that drive the acoustic analysis in this chapter are stated in (30). These questions operationalize vowel backness as variation in the second formant (F2), since this is the primary acoustic manifestation of varying tongue body backness. Specifically, back vowels exhibit lower F2 while front vowels exhibit higher F2.

(30) Four research questions:

1. Is F2 of initial [i] – [ɯ] predictable based on consonantal context (allophony)?
2. Is F2 of non-initial [i] – [ɯ] (underlying, epenthetic, and raised) predictable based on backness of the initial-syllable vowel?
3. Is F2 of low vowels following medial [i] – [ɯ] predictable based on initial-syllable vowel backness?
4. If F2 of low vowels following medial [i] – [ɯ] is predictable based on initial-syllable vowel backness, is it also predictable based on number of high unrounded vowels?

To examine the fourth question, elicited words fell into one of two conditions, exhibiting either a short- or long-distance dependency (31). The short-distance condition involved only a single high unrounded vowel, either the epenthetic POSS.1S /-m/ or the underlyingly [+high] POSS.3S suffix, which reportedly undergoes rounding harmony as well as backness harmony (31d). In contrast, the long-distance condition involved three high unrounded vowels, the first two of which were derived via raising (29), while the third was always the third-person singular possessive suffix. Roots in the short-distance condition were always monosyllabic, and roots in the long-distance condition were always disyllabic with a root-final low vowel subject to raising (29).

(31) Short- and long-distance conditions

	Short-distance	
a.	bæl-i-dæ	'waist-POSS.3S-LOC'
b.	ʧyʃ-i-dæ	'dream-POSS.3S-LOC'
c.	bæl-im-dæ	'waist-POSS.1S-LOC'
d.	ʧyʃ-ym-dæ	'dream-POSS.1S-LOC'
	Long-distance	
e.	sælli-lir-i-dæ	'turban-PL-POSS.3S-LOC'
f.	tøpi-lir-i-dæ	'hill-PL-POSS.3S-LOC'
g.	pɑlti-lir-i-dɑ	'axe-PL-POSS.3S-LOC'
h.	χurmi-lir-i-da	'persimmon-PL-POSS.3S-LOC'

Words derived from monosyllabic roots were maximally three syllables in length, while words derived from disyllabic roots were maximally five syllables in length.

Sessions were conducted in a quiet room. Participants wore a Shure-SM10A unidirectional head-mounted microphone, and all data were recorded to a Marantz PMD 661 MKII digital recorder at a sampling rate of 44.1 kHz. Each session lasted between 45 and 90 minutes.

2.3.3 Participants

Participants were recruited through existing relational networks in Chunja, Kazakhstan. Nine Uyghur speakers (five females, mean age: 44.4 years, range: 19-63 years) participated in the study. All participants reported native fluency in the target language. Most participants also reported fluency in Kazakh and Russian.

2.3.4 Segmentation

All sound files were segmented in Praat (Boersma & Weenink 2015) . The beginning and end of each vowel was set to the onset and offset of the second formant. In cases where the second formant persisted across flanking consonants, abrupt changes in energy or formant frequencies were used to indicate vowel onset and offset.

2.3.5 Statistical analysis

After segmentation, the first three formants and duration were measured at vowel midpoint. To facilitate across-speaker comparisons, the data were z-score normalized (Lobanov 1971). All vowels in the language were included in the normalization procedure (see Chapter 4 for more on this topic). Outliers were then inspected for measurement errors. In particular, a number of errors were found with the high back vowels, where the formant tracker in Praat failed to distinguish the first two formants. In these cases formant frequencies were hand measured at the approximate vowel midpoint. The data were analyzed in R (R Core Team 2017), using the lme4 package (Bates et al. 2015). A mixed effect linear regression was used to predict normalized F2 at vowel midpoint for both high unrounded vowels and subsequent alternating vowels. Model comparisons are used to assess statistical significance.

2.4 Results

2.4.1 Root-internal /i/

The distribution of F2(z) by lexeme for the six roots compared is presented in Figure 2.1. F2 values range from less than -1z, which is comparable to [ɯ] to more than 2z, a value consistent with a very peripheral [i] vowel quality. In general, though, values tend to fall between 0.5 and -0.5z, which is more [ə] or [ɨ]-like. In addition to the global distribution of F2, there is variation between different lexemes. Observe that the highest F2 is associated with the root, /ʧiʃ/ 'tooth' while the lowest F2 is associated with /jil/ ‘year’, /pil/ ‘elephant’, and /ʃilim/ ‘paste.’ If attempting to describe the distribution of F2 by immediate consonantal context, the following three generalizations appear valid. One, F2 is relatively low preceding the lateral. Two, F2 is relatively high adjacent to a postalveolar, and three, F2 is decreased following a uvular. These three generalizations can account for good deal of variation seen below, but there is still some variation that is unaccounted for by this kind of allophonic analysis of putative /i/.

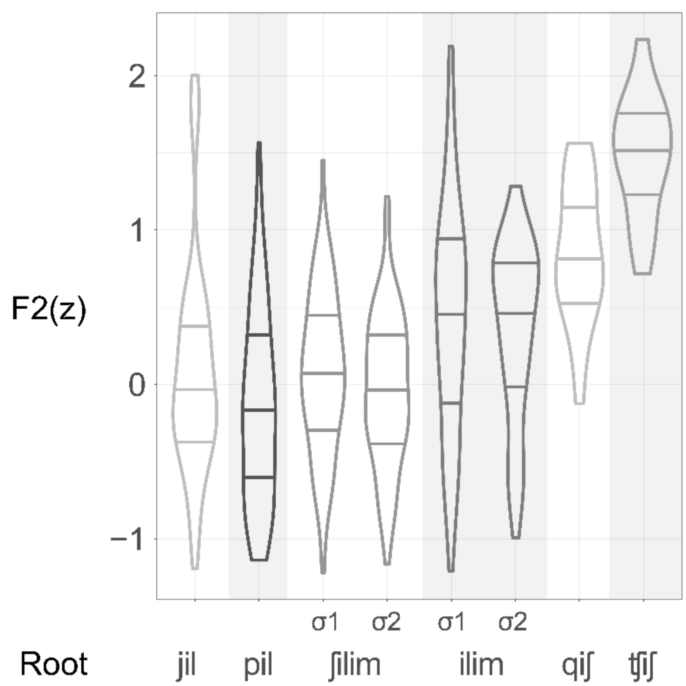

Figure 2.1: F2(z) of root internal /i/ in Uyghur by lexeme and syllable number. Within each distribution, the middle line represents the median and the two other lines indicate quartiles.

In particular, note the difference in F2 between the lexemes /ilim/ ‘science/ and /ʃilim/ ‘paste.’ In this pair, even though only the initial consonant differs, the average difference in F2 is almost 0.5z, a difference that persists over both the first and second syllables within these two roots. To be clear, there is no obvious consonantal reason for this difference, as the postalveolars appear to promote higher F2 in /qiʃ/ ‘winter’ and /ʧiʃ/ ‘tooth.’ The lexeme /ʧiʃ/ has, by far, the highest mean F2, and is flanked by two postalveolar sounds. Moreover, the following lateral does not appear to affect F2 of the initial-syllable in /ilim/ like it does the other vowels preceding a lateral in Figure 2.1. One additional difference between /ilim/ and /ʃilim/ is the range of F2 values they vary over. In /ilim/, F2 values vary wildly, spanning the entire plot. In contrast, values for /ʃilim/ exhibit more moderate variation. Based on Figure 2.1, /ilim/ actually varies between [ilim] and [ɨlɨm], while a narrower transcription for /ʃilim/ would show variation between [ʃɯlɯm] and [ʃɨlɨm].

From the data presented above, it is not clear that /i/ is allophonically distributed between a variety of high unrounded vowels, ranging from [i] to [ɯ]. Rather, it seems more likely that /i/ contrasts with a second high unrounded vowel, /ɯ/. Historical evidence clearly supports the presence of *ɯ in the language, and Hahn (1986) reports that some southern dialects maintain the synchronic distinction between /i/ and /ɯ/ (see also Yakup 2005). Although the variety of Uyghur spoken in Kazakhstan is central, and not southern, it is possible that the historical distinction has been preserved in Kazakhstani Uyghur, in part, due to contact with Kazakh, which maintains a robust backness distinction between the high unrounded vowels. Based on the evidence above, it seems most likely that the high unrounded vowels derive from two separate phonemes, /i/ and /ɯ/. This conclusion is tentative at present, but the realization of these vowels root-internally is not explainable from allophony.

If a synchronic contrast between /i/ and /ɯ/ exists in the language, then it should manifest itself in the selection of suffixes following these roots. In Figure 2.2, the percentage of [+back] vowel suffixes produced during elicitation is compared across roots and their F2 distributions. What is striking is that regardless of root F2, these vowels typically select [+back] allomorphs of alternating suffixes. There is some variation, but no root ever selects [-back] suffixes more than [+back].

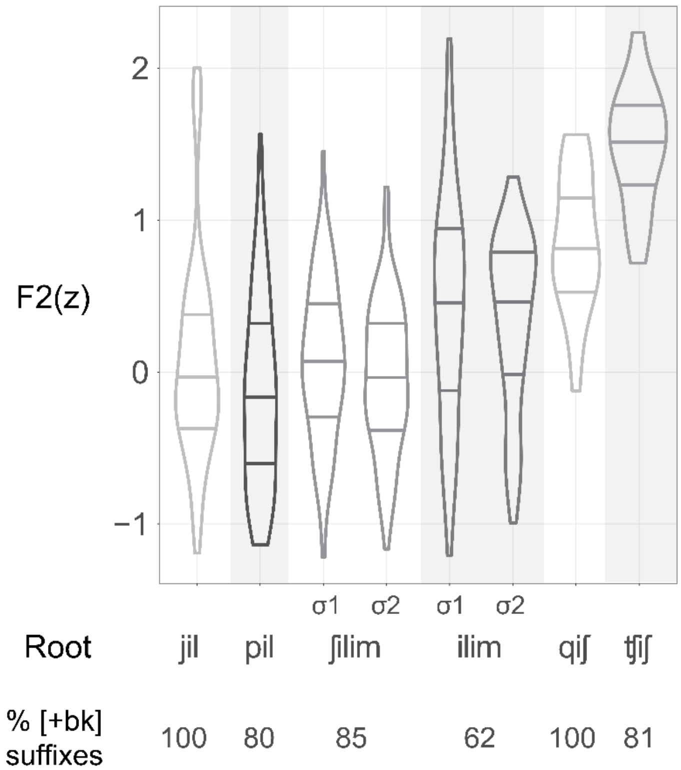

Figure 2.2: F2(z) of root internal /i/ in Uyghur by lexeme and syllable number along with percent [+back] suffixes during elicitation. Within each distribution, the middle line represents the median and the two other lines indicate the first and third quartiles

These results are compared with (27) in Table 2.2. All the monosyllabic roots are classified as [+back] in Lindblad (1990) while the two disyllabic roots are considered [-back]. More significantly, all the roots with [+back] cognates in Kyrgyz and Kazakh trigger [+back] suffixes 100% of the time, whereas roots with [-back] cognates in those languages exhibit variation in suffix selection.

Table 2.2: The phonological backness of root-internal /i/ in Uyghur

Root	Gloss	Status in Lindblad (1990)	Cognates		Count [+back] suffix	Percent [+back] suffix
			Kyrgyz	Kazakh		
qiʃ	winter	[+back]	qɯʃ	qɯs	41/41	100
ʧiʃ	tooth	[+back]	tiʃ	tɪs	39/48	81
jil	year	[+back]	ʤɯl	ʒɯl	39/39	100
pil	elephant	[+back]	pil	pɪl	28/35	80
ilim	science	[-back]	ilim	ɪlɪm ~ ʁɯlɯm	23/37	62
ʃilim	paste	[-back]	ʤelim	ʒielɪm	41/48	85

While it is clear in Table 2.2 that root-internal /i/ (at least for these six lexical items) typically triggers [+back] suffixes, there is also a lower-level phonetic effect at play. F2 of root-final /i/ also has a significant effect on the selection of suffix allomorphs. In a mixed effect logistic regression with random intercepts for speaker and lexeme, increased root-final F2 was highly predictive of a front vowel suffix ($z = 3.8$, $p < .001$). This is evident in Figure 2.3, which plots the likelihood (left axis) and counts (right axes) of a front vowel suffix as a function of the F2 of the root-final vowel (x-axis). Even though root-internal /i/ typically behaves like a back vowel, when more phonetically fronted, it tends to behave more like a front vowel.

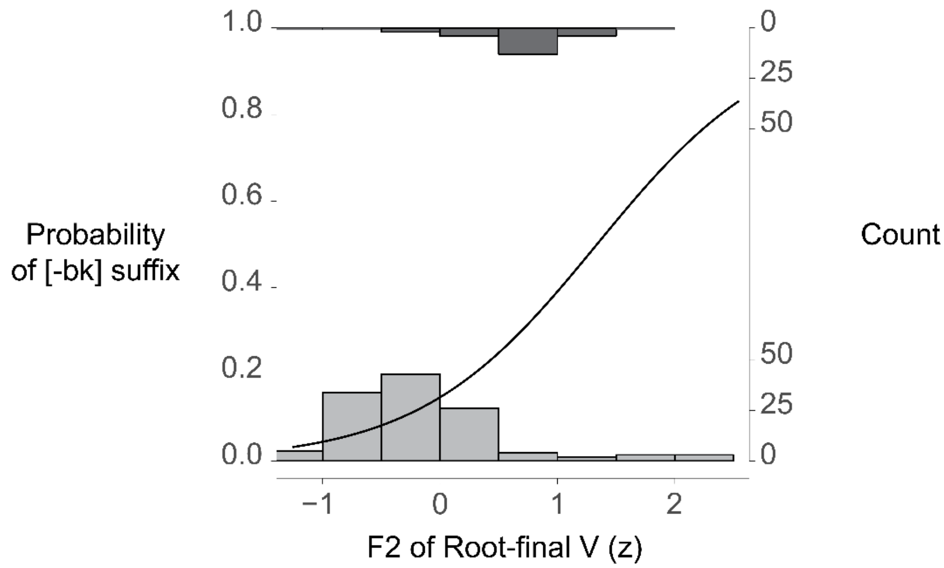

Figure 2.3: Relationship between F2(z) of root-final /i/ and suffix backness along with a histogram of counts of back (light grey) and front (dark grey) low vowel suffixes

The interaction root-internal F2 of /i/ and suffix backness also translates to suffix F2, as one would expect. When root-internal /i/ is more fronted (i.e. exhibits higher F2), the F2 of the following suffix is higher. The moderate correlation between F2 of root-internal /i/ and F2 of the following suffix is seen in Figure 2.4 (*r* =.45, p < .001). For each lexical item, there is a positive correlation between F2 of the root-final vowel and F2 of the following low vowel suffix. Note that only a few tokens of /ʧiʃ/ and /qiʃ/ are plotted below. This is because these vowel were almost always elided during production, so it was not possible to compare the backness of the root vowel with the following suffix. Also, observe that the correlation is weakest for /jil/ and /qiʃ/; there are too few tokens of /ʧiʃ/ to truly support any meaningful correlation. However, the correlation is strongest for /ilim/, /ʃilim/, and /pil/. It may be significant that all of these lexemes have [-back] cognates in related languages. It is also worth noting that the correlation seen here does not indicate the directionality of the relationship. It is possible that selecting a front .or back vowel suffix exerts an anticipatory coarticulatory effect on the root-internal vowel.

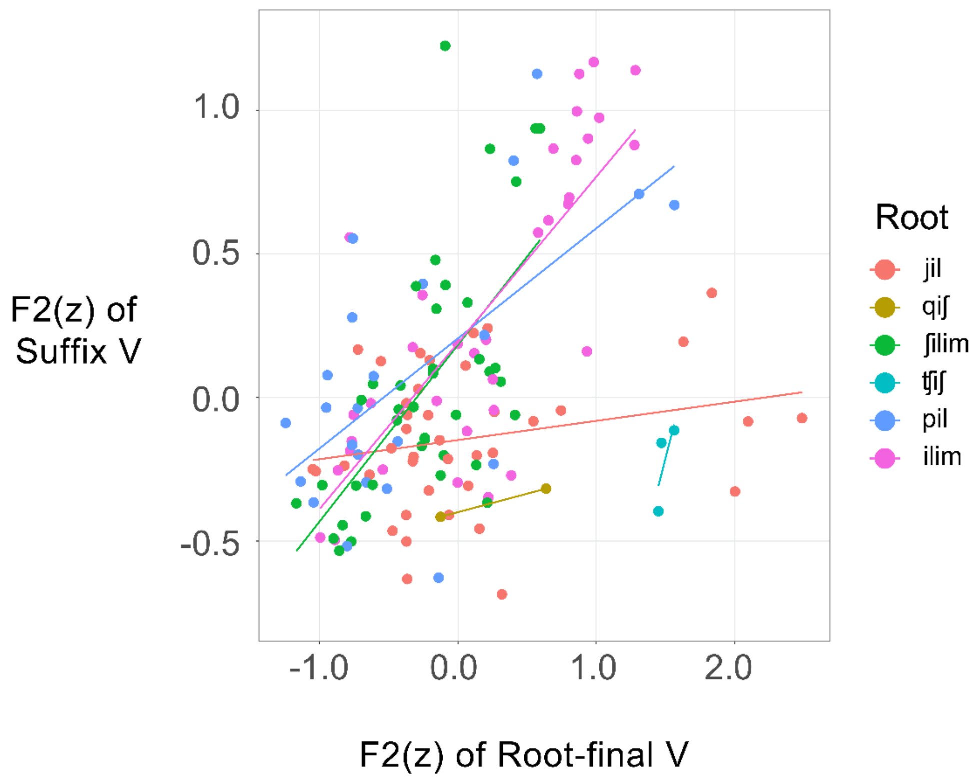

Figure 2.4: Correlations between F2(z) of root-final /i/ and suffixal /ɑ/~/æ/ by lexical item

To sum up, evidence adduced thus far is unclear on both the relationship between root-internal [i] and [ɯ], as these surface phones do not exhibit a clearly allophonic distribution. There is some evidence that phonetic backness factors into suffix selection, but not like one would expect for a typical harmonic pairing. In a typical pairing, one would expect root F2 to account for almost all variation in suffix backness. While root F2 is a significant predictor of suffix backness in the regression model above, and while there is a significant correlation between root and suffix F2, these relationships seem too weak to signal a robust phonological contrast. It is thus unclear whether /i/ has a contrastive [+back] counterpart, /ɯ/, although this seems likely. If there is a second high unrounded vowel in the language, /ɯ/, it would help explain variation between near-minimal pairs like /ilim/ [ɪlɪm] ‘science’ and /ʃɯlɯm/ [ʃɯlɯm] ‘paste.’ If one posits that there is only one high unrounded vowel phoneme in the language, it is not at all clear that it should be /i/ and not /ɯ/, since [ɯ] is much more common, and occurs with a much freer distribution (see also Hahn 1986 for discussion). When the F2 values in Figure 2 are consulted, these values suggest that vowel qualities ranging from central to back are more common than phonetically front vowels.

2.4.2 Suffixal and raised /i/

This subsection investigates whether or not /i/ alternates for harmony in suffixes and in raised vowel contexts. To that end, the subsection is broken up into two parts, one examining results from the short-distance condition, where a single /i/ intervenes between a trigger and target; the second examining results from the long-distance condition, where three /i/ vowels intervene between a harmony trigger and target.

2.4.2.1 Short-distance condition

In this condition, second-syllable high vowels from three-syllable words are compared to determine whether or not F2 varies with initial-syllable backness, as well as rounding. Rounding is considered because extant descriptions indicate that the epenthetic vowel of the first-person singular possessive suffix undergoes both backness and rounding harmony. Since both the first-person singular possessive suffix and the third-person singular possessive suffix are evaluated, this second predictor, in addition to a predictor for suffix allows the model to differentiate between general effects and effects specific to the first-person singular possessive suffix. The linear model used predicts normalized F2 at vowel midpoint from initial-vowel backness, initial-vowel roundness, suffix; with interactions between backness and roundness, and roundness and suffix, as well as a random intercept for speaker.

As seen below in Figure 2.5, root backness exerts a significant effect on F2 of putatively transparent /i/, [$\beta = -1.74$, $\chi^2(1) = 82.32$, $p < .001$]. There is an additional effect of root rounding on F2 of these high vowels [round: $\beta = -1.07$, $\chi^2(1) = 14.50$, $p < .001$]. There are two key takeaways from the short-distance condition. First, the high unrounded vowel does alternate for backness harmony, as seen in the significant reduction in F2 when following a back vowel. Thus, /i/ is not transparent to harmony in this condition, but alternates between a relatively fronted vowel, similar to [ɪ] and a more backed vowel, closer to [ɯ]. Second, the high vowel also alternates for rounding in front vowel contexts. The darker grey boxes in Figure 2.5 show decreased F2, consistent with an effect of rounding, which depresses both F2 and F3. Lastly, in the data, any effect of rounding was diminished preceding the bilabial nasal, presumably because this sound triggers significant lowering of F2 on its own [$\beta = -1.07$, $\chi^2(1) = 9.38$, $p < .01$]. If a following /m/ triggers coarticulatory F2 depression on a preceding high vowel, the acoustic distinction between rounded and unrounded vowels diminishes.

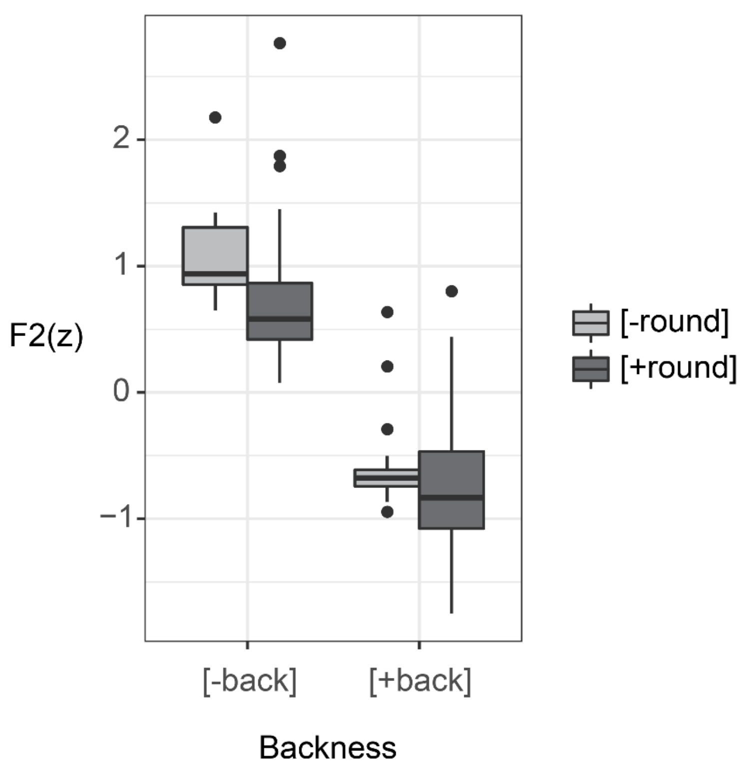

Figure 2.5: F2(z) of suffix /i/ by root backness and rounding in short-distance condition

Moving on, it is now crucial to determine whether or not regularly alternating suffixes, here the locative ([-dɑ]~[-dæ]) and plural ([-lɑr]~[-lær]) suffixes alternate for harmony after a high vowel. Normalized vowel F2 is thus predicted from initial vowel backness and rounding, as well as suffix. Additionally, the model includes an interaction term for backness and rounding. No interactions for suffix type are considered, so I have no a priori hypotheses concerning the locative or plural suffix's effect on backness or rounding.[4]

In Figure 2.6 it is clear that the low vowels alternate for backness harmony. Statistically, this is manifested in a significant effect of initial vowel backness on the F2 of the locative and plural suffixes [$\beta = -0.89$, $\chi^2(1) = 13.25$, $p < .001$]. Unlike the high vowels, F2 of the low vowels was not significant affected by rounding harmony [$\beta = 0.02$, $\chi^2(1) = 0.01$, $p = .91$], which is consistent with descriptions of the central dialect, in which rounding harmony targets high vowels only. Neither suffix nor the interaction between backness and rounding were significant predictors of low vowel F2 [suffix: $\beta = 0.18$, $\chi^2(1) = 1.18$, $p = .28$; round : back: $\beta = 0.05$, $\chi^2(1) = 0.04$, $p = .85$].

[4] One would likely predict phonetic effects of consonantal context, but these are not germane for assessing whether these suffixes alternate between /ɑ/ and /æ/ after high vowels. If these suffixes do alternate for harmony, the effect should be large enough to ignore smaller phonetic effects of context.

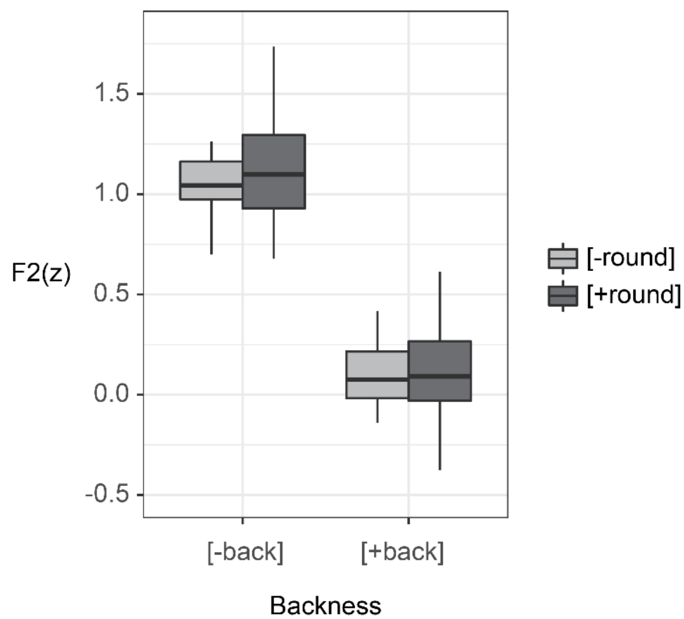

Figure 2.6: F2(z) of suffix low vowels by root backness and rounding in short-distance condition

To summarize this subsection, the key finding is that high vowels alternate for both backness and rounding. The effect is clearer for backness, but that is not surprising since the acoustic effects of rounding are typically much smaller than the effects of vowel backness. Given that the high unrounded vowels alternate for backness harmony, it is unsurprising that following low vowels also alternate for harmony. In short, results from the short-distance condition support an analysis whereby both high and low vowels regularly undergo backness harmony in Uyghur.

2.4.2.2 Long-distance condition

In this condition, the medial three syllables of five-syllable words consisted of high vowels derived from vowel raising as well as an underlying high vowel. The raised vowels occurred in syllables 2 and 3, and the underlying high vowel of the third-person singular possessive suffix occurred in the fourth syllable. As in Section 2.3.2.1, a linear mixed effects model was used to predict normalized F2 at vowel midpoint from initial-vowel backness and rounding. An additional fixed effect of syllable number was included in the model to determine if the second and third high vowels in these sequences pattern like the first high vowel. Two-way interactions between backness and rounding, backness and syllable number, and rounding and syllable number along with a random intercept for speaker were incorporated into the model.

As in the short-distance condition, initial vowel backness was a significant predictor of high vowel F2 [$\beta = -2.52$, $\chi^2(1) = 236.22$, $p < .001$]. In this condition, though, rounding was not a significant predictor of vowel F2 [$\beta = -2.52$, $\chi^2(1) = 0.96$, $p = .33$]. As seen in Figure 2.7, F2 of front vowel decreased by syllable [$\beta = -0.18$, $\chi^2(1) = 11.63$, $p < .001$], while F2 of back vowels increased by syllable [$\beta = 0.65$, $\chi^2(1) = 102.24$, $p < .001$]. In other words, F2 of front vowels decreased by 0.18z across syllables while F2 of back vowels increased by 0.47z (0.65 – 0.18 = 0.47) across syllables. The magnitude of these positional shifts is quite noticeable below. This asymmetry will be discussed in detail in Chapter 4. The interactions between backness and rounding [$\beta = 0.06$, $\chi^2(1) = 0.35$, $p = .55$], as well as syllable number and rounding, were not significant [$\beta = -0.01$, $\chi^2(1) = 0.05$, $p = .82$].

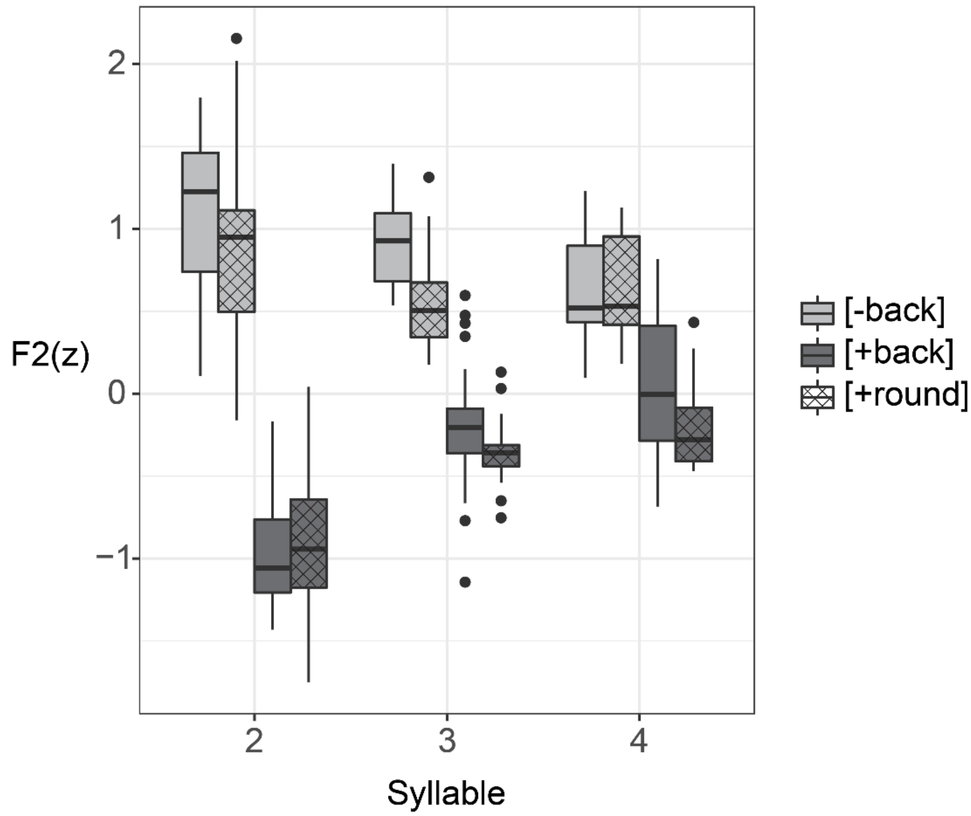

Figure 2.7: F2(z) of /i/ by root backness and rounding in long-distance condition

To determine whether or not low vowels alternate for harmony in the long-distance condition, fifth-syllable low vowels in the locative suffix ([-dɑ]~[-dæ]) were examined. Initial vowel backness and rounding, as well as their interaction were used to predict low vowel F2. As above, a random intercept for speaker was also included in the model.

As in the short-distance condition, initial-vowel backness significantly affected low vowel F2 [$\beta = -1.18$, $\chi^2(1) = 96.88$, $p < .001$]. In other words, low vowels alternate for backness harmony in the long-distance condition, shown in Figure 2.8. Initial-vowel rounding did not significant affect low vowel F2 [$\beta = -0.19$, $\chi^2(1) = 2.55$, $p = .11$], nor was the interaction between backness and rounding significant [$\beta = 0.03$, $\chi^2(1) = 0.03$, $p = .85$], although Figure 2.8 suggests a tendency toward diminished F2 after round vowels.

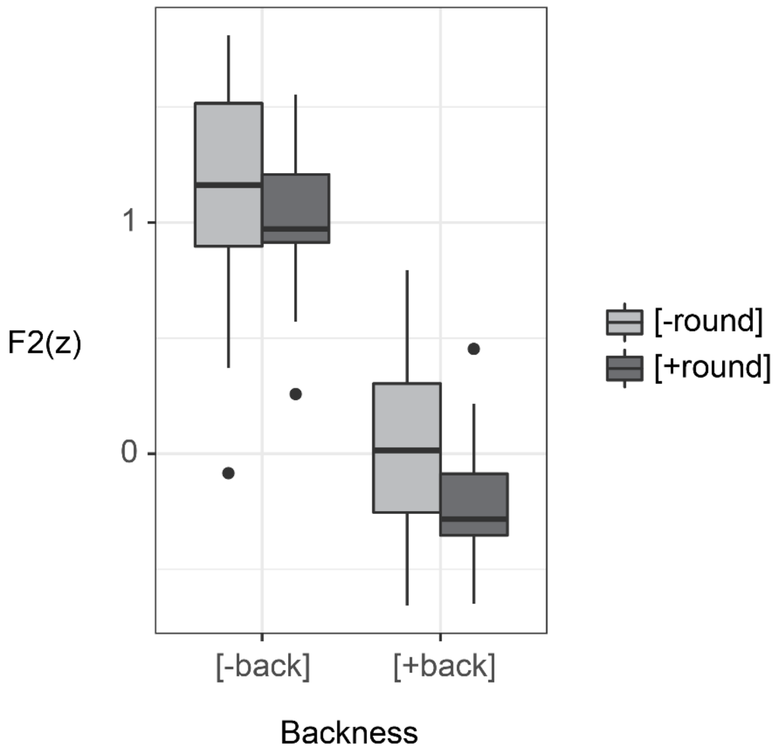

Figure 2.8: F2(z) of suffix low vowels by root backness and rounding in long-distance condition

In summary, this subsection has shown that in the long-distance condition, reportedly transparent vowels actually alternate for backness harmony, with no significant effects of rounding harmony. Like the high vowels, low vowels following the high vowels also alternate for harmony. These results mirror results from the short-distance condition, confirming that high vowels do in fact alternate for harmony in non-initial syllables. Moreover, high vowels in epenthetic, underlying, and raising contexts all alternate for harmony. These results indicate quite strongly that non-initial high vowels exhibits backness alternations, and at least some also exhibit rounding alternations, too.

2.5 Discussion

The findings from the previous section pose intriguing problems for the analysis of harmony in Uyghur, but also provide some clarity on backness harmony in the language. To see the concerns that arise from the data, particularly the distribution of root-internal high unrounded vowels, consult Table 2.3 below. The real issue is how to analyze [i] and [ɯ] within roots. Some lexically-specific patterns above, e.g. the difference in F2 for pairs like /ilim/ 'science' and /ʃilim/ 'paste', suggests that these two surface phones are not simply allophones of a single phoneme. However, the general distribution of these sounds in the lexemes tested does not convincingly support a contrastive relationship between these two sounds. If the high unrounded vowels are contrastive, then Uyghur parallels Moro, exhibiting a relatively subtle contrast that manifests itself both in triggers and targets of harmony. However, if the high unrounded vowels are not contrastive within roots, they would only exhibit phonetic alternations in suffixes. If this is the case, then the language resembles Kinande, possessing a surface inventory that is larger than its underlying inventory. It is clear, though, that high vowels are not transparent to harmony, and although there is a reduction in acoustic contrast across syllables, there is no evidence to support a Hungarian-like distance effect in the language.

Table 2.3: Schema of possible results compared against actual results

	Phonological alternations	Phonetic alternations	Transparency	Distance-delimited pseudo-transparency	Actual results
Initial [i]-[ɯ]	/i/ and /ɯ/ are contrastive	[i] and [ɯ] are allophonic	[i] and [ɯ] are allophonic	[i] and [ɯ] exhibit imperceptible allophonic differences	unclear
Non-initial [i]-[ɯ]	alternate for [back]	phonetically alternate for [back]	do not alternate for [back]	exhibit imperceptible allophonic differences	alternate for [back]
Following vowels	alternate for [back]	alternate for [back]	alternate for [back]	alternate for [back] if number of /i/ < *x*	alternate for [back]
Comparison language	Moro (Ritchart & Rose 2017)	Kinande (Gick et al. 2006)	Wolof (Dye 2015)	Hungarian (Benus & Gafos 2007; Hayes & Londe 2006)	

More generally, the distribution of root-internal high unrounded vowels forces questions like, what is contrast? Are there degrees of contrast? Kiparsky (2015) argues that structural contrast should be divorced from perceptual distinctiveness, producing a four-way typology of sounds, shown below in Table 2.4 (see also Hall 2009, 2013 on intermediate contrasts). If contrast and distinctiveness are not yoked, then two sounds may be distinctive in very perceivable ways without exhibiting actual linguistic contrast. Kiparsky (2015) points to Russian [i] and [ɨ] as examples of such a salient but non-contrastive relationship. On the other hand, some sounds may exhibit structurally different phonological behavior while being characterized by the same acoustic and/or articulatory properties. Kiparsky calls this near-merger, but to distinguish this category from the typical use of that term (e.g. Labov et al. 1991; Yu 2007), I refer to this as abstract contrast. The vowels [ɛ] and [ɔ] in Tutrugbu illustrate this sort of relationship. The front vowel [ɛ] is the surface output of both /ɪ/ and /ɛ/, while [ɔ] is the surface output of /ʊ/ and /ɔ/ (McCollum et al. 2019). Despite their surface neutralization, these sounds still maintain their abstract structural contrast for [high] for both rounding and ATR harmony.

Table 2.4: A typology of contrast and distinctiveness (Kiparsky 2015)

	Contrastive	Non-contrastive
Distinctive	phoneme (English /t/ and /d/)	quasi-phoneme (Russian [i] and [ɨ])
Non-distinctive	abstract contrast (Tutrugbu /ɛ/-/ɪ/ and /ɔ/-/ʊ/)	allophone (English [t] and [tʰ])

Despite the appeal of Kiparsky's proposal, Uyghur [i] and [ɯ] don't fit well into these categories. The acoustic differences between these two sounds are not terribly salient, as attested by the number of grammars that describe only /i/ in the language. However, the behavior of these two surface sounds is not equivalent to the abstract contrasts in Tutrugbu, /ɪ/-/ɛ/ and /ʊ/-/ɔ/, since clear acoustic differences are maintained. One possibility is that the contrast between these two sounds simply isn't categorical, but rather contrast and allophony exist on a continuum. This is the very point argued in Hall (2009; see also Goldsmith 1995). Hall's proposal offers a way to compare degree of contrast between these two vowels, but also for the comparison of contrast across languages.

Another possibility is that both the Uyghur system and other outliers, like Russian [i] and [ɨ], derive their behavior from other factors. For instance, Russian [i] and [ɨ] are represented orthographically, and figure meaningfully in Russian pedagogy. In contrast, Uyghur [i] and [ɯ] are not represented orthographically, save for a brief time in Soviet Kazakhstan (Hahn 1991). Uyghur in both contemporary Xinjiang and Kazakhstan is represented by scripts that do not convey a distinction between the two. If non-linguistic factors may play a role in the psychological reality of a sound pair, then perhaps orthography offers some explanatory power. Additionally, what role does a second language play in the maintenance or loss of contrast? Kazakhstani Uyghurs speak both Uyghur and Kazakh, and the genetic and structural similarities of the two languages may influence one another. As noted above, the backness distinction between the high unrounded vowels in Kazakh is quite robust, and given the lexical similarities between the two, it is plausible that a contrast in Kazakh may help to maintain a contrast in Kazakhstani Uyghur. Such a proposal does, though, presuppose cultural and linguistic affinity that may not actually exist. It is well known that some speech communities enhance their linguistic distinctives in order to separate themselves from a related group, so it is not at all clear that structural similarities would play a contrast-preserving role in this case.

Moreover, root-internal [i] and [ɯ] raise the question, how does one investigate structural contrast from a phonetic point of view? In the clearest of cases, the answer seems simple. One examines the distribution of sounds along some continuum to determine whether their distributions are similar or sufficiently distinct. Pairing such a production-based approach with perception testing should, in many cases, presumably offer a relatively clear picture of the relationship between two sounds. However, other factors, particular lexical factors, are known to play a significant role in production and perception. Issues like neighborhood density (Scarborough & Zellou 2013), contextual predictability (Seyfarth 2014), morphological constituency (Plag et al. 2017), and intra-paradigmatic relationships (Seyfarth et al. 2018) all affect the realization of segments, and if the contrast is relatively subtle, then it could become quite difficult to tease apart the different possible effects inherent within a set of data points. Does contrast really boil down to minimal pairs and semantic differences? If so, then Uyghur [i] and [ɯ] cannot be considered contrastive, since they do not have clear-cut (near-)minimal pairs, although perhaps the slight differences in [ilim]~[ɨlɨm] 'science' and [ʃɨlɨm]~[ʃɯlɯm] 'paste' suggest the possibility of such pairs.

In OT, contrast falls out from constraint rankings rather than restrictions on the underlying inventory (see Lubowicz 2012 for more on contrast in OT). In this framework, some allophony-driving constraints may compel the distributional facts surveyed above. For instance, a highly ranked constraint against [ɯ] preceding postalveolars and palatals, *ɯʃ, could ensure that high unrounded vowels surface as front in this context. Likewise, a constraint banning [i] before the lateral, *il, could account for the lowered F2 of the high unrounded vowels before [l]. With a set of these markedness constraints, one might be able to account for much of the contextual distribution of the high unrounded vowels. However, since harmony produces a reliable alternation between [i] and [ɯ], whatever constraint drives harmony in the language, e.g. AGREE[back], must outrank the markedness constraints governing the distribution of these two sounds generally. As an example, AGREE[back] must outrank *il, since harmony produces [i]-like vowel qualities in raised vowel and suffixal contexts. At the same time, AGREE[back] must be outranked by some constraint on these sounds within roots. Stated differently, whatever contrast exists between these two sounds is most robustly attested not in roots, the typical prominent position in which contrasts should be preserved, but rather in suffixes.

Evidence from Esimbi, a Grassfields Bantu language, indicates that a language may displace contrast (Stallcup 1980; Hyman 1988; Walker 2011; Kaplan 2015). In Esimbi, root height contrasts shift to the initial syllable, which is almost always a prefix. Walker (2011) and Kaplan (2015) argue that the Esimbi pattern falls out from licensing. In essence, the initial syllable licenses height contrasts that are illicit elsewhere in the word. The challenge for a licensing account of Uyghur is the licensing position.

In Uyghur, it appears that contrast is typically displaced to non-initial syllables, including both medial and final syllables. Medial syllables, though, are subject to reduction via vowel raising, so it would be surprising if they are targets for contrast displacement. Additionally, word-finally all high vowel contrasts are neutralized to a very peripheral [i] in the language (see Chapter 4 for more). Thus, a licensing account would have to motivate displacement to non-initial syllables, provided that the licensor is not word-final, which also happens to be where stress falls in the language. Thus, stress placement cannot motivate licensing, since high vowel are neutralized to [i] in absolute word-final (stressed) position. In Esimbi, the licensor is the cross-linguistically prominent initial syllable, but the licensing position is not so obviously definable or prominent in Uyghur.

Other work has revealed a similar pattern among some speakers of Crimean Tatar. For these speakers, /y/ surfaces as [u] in roots unless it is flanked by a coronal, which appears to enhance the frontness of the vowel (see also Flemming 2001, 2003). However, in suffixes, /y/ is produced as [y] regardless of consonantal contexts. These types of patterns defy McCarthy & Prince's (1995) claim that root faithfulness universally outranks affix faithfulness. Problems with McCarthy & Prince's (1995) FAITH-ROOT >> FAITH-AFFIX metaconstraint are evident in languages with dominant-recessive harmonies, or languages in which the intonational or tonal content of affixes overrides that of less peripheral morphemes (e.g. Bennett et al. 2018; Rolle 2018). Affix and root discrepancies like those in Uyghur suggest continued work on contrast, its psychological and cognitive foundations, as well as its differential realization in spoken language.

In spite of the various challenges posed by [i] and [ɯ], their alternation does provide a much simpler analysis of the harmony pattern than those presented in Lindblad (1990) and Vaux (2000). In fact, existing analyses of Uyghur have recently been critiqued for their computational complexity. Mayer & Major (2018) contend that if both consonantal and vocalic features define the distribution of the high unrounded vowels and the harmony pattern more generally, then harmony in Uyghur is more complex than tier-based strictly local patterns (e.g. consonant harmony), contradicting claims made in recent work on the computational complexity of phonology (e.g. Heinz 2018). The fact that high vowels all alternate for [back] simplifies the analysis of harmony in Uyghur. There is no need for theoretical devices like absolute neutralization rules, or for a less constrained harmony driver, like SPREAD (Finley 2008). So, no derivational ordering or transparency-producing constraints are necessary, making the analysis much more straightforward- the constraint driving harmony outranks the relevant markedness and faithfulness constraints, compelling backness alternations in suffixes. Typologically, this result falls nicely in line with previous work investigating putative transparency in harmony (e.g. Gick et al. 2006; Benus & Gafos 2007; Ritchart & Rose 2017), showing that harmony applies more locally than previous accounts suggest. So, while this investigation of transparency in Uyghur has not yielded the most sanitized view of contrast, it has corroborated a number of studies on other languages that support a decidedly local conception of harmony.

Recall the questions driving this study. One, do [i] and [ɯ] contrast in Uyghur? Two, do [i] and [ɯ] alternate for harmony in non-initial positions? Three, does Uyghur exhibit transparency, and if so, is transparency constrained by distance? The answer to the first is unclear. It is plausible that these two sounds exhibit some intermediate contrast (Hall 2009) between full contrast and allophony. As for the second question, the answer is clear – harmony dictates the realization of [i] and [ɯ], as well as all other vowels in the language. Finally, since Uyghur does not exhibit transparency, the issue of distance is moot. From the results in the long-distance condition, though, it is clear that there is a relationship between vowel F2 and distance from the trigger, which is discussed at length in Chapter 4.

This chapter has demonstrated that the high unrounded vowels alternate for harmony in Uyghur, contrary to a number of previous descriptions. As a result, the analysis of Uyghur is far simpler than previously thought. The next chapter continues the investigation of locality to determine whether or not consonants also undergo harmony alternations in Kyrgyz, Kazakh, and Uyghur, and the implications of consonantal participation in harmony for various theories of locality.

Chapter 3: Strict locality in harmony

3.1 Introduction

The previous chapter observed that, despite a number of claims that vowel harmony in Uyghur skips vowels, harmony is local, triggering alternations on all vowels within the word. As discussed in the previous chapter, locality offers significant formal benefits. Problematically, there is no single definition of what the term 'local' actually means in the literature. At one end of the spectrum, Nevins (2010) proposes a harmony-driving search function that is relativized to ignore certain classes of vowels, which is capable of producing long-distance transparency effects (see also Calabrese 1995; Vaux 2000). At the other end of the spectrum, a body of work (e.g. Walker 1998; Gafos 1999; Ní Chiosáin & Padgett 2001) contends that harmony is 'strictly local,' arguing that all segments either participate in or block harmony. Under this proposal, all segments either undergo harmony or block harmony, ruling out all transparency. To continue the investigation of locality in Turkic, this chapter examines the effects of harmony on intervening consonants to better understand how strict locality is in the four languages under study. Specifically, this chapter examines the realization of the sibilants as well as the lateral in each language to determine if these segments are realized in accordance with backness and rounding harmony, as is claimed by Dzhunisbekov (1980, 1991) and Abuov (1994).

This chapter is organized as follows. Section 3.2 discusses two proposals concerning locality, trans-consonantal locality and strict locality, in addition to a minor variant of the strict locality account proposed in Smith (2018). Section 3.3 details the methods used to collect, measure, and analyze the acoustic data collected, and Section 3.4 presents results from each language. Section 3.5 relates the study's findings back to the larger question of locality, and presents a brief Optimality Theoretic analysis of locality in Turkic; Section 3.6 concludes the chapter.

3.2 Locality and the phonology-phonetics interface

3.2.1 Locality

3.2.1.1 Trans-consonantal locality

The prevailing conception of vowel harmony is largely agnostic to the role of consonants. Unless consonants interact directly with the pattern, like the palatalized lateral in Turkish (e.g. Clements & Sezer 1982), most work assumes very little about consonantal participation in harmony. Deriving largely from autosegmental models of phonology, most work assumes that vowels may be projected onto their own tier, and locality holds only within the tier (Goldsmith 1976; Clements 1981; Odden 1991). Various instantiations of this general claim exist throughout the literature: projections, feature-bearing units, and formal language-theoretic tiers (e.g. Halle & Vergnaud 1981; Clements & Sezer 1982; Archangeli & Pulleyblank 1994; Krämer 2003; Hayes & Wilson 2008; Heinz 2018). The various proposals are remarkably consistent in their distinction between vowel and consonant, relevant and irrelevant for harmony (though see Clements & Hume 1995 for a feature-geometric approach to consonant-vowel interactions). Throughout the chapter I will refer to this view as trans-consonantal locality. Trans-consonantal locality is exemplified below in (32). In (32a), the first-syllable vowel spreads its value for [F] to the second-syllable vowel. This type of syllable-adjacent interaction is grammatical, while non-local interactions, like (32b) are forbidden. In (32b), the initial-syllable vowel spreads its value for [F] directly to the third-syllable vowel without harmonizing the intervening vowel. This type of transparent interaction is ruled out under trans-consonantal locality.

(32) Trans-consonantal locality

Grammatical		Ungrammatical
a. [F]		b. [F]
V V V	[vowel tier]	V V V
C V C V C V C	[segmental tier]	C V C V C V C

When it comes to the question of intervening consonants and their participation in harmony, the typically tacit assumption is that consonants simply don't matter. In Calabrese's (1995) terms, they typically are not specified for the harmonic feature, and as a result are not visible to harmony (see also Vaux 2000; Nevins & Vaux 2006; Nevins 2010). This may very well be the state of affairs in many languages with vowel harmony. In these cases, if intervening consonants are ignored by the harmony pattern, then their realization is effectively a phonetic issue.

2.1.2 Strict locality

As discussed in the previous chapter, a number of languages with reportedly transparent segments have been shown to be subject to a stricter definition of locality, like Hungarian, Turkish, Kinande, and Moro (Boyce 1990; Gick et al. 2006; Benus & Gafos 2007; Ritchart & Rose 2017). In Hungarian, Kinande, and Moro, vowels that were thought to be transparent were subsequently shown to exhibit harmonic alternations, and in the case of Moro, alternating vowels were shown to be contrastive. If vowels that were thought to be transparent have since been shown to actually participate in harmony, then the possibility that consonants may exhibit similar patterns of alternation. Some evidence for this comes from Boyce (1990), which finds that rounding harmony affects both vowels and intervening consonants, in Turkish, as lip rounding persists across all segments within the domain of harmony. As a result, it is likely that consonants serve as targets for harmony in Turkish rounding harmony. Kaun (1995, 2004) formalizes this as a universal preference for 'gestural uniformity', a prolonged extension of a single articulatory gesture across a given domain. This notion of gestural uniformity is in many ways equivalent to what has been labelled 'strict locality' in works like Walker (1998), Gafos (1999), and Ní Chiosáin & Padgett (2001). These authors argue that all segments either participate in or block harmony. In effect, all consonants and vowels become targets for harmony, shown in (33). In (33a), the initial-syllable vowel spreads its value for [F] to every vowel and consonant within the word, obeying this stricter definition of locality. In (33b), however, harmony skips all consonants, and is thus ungrammatical under strict locality.

(33) Strict locality

	Grammatical		Ungrammatical
a.	[F]	b.	[F]
	C V C V C V C		C V C V C V C

Strict locality is formalized in several different ways. Walker (1998) and Ní Chiosáin & Padgett (2001) build locality into a theory of Gen that eliminates all candidates in which harmony is non-local. Gafos (1999), and more recently Smith (2018), encode strict locality as a result of a gestural representational system. Gafos (1999) and Smith (2018) argue that harmony is the persistence of a single gesture across the entire domain of harmony. In both the gestural and featural representations employed, the overarching claim is the same: harmony targets all segments within the word. While the trans-consonantal locality account leaves phonetics to determine the realization of intervening consonants, the strict locality account makes a much stronger claim. The strict locality claim not only obligates phonological harmony to target all segments, it further obligates the phonetic module of the grammar to faithfully implement the commands from phonology. The prediction is simple and readily falsifiable- if harmony fails to trigger an alternation on some segment or class of segments, then harmony is not strictly local.

To see the restrictiveness of the strictly local account, consider a suffix with the shape CVC, like the nominalizing suffix /-lik/ in Kyrgyz. If we assume for the moment, in conformity with the strictly local analysis that there are two possible surface outputs for backness harmony for each segment in the syllable,, [l]-[ɫ], [i]-[ɯ], and [k]-[q]. Given that there are three binary oppositions here, there are eight ($2^3 = 8$) possible syllables derivable from /-lik/. The restrictiveness of each analysis is evident in the number of forms that are grammatical under each analysis. In Table 3.1, the set of grammatical strings attached to a [±back] root are compared. Of the eight possible strings, the trans-consonantal locality allows four for each value for [back], permitting all strings in which the vowel's backness is consistent with the backness of the root. Trans-consonantal locality does not demand that consonants surface in conformity with adjacent vowels. As a result, it is possible that these sounds freely vary with one another. In constraint terms, neither [ʤɑqʃɯ-lɯq] "goodness" nor the possible [ʤɑqʃɯ-lɯk] doesn't violate a trans-consonantal harmony driver, since the constraint does not pay attention to the consonant-vowel sequence [lɯ] and the vowel-consonant sequence [ɯk]. Vowels are projected onto a separate tier, and the segmental tier is ignored. In contrast, the strict locality analysis permits only a single output for each backness value, a string that contains only syllables that all agree with the root for [back].

Table 3.1: Comparison of different locality analyses

Context	Analysis	
	Trans-consonantal locality	Strict locality
[-back] stem e.g. izgi 'good'	lik, liq, ɫik, ɫiq	lik
[+back] stem e.g. ʤɑqʃɯ 'good'	ɫɯq, ɫɯk, lɯq, lɯk	ɫɯq

Evidence to-date suggests that backness and rounding harmony in Turkic are strictly local. Dzhunisbekov (1980, 1991) argues that each consonant undergoes four-way alternations for backness and rounding in Kazakh. He goes on to contend that all Turkic languages are governed by this same pattern of harmony, making strictly local harmony a distinguishing feature of the language family (see also Abuov 1994). Limited phonetic evidence is marshalled in Ní Chiosáin & Padgett (2001) and McCollum (2015) to support the putative locality of rounding harmony, but no systematic phonetic work has been done on consonantal participation in backness harmony.

With these differences in mind, this chapter attempts to determine whether the realization of consonants is more consistent with the trans-consonantal locality account or the strict locality account. Since the trans-consonantal locality account gives phonetics control of intervening consonants, it is key to understand the role phonetic coarticulation may play in the production of consonants. This topic is discussed below, including Smith's (2018) nuanced perspective on locality.

3.2.2 The phonology-phonetics interface

If the production of consonants in a language with vowel harmony is determined by phonetic forces, what possibilities would one predict for these interveners? Three possibilities seem most worthy of discussion: assimilation, no assimilation, and restricted assimilation. I will briefly discuss these in order.

A large body of literature has shown that both vowels and consonants are affected by flanking segments (e.g. Hardcastle & Hewlett 2006; Farnetani & Recasens 1997). Coarticulatory forces have even been shown to persist across a number of segments. For instance, Öhman (1966) shows that vowels may affect one another across an intervening consonant (see also Magen 1997). In much the same way, vowel articulations are known to affect intervening consonants (e.g. Recasens 1985, 1999). Given the extent of vowel coarticulatory forces and cross-linguistic reports of vowel-to-consonant coarticulation, one possibility is that all consonants undergo phonetic alternations consistent with the phonological harmony pattern due to coarticulation. For instance, given a sequence like /ulu/ versus /ili/, one might expect the lateral to undergo backing in the context of /u/ as well as potential lip rounding (Daniloff & Moll 1968; Boyce 1990). If vowel-to-consonant coarticulation occurs throughout the harmonic domain, the harmonic feature will be present to varying degrees on all segments within the word. Therefore, if trans-consonantal locality is implemented in a phonetics with widespread vowel-to-consonant coarticulation, the resultant analysis is effectively equivalent to the strictly local analysis in Turkic – all segments bear the harmonic feature.[5] The two analyses differ in where they place the locus of explanation. Trans-consonantal locality plus widespread coarticulation relies on a relatively independent conception of phonology and phonetics, whereas the strictly local account assumes that phonological considerations dictate the production of consonants. From an empirical standpoint, though, trans-consonantal locality plus widespread coarticulation is indistinguishable from strict locality, and the two are treated as strict locality below.

Despite the prevalence of coarticulation in connected speech, some research has found that consonants may resist coarticulation for at least two reasons- perceptibility and articulatory constraints. In some cases, the extension of a particular gesture into an adjacent segment may obscure some crucial aspect of one segment's identity, and in these cases coarticulation may be minimized. Butcher (1999) reports that anticipatory nasalization is only minimal in Australian languages, in contrast to languages like English, where velum lowering begins well before the onset of the nasal consonant. Moreover, Stoakes et al. 2019) argues that limited anticipatory nasal coarticulation maintains a perceptually weak place of articulation contrast among the triggering nasals in Bininj Kunwok. One can generalize from this finding that nasal coarticulation is limited in Australian languages due to language- and segment-specific

[5] While widespread coarticulation plus trans-consonantal harmony may be difficult to distinguish from strict locality in Turkic, one can imagine scenarios where they make divergent predictions. Specifically, consider a language with some progressive vowel harmony but with regressive vowel-consonant coarticulation. The opposing directionality of these patterns would likely produce different predictions.

perceptibility concerns. In this case, perception of the triggering segment militates against coarticulation, but it is also possible that the target segment may also resist coarticulation for perceptual reasons. This reasoning is leveraged in Manuel (1990) to explain why languages with larger and more crowded vowel inventories exhibit smaller patterns of consonant-to-vowel coarticulation. Manuel argues that vowel contrasts are more likely to be blurred under coarticulation, and so potential targets may resist physiologically motivated coarticulation in order to maintain phonological contrasts.

When perception-based explanations of limited vowel-to-consonant coarticulation are contextualized to Turkic, though, it is unclear what contrasts might be preserved by resisting coarticulation. Turkic does not typically employ a secondary palatalization or labialization contrast (cf. Berta 1998; Csató 1999; Kavitskaya 2010 on palatalization in West Kipchak languages). In addition, very few minor place of articulation differences are exploited. There are no contrastive bilabial-labiodental, dental-alveolar, or alveolar-postalveolar contrasts that are likely to be lost under coarticulation from backness or rounding harmony. In fact, the only consonants that are reported to alternate in most Turkic languages are the velars and uvulars, which are not contrastive. These two classes of sounds have developed a contrast only in languages with extensive borrowing, like Kazakh or Uyghur, but have maintained an allophonic relationship in much of the language family. Thus, although perceptibility factors likely limit coarticulation in some languages, there is no clear reason to expect perception to significantly limit consonantal participation in Turkic vowel harmony.

Despite the absence of perceptual factors motivating consonantal resistance to harmony, there is some articulatory reason to consider consonants immune to harmony. As argued in Recasens (1985) and Recasens et al. (1997), among other works, vowel-to-consonant coarticulation can be predicted based on the articulatory properties of the target consonant. Some consonants involve a higher degree of articulatory control, e.g. palatal consonants, while others, e.g. labials, require less stringent articulatory configurations. Findings in Recasens (1989) further suggest that carryover (left-to-right) coarticulatory patterns may be affected by antagonistic consonantal gestures more readily than anticipatory (right-to-left) coarticulation. Since harmony in Turkic is progressive (left-to-right), the role that consonantal gestures may play in their ability to resist backness and rounding coarticulation may be more significant. Thus, it is possible that consonants resist coarticulatory effects in Turkic due to some intrinsic properties of the articulations they require as well as the directionality of spreading.

While intervening consonants may (relatively) uniformly undergo or resist coarticulation from flanking vowels, it is also plausible that vowel-to-consonant coarticulation may be mediated by consonant-specific gestural factors. Consonantal participation in harmony may thus not be uniform, but depend on the specific articulatory properties of each consonant in question. This restricted coarticulatory pattern lines up very neatly with Recasens' (1985, 1987, 1989) coarticulatory resistance argument. As noted just above, consonants that require more specific articulatory configurations are predicted to resist coarticulation more than other consonants. Recasens (1985) suggests that consonants involving more tongue body contact as well as tongue body-tongue tip coupling are the most resistant to effects from flanking segments. Based on his findings, both the sibilants, particularly the postalveolar sibilants, and the lateral should be less prone to coarticulatory effects than other consonants. Tabain (2001) corroborates Recasens' claim for the sibilants, finding that the sibilants in Australian English are less prone to coarticulatory variation than other consonant. She argues that this coarticulatory resistance falls out from their specific articulatory configuration. The sibilants involve a high degree of tongue body constriction that must also channel the airstream toward the bottom teeth. This complexity of articulation is thus predicted to make them less affected by neighboring segments (cf. Gafos 1999). Similarly, the lateral requires controlling both the tongue tip for apical contact and the constriction of the dorsum to produce lateralization of the airstream. The lateral involves coupling the tongue tip and dorsum, which also is the lingual region most relevant for vowel articulations. These facts suggest that the lateral, like the sibilants involves a relatively complex set of articulatory gestures. A segment or class of segment's unique articulatory properties may therefore distinguish them from other consonants in the language, producing a more restricted patterns of coarticulation or consonantal participation in harmony.

A third analytical possibility is thus derivable from coarticulatory resistance, what I will call violable strict locality. Under this model of locality, consonants surface in accordance with the harmony pattern so long as they are not produced with an articulatory configuration that is antagonistic to the harmonic feature. This idea is discussed extensively in Smith (2018), although she focuses on vowel articulations and their more obvious role in harmony. She argues that the spreading gesture extends throughout transparent segments, but that its effect is masked due to other articulatory constraints. If coarticulatory resistance, as in Recasens' work, may hold the same status as the vowel-specific articulatory factors that shape the typology discussed in Smith (2018), articulatory antagonism provides a more nuanced framework within which to understand consonantal participation in harmony. Instead of uniform consonantal participation or resistance to harmony, consonants may exert varying degrees of resistance, in accordance with their own articulatory specifications.

The three analyses evaluated and discussed throughout the rest of the chapter are thus trans-consonantal locality, violable strict locality, and strict locality. The next section details the specific predictions of these analyses and how they relate to the sibilants and lateral in Kyrgyz, Kazakh, Uyghur, and Uzbek.

3.3 Methods

3.3.1 Research questions

To concretize the larger research question of consonantal participation in harmony, this chapter examines the realization of two classes of segments, the sibilants and the lateral. The larger research question can thus be broken down to consider the realization of each of these classes of segments for both backness and rounding harmony. These questions are explicitly formulated in (34).

(34) Research questions

1. Do the sibilants (i.e. /s z ʃ ʒ ʧ ʤ) alternate based on root backness?
2. Do the sibilants (i.e. /s z ʃ ʒ ʧ ʤ) alternate based on root rounding?
3. Does the lateral alternate based on root backness?
4. Does the lateral alternate based on root rounding?

Predictions from each of the analyses described above are listed in Table 3.2. For each of the questions in (34), trans-consonantal locality is agnostic to phonetic alternations among the sibilants and lateral. The trans-consonantal locality analysis leaves consonantal realization up to the phonetics of each language, considering only vowels for the harmony pattern. The violable strict locality account predicts that the sibilants and the lateral should alternate for rounding harmony since the lips are independent of the tongue body. Similar to the trans-consonantal locality account, the violable strict locality account makes no explicit predictions regarding backness harmony. Given the large differences between the European languages discussed within work on coarticulatory resistance and the Turkic languages under study here, it is unclear if the trends in Recasens (1985, 1987, 1989) and Tabain (2001) will translate directly into Turkic. Unlike the first two analyses, the strict locality account predicts that the sibilants and lateral should alternate for harmony.

Table 3.2: Predictions concerning the four research questions in (34)

<table>
<tr><th></th><th>Alternation</th><th>Trans-consonantal locality</th><th>Violable strict locality</th><th>Strict locality</th></tr>
<tr><td rowspan="2">Sibilants</td><td>[back]</td><td rowspan="4">Perhaps</td><td>Perhaps</td><td rowspan="4">Yes</td></tr>
<tr><td>[round]</td><td>Yes</td></tr>
<tr><td rowspan="2">Lateral</td><td>[back]</td><td>Perhaps</td></tr>
<tr><td>[round]</td><td>Yes</td></tr>
</table>

While Kyrgyz, Kazakh, and Uyghur demonstrate consistent vowel harmony patterns, vowel alternations are largely absent in contemporary Uzbek. Limited rounding harmony is still operational in the Osh dialect (the dialect under study), but backness alternations are unattested. Therefore, Uzbek provides a control against which to compare the other three languages in at least two respects. First, what alternations are statistically significant? Second, how do the magnitudes of significant alternations compare? A priori, Uzbek should exhibit fewer and smaller backness and rounding alternations on consonants, since harmony is no longer at work in the language. It should be noted that the absence of phonological harmony does not preclude the possibility for continued phonetic effects on the consonants. Nonetheless, I predict that the number of statistically significant alternations should be fewer and that their magnitude should be smaller than comparable alternations in the other three languages.

3.3.2 Stimuli

The elicited data set derives from a word list composed of monosyllabic and disyllabic roots, presented in Table 3.3. All monosyllabic roots ended in a sibilant or a liquid (e.g. Kyrgyz /bɑʃ/ 'head' and /bɑl/ 'honey'). Disyllabic stimuli were composed of two vowels that agree in height with an intervening sibilant or liquid (e.g. Kyrgyz /bɑltɑ/ 'axe' and /ɯrɯm/ 'superstition'). Stimuli were presented pictorially on a laptop computer, prompting each participant's response. If a participant was unable to guess the appropriate word, I provided synonyms in the target language or a Russian translation of the appropriate word. All words were produced in isolation to the pictorial prompts. Each word was prompted only once, but if a speaker repeated a word, all repetitions were included in the analysis. Note that there are a number of gaps in the wordlist in Table 3.3. Some of these gaps are due to inventory gaps, and these are notated with "N/A". Other gaps are due to the restricted distribution of a given phoneme, like Kyrgyz /ʧ/ and Kazakh /ʃ/. Since the affricates are composed of a stop + frication sequence, the analysis groups the two sounds together, which removes some of the gaps below. See Chapter 4 for more on the word list.

Table 3.3: Stimulus lexemes for all four languages

		Kyrgyz		Uzbek		Kazakh		Uyghur	
		[-bk]	[+bk]	[-bk]	[+bk]	[-bk]	[+bk]	[-bk]	[+bk]
/s/	[-rd]	es		es	ɑsɑl, bɑrs	ies, tɪs, bæs, sældie	bɑs, qɯs	æs	
	[+rd]		dos		dos	tʏs	dos		
/z/	[-rd]	tezek							
	[+rd]	køz, ʤyzym	muz	køz, jyzym	muz	køz, ʒʏzʏm	mʊz	yzym	muz
/ʃ/	[-rd]	tiʃ	bɑʃ, qɯʃ	tiʃ, qiʃ	dɑʃt			ʧiʃ	bɑʃ
	[+rd]	tøʃøk, tyʃ		tøʃɑk, tyʃ	boʃ			tøʃæk, ʧyʃ	
/ʒ/	[-rd]	N/A		N/A			ʒɯl	N/A	
	[+rd]					ʒʏzʏm	ʒol		
/ʧ/	[-rd]			ʧiʃ		N/A		ʧiʃ	ʧɑlmɑ
	[+rd]								
/ʤ/	[-rd]		ʤɯl			N/A			
	[+rd]	ʤyzym	ʤol						
/l/	[-rd]	bel, selde, kelte, pil, ilim	bɑl, bɑltɑ, ʤɯl	bel, selle, ilim	ɑsɑl	biel, sældie, ɪlɪm	bɑl, bɑltɑ	bæl, sællæ	bɑl, pɑltɑ
	[+rd]	køl, gyl	ʤol, moldo, qul	køl	jol, boltɑ, moldɑ	køl	ʒol, moldɑ	køl	jol, moldɑ

In addition to these lexemes, the lateral occurs in some allomorphs of the plural suffix for all four languages, and was included for analysis. In Uyghur and Uzbek, the lateral is the initial segment of the plural (e.g. Uyghur [lɑr]-[lær]) regardless of preceding segment, but in Kazakh and Kyrgyz, the initial segment of the plural varies according to the stem-final segment, and only vowel-final stems trigger the lateral-initial allomorph of this suffix (see Gouskova 2004 for an analysis).

3.3.3 Participants

Participants were recruited from existing relationship networks in each of the field sites: Taldykorgan and Chunja, Kazakhstan; and Bishkek and Osh, Kyrgyzstan. In Taldykorgan, nine Kazakh speakers (seven females, mean age: 33.4 years, range: 19-49 years), and in Chunja, nine Uyghur speakers (five females, mean age: 44.4 years, range: 19-63 years) participated in the study. In Bishkek, thirteen Kyrgyz speakers (eleven females, mean age: 35.0 years, range: 18-57 years), and in Osh, nine Uzbek speakers (four females, mean age: 28.8 years, range:19-45 years) participated in the study. All participants reported native fluency in the target language. Many participants reported additional fluency in at least one neighboring Turkic language, as well as Russian.

3.3.4 Measurement and statistical analysis

Sessions were conducted in a quiet room. Sound files were recorded to a Marantz PMD 661 MKII digital recorder, using a Shure SM-10A head-mounted microphone. All sound files were segmented in Praat (Boersma & Weenink 2015) . The beginning and end of each sibilant was determined to correspond to the beginning and end of high-frequency frication noise in the spectrogram. For the affricates, the stop portion was ignored and only the frication release was considered. Abrupt shifts in energy from flanking vowels along with changes in formant structure were used to segment the lateral from adjacent vowels. In some cases, the lateral immediately preceded an obstruent, and in these instances, the end of the second formant was aligned to the end of the lateral.

For statistical analysis, the alternation of the sibilants and lateral respectively are operationalized as differences in center of gravity and the second formant (F2), respectively. It has been shown that sibilant place of articulation is cued by the distribution of energy in their spectra (e.g. Jongman et al. 2000). For the sibilants, this translates into a difference in center of gravity, with more posterior sibilants exhibiting lower center of gravity. Additionally, Ní Chiosáin & Padgett (2001) and McCollum (2015) demonstrate that some fricatives show parallel differences in center of gravity based on root roundness. Consider the sibilants in Figure 3.1 below. In the two panels, waveforms and spectrograms for the Kazakh words /tɪs/ 'tooth' and /tʏs/ 'dream' are shown. Observe the distribution of energy during the two fricatives. More intense energy is distributed lower in the spectrum in /tʏs/ due to preceding vowel rounding. Lip rounding extends the distance from the point of constriction to the end of the vocal tract, depressing spectral center of gravity. Thus, differences in place of articulation and lip rounding result in the same acoustic consequences, lowered fricative center of gravity.

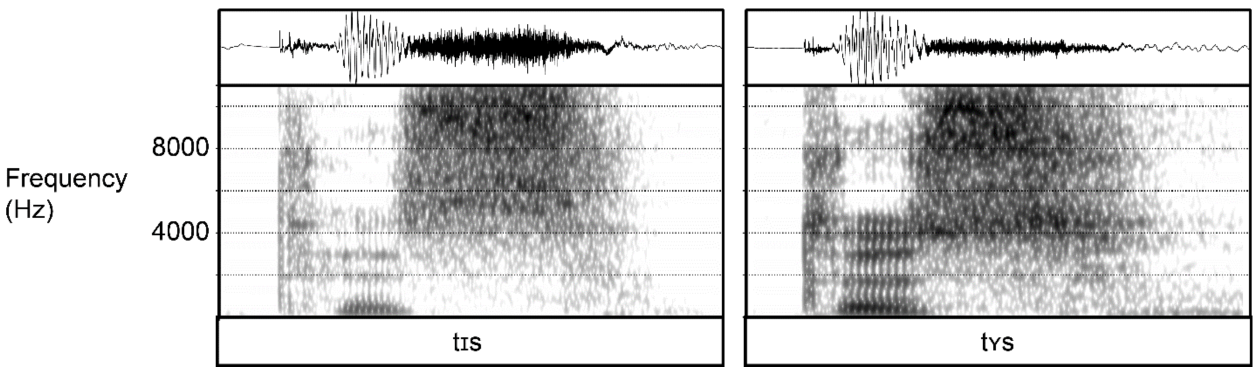

Figure 3.1: Waveforms and spectrograms for /tɪs/ 'tooth' (left) and /tʏs/ 'dream' from a female Kazakh speaker

The same difference in center of gravity is also visible in spectral slices from a 15 millisecond window at the midpoint of the sibilant in /tɪs/ and /tʏs/, shown below in Figure 3.2. In the right panel there is a peak in spectral energy at around 4,000 Hz that is notably absent from the left panel. The increase in energy from around 1,000 Hz to around 5,000 Hz after the round vowel in /tʏs/ results in a lower center of gravity. Center of gravity for the 15 msec. window sampled is 8113 Hz for /tɪs/ and 6603 Hz for /tʏs/.

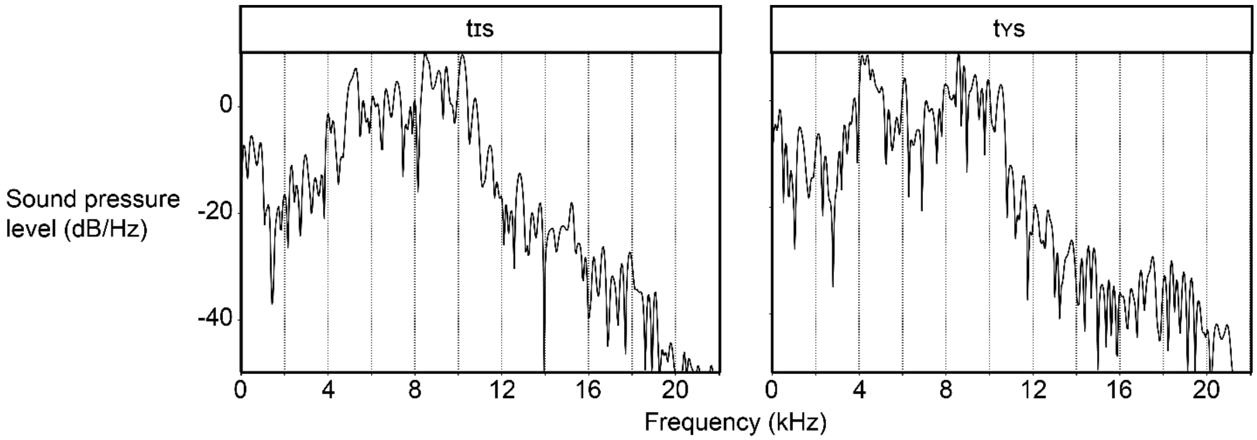

Figure 3.2: Spectral slices from 15 msec. windows at the midpoint of the sibilants in /tɪs/ 'tooth' (left) and /tʏs/ 'dream' from a female Kazakh speaker

As for the lateral, clear [l] and dark (or velarized) [ɫ] differ in the distribution of acoustic energy, with dark [ɫ] having a lower second formant (F2; Recasens & Espinosa 2005; Carter & Local 2007). If the lateral differs in tongue body position according to backness, then similar changes in the second formant should occur. Observe the difference in two tokens of /l/ in Kazakh, shown in Figure 3.3. After the back vowel /ɯ/, F2 of the lateral is very low in the left panel, around 1000 Hz. However, in the right panel, F2 of the lateral is approximately 2000 Hz after the front vowel /ɪ/. Also, since lip rounding decreases F2 generally, F2 of the lateral in round vowel contexts should be lower than in unrounded contexts.

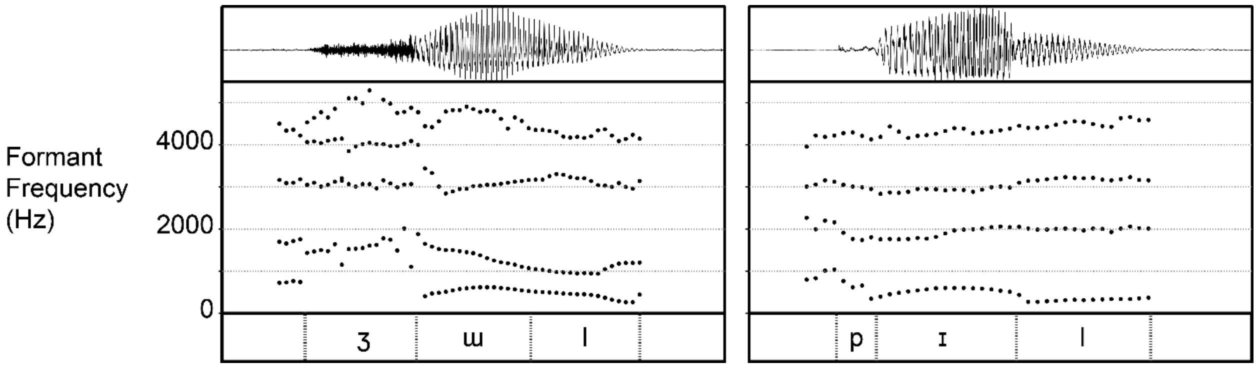

Figure 3.3: Waveforms and spectrograms for /ʒɯl/ 'year' (left) and /pɪl/ 'elephant' from a female Kazakh speaker

Therefore, the four research questions from (34) can be restated in testable, acoustic terms, as in (35).

(35) Research questions

1. Do the sibilants (i.e. /s z ʃ ʒ ʧ ʤ) differ in center of gravity based on root backness?
2. Do the sibilants (i.e. /s z ʃ ʒ ʧ ʤ) differ in center of gravity based on root rounding?
3. Does the lateral differ in F2 based on root backness?
4. Does the lateral differ in F2 based on root rounding?

After segmentation, center of gravity was measured over a 15 millisecond window at the midpoint of each sibilant. For the affricates /ʧ ʤ/, the midpoint of the fricative release rather than the midpoint of the entire affricate was used for the analysis. For the lateral, the second formant was measured at the midpoint of each token. The data were analyzed in R (R Core Team 2017), using the lme4 package (Bates et al. 2015).

Outliers were then inspected for measurement errors. To facilitate across-speaker comparisons, F2 data were z-score normalized (Lobanov 1971). Four tokens of each vowel were taken from monosyllabic words from each speaker's data. If four tokens of a given vowel were not present in monosyllables, then the remaining tokens were taken from the initial syllable of disyllabic words. Thus, the normalization compares the lateral against the normalized vowel space for each speaker.

A mixed effect linear regression was used to predict center of gravity based on root backness and rounding as well as place of articulation (alveolar or post-alveolar), The model for the laterals predicted F2 based on root backness and rounding. In Uzbek, since vowel harmony is almost entirely lost in Uzbek, preceding vowel backness and rounding were used instead of root backness and rounding to predict F2 of the lateral. In both Uyghur and Uzbek, the lateral of the plural suffix may attach to stems ending in a lateral, creating a derived geminate. In these cases, the geminate was treated as root-internal. In the plots below for root-internal and suffixal laterals, these tokens are grouped together with root-internal forms. Also, since in these derived geminates any transition from the root-internal to suffixal lateral is possible, the F2 at the midpoint of the entire geminate was measured for the analysis. Both the regression models for the sibilants and lateral included by-speaker random slopes for backness and rounding. In all cases below, statistical significance was assessed using model comparison. Specifically, to determine whether a given effect (or interaction) is significant, a chi-squared test was used to compare the full model with another model that is identical to the full model except the relevant effect is removed. By removing only the effect under consideration, this test assesses the contribution of the relevant effect.

3.4 Results

3.4.1 Sibilants

3.4.1.1 Kyrgyz

In Kyrgyz, backness, rounding, and place of articulation are all significant predictors of sibilant center of gravity. Back vowels [$\beta = -625$, $\chi^2(1) = 26.69$, $p < .001$] and round vowels [$\beta = -1000$, $\chi^2(1) = 31.96$, $p < .001$] both depress center of gravity, in line with predictions from previous work. Additionally, posterior sibilants exhibit significantly lower center of gravity [$\beta = -2257$, $\chi^2(1) = 1102.9$, $p < .001$]. These results, which are shown in Figure 3.4, are consistent with claims in Ní Chiosáin and Padgett (2001) and McCollum (2015), suggesting that sibilants in Kyrgyz exhibit alternations based on both backness and rounding. A slight difference based on backness and rounding is evident in the upper right panel, in /ʃ/. Differences in rounding are also observable in the lower two panels, in /z/ and /ʤ/. In summary, Kyrgyz sibilants exhibit lower center of gravity after back vowels and round vowels, which supports the claim that the sibilants are produced with more posterior constrictions in back vowel contexts, and extended lip rounding throughout the consonant gesture. (see McCollum 2015 for some video evidence for Kazakh). The size of the rounding effect is 60% larger than the backness effect (1000 Hz vs. 625 Hz), but both are still relatively small. The perceptibility of these alternations is discussed below.

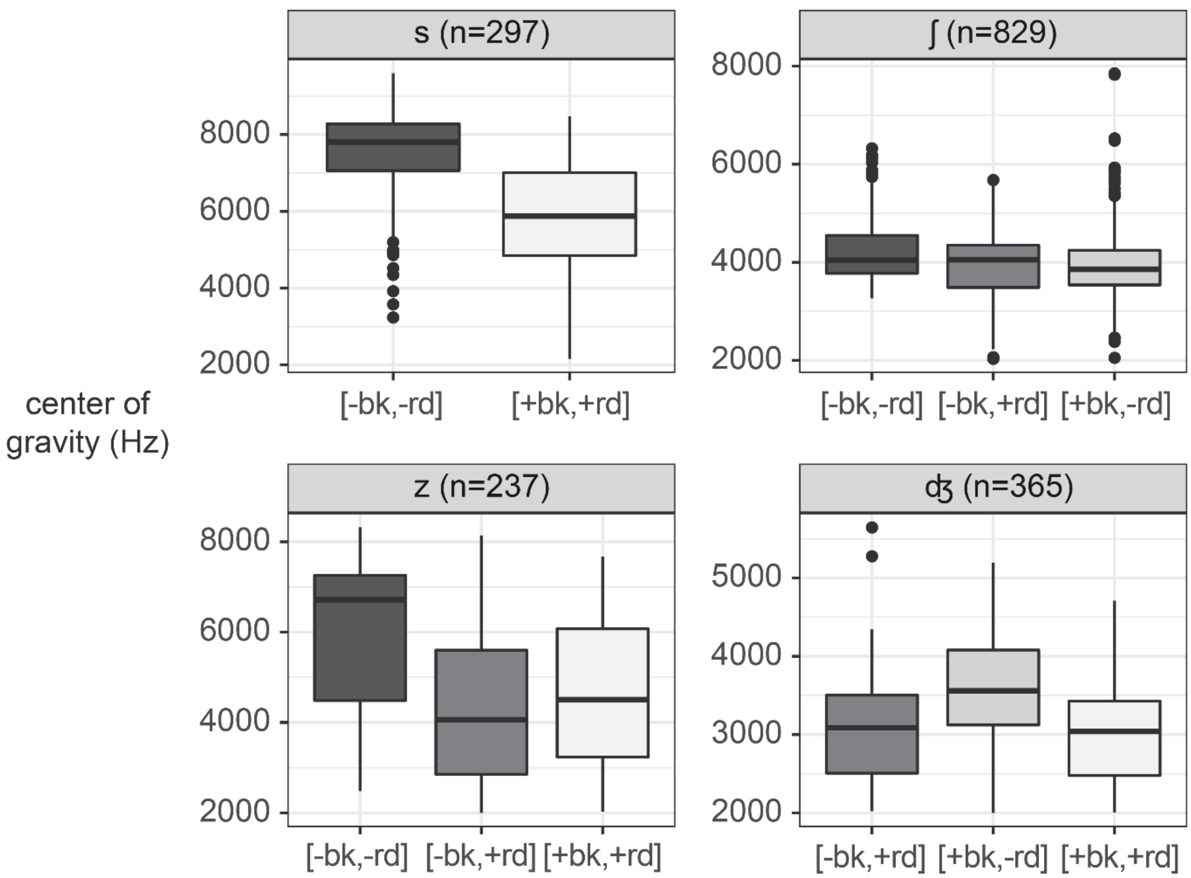

Figure 3.4: Center of gravity in Hz by backness and rounding for the Kyrgyz sibilants (frication portions of the affricates /ʤ/ are grouped with /ʒ/)

3.4.1.2 Kazakh

Kazakh sibilants are produced with significantly lower center of gravity in round vowel contexts [β= -1778, $\chi^2(1)$= 33.55, p < .001]. In contrast, Kazakh sibilants exhibit no significant phonetic differences based on vowel backness [β= -81, $\chi^2(1)$= 0.80, p = .37]. In other words, vowel rounding affects sibilant center of gravity but vowel backness does not. Additionally, more posterior sibilants (post-alveolar) are produced with significantly lower center of gravity [β= -2858, $\chi^2(1)$= 464.93, p < .001]. These effects can be seen in Figure 3.5 below. All four combinations of [back] and [round] were produced with /s/, with both [+round] contexts producing significantly lower center of gravity. In contrast, backness has no effect on /s/. Like /s/, backness has no effect on the center of gravity of /z/ either. The low center of gravity for all /z/ tokens may derive from their rounding context. All tokens of /z/ occurred after round vowels, although this different is much larger than the difference between round and unrounded contexts for /s/. Finally, there are smaller differences between unrounded and round contexts for /ʒ/, but no difference by backness context.

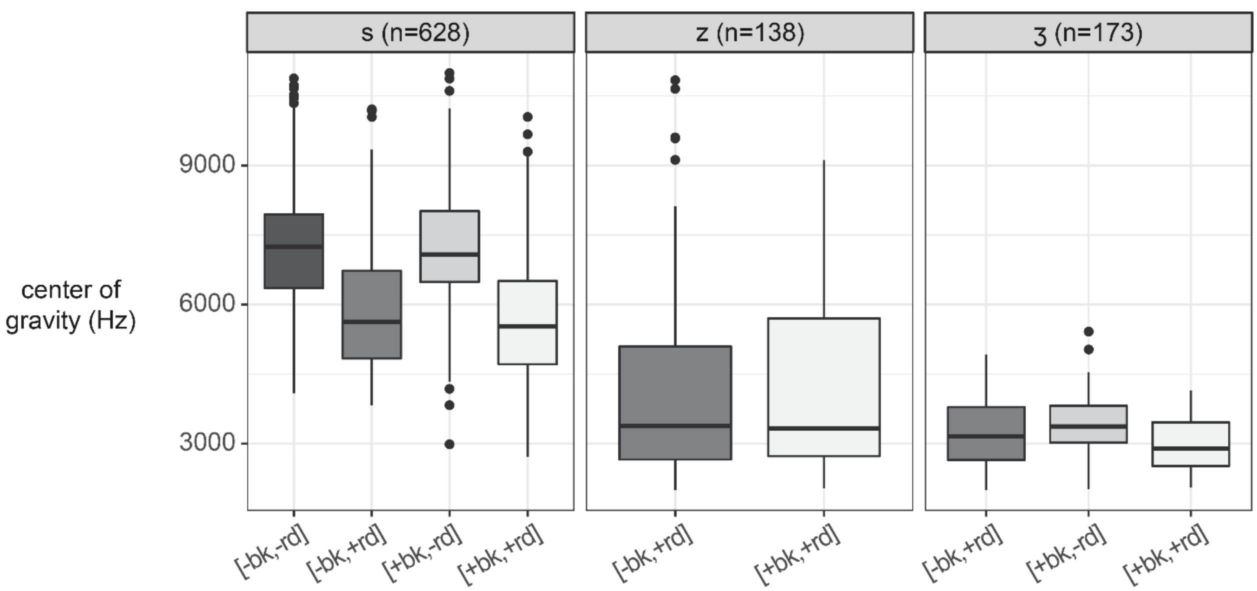

Figure 3.5: Center of gravity in Hz by backness and rounding for the Kazakh sibilants

3.4.1.3 Uyghur

Sibilants in Uyghur are produced with significantly lower center of gravity in back vowel contexts [β= -507, $\chi^2(1)$= 20.73, p < .001] and significantly lower center of gravity in round vowel contexts [β= -1258, $\chi^2(1)$= 31.59, p < .001]. The effect of vowel backness and rounding are consistent with findings in Kyrgyz detailed above. In addition, post-alveolar sibilants are produced with a significantly lower center of gravity than their alveolar counterparts [β= -2275, $\chi^2(1)$= 878.09, p < .001]. These effects are shown in Figure 6. The voiceless alveolar sibilant exhibits the highest center of gravity. All other sibilants exhibit a much lower center of gravity, although the voiced alveolar sibilant was only produced in round vowel contexts, which may depress the center of gravity plotted for /z/ in Figure 3.6. The post-alveolar sibilants occurred with a greater range of [back] and [round] contexts. The voiceless post-alveolar sibilant shows lowered center of gravity following a round vowel, which is consistent with the statistical model. Before moving on, it is worth noting that the significance of backness and rounding on sibilant center of gravity in Uyghur may result from the uniquely high center of gravity of /s/ relative to the other sibilants under examination, as well as the fact that /s/ occurred only in [-back, -round] contexts. Future work utilizing a more balanced design should resolve this potential interpretation issue.

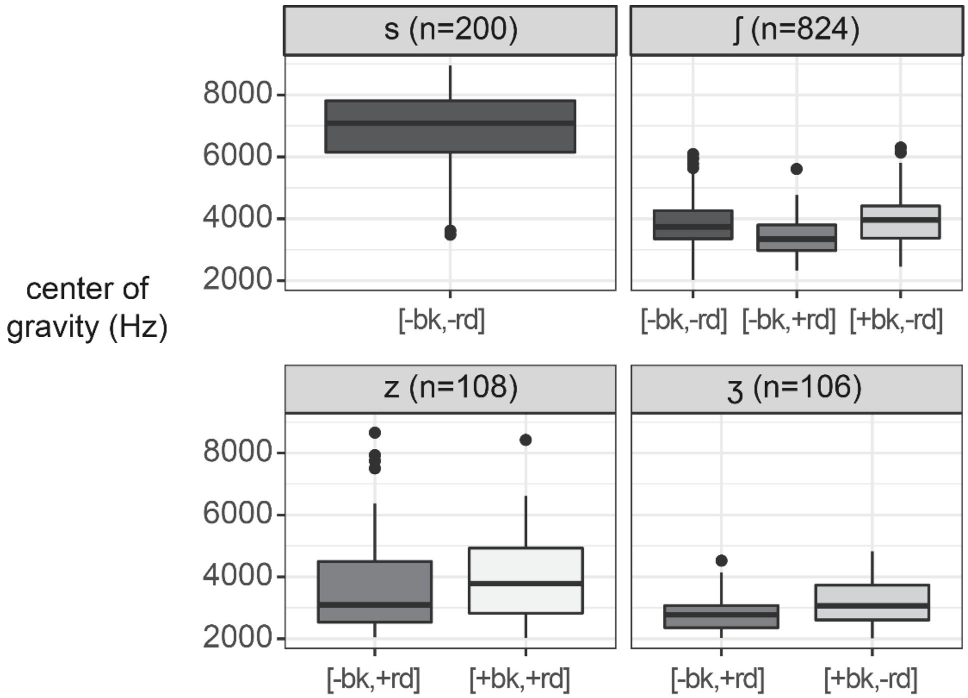

Figure 3.6: Center of gravity in Hz by backness and rounding for the Uyghur sibilants (frication portions of the affricates /ʧ/ are grouped with /ʃ/)

3.4.1.4 Uzbek

Uzbek sibilants are marked by significantly lower center of gravity after round vowels [β= -1101, $\chi^2(1)$= 17.85, $p < .001$], like the other three languages under study. Also, like Kazakh, the center of gravity of Uzbek sibilants is not significantly decreased by the backness of the preceding vowel [β= -74, $\chi^2(1)$= 0.28, $p = .59$]. These results, presented in Figure 3.7, suggest that lip rounding extends throughout the target vowel and the following sibilant. However, there is no evidence that tongue body backness differences in vowels persist into following sibilants. As for place of articulation, the center of gravity of the post-alveolar sibilants is, as expected, significantly lower than the center of gravity of alveolar sibilants [β= -1864, $\chi^2(1)$= 563.71, $p < .001$].

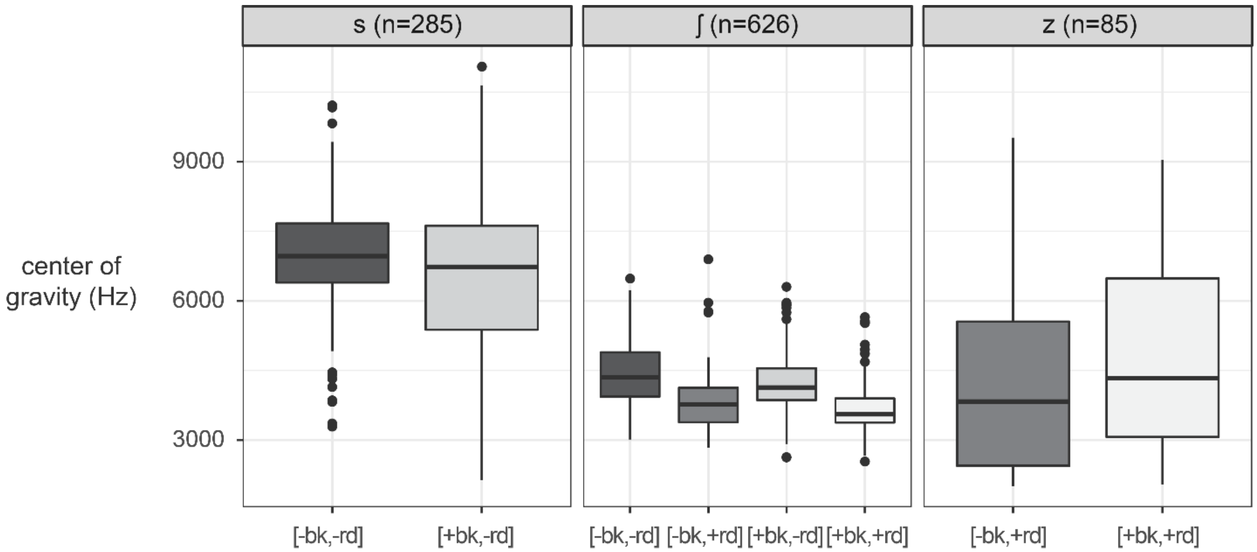

Figure 3.7: Center of gravity in Hz by backness and rounding for the Uzbek sibilants (frication portions of the affricates /ʧ/ are grouped with /ʃ/)

3.4.2 Lateral

3.4.2.1 Kyrgyz

Recall from Section 3.4.1.1 above that the sibilants undergo phonetic alternations for backness and rounding in Kyrgyz. Like the sibilants, the lateral exhibits significant phonetic differences based on backness and rounding context. Specifically, the F2 of the lateral decreases in back vowel [β= -1.59, $\chi^2(1)$= 59.25, $p < .001$] and round vowel [β= -0.25, $\chi^2(1)$= 22.28, $p < .001$] contexts. F2 of the lateral within roots is presented in Figure 3.8, while F2 of the lateral in the plural suffix is shown in Figure 3.9. In both, back vowel contexts trigger a substantial drop in lateral F2 while round vowel contexts trigger a systematic, but smaller decrease in F2. These findings support an analysis whereby the lateral is produced with a more posterior lingual constriction around back vowels, and with lip rounding in round vowel contexts.

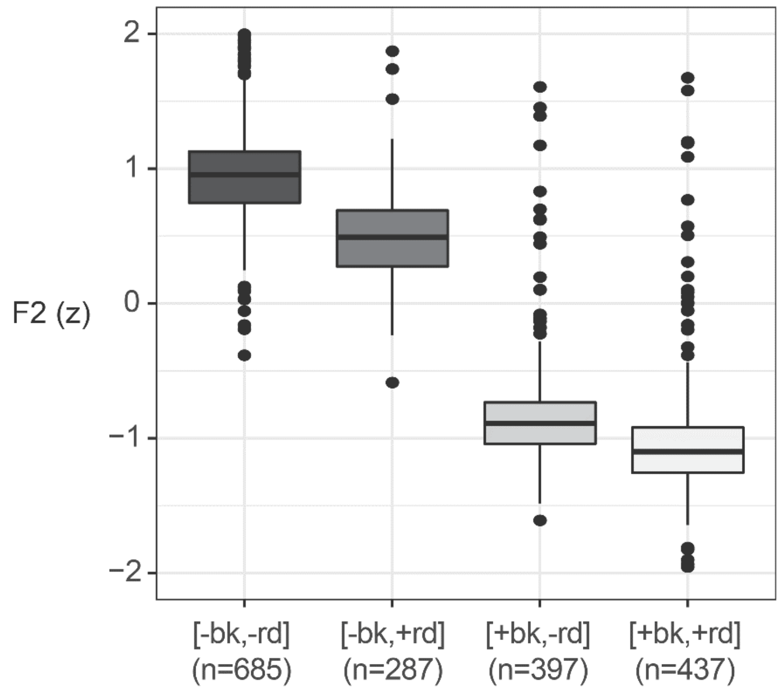

Figure 3.8: F2(z) at midpoint of the Kyrgyz lateral (root-internal) based on backness and rounding context

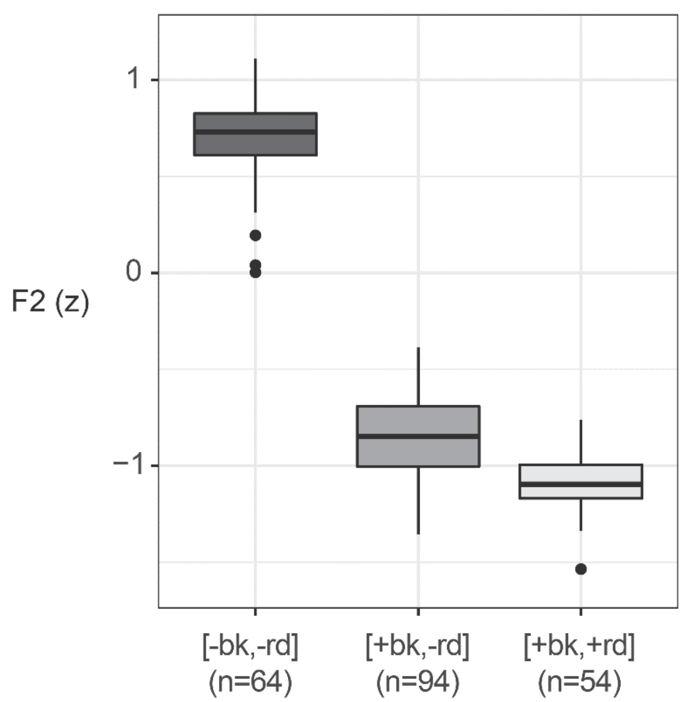

Figure 3.9: F2(z) at midpoint of the Kyrgyz lateral (plural suffix) based on backness and rounding context. No lateral-initial allomorphs of the plural suffix were produced in [-bk,+rd] contexts.

3.4.2.2 Kazakh

Although Kazakh sibilants undergo alternations only for [round], the lateral in Kazakh undergoes alternations for both [back] and [round]. Specifically, F2 of the lateral is significantly lower in back vowel [β= -1.85, $\chi^2(1)$= 31.57, $p < .001$] and round vowel [β= -0.27, $\chi^2(1)$= 13.65, $p < .001$] contexts. The Kazakh pattern of lateral alternations is evident in Figures 3.10-3.11, which show root-internal and suffixal variants of /l/. Within roots, the lateral undergoes clear alternations for both backness and rounding. The magnitude of these differ a great deal, with back vowels triggering a large decrease in F2 while round vowels trigger a more modest reduction in F2. In the plural suffix though, no rounding alternations are evident. In fact, F2 of laterals after initial round vowels is slightly higher than F2 of unrounded vowels. This lack of [round] alternations in the plural suffix is explained by the restrictions on rounding harmony in the language. Rounding harmony in contemporary Kazakh does not typically target [-high] vowels (McCollum 2018). Since the lateral-initial allomorph of the plural was only produced after vowel-final roots (Gouskova 2004), and all vowel-final roots were [-high], no lateral-initial plural suffixes were immediately preceded by a round vowel. Thus, although the initial-syllable vowel was [+round] in words like /kørpie/ 'sleeping mat' or /moldɑ/ 'mullah', the second-syllable vowel, which immediately precedes the plural suffix, was never rounded.

In sum, the Kazakh lateral alternates for both [back] and [round], as manifested by the decreases in F2 for laterals after back vowels and round vowels. However, rounding alternations were evident only within roots due to the lexemes elicited and the restrictions on rounding harmony in the language. Unlike the sibilants, which alternate for [round] only, the lateral appears to vary according to both harmony patterns in Kazakh.

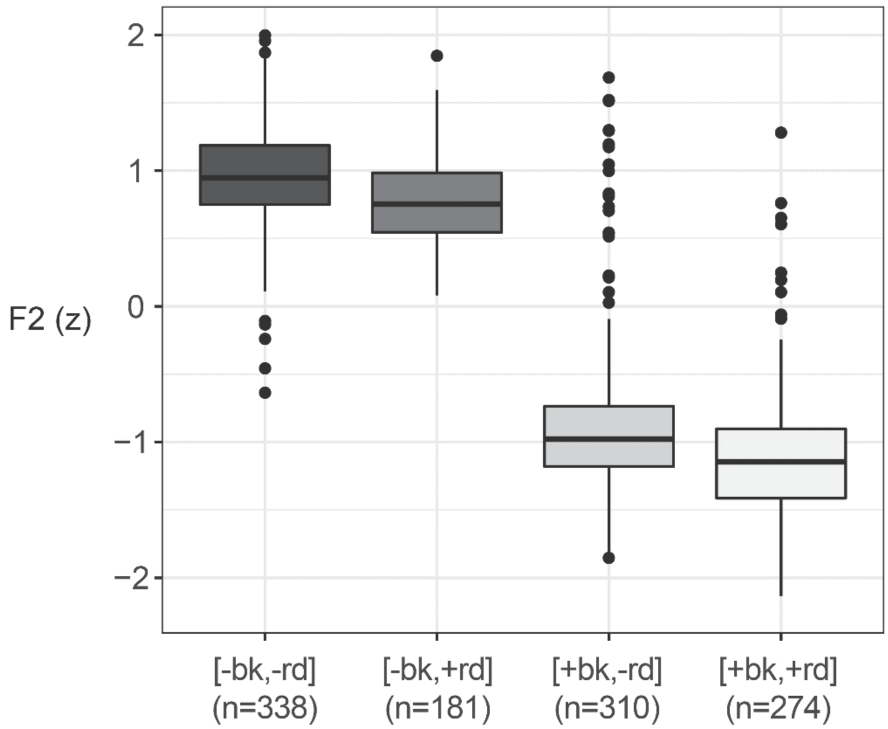

Figure 3.10: F2(z) at midpoint of the Kazakh lateral (root-internal) based on backness and rounding context.

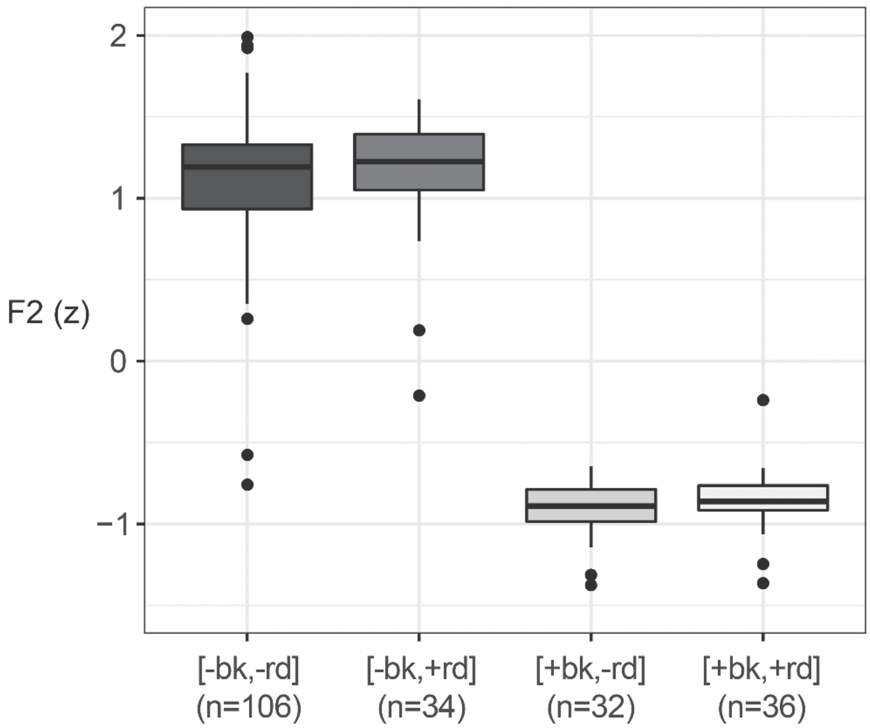

Figure 3.11: F2(z) at midpoint of the Kazakh lateral (plural suffix) based on backness and rounding context. Since all plural suffixes occurred after a [-high] vowel and rounding harmony does not target [-high] vowels, no immediately preceding vowels were actually round.

3.4.2.3 Uyghur

The lateral in Uyghur is produced with significantly lower F2 in back vowel contexts [β= -1.11, $\chi^2(1)$= 42.69, $p < .001$]. This substantial decrease in lateral F2 is seen in Figures 3.12-3.13. This decrease in F2 is consistent both within roots and in the initial segment of the plural suffix. In the context of round vowels, the lateral is produced with significantly higher F2 [β= 0.25, $\chi^2(1)$= 11.47, $p < .001$], which again is shown to be consistent in both roots and in the plural suffix. This increase is unexpected, since lip rounding, all else being equal, should lower F2. When one examines Figures 3.12 and 3.13 below, this difference in F2 in [round] contexts appears to be driven by the front rounded vowels. In back vowel words, the round vowels appear to demonstrate the same pattern seen in the other languages investigated, with round vowels triggering a slight lowering of F2. Within roots, lowering of F2 in the lateral appears greater while differences between back round and back unrounded vowels are diminished in the plural suffix.

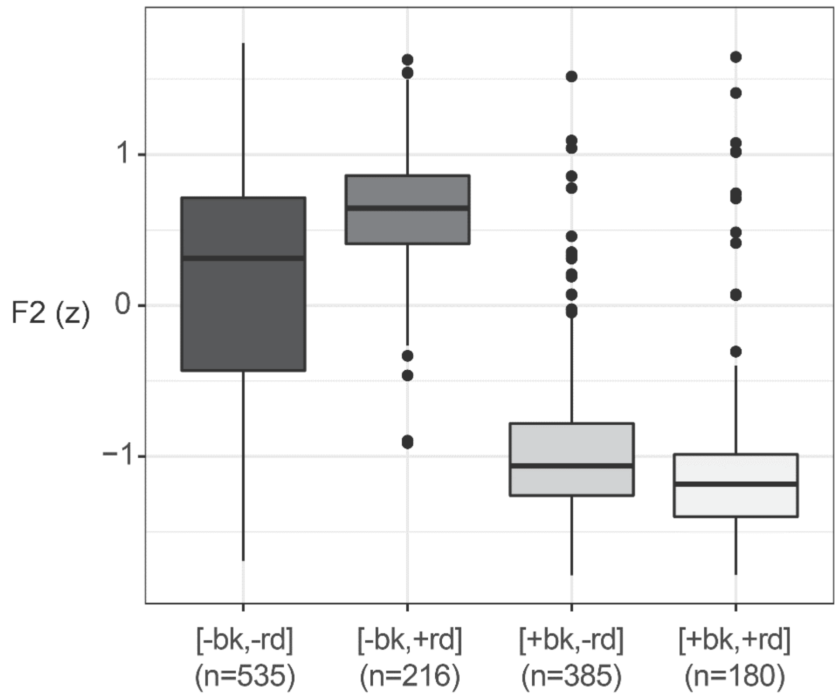

Figure 3.12: F2(z) at midpoint of the Uyghur lateral (root-internal) based on backness and rounding context

In sum, the lateral in Uyghur exhibits alternations for vowel backness. This alternation does not appear to be morphologically constrained, as it holds both within roots and across morpheme boundaries. Yet, the lateral does not appear to alternate for vowel rounding. In fact, F2 of the lateral is higher in round vowel contexts, although this effect appears to be driven by the front rounded vowels, with the lateral in back rounded contexts showing a pattern similar to Kyrgyz and Kazakh, though smaller in magnitude.

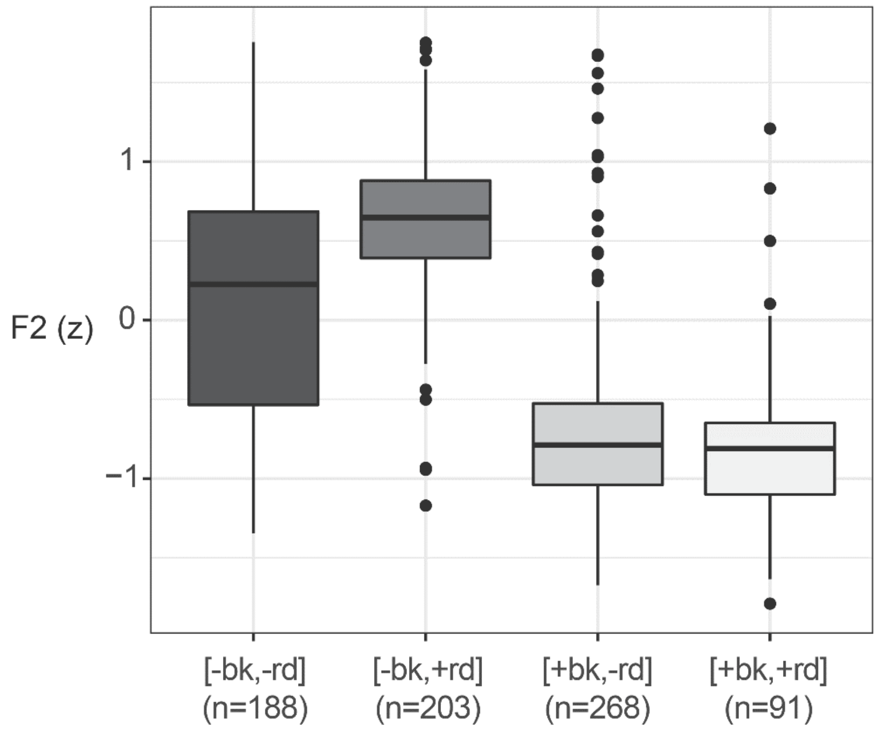

Figure 3.13: F2(z) at midpoint of the Uyghur lateral (plural suffix) based on backness and rounding context

3.4.2.4 Uzbek

In Uzbek, F2 of the lateral significantly varies according to both preceding vowel backness [β= -0.73, $\chi^2(1)$= 30.87, p < .001] and rounding [β= -0.23, $\chi^2(1)$= 10.86, p < .001]. The lateral is realized with lower F2 after back vowels and round vowels, presumably due to a more posterior tongue body position in back vowel contexts and continued lip rounding in round vowel contexts. Two patterns are observable in Figures 14-15, which show F2 of root-internal and suffixal laterals. Within roots, the lateral alternates for both backness and rounding, as in Figure 3.14. However, the lateral of the plural suffix ([-lɑr]) exhibits much smaller phonetic differences based on backness and rounding, seen in Figure 3.15. A second regression model was run to assess the significance of suffixal lateral alternations. Results indicate that vowel backness has a significant effect on F2 of the lateral of the plural suffix [β= -0.28, $\chi^2(1)$= 10.86, p < .001], although rounding does not [β= -0.01, $\chi^2(1)$= 0.04, p = .84].

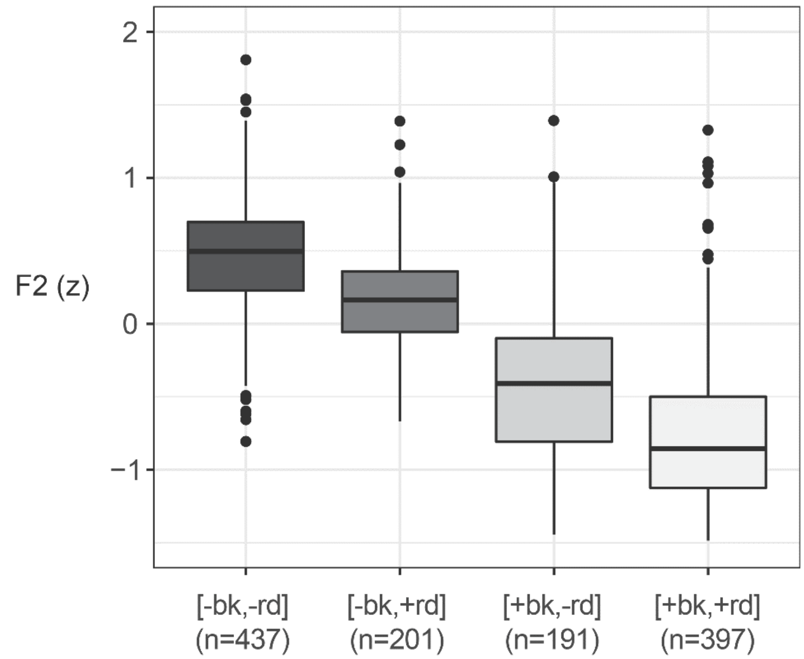

Figure 3.14: F2(z) at midpoint of the Uzbek lateral (root-internal) based on backness and rounding context.

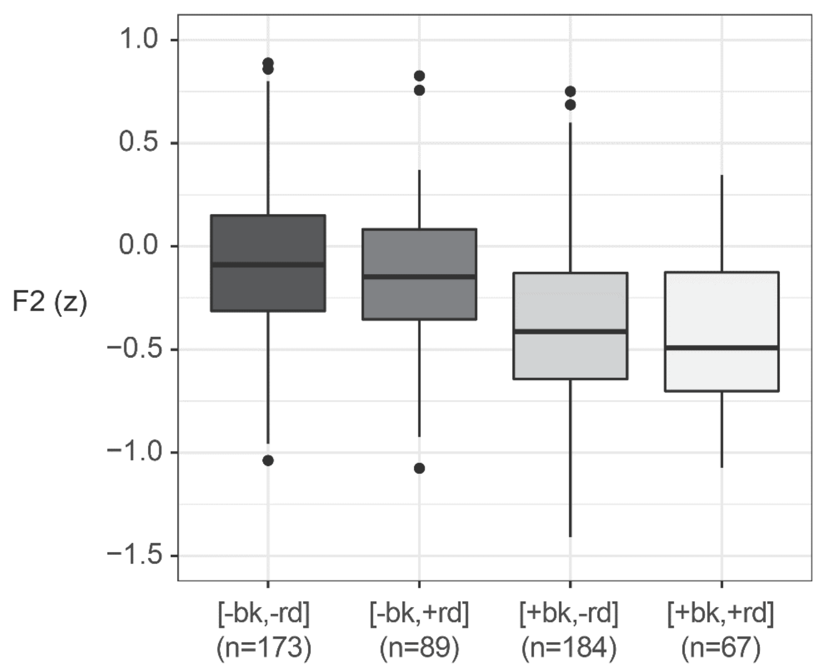

Figure 3.15: F2(z) at midpoint of the Uzbek lateral (plural suffix) based on backness and rounding context. Since backness and rounding harmony are almost non-existent, preceding vowel backness and rounding were used to determine context.

To summarize, the lateral in Uzbek alternates based on preceding vowel backness and rounding. In suffixes, though, this alternation is smaller, being significant only for vowel backness. These findings are interesting because they suggest that at least some consonants exhibit low-level alternations for vowel backness and rounding despite the fact that vowel harmony is not active in the contemporary language. Nonetheless, the lateral exhibits a pattern of vowel quality-dependence similar to those found in neighboring languages with robust harmony patterns.

3.4.3 Summary

In Kyrgyz, both the sibilants and the lateral display significant phonetic differences for [back] and [round]. In Kazakh, the lateral exhibits significant alternations for both features, while the sibilants exhibit a significant alternation for [round] only. In Uyghur, the sibilants alternate for [round] and [back] while the lateral alternates only for [back]. Finally, in Uzbek, the lateral alternates for both features, but the sibilants alternate for [round] only, similar to Kazakh. In general, all sibilants show alternations for [round], while only Kyrgyz and Uyghur sibilants show an alternation for [back]. In contrast, in all four languages the lateral alternates for [back]. In every language except Uyghur, the lateral also displays significant systematic differences for [round]. Results are summarized in Table 3.4.

Table 3.4: Summary of significant consonant alternations across the four languages

	Context	Language			
		Kyrgyz	Kazakh	Uyghur	Uzbek
Sibilants	[back]	✓		✓	
	[round]	✓	✓	✓	✓
Lateral	[back]	✓	✓	✓	✓
	[round]	✓	✓		✓

One question yet to be addressed is the perceptual significance of these differences. Ghosh et al. (2010) reports a wide range of just noticeable differences (JNDs) across speakers in a task discriminating between /s/ and /ʃ/. Their reported JNDs range from around 300 Hz for some speakers to over 1000 Hz for others. Generally, the participants in Ghosh et al (2010) were able to distinguish between voiceless sibilants differing in a center of gravity of approximately 600-700 Hz. Comparing Ghosh et al's (2010) finding with the mean differences from the statistical model, which are reported in Table 3.5, most of the differences among the sibilant may be perceivable, although the Uyghur and Kyrgyz differences for [back] (507 Hz and 625 Hz) are less likely to be noticeable than the other differences.

Table 3.5: Mean significant center of gravity differences (Hz) for the sibilants based on vowel context

	Context	Language			
		Kyrgyz	Kazakh	Uyghur	Uzbek
Mean significant difference (Hz)	[back]	625		507	
	[round]	1000	1778	1258	1101

Although the perception of the alveolar versus the velarized lateral is not well known, the magnitude of the differences in the lateral based on [back] and [round] context is shown below in Table 3.6. Note that z-score values are language-specific, but the similarities of the vowel inventories between these languages makes tentative comparison possible. Of the four languages, Kyrgyz and Kazakh exhibit the largest difference in lateral F2 for backness, with smaller differences for Uyghur and Uzbek. Differences in lateral F2 after round vowels are relatively consistent across languages, and as expected, are much smaller than backness differences.

Table 3.6: Mean significant differences in F2 (z) at lateral midpoint based on vowel context

	Context	Language			
		Kyrgyz	Kazakh	Uyghur	Uzbek
Mean significant difference (z)	[back]	1.59	1.85	1.11	0.73
	[round]	0.25	0.27		0.23

3.5 Discussion and analysis

Returning to the predictions of the three analyses, it is clear that some of the consonants examined do not alternate for both harmony patterns in Turkic. That being said, the Kyrgyz pattern of consonant alternations most clearly exemplifies what strictly local harmony should look like. Based on the findings above, sibilants and the lateral exhibit phonetic differences consistent with both backness and rounding harmony, suggesting that tongue body backing and lip rounding extend throughout intervening consonants. The other three languages each exhibit three significant alternations. Kazakh and Uzbek lack only a [back] alternation for the sibilants, and Uyghur lacks only a [round] alternation for the lateral. The fact that Uzbek exhibits comparable alternations to the other languages is somewhat surprising given that vowel harmony no longer robustly operates in the language.

Results suggest that vowel harmony is not local in the strict sense, since some consonants fail to alternate for harmony in three of the four languages. If phonology dictates phonetic content, as in the strictly local analysis, one must explain why some consonants fail to alternate. Ní Chiosáin & Padgett (2001:119) presents strict locality as a stipulative constraint on human grammar. Effectively, they claim that phonology exerts authoritarian force over phonetics in spreading. They write:

> We assume that locality holds strictly in two senses of "strict." First, spreading respects segmental adjacency, as proposed by the references above. An essential result of this view is that segments are either blockers or participants in spreading; there is no transparency or skipping. Second, segmentally strict locality is inviolable; in Optimality Theoretic term, Gen does not produce structures in which segments are skipped in a spreading domain.

If segmental locality is inviolable, as they claim, phonological computation cannot be affected by the segment-specific articulatory or auditory properties of segments that undergo harmony. These phonetic facts must be overcome by harmony, or else the possibility for skipped segments re-emerges. In their proposal, harmony must dictate the realization of intervening consonants, even if there are phonetic forces that militate against the realization of the harmonic feature on a given consonant. As noted above, this predicts that the sibilants and the lateral will undergo phonetic alternations consistent with harmony. Yet, in three of the four languages examined, some segment fails to undergo for at least one harmony pattern.

At the same time, the fact that so many do alternate for harmony is problematic for the trans-consonantal locality approach. Under the trans-consonantal locality analysis, the phonological and phonetic modules of the grammar both trigger backness and rounding alternations, one via harmony and the other via coarticulation, which poses a duplication problem. The violable strict locality account offers an explanation for why the sibilants in Kazakh fail to undergo harmony, namely their class-specific articulatory properties, which militate against syntagmatic effects on the place of constriction in the oral cavity. Under this account, phonology neither dominates phonetics nor are the two completely independent.

Consider why the sibilants and lateral typically alternate in accordance with adjacent vowels. If the phonological and phonetic modules are completely independent, there is no reason to presume that they would. One might expect some language's phonology to make all vowels [+back] within a given word while that same language's phonetics may have reason to compel consonants in that word toward more anterior gestural configurations. Or perhaps more realistically, one might expect no obvious differences in consonantal articulation or acoustics if vowel harmony and the phonetics of intervening consonants are not tied together somehow. Yet, it is clear that phonology and phonetics are tied together in many ways (see Kingston 2007 for discussion). If for no other reason, coarticulatory pressures should compel intervening consonants to become more like the phonologically-dictated [back] vowels that flank them.

If one establishes that phonology and phonetics are not wholly independent, two questions arise. First, how related are they? Some have argued that they are so related that they are indistinguishable (e.g. Pierrehumbert 1994; Pierrehumbert et al. 2000). This is a difficult question to answer, although subsequent chapters take up the topic of the phonology-phonetics interface, arguing for a distinction between the two, and discussing some diagnostics for classifying sound patterns as either phonological or phonetic. In addition to the extent of their relationship, the second question that arises relates to the nature of their relationship. How do phonology and phonetics relate to each other? Early proposals argued that phonetics implements phonological commands, translating abstract phonological symbols into space and time (see Cohn 2006; Kingston 2007 for discussion). As Kingston (2007) notes, even in simple cases, phonetic and phonological forces may compete and conflict, yielding outputs that are tenuous resolutions of both phonological and phonetic concerns. This conception of phonology and phonetics, which is consistent with the violable strict locality account, offers an intermediate way forward for the analysis of (non-)alternating consonants in these languages.

To formalize violable strict locality is simply to define a set of harmony-driving constraints that is able to both force consonantal alternations and ignore consonants. To this end, I propose a version of AGREE (36; Lombardi 1999; Baković 2000), that imposes strict segmental locality, as advocated by Gafos (1999) and Ní Chiosáin & Padgett (2001), among others. This strictly local harmony driver works in tandem with a second AGREE constraint that compels vowels to agree in backness, (37).

(36) AGREE-CV-[BK]: assign a violation to every adjacent pair of segments that do not agree for the feature [back]

(37) AGREE-VV-[BK]: assign a violation to every syllable-adjacent pair of vowels that do not agree for the feature [back]

This first constraint assumes that all segments have a possible [+back] or [-back] variant. For some sounds this is easily represented, e.g. [l] vs. [ł], but for most consonants I employ the advanced and retracted IPA diacritics (combining subscript + indicates advanced; combining subscript - indicates retracted) to mark the relative backness of the consonant. I illustrate the claim below with several tableaux. In (38), strict locality, as in Kyrgyz is demonstrated. Underlyingly, the only [+back] segment in the word is the initial-syllable vowel. If we assume a set of positional faithfulness constraints that privilege the initial-syllable vowel over other segments, then due to the ranking of the two AGREE constraints over the relevant faithfulness and markedness constraints, /ɑ/ drives harmony, forcing all other segments to be produced with a more posterior gesture. In (38a), the faithful output incurs a violation of AGREE-VV because the two vowels do not agree in backness. Additionally, (38a) incurs two violations of AGREE-CV because the two segments in the substrings [b̟ɑ] and [ɑl] disagree in backness. In (38b), the vowels disagree for backness, incurring a fatal violation of AGREE-VV. However, in this output candidate, the vowel triggers assimilation of /b̟/ to [b̠] and /l/ to [ł]. As a result, only one bigram disagrees in backness, [łd̟], so AGREE-CV is violated only once. The third candidate, (38c), satisfies Agree-VV by mapping input /e/ to [ɑ], but critically violates AGREE-CV by failing to map /d̟/ to [d̠]. Harmony thus skips transparent [d̟] in (38c), incurring two violations of AGREE-CV, one for the bigram [łd̟] and one for [d̟ɑ]. The winning candidate (38d) satisfies both AGREE-VV and AGREE-CV by outputting all segments in conformity to the backness of the initial-syllable vowel.

(38)

		/b̟ɑl-d̟e/	AGREE-VV-[BK]	AGREE-CV-[BK]	ID-IO[BK]	*[+BACK]
	a.	b̟ɑld̟e	*!	**		*
	b.	b̠ɑłd̟e	*!	*	**	***
	c.	b̠ɑłd̟ɑ		*!*	***	****
☞	d.	b̠ɑłd̠ɑ			****	*****

The addition of a specific markedness constraint banning alternations on a given segment can build on the ranking instantiated above to produce a pattern like that in Kazakh. In Kazakh, results above indicate that the sibilants do not alternate for backness harmony. This finding makes sense given several findings from previous work on sibilants, specifically findings in Recasens (1985) and Tabain (2001). Since sibilants are resistant to coarticulatory forces, they should also be resistant to harmonic alternations, whether these alternations be construed as the production of a phonological rule targeting the sibilant or as coarticulation from flanking vowels. To formalize this phonetic insight, I introduce a constraint against backness alternations for the sibilants *S̱/Ṣ, in (39). Note that this constraint penalizes both fronted and backed sibilants. I have no reason at present to stipulate that the sibilants are produced with an intrinsically fronted or backed articulatory configuration. Instead, I simply assume that the sibilants are produced with gestures orthogonal to backness harmony. To emphasize this point, this constraint treats a third, harmony-neutral articulatory configuration as least marked. This may be incorrect, and future work may demonstrate that sibilants are by default produced with gestural coordination that matches some value of the harmonic feature, but for the present I will assume that the specific articulatory configuration for the sibilants is determined independent of harmony.

(39) *S̱/Ṣ — assign a violation to every sibilant that is produced with either phonetic fronting (S̱) or backing (Ṣ).

If *S̱/Ṣ outranks the strictly local AGREE-CV constraint, then the transparency of the Kazakh sibilants to backness harmony receives an analysis, shown below in (40). Candidates (40a-c) all incur a crucial violation of AGREE-VV due to the failure of harmony on underlying /ie/. Candidate (40d) exhibits strictly local harmony, triggering harmony on both /s/ and all other segments to force them to agree with initial-syllable /ɑ/. Candidate (40e) is optimal under this ranking because it assimilates the second-syllable vowel and all consonants except the sibilant to the backness of the initial-syllable vowel. Under this analysis, the sibilants are neither specified for a fronted or backed articulation, highlighting the fact that their articulatory configuration is completely orthogonal to harmony.

(40)

		/b̟ɑs-t̟ie/	AGREE-VV-[BK]	*S̠/S̟	AGREE-CV-[BK]	ID-IO[BK]	*[+BACK]
	a.	b̟ɑst̟ie	*!		***		*
	b.	b̠ɑst̟ie	*!		*		
	c.	b̠ɑs̠t̠ie	*!	*	*	***	****
	d.	b̠ɑs̠t̠ɑ		*!		****	*****
☞	e.	b̠ɑst̠ɑ			**	***	****

Under this account of locality, AGREE-CV attempts to dictate the realization of every segment in the word, consistent with the conceptions of strict locality developed in the literature. The fact that this constraint is violable introduces the possibility that other factors override local harmony. In the case just illustrated in (40), the sibilant's inherent resistance to coarticulation is encoded in the phonology as a markedness constraints against its alternation. Note the importance of AGREE-VV for an adequate account of vowel-consonant interactions. If AGREE-VV is removed from the constraint set, then the specific markedness constraint *S̠/S̟ will block harmony on the locative suffix above due to the markedness of [+back] segments. As described at length in Bakovic (2000), AGREE constraints always favor blocking over transparency. Thus, AGREE-VV is necessary to generate transparency across the non-participating sibilant.

One may also wonder if the specific markedness constraint above, *S̠/S̟, is truly phonological or phonetic. Given that the space of OT constraints is constrained (at least partially) by the space of potential phonetic effects, then there is no obvious answer to this. If harmony is the phonologization of vowel-to-vowel coarticulation (Öhman 1966) or misperception (Ohala 1994), then the same phonetic forces driving coarticulation are present in the form of phonological constraints. Thus, there is clear evidence that phonetic modulation of the sibilants is dispreferred. I suggest that this dispreference, like the active preference for adjacent vowels to be similar, can be encoded in the phonological grammar as a markedness constraint.

The analysis developed here is similar in several ways to the gestural theory developed in Smith (2018). Smith formalizes harmony in terms of gestural representations rather than harmony-driving markedness constraints, but similarly contends that harmony is strictly local. Smith's analysis, like mine, is more nuanced than previous analyses because she allows transparency on a restricted scale, while authors like Gafos (1999) do not. In her theory, segments that are produced with gestures intrinsically antagonistic to harmony mask the propagation of the harmonic feature. Thus, segments like the sibilants balance participation in harmony with maintaining their gestural specification via Task Dynamics (Saltzman & Munhall 1989), producing a gestural blend. In these cases, a gestural blend might be equivalent to the low-level phonetic effects reported in Benus & Gafos (2007). In their work on transparent vowels in Hungarian, the transparent vowels are never backed like a [+back] vowel, but exhibit small articulatory and acoustic differences that Benus & Gafos (2007) use to argue that harmony obeys locality. It is possible that a gestural blend between a sibilant and [+back] vowel may differ somewhat from the blend between a sibilant and a [-back] vowel. If so, Smith's account appears to predict low-level phonetic alternations in these cases. The results from this chapter do not support this interpretation, at least for the sibilants and backness harmony in Turkic. The sibilants don't appear to alternate at all (except in Kyrgyz), precluding the need for conflict resolution via gestural blending. Also, Smith's account is reminiscent of early generative accounts, like Vago (1973, 1976), wherein transparent segments undergo harmony at one level of representation, only to be subject to a later rule returning them to their original feature value. Smith's analysis accounts for both strictly local spreading and transparency from the interplay between gestural spreading and antagonistic gestures, which is similar to rule-ordering analyses like Vago (1973, 1976). Like in a rule-ordering analysis, Smith upholds that harmony affects the transparent segment even though the segment in question does not surface in accordance with harmony.

Unlike Smith (2018), I use two AGREE constraints to model limited transparency in harmony. I assume that the general AGREE-VV constraint is necessary to account for languages with real trans-vocalic transparency (see Dye 2015), and AGREE-CV is necessary to account for strictly local harmony. Since transparency is reported to exist, the two constraints cannot be conflated, and their mutual existence entails that they may cooperate when ranked appropriately, as in the sample tableau above for Kazakh. Within OT, other constraints can model the same behavior, like Finley's (2008) definition of SPREAD or McCarthy's (2004) headed spans. Each of these constraints come with their own drawbacks and seemingly pathological predictions. My aim is thus not to adjudicate between the various constraints available in the literature, but to argue that strict locality can be reasonably accounted for if we remove Ní Chiosáin & Padgett's (2001) stipulation on Gen, and formalize strict locality the same way that many other forms of locality are encoded in OT, as a constraint.

3.6 Conclusion

This chapter has examined the claim that harmony in Turkic is strictly local, affecting all segments within the harmonic domain. Findings from an acoustic study indicate that the sibilants vary according to [round] context in all four languages but are, except in Kyrgyz, invariant for [back]. As for the lateral, findings suggest that the lateral generally alternates for both [back] and [round] in each language. Results also support the common view of harmony across these languages, with harmony being the strictest in Kyrgyz, only phonetic in Uzbek, with Kazakh and Uyghur somewhere in between. As it relates to locality, this chapter supports a violable conception of strictly local spreading, deriving both strictly local and transparent patterns from constraint ranking rather than a stipulation on the grammar.

This chapter as well as the previous chapter have dealt with locality and transparency in Turkic, marshalling experimental data to inform the analysis of harmony in Uyghur, and the effects of harmony on intervening consonants in Kyrgyz, Kazakh, Uyghur, and Uzbek. At this point, the dissertation pivots from investigating locality to examine gradience in these languages, developing a larger analysis of gradience in vowel harmony that extends beyond Central Asian Turkic.

Chapter 4: Gradience in backness harmony: A production study

The previous two chapters focused on locality issues in harmony, but this chapters shifts to examine a different theoretical topic, gradience. In fact, the rest of the dissertation continues the investigation of gradience, moving from acoustic data in this chapter, to perceptual data in Chapter 5, the formal analysis in Chapter 6, and finally to a larger discussion of gradience in Chapter 7.

4.1 Background

In Structuralist and early Generative work, all grammatical knowledge was assumed to be categorical and discrete (e.g. Bloomfield 1933; Chomsky & Halle 1968). As a result, phonetics, being inherently gradient, was regarded as extra-grammatical, as reflected in de Saussure's distinction between speech [parole], which has no linguistic significance, and language [langue], the purview of linguistic theory, which includes phonology, syntax, and semantics. The core tenets of this view are evident in de Saussure's (1993) lectures – phonology is part of linguistic competence while phonetics is not, being constrained by physiology only. "The physiology of speech sounds [sons de la parole] is not part of linguistics" (53), and "The speech part [parole] of language has no essential connexions with the language [langue] part (73)." Further, de Saussure (1993:72) likens language to Morse code. Language is the code, while phonetics is the trivial collection of electrical components used to express the code. "Phonation might appear to command an important place among the phenomena of language; appears as inessential as the various pieces of electrical apparatus which may be used to transmit signs of the Morse alphabet."

A very similar view of phonetics is espoused by Chomsky & Halle (1968). They write, "Given the surface structure of a sentence, *the phonological rules of the language interact with certain universal phonetic constraints* to derive all grammatically determined facts about the production and perception of this sentence" (293; emphasis mine). Chomsky & Halle notes the multitudinous possible pronunciations of a given word or sentence, suggesting that this does not form a meaningful level of representation for the linguist. Instead, the linguistic surface representation is an abstraction from the detailed phonetics, which was the mainstream view in phonological theory for decades (see Pierrehumbert 1994; Kirchner 1997 for critiques). Chomsky & Halle (1968) further justify this abstraction by noting the analyst's inability to adequately transcribe all the phonetic variation encountered in language, "In addition, phonetic transcriptions omit properties of the signal that are supplied by *universal rules*. These properties include, for example, the different articulatory gestures and various coarticulation effects- the transition between a vowel and an adjacent consonant, the adjustments in the vocal tract shape made in anticipation of subsequent motions, etc" (295, emphasis mine). They suggest that, among other phonetic forces, coarticulation should exhibit universal tendencies constrained by physics and be completely independent of language-specific phonological rules.

However, while Chomsky & Halle advocated for abstract sound representations, a body of phonetics research has discovered that phonetic patterns, in particular coarticulatory patterns, are language-specific. Therefore, phonetic patterns cannot be universal, and must form part of a speaker's knowledge of their language (e.g. Öhman 1966; Purcell 1979; Manuel & Krakow 1984; Keating 1985a; Maddieson & Emmorey 1985). As an example, Cohn (1993) finds that nasal consonants trigger gradient nasalization on a preceding vowel in English, but only minimally affect a preceding vowel in French (for similar patterns of nasal coarticulation, see Butcher 1999; Stoakes et al. 1999). If gradient phonetic patterns are not universal, but exhibit the same kind of language-specificity found in phonological patterns, gradient phonetics must be part of the linguistic grammar, and not simply a peripheral component of language. In addition to supporting the grammatical status of phonetics, such findings also suggest that all grammatical knowledge is not, by definition, categorical because phonetic patterns are gradient.

Irrespective of the ontological status of grammatical knowledge, the prevailing assumption is that phonological knowledge is categorical. If this is assumed, then gradient patterns are, by default, not phonological. In this way, gradience is typically used to differentiate phonetic from phonological patterns, in effect, further entrenching the assumed categoricality of phonology (e.g. Cohn 1993, 2006; Zsiga 1997). To be clear, I am not using the term to refer to optionality. Instead, I use the term gradience to refer to degrees of phonetic effect, sometimes called subphonemic effects. The commonly held view is that phonological patterns involve substitution of abstract symbols, producing outputs that are indistinguishable from their non-derived counterparts. In contrast, phonetics involves the translation of phonological abstractions into continuous space and time.

Gradience as the primary diagnostic of phonetic versus phonological status has encountered one recurrent challenge, incomplete neutralization. In many cases, putatively phonological patterns, like word-final devoicing, have been shown to exhibit subphonemic gradience. In other words, these patterns produce sounds that are not acoustically or articulatorily identical to their non-derived counterparts (e.g. Port & O'Dell 1985; Warner et al. 2004; Ernestus 2011; Braver 2013). One interesting generalization concerning attested cases of incomplete neutralization is that they have all, at some time or another, been analyzed as post-lexical. Within Lexical Phonology (Mohanan 1982; Kiparsky 1985), the core of phonology is categorical and structure-preserving. Post-lexical phonology, on the other hand, is not constrained by Structure Preservation. As a result, post-lexical phonology may be gradient, a point discussed in Liberman & Pierrehumbert (1984) and Kiparsky (1985). Thus, in the traditional Generative view, gradience is confined to phonetics, or in Lexical Phonology, phonetics and post-lexical phonology. Under both views, though, the categoricality of morphophonology is preserved. If all attested cases of subphonemic gradience are either post-lexical or phonetic, the question is thus, may morphophonological patterns be gradient? To date, the answer has been no. Cohn (2006:36) states the issue thusly: "Morphophonemic alternations are at the very core of what most phonologists think of as phonology . . . If these sorts of cases are shown to involve gradience, this would strike at the core of our understanding of the phonology, since these are the least disputable candidates for 'being phonology.'" Work like Zsiga (1995,1997) demonstrates key differences between phonetic and phonological patterns from phonetic data. Although it can be challenging to distinguish the phonetic from the phonological from surface phonetics alone, previous work offers a range of diagnostics for evaluating whether a given pattern falls out from phonological or phonetic forces (e.g. Cohn 1993). With diagnostics from previous literature, discovering that phonology is gradient entails finding gradience in surface phonetic forms that is not derivable from known phonetic forces. If surface gradience is inconsistent with well-described phonetic forces like coarticulation or reduction, then the source of gradience may plausibly be phonological.

In this chapter, I examine acoustic data from vowel harmony in three languages, showing that Kazakh and Uyghur vowels are asymmetrically fronted in non-initial syllables in a manner not derivable from known phonetic forces of reduction or interpolation across unspecified segments. This contrasts with results from Kyrgyz and Uzbek, where vowels exhibit relatively symmetrical centralization, which is derivable from phonetic reduction. Building from these results, as well as evidence from elsewhere in the phonologies of these languages, I contend that phonology is not ontologically categorical, but may exhibit both categoricality and gradience. For background on vowel harmony in these languages, consult Chapter 1.

4.2 Methods

4.2.1 Participants

Participants were recruited from existing relationship networks in each of the field sites: Taldykorgan and Chunja, Kazakhstan; and Bishkek and Osh, Kyrgyzstan. In Taldykorgan, nine Kazakh speakers (seven females, mean age: 33.4 years, range: 19-49 years), and in Chunja, nine Uyghur speakers (five females, mean age: 44.4 years, range: 19-63 years) participated in the study. In Bishkek, thirteen Kyrgyz speakers (eleven females, mean age: 35.0 years, range: 18-57 years), and in Osh, nine Uzbek speakers (four females, mean age: 28.8 years, range:19-45 years) participated in the study. All participants reported native fluency in the target language. Many participants reported additional fluency in at least one neighboring Turkic language, as well as Russian.

4.2.2 Stimuli

The elicited data set derives from a word list composed of monosyllabic and disyllabic roots exhibiting all (short) vowel phonemes in each language. All monosyllabic roots ended in a sibilant or a liquid (e.g. Kyrgyz /bɑʃ/ 'head' and /bɑl/ 'honey'). Disyllabic stimuli were composed of two vowels that agree in height (e.g. Kyrgyz /bɑltɑ/ 'axe' and /ɯrɯm/ 'superstition'; Kazakh /moldɑ/ 'mullah'; Uyghur /tøʃæk/ 'sleeping mat'). See Table 4.2 for the full list of stimuli.

Each lexical item was prompted in thirteen different grammatical contexts, which are illustrated in Table 4.1 below for the Kyrgyz root /bɑʃ/ 'head'. With monosyllabic roots, this created words of up to three syllables in length, and for disyllabic roots, words of up to four syllables in length. As part of the investigation of locality in Uyghur (see Chapter 2), words of up to five syllables in length were elicited by adding the POSS.3S suffix to plural forms (e.g. /sællæ-lær-i-dæ/ 'turban-PL-POSS.3S-LOC').

Table 4.1: Grammatical prompts for the Kyrgyz root /bɑʃ/ ‘head’

Singular		Plural	
Word	Gloss	Word	Gloss
bɑʃ	NOM	bɑʃ-tɑr	(NOM) PL
bɑʃ-tɑ	LOC	bɑʃ-tɑr-dɑ	PL-LOC
bɑʃ-tɑn	ABL	bɑʃ-tɑr-dɑn	PL-ABL
bɑʃ-tɯ	ACC	bɑʃ-tɑr-dɯ	PL-ACC
bɑʃ-ɯ	POSS.3S		
bɑʃ-ɯm	POSS.1S		
bɑʃ-ɯ-nɯ	POSS.3S-ACC		
bɑʃ-ɯm-dɑ	POSS.1S-LOC		
bɑʃ-ɯm-dɯ	POSS.1S-ACC		

In Kyrgyz, both backness and rounding harmony operate throughout the word, so all eight phonemic vowels occurred in all four syllables elicited (Hebert & Poppe 1963; Kirchner 1998b). In the other languages, though, the distribution of the round vowels is more restricted. In all three languages, the mid round vowels, /o ø/, occurred in initial syllables only. Also, in Uzbek, Kazakh, and Uyghur, the high round vowels occurred less frequently and with reduced lip rounding in non-initial syllables (Balakayev 1962; Ibrohimov 1967; Hahn 1991; Kirchner 1998a; McCollum 2018).

Table 4.2: Stimulus lexemes for all four languages

Root type	Gloss	Kyrgyz	Uzbek	Kazakh	Uyghur	Vowel
Monosyllabic	head	bɑʃ	bɑʃ ~ boʃ	bɑs	bɑʃ	ɑ
	honey	bɑl	bɑl ~ ɑsɑl	bɑl	bɑl	
	friend	dos	dos	dos		o ɔ
	road	ʤol	jol	ʒol	jol	
	winter	qɯʃ	qiʃ	qɯʃ	qiʃ	ɯ i
	year	ʤɯl	jil	ʒɯl	jil ~ jɯl	
	ice	muz	muz	mʊz	muz	u ʊ
	slave	qul	qul	qʊl	qul	
	memory	es	es	ies	æs	æ e ie
	waist, lower back	bel	bel	biel	bæl	
	eye	køz ~ gøz	køz ~ gøz	køz	køz	ø
	lake	køl ~ gøz	køl ~ gøz	køl	køl	
	tooth	tiʃ	tiʃ ~ ʧiʃ	tɪs	ʧiʃ	i ɪ
	elephant	pil	fil	pɪl	pɯl ~ pil	
	dream	tyʃ	tyʃ	tʏs	tyʃ	y ʏ
	flower	gyl	gyl	gʏl	gyl	
	bet			bæs		æ
	nutrition			nær		
Disyllabic	axe	bɑltɑ	boltɑ	bɑltɑ	pɑltɑ	ɑ
	curtain	pɑrdɑ ~ pɑrdo	pɑrdɑ	pierdie	pɑrdɑ	
	mother				ɑpɑ / ɑnɑ	
	frog				pɑqɑ	
	turban				ʧɑlmɑ	
	mullah	moldo	moldo	moldɑ	moldɑ	o
	persimmon				χormɑ	
	superstition	ɯrɯm	irim	ɯrɯm		ɯ
	paste				ʃɯlɯm	
	soot		qurum	qʊrʊm	qurum	u
	turban	selde ~ selle		sældie	sællæ	e æ
	a type of old gun	kelte				
	sleeping mat	kørpø	kørpɑ	kørpie	kørpæ	ø
	bed	tøʃøk	tøʃɑk		tøʃæk	
	science	ilim	ilim	ɪlɪm	ilim	i
	grape	ʤyzym	jyzym	ʒʏzʏm	yzym ~ yʒym	y

The allomorphs of the suffixes elicited are shown below in Table 3 for all four languages. Due to the robustness of rounding harmony in Kyrgyz, Kyrgyz exhibits the largest number of allomorphs of each suffix, while Uzbek typically exhibits the least. Note several differences between Uyghur and the other three languages that will be relevant throughout. First, the plural suffix in Uyghur may undergo raising, which results in a six-way alternation for backness and rounding. Second, the ablative suffix in Uyghur is realized with a high vowel, unlike the other languages, in which the ablative is always non-high. This difference in vowel height results in a four-way alternation for vowel backness and rounding, like for the other underlyingly high vowel suffixes

Table 4.3: Elicited suffix allomorphs

	Kyrgyz	Uzbek	Kazakh	Uyghur
PL	-lɑr~lor~ler~lør, -dɑr~dor~der~dør, -tɑr~tor~ter~tør	-lɑr	-lɑr~lier, -dɑr~dier, -tɑr~tier	-lɑr~lær, plus raised variants (-lɯr~lur~lir~lyr)
LOC	-dɑ~do~de~dø, -tɑ~to~te~tø	-dɑ, -tɑ	-dɑ~die -tɑ~tie	-dɑ~dæ -tɑ~tæ
ABL	-dɑn~don~den~døn, -tɑn~ton~ten~tøn	-dɑn, -tɑn	-dɑn~dien, -tɑn~tien	-dɯn~dun~din~dyn, -tɯn~tun~tin~tyn
ACC	-nɯ~nu~ni~ny, -dɯ~du~di~dy, -tɯ~tu~ti~ty	-nɯ	-nɯ~nu~ni~ny, -dɯ~du~di~dy, -tɯ~tu~ti~ty -n	-ni
POSS.3	-ɯ~u~i~y, -sɯ~su~si~sy	-ɯ, -sɯ	-ɯ~u~i~y, -sɯ~su~si~sy	-ɯ~u~i~y, -sɯ~su~si~sy
POSS.1S	-ɯm~um~im~ym, -m	-ɯm~um~im~ym, -m	-ɯ~u~i~y, -m	-ɯm~um~im~ym, -m

Uzbek differs from the other three languages in its lack of backness harmony, as evident in the lack of backness alternations above in Table 4.3. Although its genetic predecessor, Chagatai, had productive backness harmony, the standard variety of Uzbek no longer exhibits harmony (Bodrogligeti 2001; Eckmann 2017; Sjoberg 1963; Tursunov et al. 1992). Non-standard varieties of the language exhibit backness and rounding harmony to varying degrees (Ibrohimov 1967; Razhabov 1996), but the Osh dialect recorded during fieldwork displayed only lexeme-specific rounding harmony with no backness harmony. For the present study, Uzbek speakers serve as a control group with which to compare results from the other three languages, as well as the general validity of the methods employed.

4.2.3 Task

Each session was divided into training and recording phases. During the training phase, participants were taught a small set of pictorial-grammatical correspondences involving number, case, and possession. This phase typically lasted less than 5 minutes. After participants completed the training, the recording phase began. Throughout each session, participants were presented images on a laptop computer screen that prompted both a lexical item paired with at least one suffix from the training phase. Example prompts for the target Kyrgyz words /bɑltɑ-dɑ/ 'axe-LOC' and /bɑltɑ-lɑr-dɑ/ 'axe-PL-LOC' are shown in Figure 4.1. All target words were prompted individually, and the order of prompts was randomized. When speakers were unable to guess the target word from the prompt, they were given either the equivalent Russian word or a paraphrase in the target language. Sessions were conducted in a quiet room. Participants wore a Shure-SM10A unidirectional head-mounted microphone, and all data were recorded to a Marantz PMD 661 MKII digital recorder at a sampling rate of 44.1 kHz. Each session lasted between 45 and 90 minutes.

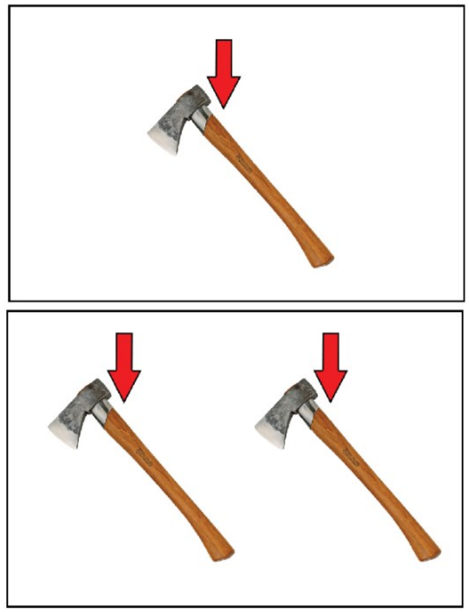

Figure 4.1: Example prompts for the target words /bɑltɑ-dɑ/ 'axe-LOC' and /bɑltɑ-lɑr-dɑ/ 'axe-PL-LOC.' The axe indicates the target root, with the number of axes indicating grammatical number. The downward arrow indicates locative case.

4.2.4 Measurement and statistical analysis

All sound files were segmented in Praat (Boersma & Weenink 2015) . The beginning and end of each vowel was set to the onset and offset of the second formant. In cases where the second formant persisted across flanking consonants (i.e. sonorants), abrupt shifts in the amount and distribution of spectral energy were used to indicate vowel onset and offset.

After segmentation, the first two formants (F1-F2) and vowel duration were measured. The first formant is a good indicator of vowel height, while the second formant is a reliable indicator of both lip rounding and vowel backness. Formants were measured at the 25, 50, 75% point in each vowel using a modified Praat script (Crosswhite 2003). The data were analyzed in R (R Core Team 2017), using the lme4 package (Bates et al. 2015). Formant measurements from the midpoint of each vowel were used in the analysis.

Outliers were then inspected for measurement errors. In particular, a number of errors were found with the high back vowels, where the formant tracker in Praat failed to distinguish the first two formants. In these cases, formant frequencies were hand measured at the approximate vowel midpoint. To facilitate across-speaker comparisons, the data were z-score normalized (Lobanov 1971). The data for normalization consisted of four tokens of each vowel and were taken from monosyllabic words. If four tokens of a given vowel were not present in monosyllables, then the remaining tokens were taken from the initial syllable of disyllabic tokens.

A mixed effect linear regression was used to predict normalized F2 for each language based on the following fixed effects: initial vowel backness, target vowel height, syllable number (counting from the left), root type (mono- or disyllabic), preceding consonant place of articulation, and following consonant place of articulation. Since laterals alternate for vowel backness in all four languages, laterals flanking back vowels were coded as dorsal while laterals flanking front vowels were coded as coronal. Additionally, since backness harmony does not operate in Uzbek, vowels were coded for target backness instead of initial vowel backness. The model incorporated two-way interactions between: preceding consonant place of articulation and target vowel height, following consonant place of articulation and target vowel height, initial vowel backness and target vowel height, initial vowel backness and syllable number, initial vowel backness and root type, and syllable number and root type. The model also incorporated a three-way interaction for initial vowel backness, syllable number, and root type. The model included random intercepts for speaker and target vowel, as well as by-syllable random slopes for speaker and target vowel. More elaborated random effect structure resulted in models that failed to converge. Statistical significance was assessed using likelihood ratio tests.

The key predictors to observe are the main effect of syllable number and the interaction between initial vowel backness and syllable number. Positional shifts for front vowels are observable from the main effect alone, whereas positional shifts for back vowels require examining both the main effect and interaction terms. Throughout the results section and analysis, harmonic pairings are compared. Each language's harmonic pairings are summarized in Table 4.4. Kyrgyz exhibits four harmonic pairings. Uzbek exhibits only one, for rounding harmony. Kazakh exhibits three pairings, and Uyghur exhibits three generally. During data collection, raised vowels (see Chapter 2 for more details) were not impressionistically identical to underlying high vowels, so raised vowel alternations are analyzed separately. Raised vowels are indicated by a subscript.

Table 4.4: Language-specific harmonic pairings

	Kyrgyz	Uzbek	Kazakh	Uyghur
[-high, -round]	e-ɑ	N/A	ie-ɑ	æ-ɑ
[-high, +round]	ø-o	N/A	N/A	N/A
[+high, -round]	i-ɯ	N/A	ɪ-ɯ	i-ɯ, i$_R$-ɯ$_R$
[+high, +round]	y-u	y-u	ʏ-ʊ	y-u, y$_R$-u$_R$

4.2.5 Predictions

Little work has been done on the positional realization of harmonic vowels, but Zsiga (1997:234-235) predicts that alternating vowels should be acoustically indistinguishable from non-alternating vowels. If the targets of harmony, i.e. non-initial syllables, are acoustically indistinguishable from initial-syllable trigger vowels, we can conclude that harmony operates categorically in the language, as is found for Turkish (Lanfranca 2012; Sabev & Payne 2018). Alternatively, it is possible that both classes of [back] vowels exhibit symmetric shifts toward a more central F2 in non-initial positions. This type of centralization is found in many languages (Vayra & Fowler 1992; Johnson & Martin 2001). Lastly, it is possible that only one class of [back] vowels exhibits a positional shift in F2. This type of asymmetric shift is attested in related Kazakh (McCollum 2015; McCollum & Chen 2019). Each of these patterns and their phonetic and phonological analyses are schematized in Table 4.5.

Table 4.5: Schema for possible patterns and their analysis. Each pattern is paired with at least one analysis, which involves a phonological and phonetic component. Dashed lines indicate front vowel shifts, while solid lines indicate back vowel shifts.

	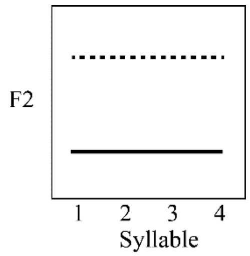	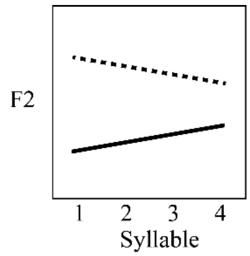	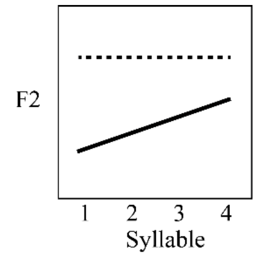		
Pattern	**No positional shifts**	**Symmetric shifts**	**Asymmetric shifts**		
Analysis	Categorical harmony	Gradient centralization	Under-specification	Adaptive hypo-articulation	Gradient phonology
Phonology	Categorical harmony	Categorical harmony	No phonological harmony	Categorical harmony	Gradient harmony
Phonetics	No reduction	Centralization	Interpolation across underspecified syllables	Reduction of effort in predictable contexts	No reduction

Of the three possible patterns, analyzing the first two is straightforward– if no positional shifts occur, then we can conclude that Zsiga's prediction is borne out. In this scenario, harmony operates categorically, also blocking phonetic reduction (Pearce 2008). If non-initial vowels undergo relatively symmetric shifts toward a central F2, such a pattern is analyzable as categorical phonology with gradient phonetic implementation in the form of centralization. For the third pattern, three possible analyses are sketched out below.

One account of asymmetric positional shifts involves interpolation across underspecified segments or syllables (see Cohn 1993 for an analysis). Under this analysis, speakers gradiently transition from [+back] initial-syllable vowels to a [-back] default articulatory setting at the end of the word. This analysis depends on two claims. First, phonological specifications are only provided for the initial syllable and the default articulatory setting. As a result, all non-initial syllables do not receive a phonological specification for [back]. Default articulatory settings have been found at the ends of larger prosodic units (Gick et al. 2004; Wilson & Gick 2006)., and since words were recorded in isolation, the end of the word is also the end of the utterance Also, studies on languages with vowel harmony have found that the default articulatory settings in those languages typically correspond to the unmarked feature value in the harmony pattern (Hudu 2010; Allen et al. 2013). Under this analysis, harmony is not phonological, but a pattern of phonetic implementation (Keating 1988, 1990; Cohn 1993).

A second possible account involves what Lindblom (1990) calls adaptive hypoarticulation. Under this analysis, speakers utilize listener expectations to economize speech production. Since these languages, excluding Uzbek, exhibit harmony, a listener should expect the backness of the initial vowel to predict the backness of all following vowels. If a speaker uses this expectation, then a speaker can reduce more effortful articulations without losing efficient transmission of information. Since more effortful articulations are reduced, this analysis can predict reduction of one class of [back] vowels, assuming that they are in some sense more effortful, or in more phonological terms, more marked (see Harrington et al. 2011 for evidence supporting the effortfulness of /u/).

Finally, a third way to analyze asymmetric gradience is to allow the gradience to fall out from phonological rather than phonetic factors. Under this analysis, [back] spreading subphonemically peters out over the course of the word, leaving each subsequent vowel slightly less assimilated to the active feature value. A number of impressionistic descriptions suggest that harmony may be gradient, like Stegen (2002:137), which reports that ATR harmony in Rangi “seems to be gradient, i.e. having diminished effect with increase distance from the [+ATR] spreading vowel” (see Kirchner 1998:320-321; Casali 2002:25 for other similar descriptions).

4.3 Results

4.3.1 Kyrgyz

Vowel quality by position is presented in Figure 4.2 and Tables 4.5-4.6 below. Figure 4.2 presents mean vowel quality by position and Tables 4.5-4.6 provides additional information, including standard deviation, and token count (total n = 9,374).

In Figure 4.2, the vowel space contracts by position, both in terms of F1 and F2. Since F2 is the main acoustic correlate of tongue backness, we focus on that here. F2 of front vowels decreases by position [Syllable: β= -0.07, $\chi^2(1)$= 4.63, p= .03] while F2 of back vowels increases by position [Backness : Syllable: β= 0.11, $\chi^2(1)$= 5.52, p= .02]. The magnitude of each of these shifts is relatively small, -0.07z for front vowels and 0.04z for back vowels (main effect + interaction, -0.07z + .11z). Despite exhibiting relatively small shifts, these positional shifts in F2 are significant due to the systematicity of each vowel's positional shift, which is seen in Tables 4.6 and 4.7.

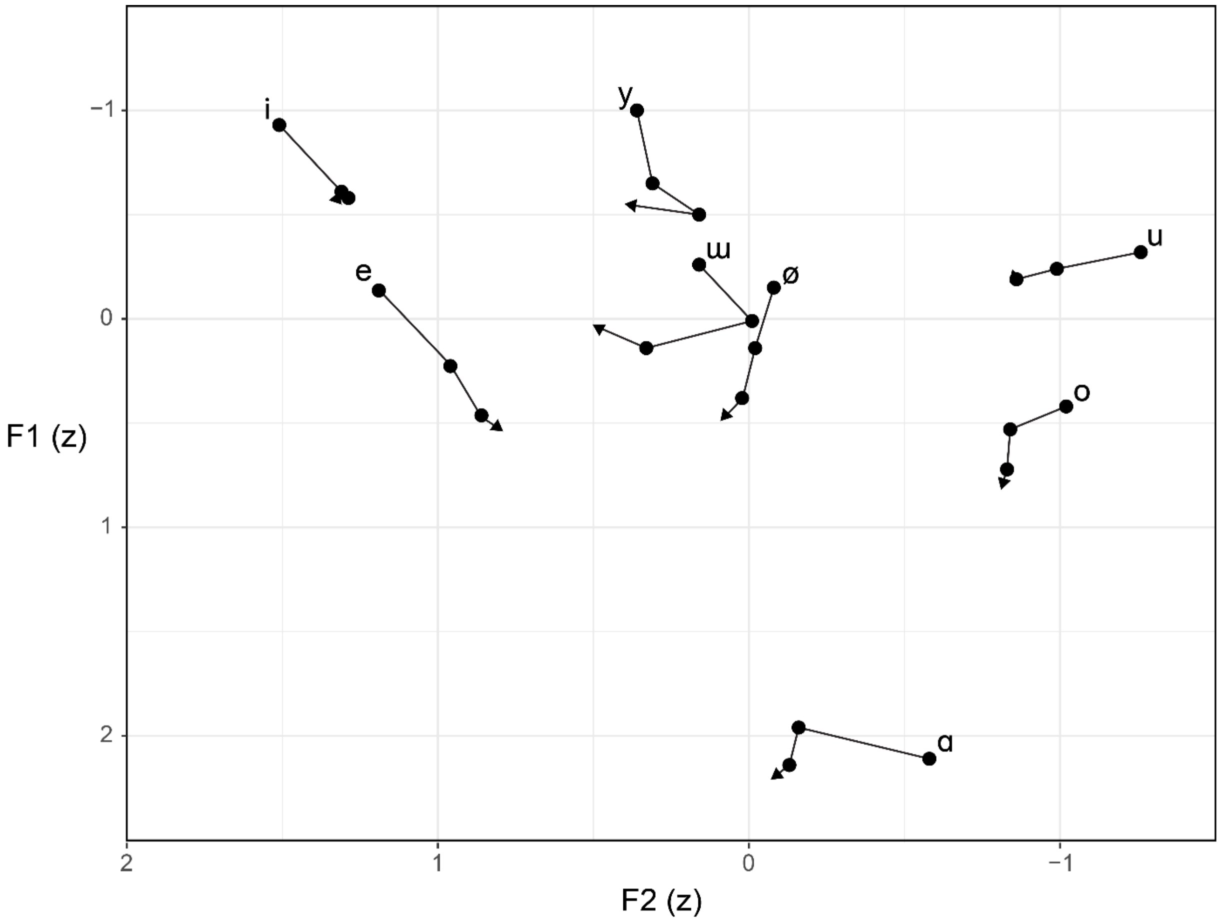

Figure 4.2: Kyrgyz mean vowel F1-F2 (z) by position. Lines connect different mean F1-F2 for a particular vowel. The initial-syllable mean is closest to the transcription and the arrow indicates the final mean.

Table 4.6: Mean F1 (in z, with SD) for Kyrgyz

	Syll. 1		Syll. 2		Syll. 3		Syll. 4	
	F1	n	F1	n	F1	n	F1	n
ɑ	2.11 (0.50)	514	1.96 (0.55)	644	2.14 (0.52)	420	2.21 (0.49)	62
o	0.42 (0.48)	438	0.53 (0.44)	327	0.72 (0.41)	156	0.82 (0.31)	27
ɯ	-0.26 (0.42)	587	0.01 (0.50)	496	0.14 (0.59)	231	0.03 (0.31)	38
u	-0.32 (0.45)	276	-0.24 (0.56)	239	-0.19 (0.49)	106	-0.22 (0.39)	13
e	-0.14 (0.33)	550	0.22 (0.42)	535	0.47 (0.37)	341	0.54 (0.36)	64
ø	-0.15 (0.30)	423	0.14 (0.31)	406	0.38 (0.39)	254	0.49 (0.38)	65
i	-0.93 (0.38)	423	-0.58 (0.41)	385	-0.57 (0.45)	224	-0.61 (0.34)	38
y	-1.00 (0.37)	425	-0.65 (0.37)	393	-0.50 (0.40)	228	-0.55 (0.29)	46
Total		3,636		3,425		1,960		353

Table 4.7: Mean F2 (in z, with SD) for Kyrgyz

	Syll. 1		Syll. 2		Syll. 3		Syll. 4	
	F2	n	F2	n	F2	n	F2	n
ɑ	-0.58 (0.30)	514	-0.16 (0.28)	644	-0.13 (0.25)	420	-0.07 (0.22)	62
o	-1.02 (0.32)	438	-0.84 (0.23)	327	-0.83 (0.23)	156	-0.81 (0.25)	27
ɯ	0.16 (0.40)	587	-0.01 (0.43)	496	0.33 (0.55)	231	0.50 (0.36)	38
u	-1.26 (0.36)	276	-0.99 (0.35)	239	-0.86 (0.35)	106	-0.85 (0.29)	13
e	1.19 (0.32)	550	0.96 (0.30)	535	0.86 (0.37)	341	0.79 (0.48)	64
ø	-0.08 (0.29)	423	-0.02 (0.22)	406	0.02 (0.21)	254	0.09 (0.20)	65
i	1.51 (0.32)	423	1.29 (0.34)	385	1.35 (0.40)	224	1.31 (0.30)	38
y	0.36 (0.63)	425	0.31 (0.31)	393	0.16 (0.48)	228	0.40 (0.18)	46
Total		3,636		3,425		1,960		353

Figure 4.3 compares positional F2 by harmonic pairing. The /a-e/ pairing shows the clearest pattern of symmetrical centralization. The trends in the other three pairings are less obvious. The /ɯ-i/ pairing appears to exhibit centralization, albeit to a lesser extent and with more positional variation of /ɯ/. The round vowel pairings, /o-ø/ and /u-y/, show a slight tendency toward higher F2 for both front and back vowels.

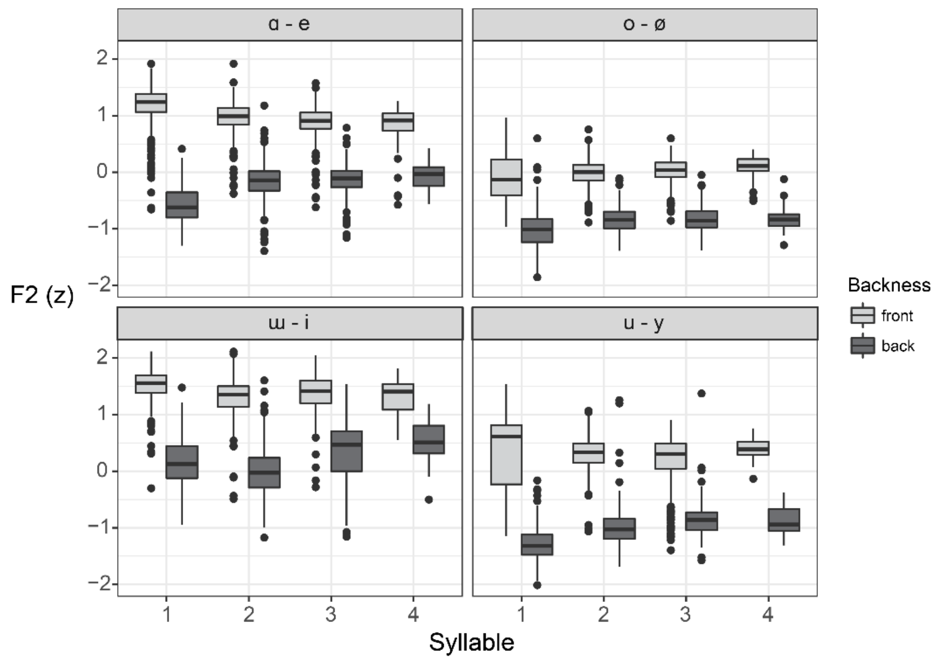

Figure 4.3: Kyrgyz F2 (z) by position and harmonic pairing

In addition to position, morphology exerts a significant effect on vowel F2 in the data. Front vowels exhibit higher F2 in words derived from disyllabic roots [Root type: β= 0.33, $\chi^2(1)$= 144.07, p< .001], while back vowels exhibit lower F2 in words derived from disyllabic roots [Backness : Root type: β= -0.45, $\chi^2(1)$= 148.0, p< .001]. This is evident in Figure 4.4 below. In Figure 4.4, the second and third syllables in words derived from monosyllabic roots are suffixes. In words formed from disyllabic roots, though, the second-syllable vowel is root-internal. This effect of morphology is most evident in the high front vowels as well as the non-high back vowels, which exhibit more peripheral F2 values than their counterparts in words derived from monosyllabic roots. However, more peripheral vowel qualities in words formed from disyllabic roots is localized toward the left edge of the word, diminishing in non-initial syllables for both front and back vowels [Syllable : Root type: β= -0.07, $\chi^2(1)$= 35.22, $p< .001$; Backness : Syllable : Root type: β= 0.10, $\chi^2(1)$= 36.96, $p< .001$]. In other words, the initial syllables of words formed from disyllabic roots are produced with more peripheral vowel qualities. Consonantal context exerts a range of effects, with the model predicting that flanking consonants typically depress vowel F2. See Table 4.8 for full output.

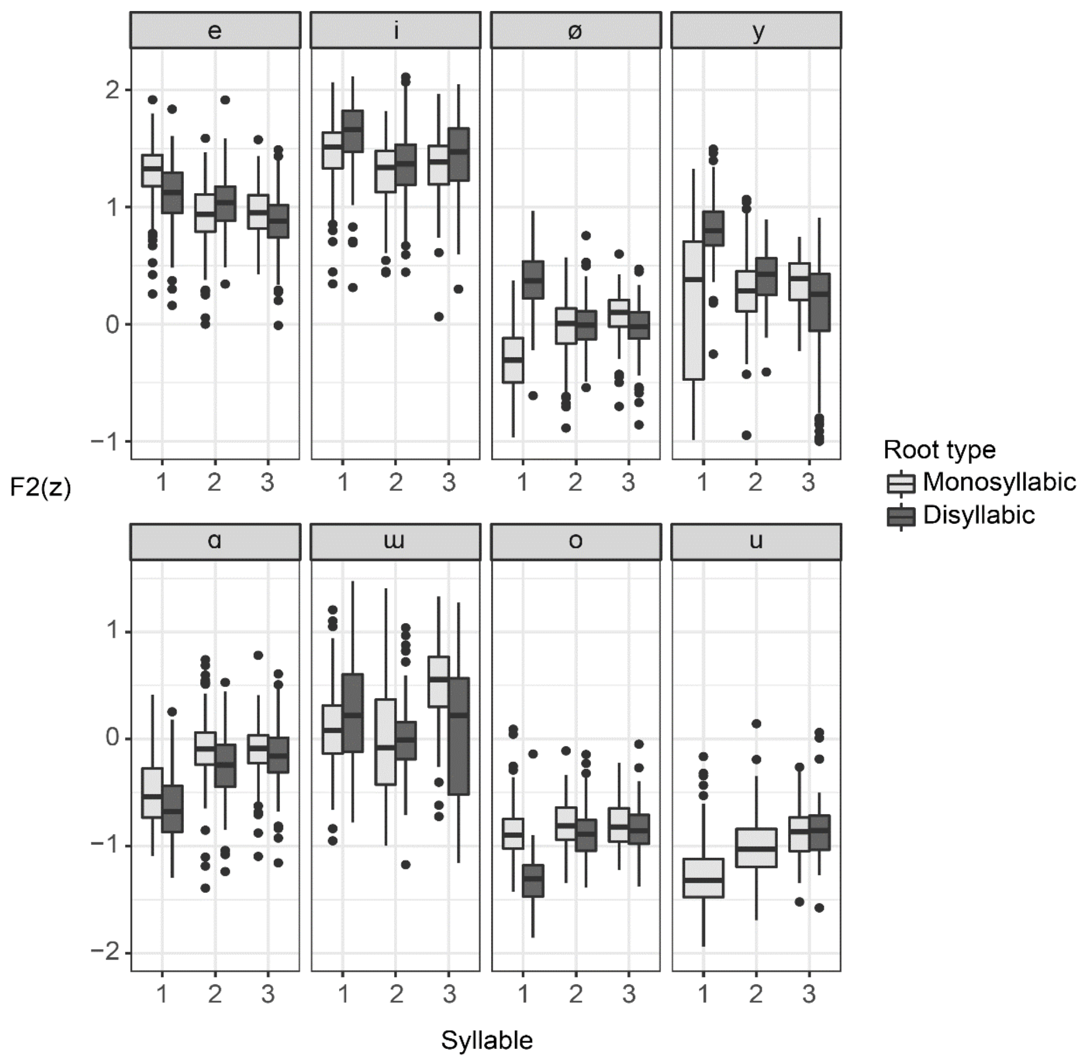

Figure 4.4: F2 (z, with SD) in syllables 1-3 based on root type in Kyrgyz

Table 4.8: Model output for Kyrgyz

Random effects					
			Variance	SD	
	Speaker	(Intercept)	0.004	0.06	
		Syllable	0.001	0.03	
	Target vowel	(Intercept)	0.59	0.77	
		Syllable	0.00	0.06	
	Residual		0.1	0.31	
Fixed effects					
		β	SE	df	t-value
	(Intercept)	0.90	0.48	6	1.88
	Backness	-0.88	0.67	6	-1.31
	Target height	-0.20	0.56	4	-0.36
	Syllable	-0.07	0.03	8	-2.11
	Root type	0.33	0.03	9286	12.03
	Prec.C.Bilabial	-0.32	0.03	9273	-9.81
	Prec.C.Coronal	-0.23	0.03	8848	-7.10
	Prec.C.Dorsal	-0.48	0.04	8242	-13.63
	Foll.C.Bilabial	-0.01	0.03	9321	-3.30
	Foll.C.Coronal	-0.03	0.01	9207	-1.89
	Foll.C.Dorsal	-0.22	0.02	9212	-11.98
	Backness : Target height	-0.71	0.79	4	-0.90
	Backness : Syllable	0.11	0.05	7	2.41
	Syllable : Root type	-0.07	0.01	9307	-5.93
	Backness : Root type	-0.45	0.04	9290	-12.19
	Target height : Prec.C.Bilabial	0.47	0.04	8961	10.67
	Target height : Prec.C.Coronal	0.58	0.04	8935	14.06
	Target height : Prec.C.Dorsal	0.25	0.05	8452	5.63
	Target height : Foll.C.Bilabial	-0.18	0.03	8107	-5.31
	Target height : Foll.C.Coronal	0.04	0.03	1831	1.26
	Target height : Foll.C.Dorsal	0.09	0.04	2625	2.54
	Backness : Syllable : Root type	0.10	0.02	9297	6.06

In sum, Kyrgyz vowels are realized with more centralized acoustic qualities in non-initial positions. Positional variation also interacts with morphology, as vowel qualities are more peripheral in words derived from disyllabic roots.

4.3.2 Uzbek

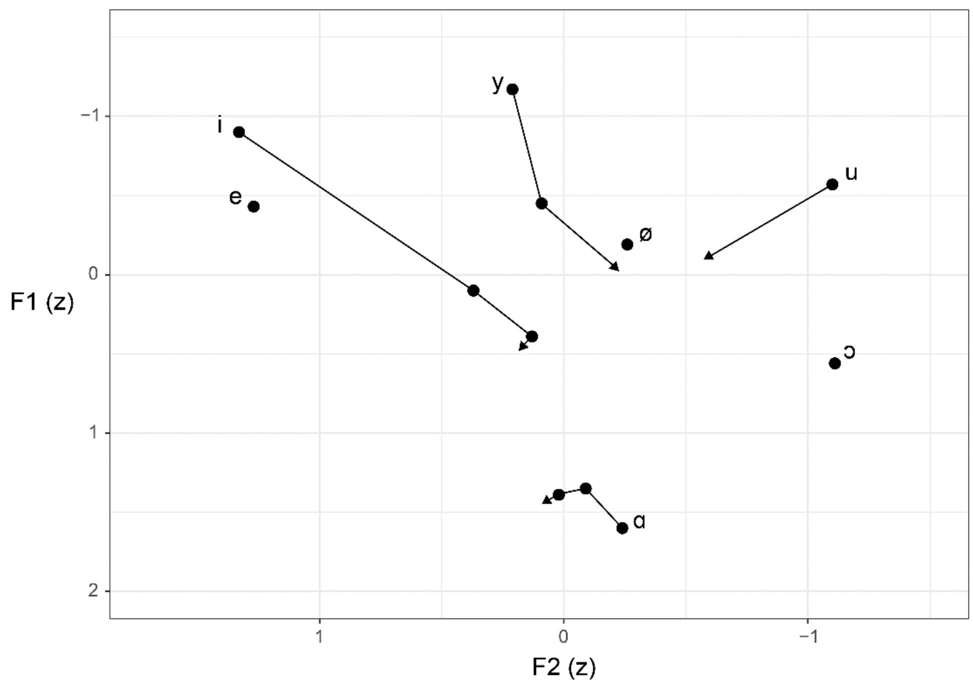

Figure 4.5: Uzbek mean vowel F1-F2 (z) by position. Lines connect different mean F1-F2 for a particular vowel. The initial-syllable mean is closest to the transcription and the arrow indicates the final mean.

Figure 4.5 and Tables 4.9-4.10 present positional vowel quality in Uzbek. Figure 4.5 presents mean vowel quality by position and Tables 4.9-4.10 provide additional information, including standard deviation, and token count (total n = 5,634). As in Kyrgyz, there is a noticeable increase in F1 by position, but also a tendency toward more central F2 values in non-initial syllables. This shift toward more central F2 is not significant for the front vowels, but the back vowels exhibit a significant pattern of fronting [Syllable: β= -0.01, $\chi^2(1)$= 0.02, p= .90; Backness : Syllable: β= 0.21, $\chi^2(1)$= 57.77, p< .001].

This appears to suggest that Uzbek vowels are asymmetrically fronted in non-initial syllables, but when Figures 4.5-4.6, along with Tables 4.11-4.12 are consulted, a different generalization emerges.

Table 4.9: Mean F1 (in z, with SD) for Uzbek

	Syllable 1		Syllable 2		Syllable 3		Syllable 4	
	F1	n	F1	n	F1	n	F1	n
ɑ	1.60 (0.57)	480	1.35 (0.49)	1,232	1.39 (0.46)	835	1.43 (0.48)	153
ɔ	0.56 (0.51)	398						
u	-0.57 (0.46)	224	-0.10 (0.50)	129				
e	-0.43 (0.32)	213						
ø	-0.19 (0.28)	302						
i	-0.90 (0.46)	548	0.10 (0.64)	800	0.39 (0.65)	609	0.47 (0.58)	115
y	-1.17 (0.37)	240	-0.45 (0.53)	232	-0.04 (0.63)	37		
Total		2,405		2,393		1,481		268

Table 4.10: Mean F2 (in z, with SD) for Uzbek

	Syllable 1		Syllable 2		Syllable 3		Syllable 4	
	F2	n	F2	n	F2	n	F2	n
ɑ	-0.24 (0.39)	480	-0.09 (0.29)	1,232	0.02 (0.29)	835	0.07 (0.32)	153
ɔ	-1.11 (0.27)	398						
u	-1.10 (0.66)	224	-0.58 (0.48)	129				
e	1.27 (0.46)	213						
ø	-0.26 (0.35)	302						
i	1.33 (0.58)	548	0.37 (0.50)	800	0.13 (0.59)	609	0.18 (0.73)	115
y	0.21 (0.56)	240	0.09 (0.32)	232	-0.22 (0.44)	37		
Total		2,405		2,393		1,481		268

First, the front vowels exhibit far more variation in F2 than the back vowels, as seen in Table 4.10 above. The standard deviations for F2 of front vowels exceed those found among the back vowels in every position. This variation prevents any systematic generalization from emerging for the front vowels. Second, this variation interacts significantly with morphology in the language. Among the front vowels, words derived from disyllabic roots are produced with significantly lower F2 [Root type: β= -0.42, $\chi^2(1)$= 61.41, $p< .001$; Syllable: Root type: β= -0.18, $\chi^2(1)$= 65.4, $p< .001$]. Root type exerts no effect on back vowels [Backness : Root type: β= 0.40, $\chi^2(1)$= 27.30, $p< .001$], though, since the interaction of backness and root type nullifies the main effect of root type (-0.40 + 0.40 = -0.02). For back vowels, F2 of vowels derived from disyllabic roots was lower than F2 of vowels derived from monosyllabic roots in non-initial syllables [Backness : Syllable : Root type: β= -0.16, $\chi^2(1)$= 26.91, $p< .001$].

The model predicts that F2 of front vowels does not decrease by position, but F2 of back vowels significantly increases by position. However, F2 of front vowels decreases significantly in words derived from disyllabic roots, while F2 of back vowels derived from disyllabic roots exhibit no noticeable shifts. These are evident in Table 4.10, which presents the full statistical model for Uzbek. In one case, there is asymmetric fronting and in the other there is asymmetric backing. In Figure 4.6, though, the distribution of F2 in terms of root type suggests only small effects of morphology. I conjecture that this particular result relates specifically to the status of /i/ in the language. Consider the realization of /i/ in mono- and disyllabic roots, shown in Figure 4.7. This vowel typically exhibits higher F2 in monosyllabic roots, even in consonantal contexts that are very similar to comparable disyllabic roots. In second syllables of disyllabic roots, F2 of /i/ is drastically reduced, with a vowel quality approximating [ə] in most forms, with variation across forms that is challenging to account for. The realization of /i/ in these syllables, along with its general backing to [ə] in non-initial positions may indicate that the historical loss of *ɯ is overstated in the literature (e.g. Harrison et al. 2002). Perhaps, the complex set of allophones described in Sjoberg (1963) reflect an incomplete merger of *ɯ and *i. If /ɯ/ is still present in the language, even in a reduced capacity, this would significantly affect the interpretation of the Uzbek data.

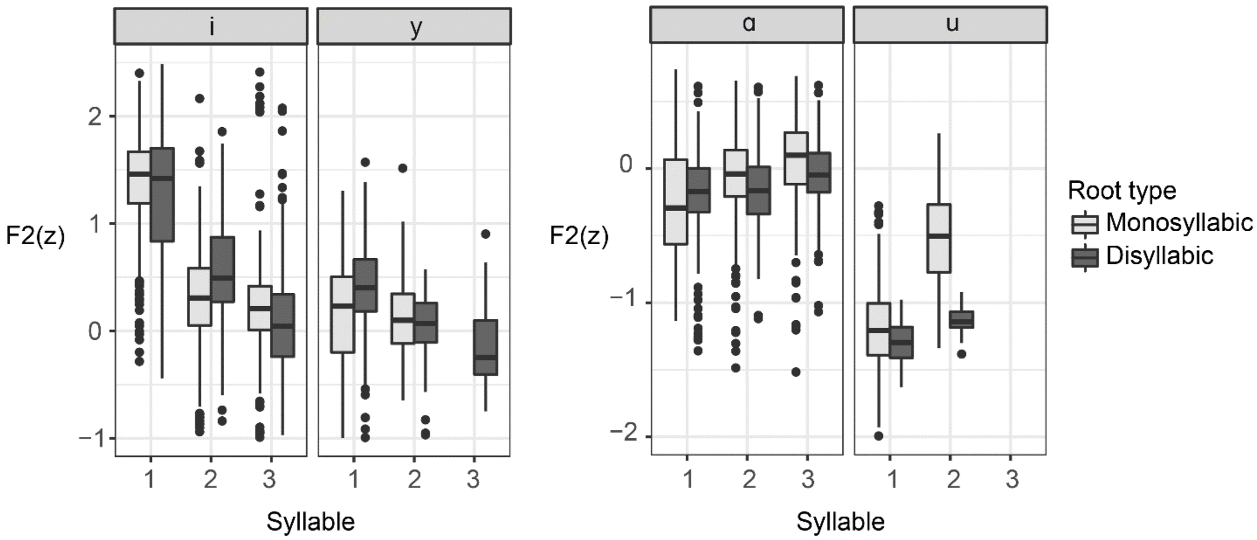

Figure 4.6: F2 (z, with SD) in syllables 1-3 based on root type in Uzbek

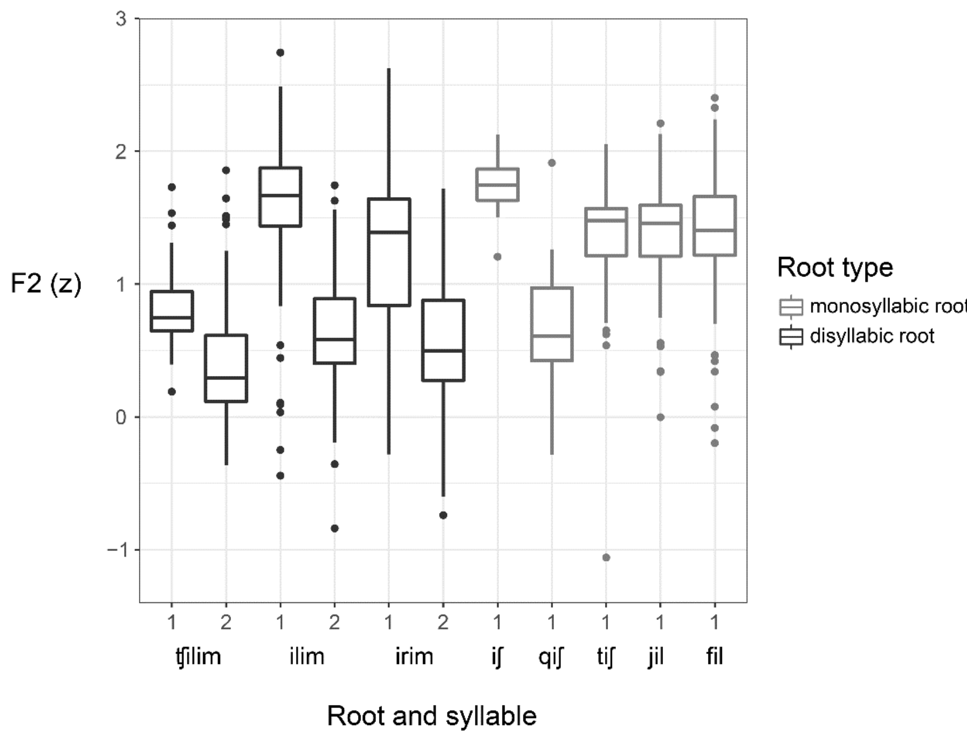

Figure 4.7: F2 (z) of root-internal /i/ in Uzbek by lexeme and syllable

Table 4.11: Model output for Uzbek

Random effects					
			Variance	SD	
	Speaker	(Intercept)	0.02	0.13	
		Syllable	0.00	0.06	
	Target vowel	(Intercept)	0.94	0.97	
		Syllable	0.11	0.33	
	Residual		0.18	0.42	
Fixed effects					
		β	SE	df	t-value
	(Intercept)	0.08	0.50	3	0.16
	Backness	-0.55	0.07	4388	-8.36
	Target height	0.47	0.10	2	4.52
	Syllable	-0.01	0.16	3	-0.05
	Root type	-0.42	0.05	5186	-7.74
	Prec.C.Bilabial	-0.27	0.06	5595	-4.97
	Prec.C.Coronal	-0.03	0.06	5599	-0.59
	Prec.C.Dorsal	-0.18	0.06	5599	-3.02
	Foll.C.Bilabial	-0.30	0.05	5591	-5.67
	Foll.C.Coronal	-0.12	0.02	5594	-5.44
	Foll.C.Dorsal	-0.32	0.04	5598	-8.96
	Backness : Target height	-0.21	0.03	5416	-6.93
	Backness : Syllable	0.21	0.03	4052	7.47
	Syllable : Root type	0.18	0.02	5357	7.99
	Backness : Root type	0.40	0.07	5400	5.51
	Target height : Prec.C.Bilabial	-0.29	0.07	3456	-4.10
	Target height : Prec.C.Coronal	-0.31	0.07	5527	-4.68
	Target height : Prec.C.Dorsal	-0.38	0.08	5575	-4.95
	Target height : Foll.C.Bilabial	0.26	0.06	4227	4.55
	Target height : Foll.C.Coronal	0.51	0.05	5584	11.04
	Target height : Foll.C.Dorsal	0.38	0.08	4810	4.66
	Backness : Syllable : Root type	-0.16	0.03	5253	-5.10

To investigate what effects fall out from the variation in /i/, I excluded all root-internal /i/ and reran the model (779 vowels; new n= 4,855), which is reported in Table 4.12. All other aspects of the model were held constant. In this model, both classes of [back] vowels exhibit non-significant shifts in vowel quality [Syllable: β= 0.05, $\chi^2(1)$= 0.30, p= .58; Backness : Syllable: β= 0.02, $\chi^2(1)$= 0.64, p= .42]. Moreover, in this model morphology plays no obvious role in vowel backness. This holds for both front vowels [Root type: β= -0.00, $\chi^2(1)$= 0.002, p= .96; Syllable: Root type: β= 0.01, $\chi^2(1)$= 0.05, p= .83], and back vowels [Backness : Root type: β= -0.04, $\chi^2(1)$= 0.33, p= .56; Backness : Syllable : Root type: β= 0.02, $\chi^2(1)$= 0.18, p= .67].

Where the original model predicted asymmetric fronting of back vowels and asymmetric backing of front vowels in words derived from disyllabic roots, this model predicts no significant effect of position or morphology. The numerous differences between these two models suggest that root-internal /i/ played a really large role in the conflicting generalizations in Table 4.10. When the general trends in Figure 4.4 are considered, the data strongly suggests that back vowels are fronted and front vowels are backed in non-initial positions. From this point on, I will interpret the Uzbek results as reflecting a general tendency towards centralization. It is worth noting that the status of /i/ and historical *ɯ may bear on the analysis of Uzbek, but without more data, I leave that to future work.

Table 4.12: Model output for Uzbek with no root-internal /i/

Random effects					
			Variance	SD	
	Speaker	(Intercept)	0.01	0.09	
		Syllable	0.00	0.06	
	Target vowel	(Intercept)	0.49	0.70	
		Syllable	0.04	0.20	
	Residual		0.15	0.39	
Fixed effects					
		β	SE	df	t-value
	(Intercept)	-0.06	0.32	4	-0.18
	Backness	-0.10	0.07	4212	-1.50
	Target height	0.14	0.09	2835	1.59
	Syllable	0.05	0.10	5	0.50
	Root type	-0.00	0.09	4820	-0.04
	Prec.C.Bilabial	-0.29	0.05	4820	-5.55
	Prec.C.Coronal	-0.02	0.05	4822	-0.31
	Prec.C.Dorsal	-0.19	0.06	4822	-3.28
	Foll.C.Bilabial	-0.28	0.05	4814	-5.52
	Foll.C.Coronal	-0.10	0.02	4819	-4.50
	Foll.C.Dorsal	-0.30	0.03	4823	-8.82
	Backness : Target height	-0.12	0.03	4275	-3.75
	Backness : Syllable	0.02	0.03	4182	0.79
	Syllable : Root type	0.01	0.03	4756	0.20
	Backness : Root type	-0.06	0.01	4818	-0.59
	Target height : Prec.C.Bilabial	-0.09	0.09	3632	-1.01
	Target height : Prec.C.Coronal	-0.00	0.08	4635	-0.13
	Target height : Prec.C.Dorsal	0.03	0.09	4799	0.22
	Target height : Foll.C.Bilabial	0.21	0.06	4674	3.95
	Target height : Foll.C.Coronal	0.34	0.05	4594	7.11
	Target height : Foll.C.Dorsal	0.43	0.08	3782	0.79
	Backness : Syllable : Root type	0.02	0.04	4807	0.20

4.3.3 Kazakh

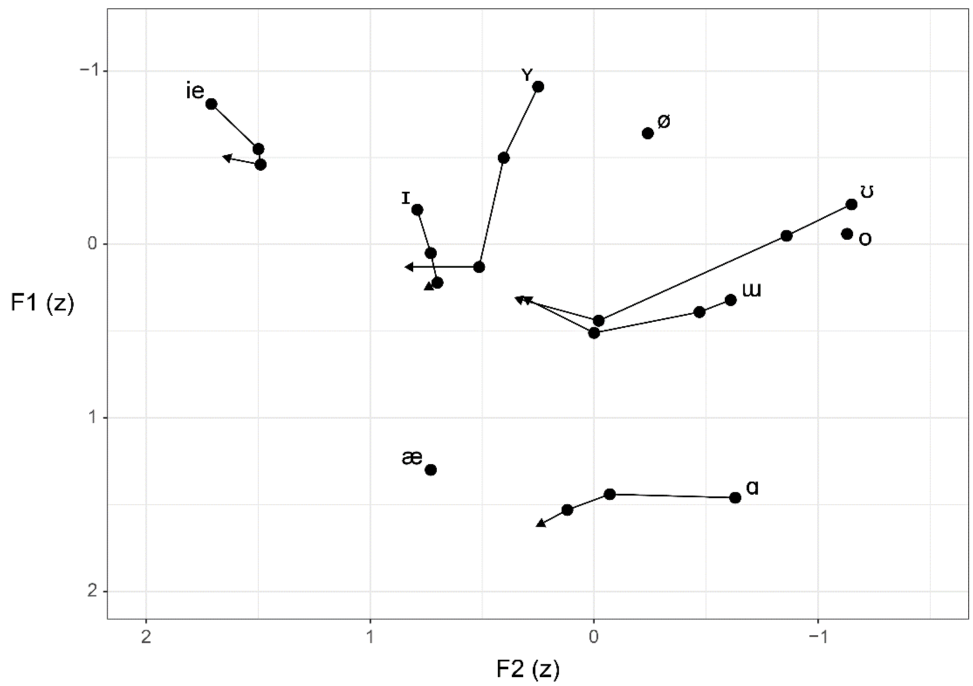

Figure 4.8: Kazakh mean vowel F1-F2 (z) by position. Lines connect different mean F1-F2 for a particular vowel. The initial-syllable mean is closest to the transcription and the arrow indicates the final mean.

Vowel quality by position is presented in Figure 4.8 above and Tables 4.13-4.14 below. Figure 4.8 presents mean vowel quality by position and Tables 4.13-4.14 provide additional information, including standard deviation, and token count (total n = 5,342). Unlike Kyrgyz and Uzbek, Kazakh vowels display an asymmetric pattern of positional variation. The back vowels are marked by increasing F2 by position while the front vowels exhibit no obviously significant effects of position. The slight positional decrease in F2 for the front vowels is not significant [Syllable: β= -0.05, $\chi^2(1)$= 0.87, p= .35]. However, the back vowels exhibit a much larger shift, which is statistically significant [Backness : Syllable: β= 0.28, $\chi^2(1)$= 7.39, p= .007]. Thus, positional changes in F2 are asymmetric in Kazakh, with the magnitude of the back vowel shift, 0.23z (main effect + interaction, -0.05z + 0.28z) being over four times larger than the magnitude of the front vowel shift, -0.05z.

Table 4.13: Mean F1 (in z, with SD) for Kazakh

	Syllable 1		Syllable 2		Syllable 3		Syllable 4	
	F1	n	F1	n	F1	n	F1	n
ɑ	1.46 (0.50)	277	1.44 (0.50)	519	1.53 (0.40)	345	1.62 (0.39)	64
o	-0.06 (0.46)	276						
ɯ	0.32 (0.79)	264	0.39 (0.60)	231	0.51 (0.54)	167	0.31 (0.41)	35
ʊ	-0.23 (0.47)	264	-0.05 (0.57)	236	0.44 (0.36)	61	0.31 (0.42)	6
ie	-0.81 (0.29)	258	-0.55 (0.36)	724	-0.46 (0.53)	464	-0.50 (0.33)	83
ø	-0.64 (0.26)	277						
ɪ	-0.20 (0.56)	212	0.05 (0.50)	293	0.22 (0.43)	209	0.24 (0.39)	53
ʏ	-0.91 (0.29)	234	-0.50 (0.54)	258	0.13 (0.59)	76	0.13 (0.29)	9
æ	1.30 (0.49)	281						
Total		2,343		2,261		1,322		250

Table 4.14: Mean F2 (in z, with SD) for Kazakh

	Syllable 1		Syllable 2		Syllable 3		Syllable 4	
	F2	n	F2	n	F2	n	F2	n
ɑ	-0.63 (0.39)	277	-0.07 (0.35)	519	0.12 (0.29)	345	0.25 (0.18)	64
o	-1.13 (0.47)	276						
ɯ	-0.61 (0.41)	264	-0.47 (0.47)	231	0.00 (0.54)	167	0.31 (0.20)	35
ʊ	-1.15 (0.30)	264	-0.86 (0.49)	236	-0.02 (0.58)	61	0.35 (0.12)	6
ie	1.71 (0.38)	258	1.50 (0.40)	724	1.49 (0.50)	464	1.65 (0.30)	83
ø	-0.24 (0.41)	277						
ɪ	0.79 (0.34)	212	0.73 (0.38)	293	0. 70 (0.39)	209	0.73 (0.42)	53
ʏ	0.25 (0.59)	234	0.44 (0.30)	258	0.52 (0.46)	76	0.84 (0.24)	9
æ	0.73 (0.41)	281						
Total		2,343		2,261		1,322		250

Figure 4.9 presents result by harmonic pairing. In all three harmonic pairings, the F2 of back vowels increases, and for the high vowels almost approximates that of their front vowel counterparts. In all three pairing, positional shifts are asymmetric, in contrast to the shifts found in Kyrgyz. The full statistical model is reported in Table 4.15.

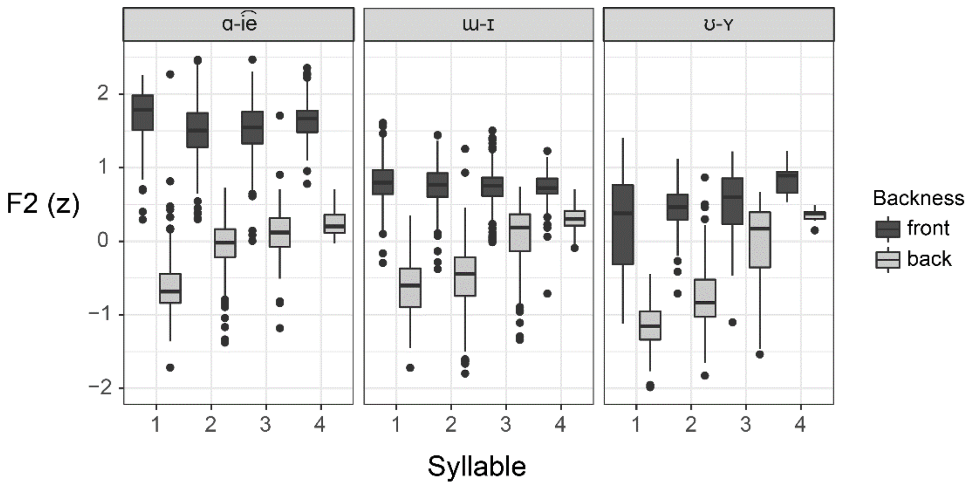

Figure 4.9: Kazakh F2 (z) by position and harmonic pairing

Morphology plays a relatively minor role in the realization of backness in Kazakh. F2 is slightly but not significantly higher for both front and back vowels in words derived from disyllabic roots [Root type: β= 0.07, χ2(1)= 3.37, p= .07; Backness : Root type: β= 0.06, χ2(1)= 0.86, p= .35]. Similarly, root type has little effect on the realization of non-initial vowels [Syllable : Root type: β= -0.01, χ2(1)= 0.25, p= .61]. However, there is a positional effect of root type on [+back] words. Although the first-syllable vowel is more centralized in disyllabic roots with [+back] vowels, subsequent vowels are more peripheral Backness : Syllable : Root type: β= -0.08, χ2(1)= 9.26, p= .002]. F2 by root type for the first three syllables is presented in Figure 4.10.

Table 4.15: Model output for Kazakh

Random effects					
			Variance	SD	
	Speaker	(Intercept)	0.01	0.08	
		Syllable	0.00	0.07	
	Target vowel	(Intercept)	0.13	0.35	
		Syllable	0.01	0.11	
	Residual		0.12	0.35	
Fixed effects					
		β	SE	df	t-value
	(Intercept)	2.13	0.22	5	9.75
	Backness	-2.08	0.30	4	-6.99
	Target height	-1.40	0.08	6	-16.98
	Syllable	-0.05	0.07	5	-0.78
	Root type	0.07	0.04	5304	1.80
	Prec.C.Bilabial	-0.43	0.05	5306	-8.94
	Prec.C.Coronal	-0.39	0.05	5270	-8.03
	Prec.C.Dorsal	-0.55	0.07	5271	-7.79
	Foll.C.Bilabial	-0.21	0.04	5303	-5.54
	Foll.C.Coronal	-0.24	0.02	5269	-12.80
	Foll.C.Dorsal	-0.73	0.04	5239	-19.75
	Backness : Target height	0.61	0.08	1	7.72
	Backness : Syllable	0.28	0.09	4	3.06
	Syllable : Root type	-0.01	0.02	5305	-0.45
	Backness : Root type	0.06	0.06	5303	0.93
	Target height : Prec.C.Bilabial	0.23	0.06	5306	3.78
	Target height : Prec.C.Coronal	0.62	0.06	5271	10.59
	Target height : Prec.C.Dorsal	0.18	0.08	5269	2.22
	Target height : Foll.C.Bilabial	-0.19	0.04	5146	-4.40
	Target height : Foll.C.Coronal	0.15	0.06	4442	2.41
	Backness : Syllable : Root type	-0.08	0.03	5294	-3.08

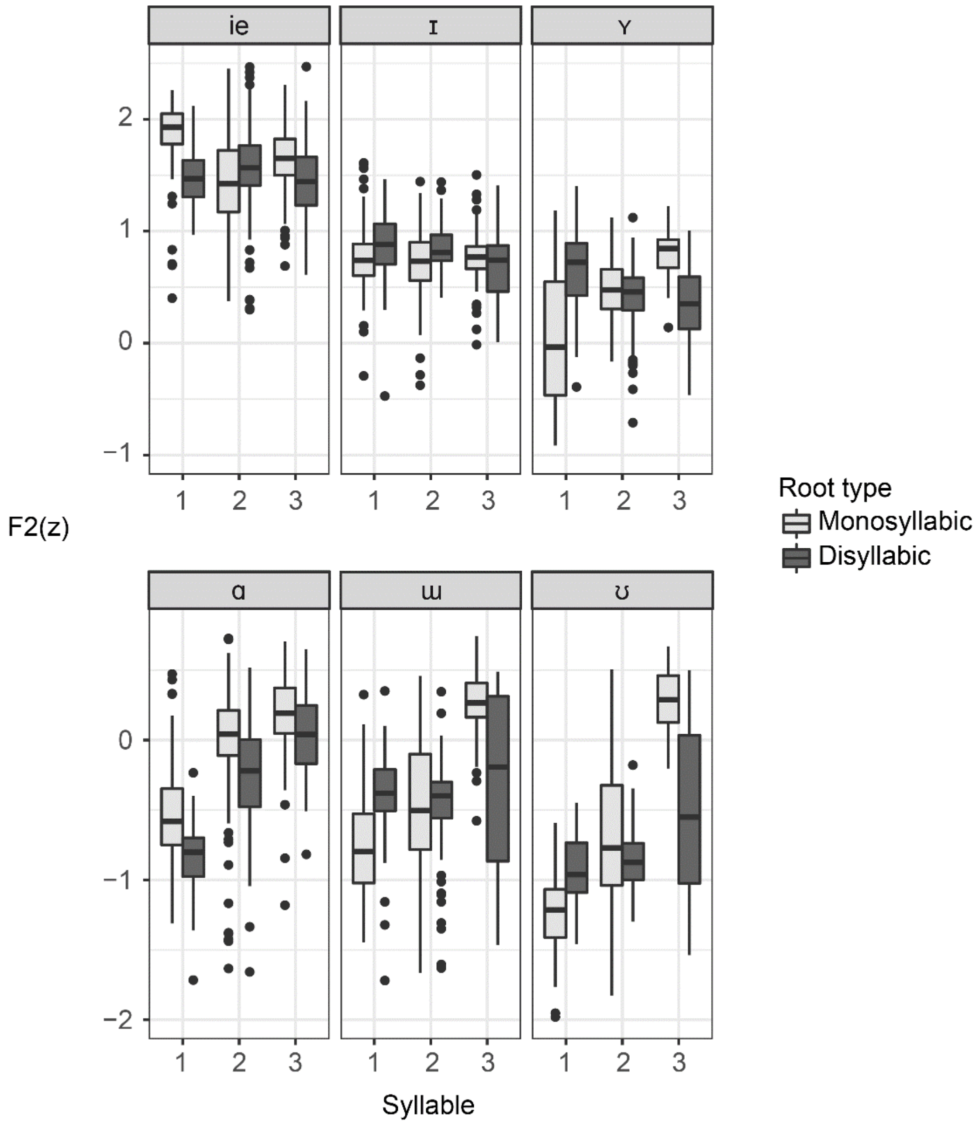

Figure 4.10: F2 (z, with SD) in syllables 1-3 based on root type in Kazakh

4.3.4 Uyghur

Vowel quality by position is presented in Figure 4.11 and Tables 4.16-4.17 below. Figure 4.11 presents mean vowel quality by position and Tables 4.16-4.17 provide additional information, including standard deviation, and token count (total n = 5,927). Syllable number does not exert a significant effect on F2 of front vowels in Uyghur [Syllable: β= 0.01, $\chi^2(1)$= 0.02, p= .88]. In contrast, back vowels are produced with significantly increasing F2 by position [Backness : Syllable: β= 0.35, $\chi^2(1)$= 32.48, p< .001]. In other words, the Uyghur vowel space undergoes an asymmetric fronting by position – back vowels are fronted in non-initial syllables while front vowels exhibit no concomitant changes by position. Thus, the Uyghur pattern strongly resembles the pattern of positional variation found in Kazakh above.

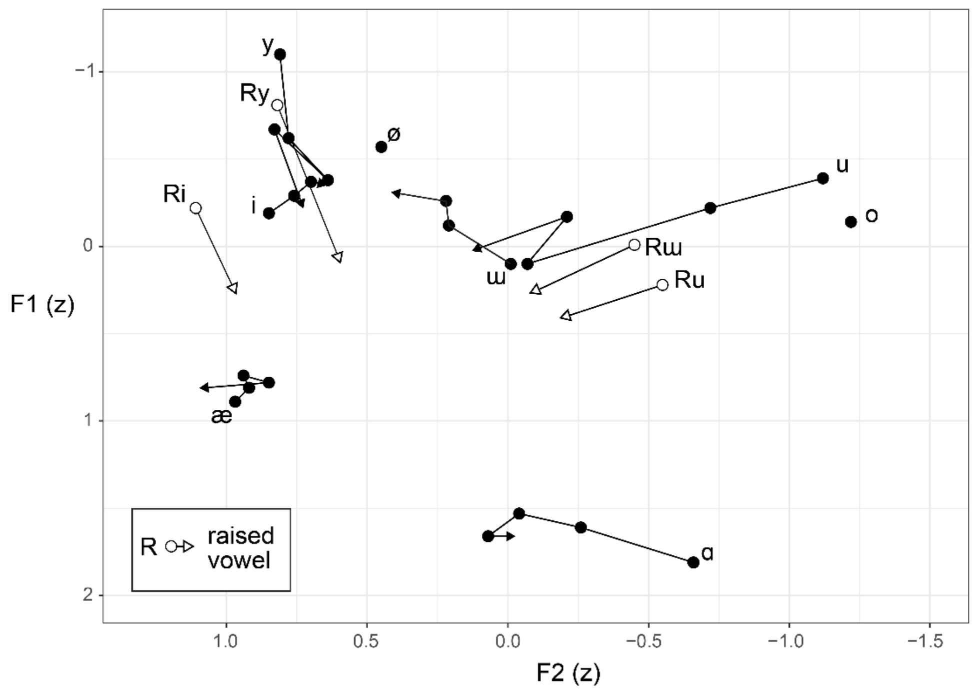

Figure 4.11: Uyghur mean vowel F1-F2 (z) by position. Lines connect different mean F1-F2 for a particular vowel. The initial-syllable mean is closest to the transcription and the arrow indicates the final mean.

Table 4.16: Mean F1 (in z, with SD) for Uyghur. Raised vowels are indicated by a subscript.

	Syllable 1		Syllable 2		Syllable 3		Syllable 4		Syllable 5	
	F1	n	F1	n	F1	n	F1	n	F1	n
ɑ	1.81 (0.53)	592	1.61 (0.53)	523	1.53 (0.55)	469	1.66 (0.45)	71	1.66 (0.41)	59
o	-0.14 (0.42)	207								
ɯ			0.10 (0.65)	114	-0.12 (0.61)	130	-0.26 (0.32)	85	-0.31 (0.38)	28
u	-0.39 (0.71)	290	-0.22 (0.81)	135	0.10 (0.48)	21	-0.17 (0.42)	12	0.02 (0.27)	7
æ	0.89 (0.76)	311	0.74 (0.49)	465	0.81 (0.48)	310	0.78 (0.43)	65	0.81 (0.32)	36
ø	-0.57 (0.26)	252								
i			-0.19 (1.13)	123	-0.29 (0.46)	92	-0.37 (0.36)	52	-0.36 (0.38)	12
y	-1.10 (0.82)	328	-0.62 (1.09)	256	-0.38 (0.41)	55	-0.67 (0.26)	6	-0.23 (0.22)	4
$ɯ_R$			-0.01 (0.59)	290	0.27 (0.36)	113				
u_R			0.22 (0.61)	123	0.41 (0.48)	35				
i_R			-0.22 (0.59)	81	0.26 (0.38)	45				
y_R			-0.81 (0.59)	105	0.07 (0.42)	25				
Total		1,980		2,215		1,295		291		146

Table 4.17: Mean F2 (in z, with SD) for Uyghur. Raised vowels are indicated by a subscript.

	Syllable 1		Syllable 2		Syllable 3		Syllable 4		Syllable 5	
	F2	n	F2	n	F2	n	F2	n	F2	n
ɑ	-0.66 (0.28)	592	-0.26 (0.32)	523	-0.04 (0.34)	469	0.07 (0.33)	71	-0.02	59
o	-1.22 (0.24)	207								
ɯ			-0.01 (0.60)	114	0.21 (0.42)	130	0.22 (0.33)	85	0.41 (0.24)	28
u	-1.12 (0.84)	290	-0.72 (0.71)	135	-0.07 (0.61)	21	-0.21 (0.29)	12	0.12 (0.44)	7
æ	0.97 (0.71)	311	0.94 (0.59)	465	0.92 (0.48)	310	0.85 (0.37)	65	1.08 (0.43)	36
ø	0.45 (0.29)	252								
i			0.85 (0.52)	123	0.76 (0.25)	92	0.70 (0.26)	52	0.66 (0.30)	12
y	0.81 (0.62)	328	0.78 (0.55)	256	0.64 (0.38)	55	0.83 (0.31)	6	0.73 (0.30)	4
$ɯ_R$			-0.45 (0.51)	290	-0.08 (0.32)	113				
u_R			-0.55 (0.53)	123	-0.19 (0.31)	35				
i_R			1.11 (0.42)	81	0.97 (0.30)	45				
y_R			0.82 (0.50)	105	0.60 (0.29)	25				
Total		1,980		2,215		1,295		291		146

Positional variation by harmonic pairing is displayed in Figure 4.12. For the three main harmonic pairings, the distribution of vowel F2 strongly resembles the Kazakh pattern. F2 of the back vowels increases by position without any obvious changes for the front vowel. Of the raised vowel alternations, the $ɯ_R$-i_R alternation exhibits the same type of behavior, but the u_R-y_R pairing appears to show more symmetrical centralization. The full statistical model is presented in Table 4.18 below.

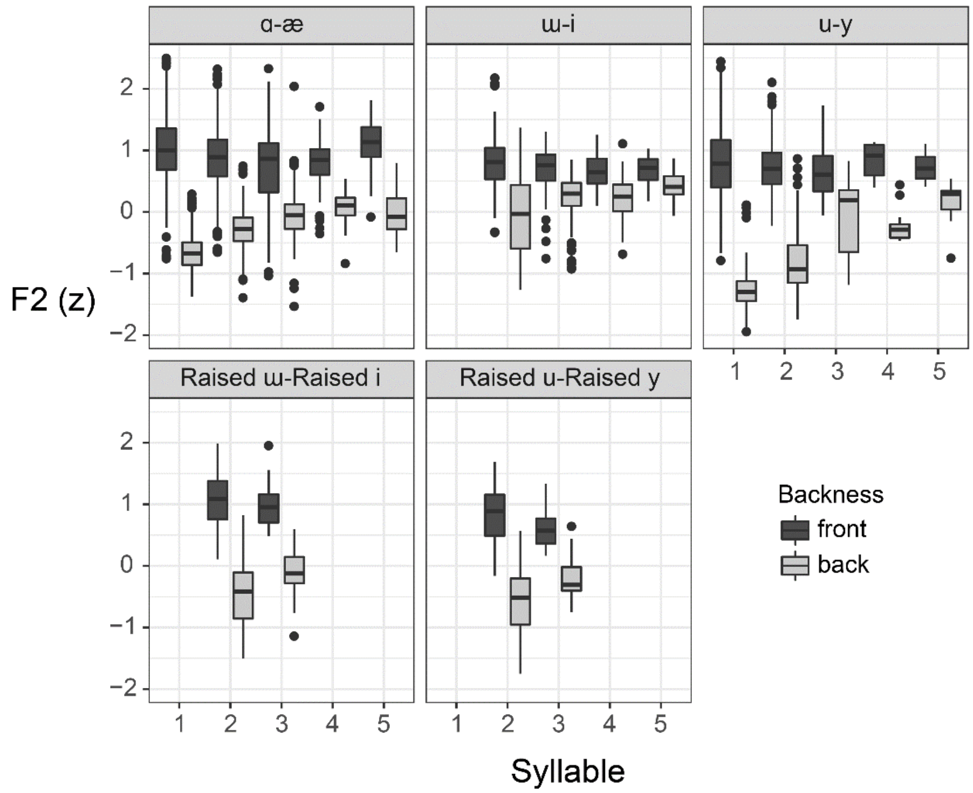

Figure 4.12: Uyghur F2 (z) by position and harmonic pairing

Similar to Kyrgyz, morphology exerts a significant effect on vowel F2 in Uyghur, with words derived from disyllabic roots being realized with higher F2 than words derived from monosyllabic roots [Root type: β= 0.33, $\chi 2(1)$= 25.98, p< .001; Backness : Root type: β= -0.25, $\chi 2(1)$= 8.09, p= .004]. This effect is strongest among front vowels, 0.33z, with only minimal effect among back vowels, 0.08z (0.33 main effect + -0.25 interaction). This fronting effect diminishes throughout the word, like in Kazakh [Syllable : Root type: β= -0.08, $\chi 2(1)$= 7.05, p= .008; Backness : Syllable : Root type: β= -0.07, $\chi 2(1)$= 2.92, p= .09]. Figure 4.12 reports vowel F2 by root type for the first three syllables in Uyghur. Since raised vowels only occurred in words formed from disyllabic roots, they are not included in Figure 4.12.

Table 4.18: Model output for Uyghur

Random effects					
			Variance	SD	
	Speaker	(Intercept)	0.03	0.17	
		Syllable	0.00	0.05	
	Target vowel	(Intercept)	0.37	0.61	
		Syllable	0.01	0.10	
	Residual		0.23	0.48	
Fixed effects					
		β	SE	df	t-value
	(Intercept)	0.58	0.35	7	1.65
	Backness	-0.97	0.14	98	-6.81
	Target height	-0.30	0.59	34	-0.51
	Syllable	0.01	0.06	10	0.11
	Root type	0.33	0.06	4938	5.09
	Prec.C.Bilabial	-0.08	0.04	5426	-1.95
	Prec.C.Coronal	-0.08	0.04	5355	-1.91
	Prec.C.Dorsal	0.01	0.05	5357	0.23
	Foll.C.Bilabial	-0.17	0.04	5428	-3.97
	Foll.C.Coronal	-0.16	0.03	5370	-6.02
	Foll.C.Dorsal	0.11	0.04	5419	2.59
	Backness : Target height	-0.09	0.09	2690	-1.09
	Backness : Syllable	0.35	0.06	110	6.46
	Syllable : Root type	-0.08	0.03	5143	-2.64
	Backness : Root type	-0.25	0.09	5117	-2.85
	Target height : Prec.C.Bilabial	-0.03	0.08	3794	-0.42
	Target height : Prec.C.Coronal	0.17	0.08	3407	2.30
	Target height : Prec.C.Dorsal	-0.13	0.08	3996	-1.72
	Target height : Foll.C.Bilabial	0.20	0.05	5422	0.43
	Target height : Foll.C.Coronal	0.34	0.48	5422	0.71
	Target height : Foll.C.Dorsal	-0.00	0.49	5422	0.00
	Backness : Syllable : Root type	-0.07	0.04	5243	-1.70

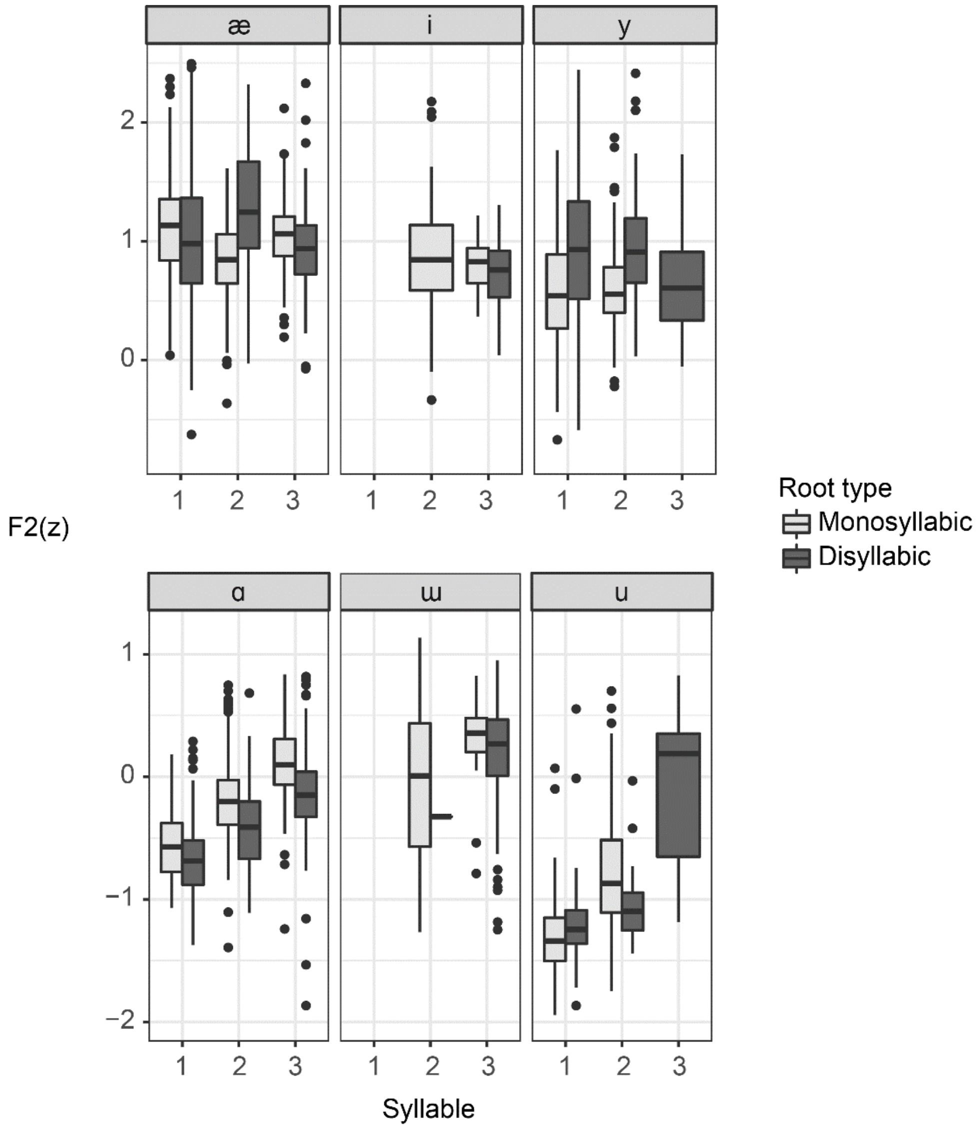

Figure 4.13: F2 (z, with SD) in syllables 1-3 based on root type in Uyghur

In sum, Uyghur vowels exhibit a pattern of asymmetric fronting, whereby F2 of back vowels increases significantly in non-initial syllables. Morphology also affects the realization of vowel backness, producing fronting in both [-back] and [+back] vowel words derived from disyllabic roots.

4.3.5 Summary

This section reported results on the realization of vowel backness across positions to investigate positional changes in F2 and their relationship to position. Returning to the predictions from Section 2.5, two of the three predicted patterns of positional changes in vowel F2 were attested. Uzbek and Kyrgyz exhibit a relatively symmetrical pattern of F2 centralization. In contrast, Uyghur and Kazakh demonstrate an asymmetrical pattern, with F2 of back vowels increasing with no comparable shift in front vowel F2.

Additionally, this section has reported on results relating morphology to positional vowel backness. Morphology exerts the most force in Kyrgyz, making both front and back vowels more peripheral in words derived from disyllabic roots. Morphology affects the realization of vowel F2 less so in Uyghur, and almost none in both Kazakh and Uzbek. The role of morphology will factor into the formalization of harmony in Chapter 6.

The rest of this chapter focuses on the analysis of the two patterns attested in this section, comparing the phonetic and phonological accounts of gradience sketched earlier in the chapter.

4.4 Analysis

This section takes the empirical results from the previous section and develops an analysis of each pattern of positional shifts in F2. Throughout the section I will compare predictions from phonetic accounts of reduction and centralization with the empirical data, using extant phonetic accounts to assess the role that phonology and phonetics play in these patterns.

4.4.1 Possible analyses

Before working out the details of the analysis, recall the possible analyses laid out in Table 4.5, replicated below in Table 4.18. Three empirical patterns and five possible analyses were described. Of these, the first empirical pattern involves non-converging F2, where F2 of both classes of [back] vowels does not change by position. This analysis relies on both a categorical phonological module and the claim in Zsiga (1997) and Pearce (2008), that harmony blocks phonetic reduction. This type of pattern is attested in Turkish (Lanfranca 2012; Sabev & Payne 2018), but does not match any of the patterns attested in this chapter.

Table 4.19: Schema for possible patterns and their analysis. Each pattern is paired with at least one analysis, which involves a phonological and phonetic component. Dashed lines indicate front vowel shifts, while solid lines indicate back vowel shifts.

	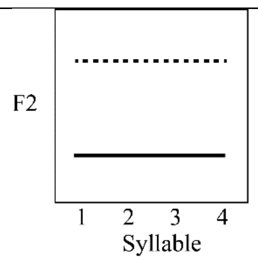	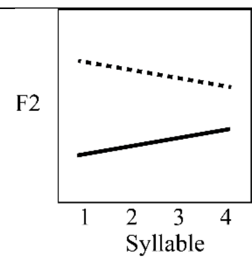	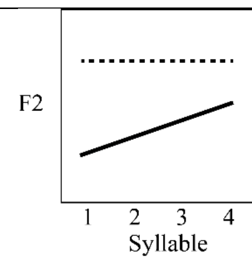		
Pattern	**No positional shifts**	**Symmetric shifts**	**Asymmetric shifts**		
Analysis	Categorical harmony	Gradient centralization	Under-specification	Adaptive hypo-articulation	Gradient phonology
Phonology	Categorical harmony	Categorical harmony	No phonological harmony	Categorical harmony	Gradient harmony
Phonetics	No reduction	Centralization	Interpolation across underspecified syllables	Reduction of effortful articulations in predictable contexts	No reduction

The second possibility is symmetric convergence. In this pattern, F2 of both front and back vowels shift toward a relatively central value in non-initial syllables. This empirical pattern is most consistent with the data found for Kyrgyz and Uzbek. The analysis of these patterns depends on gradient phonetic reduction. Under this analysis, harmony in Kyrgyz categorically associates [back] to all vowels in the word within the phonological module. In the phonetics, though, non-initial vowels are reduced due to phonetic forces. In Uzbek, no harmony operates, but phonetic reduction similarly occurs during speech production and acoustic vowel contrasts are minimized in non-initial syllables. It is not immediately germane to this chapter to discuss the details of phonetic reduction in each language, but these effects may derive from various factors, including the prominence of prosodic edges, (e.g. initial strengthening: Fougeron & Keating 1997; Cho 2005; Cho & Keating 2009; final lengthening: Wightman et al. 1992; Tabain 2003) reduction in metrically defined weak positions (Savy & Cutugno 1998; Fourakis 1991; Fourakis et al. 1999; Padgett & Tabain 2005), phrase-level patterns of articulatory reduction (Vayra & Fowler 1992; Johnson & Martin 2001; Krakow et al. 1995; Nord 1986; Herman et al. 1997), or undershoot (Lindblom (1963).

The third empirical pattern in Table 4.18 involves an asymmetric shift in the vowel space. In this pattern, F2 of one class of vowels shifts toward the other while without any concomitant shifts attested for the other class of vowels. This pattern was found for Kazakh and Uyghur. To analyze the asymmetric fronting of back vowels in these two languages, three analyses are possible, underspecification, hypoarticulation, and gradient phonology.

4.4.1.1 Underspecification

First, positional changes in vowel quality may derive from coarticulation and underspecification. Interpolatory coarticulation occurs when distinct phonological specifications are separated by segments unspecified for the relevant feature. Cohn (1993) demonstrates that regressive coarticulatory nasalization on English vowels results from interpolating from a consonant specified as [-nasal] across a vowel to another that is specified as [+nasal]. Cohn argues that interpolation is possible when the phonology of a language leaves a feature unspecified, which may then be filled in during phonetic implementation. In English, nasal interpolation may occur because nasality is not contrastive on vowels. Vowels immediately preceding a nasal are characterized by gradually increasing nasal airflow in preparation for the upcoming nasal consonant. Cohn (1993) also shows that such nasal interpolation does not occur in French, a language with phonologically contrastive nasality on vowels. She contends that this type of phonetic interpolation is constrained by phonological contrast, with interpolatory effects occurring only on segments that are not specified for the relevant phonological contrast (see also Keating 1988).

Thus, in a [+back] form, like the Kazakh word /bɑltɑ-lɑr-dɑn/ 'axe-PL-ABL', only two positions receive a phonological specification, the initial syllable and the right edge of the word, which is also the end of the utterance. The initial syllable would be underlyingly specified as [+back], while the end of the utterance would be specified as [-back], in conformity with a language-specific default articulatory setting. It has been shown that default articulatory settings (sometimes called interspeech posture; Wilson & Gick 2006; Gick et al. 2004) vary across languages. More importantly, in languages with vowel harmony, default articulatory settings are consistent with the unmarked feature value in the harmony system (Hudu 2010; Allen et al. 2013).

4.4.1.2 Adaptive hypoarticulation

Second, positional changes in vowel quality may also derive from what Lindblom (1990) calls adaptive hypoarticulation. Lindblom proposes that a speaker may reduce effortful articulations in predictable positions. In languages with vowel harmony controlled by the initial syllable, non-initial syllables are largely predictable based on the featural content of the initial-syllable vowel (Goldsmith & Riggle 2012). In a language with harmony, a speaker can utilize a listener's expectations to produce a less effortful version of the target vowel, reducing articulatory effort without compromising intelligibility. This type of analysis can generate both symmetrical shifts as well as asymmetrical shifts, depending on the relative markedness of the feature values involved. In a symmetrical pattern, this analysis could be indistinguishable from articulatory reduction. For an asymmetrical shift, this analysis would require one particular value of the relevant feature to be more marked. The less economical (i.e. more effortful) articulatory configuration is reduced, incrementally shifting toward the more economical (i.e. less effortful) configuration. Since [+back] vowels are fronted in non-initial syllables, this suggests that [+back] is the more effortful articulatory configuration in Kazakh and Uyghur.

This analysis shares many of the same underpinnings as the underspecification analysis. In both, speech production is constrained by economization of effort, with [+back] being more effortful in Kazakh and Uyghur. However, under this analysis there is no requirement that non-initial vowels be left unspecified for [back] in the phonology. Under this analysis the phonology categorically maps the feature [back] to all syllables in the word, but then speech production modulates the received specification to minimize effort while maintaining efficient information transmission.[6]

[6] One variant of this analysis (or perhaps more accurately, a hybrid of the two phonetic analyses) is also possible within Keating's window model of coarticulation (Keating 1988, 1990). In her model, Keating proposes that segments are defined not by a single articulatory target, but by a range, which allows for coarticulation from flanking segments. The window model can account for the data if the windows may shift by position, so the window for an initial-syllable /ɑ/ would be slightly different from a second- or third-syllable /ɑ/. For the analysis to work, these differences must align with the shifts attested, but herein lies the problem. The window model cannot, in and of itself, constrain the change in window position without some other mechanism. Stated differently, in principle there is nothing in Keating's model to prevent [+back] vowel windows from moving toward more posterior articulatory positions during the word, and so some mechanism must be invoked to force the windows to move toward more anterior positions. This same mechanism can likely be invoked in a general adaptive hypoarticulation analysis without the window model, and so the two are treated equivalently below.

4.4.1.3 Gradient phonology

Third and finally, gradient vowel shifts may result from phonological, and not phonetic factors. Deriving gradient vowel fronting in Kazakh and Uyghur shares one key commitment with the other two analyses- [+back] is the marked feature value. Stated simply, under this analysis [+back] spreading peters out over the course of the word. If numerical values are substituted for binary feature values, with 1 being categorically [+back], and 0 being categorically [-back], the values for [back] in the Kazakh word, /bɑltɑ-lɑr-dɑn/ 'axe-PL-ABL' would be something like /1-0.9-0.8-0.7/. McPherson & Hayes (2016) argues that the frequency of harmony may peter out throughout the word, maintaining the assumption that phonological patterns are categorical. The particular analysis advocated herein, however, supposes that phonological patterns may be gradient, allowing subphonemic patterns to fall out from the phonology (see Cohn 2006; Ernestus 2011 for overviews on gradience in phonology).

4.4.2 Predictions and independent evidence

It should be clear that all three of the analyses just outlined can derive the asymmetric gradience attested in Kazakh and Uyghur. The real question is then, which of these is most appropriate. This question hinges upon the analytical commitments each analysis makes, along with their resultant predictions. The assumptions of each analysis are listed in (41).

(41) Assumptions of each analysis

A. Underspecification

1. [+back] is the active feature value
2. Non-initial vowels are not specified for [back] in the phonology of each language
3. The production of non-initial vowels is constrained by economy, with all non-initial vowels shifting toward the default articulatory setting for the language.

B. Adaptive hypoarticulation

1. [+back] is the active feature value
2. Non-initial vowels are specified for [back] in the phonology
3. The production of non-initial vowels is constrained by economy, with all vowels shifting toward the least effortful articulatory configuration in each language

C. Gradient phonology

1. [+back] is the active feature value
2. Non-initial vowels are specified for [back] in the phonology
3. The realization of non-initial vowels is not necessarily constrained by economy

In order to evaluate these assumptions, and the larger question of analytical superiority, I consider evidence from related aspects of Kazakh and Uyghur phonology below. I first address the shared assumption that [+back] is the active feature value, and then move on to address the phonological specification of non-initial vowels and the general constraints on vowel realization that fall out from each analysis.

Each analysis argues that the specific patterns of back vowel fronting in Kazakh and Uyghur derive from the markedness of [+back]. It is easy to see that this is not the only analytical possibility, though. It is possible that [-back] is active in each language, and the Turkological term palatal harmony indeed suggests that [-back] is marked, and therefore active in the language family. However, evidence for the markedness of [+back] comes from several sources: loan word adaptation, dialectal variation, and high vowel fronting.

In Kazakh, loan words are typically borrowed with vowel pronunciations either faithful to the lending language, or with front vowels. Crucially, I know of no loan words with front vowels in the lending language that are borrowed with back vowels in Kazakh, although there are myriad examples of the reverse, some of which are shown below in Table 4.20. Words borrowed from Arabic, Persian, Russian, and Sogdian exhibit the same tendency, for words with back vowels to be borrowed with front vowels. The variety of languages borrowed from and the range of phonological adaptations are significant here. Persian and Arabic may produce the low vowel, /a/, with significant variation, from a more fronted, almost [æ]-like quality to [ɑ] or [ʌ], making it possible to account for some Kazakh borrowing patterns from the surface phonetics, but the many cases of original /a/ being adapted as /ie/, in addition to the various other adaptations found below suggests a more general pattern of loan word fronting.

Table 4.20: Loan word adaptations in Kazakh

Kazakh	Original	Vowel adaptation	Lending language	Gloss
sæliem ~ sɑlɑm	salaːm	a → æ	Arabic	peace
dʏnijie	dunja	u → ʏ	Arabic	world
æliem	ʕalaːm	a → æ	Arabic	world
æskier	ʕaskaːr	a → æ	Arabic	soldier, army
mieʃɪt	masʤid	a → ie	Arabic	mosque
miektiep	maktab	a → ie	Arabic	school
kʏnæ	gonah	o → ʏ, a → æ	Persian	sin
kieruwien	ka:rava:n	a → ie	Persian	caravan
kieriewiet	krovat^{j}	o → ie, a → ie	Russian	bed
mæʃinie (colloquial variant in southeast)	maʃina	a → æ	Russian	car
pætier	kfartira	a → æ	Russian	apartment
sømkie	sumka	u → ø	Russian	luggage
bøtielkie	butɨlka	u → ø	Russian	bottle
kient	kand	a → ie	Sogdian	town

For Uyghur, Yakup (2005:187-189) describes a tendency for loans to be adapted to front vowels, and especially for initial-syllable /a/ in Persian and Arabic words to be adapted as /e/. Some examples from Yakup (2005) are shown in Table 4.20. Uyghur thus exhibits a similar tendency for loanwords to be borrowed with front vowels, indirectly suggesting the markedness of back vowels. Kenstowicz (2005) and Kang (2011) discuss cases of loanword adaptation where the borrowing language imposes a stricter requirement on borrowed words than on the native lexical strata. They call this pattern "retreat to the unmarked." In Table 4.21, the native inventory includes sounds very similar to the sounds adapted from Persian, Arabic, and Russian, /o a u a/. Nonetheless, those phonetically more similar sounds are bypassed in adaptation in favor of the less marked feature value, [-back]. In Uyghur, both [+back] and [-back] vowels are licit, but in loans the preference for [-back] vowels is manifest in these patterns of adaptation.

Table 4.21: Loan word adaptations in Uyghur (Yakup 2005:187-189)

Uyghur	Original	Vowel adaptation	Lending language	Gloss
gøʃ	goʃt	o → ø	Persian	meat
dærgyʃ	darviːʃ	a → æ	Persian	dervish
heqi	ʔaql	a → e	Arabic	wisdom
ømyr	ʔumr	u →ø	Arabic	life
mæhællæ	mahalla	a → æ	Arabic	village
ʧøtkæ (Turfan dialect)	ʃotka	o → ø, a → æ	Russian	brush
ædijɑl	odejɑlo	o → æ	Russian	blanket cover

In addition to loans, dialectal variation supports the markedness of [+back] in both Kazakh and Uyghur. In Kazakh, speakers of eastern dialects may front words that are both written with back vowels and produced by speakers of other dialects as [+back]. For instance, (Niyazgalieva & Turganalieva 2013:78) cites a number of examples, like /tɪrniektiep/ as a variant of /tɯrnɑqtɑp/ 'scratching' in southeastern varieties, and /ʒiewdɪriep/ as a variant of /ʒɑwdɯrɑp/ 'showering' in southern varieties of the language. In Uyghur, a range of dialectal adaptations are attested. These varied adaptations show the same tendency to adapt back vowels as front. In all the loans presented in Yakup (2005:187-189), foreign /a/ is variably adapted as /ɑ/ or /æ/, as shown in Table 4.22 (see also Abdurehim 2014:79-80). While back vowels may be fronted, the reverse, where loans are backed, is to my knowledge unattested or at most, rare.

Table 4.22: Dialectal variation in Uyghur loanword adaptation (Yakup 2005:187-189)

Kashgar dialect	Turfan dialect	Standard Uyghur	Original	Gloss
ɑdɑwæt	ɑdɑgæt	ɑdɑwæt	ada:wat	contradiction
æwwæl	ɑgɑl	ɑwɑl	awwal	first
ʁælvi:	ʁalvir	ʁælvir	ʁalbi:r	sieve
	mejman	mehman	mehma:n	guest

Finally, the proposal that [+back] is marked in these languages receives significant support from the realization of high vowels in word-final position. In general, marked values should be characterized by a more restricted distribution than their unmarked counterparts. One prime example of this is word-final high vowel fronting in Kazakh and Uyghur. In Uyghur, the contrast in backness and rounding among the high vowels is completely neutralized in absolute word-final position. Word-finally, all high vowels are realized as a very peripheral [i], as seen in Figure 4.14. In the figure, observe the expected distribution for F2 in closed syllables- /i/ and /y/ have high F2 while /ɯ/ and /u/ have low F2. However, all high vowels are realized with F2 exceeding 2z in final open syllables (i.e. word-finally), which is substantially higher than F2 of any vowel in any other position in the word.

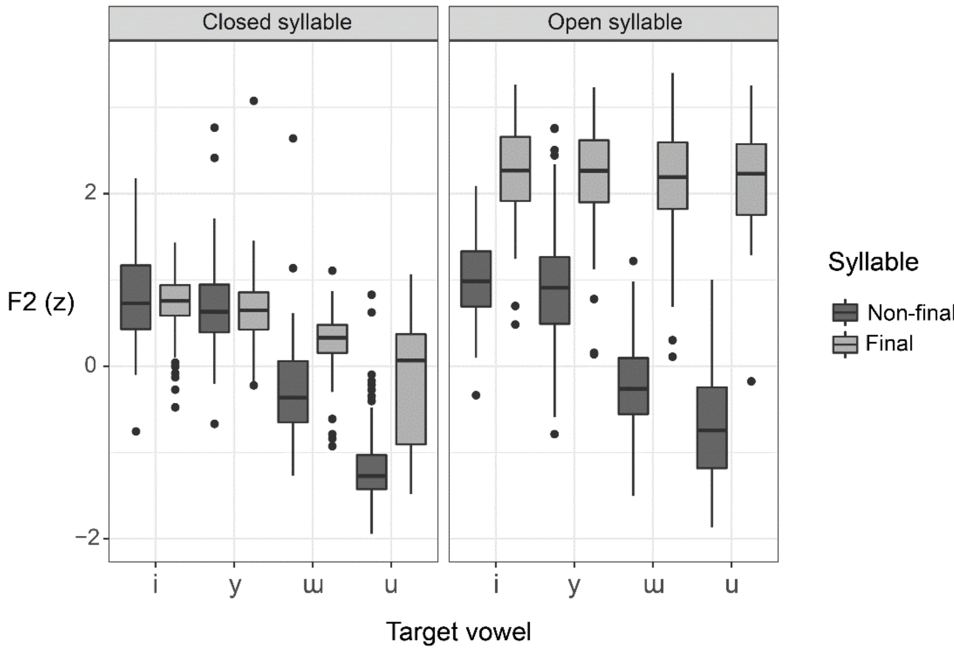

Figure 4.14: F2 (z) of high vowels by position and syllable type in Uyghur

Word-final high vowels undergo fronting in Kazakh, as well. However, note in Figure 4.15 that this process is gradient, failing to completely neutralize acoustic differences between all four high vowels. Word-final /ʏ/ is realized with F2 approximating word-final /ɪ/, but word-final /ɯ/ and /ʊ/ are produced with F2 values in between canonical front and back vowels in the language.

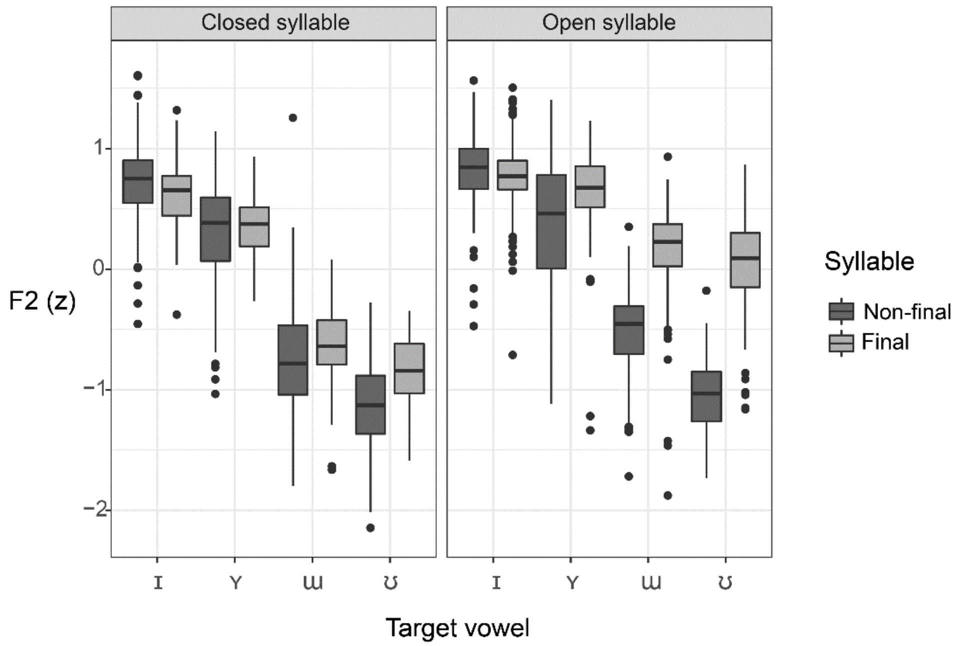

Figure 4.15: F2 (z) of high vowels by position and syllable type in Kazakh

Word-final high vowel fronting suggests that [+back] is marked among the high vowels. High back vowels and high round vowels are not permitted word-finally in Uyghur, with all contrasts for backness and rounding neutralized to a very cardinal [i]. In Kazakh, those same contrasts are incompletely neutralized, shifting toward the front unrounded [ɪ]. In tandem with loan word adaptation and dialectal variation in these two languages, high vowel fronting further supports the claim that [+back] is the marked feature value.

Markedness relations have traditionally been tied to featural activity – the marked feature value is the active feature value (Stewart 1967; Schachter 1969; Ultan 1973; Archangeli 1984; Smolensky 2006). Conversely, the unmarked feature value is inert. In OT, harmony is often formalized as two relations between markedness constraints. First, the general markedness constraint *[+F] must outrank the constraint *[-F]. Second, a harmony-motivating constraint, like AGREE[F] (Lombardi 1999; Baković 2000) must outrank *[+F], as well as the relevant faithfulness constraints. If the general markedness constraint outranks the harmony-driver, then harmony fails (see e.g. Kiparsky & Pajusalu 2003; Smolensky 2006 for discussion). The general state of affairs is aptly summarized by Baković (personal communication), "when harmony fails, markedness prevails." Among other factors discussed in Casali (2012), when harmony is blocked by some structural condition (e.g. the low vowel in Akan), the feature value of vowels further from the harmony trigger should provide evidence for the unmarked feature value in that language. In Akan, only [-ATR] vowels occur to the left of /a/. In Kyrgyz rounding harmony (Chapter 1), /ɑ/ does not assimilate to [o] after [u] (as a markedness constraint *u…o), producing forms like [qul-dɑr] 'slave-PL' rather than *qul-dor. In longer words, only unrounded vowels occur to the right of unassimilated [ɑ], e.g. quldɑrdɯ 'slave-PL-ACC.' The generalization is that blocking reveals which feature value is marked. In the cases examined above, no exceptional morphemes or structural conditions demonstrate this, but the direction of gradience does. When harmony categorically fails, the unmarked feature value emerges. In like manner, when harmony fails gradiently, a less marked phonological output occurs. Thus the gradient, incomplete failure of harmony manifests the same markedness generalizations as in Akan and Kyrgyz. For Kazakh and Uyghur, the diagnostics discussed thus support the markedness of [+back], and as a consequence, its status as the active feature value in these languages.

Despite the different pattern of positional variation in Kyrgyz, I contend that Kyrgyz, too, exhibits the same featural asymmetry, with [+back] being the marked and active feature value. Based on my own experience, Kyrgyz is more likely than Kazakh to faithfully adapt loans from other languages. Perhaps my impressions derive from the relative paucity of loans in Kyrgyz compared to Kazakh. Regardless, Kyrgyz exhibits the same tendency to front back vowels in loans, as seen below in Table 4.23.

Table 4.23: Loanword adaptations in Kyrgyz

Kyrgyz	Original	Vowel adaptation	Lending language	Gloss
tɑmeki	tabak	a → e	probably Russian (possibly French or Arabic)	tobacco
ystøl	stol	o → ø anaptyptic [y]	Russian	table
kerebet	krovat^{j}	o → e, a → e	Russian	bed
myʃtøk	mundʃtuk	u → y	probably Russian (possibly German)	bit (in horse's mouth)
bøtølkø	butɨlka	u → ø	Russian	bottle
ømyr	ʔumr	u →ø	Arabic	life
ɑsker	ʔaskar	a → e	Arabic	soldier, army

Kyrgyz dialectal variation also lends support for the markedness of [+back]. Standard Kyrgyz is based on the Northern dialect of the language, and where the Southern dialect diverges from the Northern, back vowels are fronted, especially to a ninth phonemic vowel, /æ/. Many have noted the significant influence of Uzbek on the Southern dialect of Kyrgyz, but one thing to note is that dialects of Uzbek do not have phonemic /æ/, but instead /ɑ/ (Batmanov 1940; Reshetov & Shoabdurahmonov 1978; Yunusaliev 1971; Ibrohimov 1967; Töjchiboyev & Hasanov 2004). Thus, the introduction of /æ/ and its role in the phonology of the Southern dialect is not entirely reducible to Uzbek influence. All else being equal, this dialectal difference suggests a general dispreference for [+back].

Like Uyghur and Kazakh, Kyrgyz exhibits a pattern of high vowel fronting in absolute word-final position. In closed syllables, each target vowel shows no substantial difference in F2 in non-final and final syllables. However, in final open syllables, the back vowels show a large increase in F2. In short, the back vowels are fronted (in addition to centralization) in word-final position. This pattern in evident in Figure 4.16 below. Observe the significant increase in F2 for /ɯ u/ in final open syllables. When the magnitude of this shift is compared to that found in closed syllables, on the left side of the figure, it is clear that this is fronting over and above the general centralization of non-initial vowels.

To summarize, all high vowels neutralize categorically to cardinal [i] in Uyghur, while the high vowels neutralize incompletely to [ɪ] in Kazakh. In this respect, Uyghur and Kazakh differ in the categoricality of their fronting patterns. In Kyrgyz, the pattern is gradient and only affects a subset of the high vowel inventory. The high front vowels are unaffected while the two high back vowels, /ɯ u/ undergo a phonetic fronting in Kyrgyz. Together with loanword adaptations and dialectal variation, high vowel fronting supports an analysis whereby [+back] is the feature value that is active (i.e. the marked value) for harmony in these three languages.

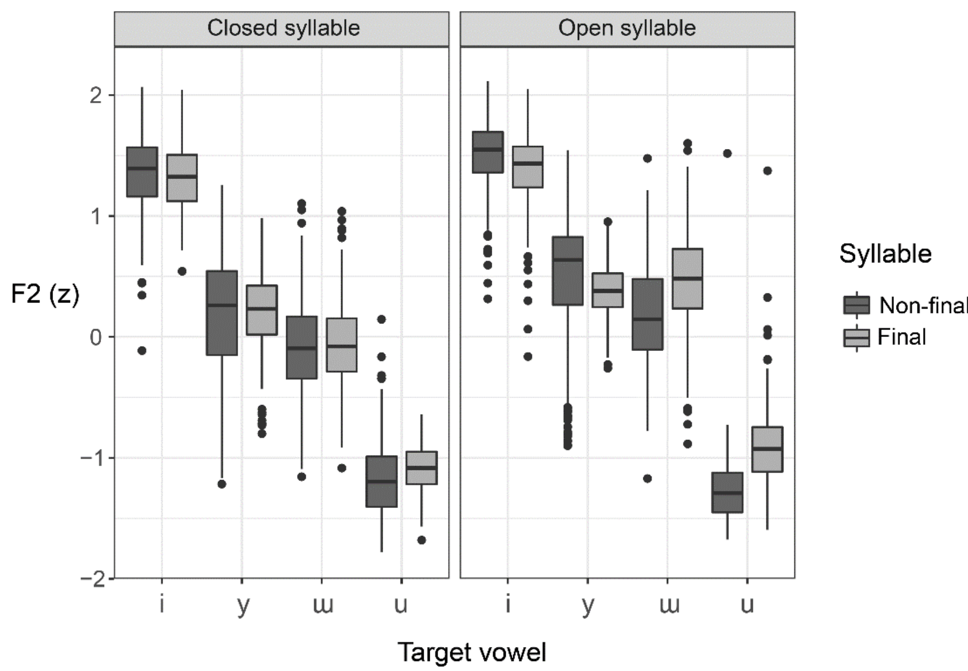

Figure 4.16: F2 (z) of high vowels by position and syllable type in Kyrgyz

Moving on to the second prediction, the underspecification analysis differs from both the hypoarticulation and gradient phonology accounts in that it predicts that non-initial vowels should be unspecified for [back] in the phonology, leaving phonetics to operate over segments unspecified for [back]. In all three languages, vowel backness triggers alternations on dorsal obstruents. When adjacent to a front vowel, either preceding or following, dorsal obstruents are produced with a velar place of articulation. In back vowel contexts, again either preceding or following, dorsal obstruents are produced with a uvular place of articulation. Some representative data are shown below in (42). In (42a), the nominative form occurs with a back vowel, indicating the underlying quality of the initial-syllable vowel. In (42b), the back vowel triggers a [+back] allomorph of the locative suffix, whose consonant is coronal, not dorsal. The first two examples in (42) indicate that [+back] is underlying in these words, and that the backness of the root determines the backness of the following suffix. In (42c-e), the dative suffix is attached to the same root, and the initial dorsal consonant of the dative suffix surfaces as uvular. In the front vowel words in (42g), a front vowel root likewise triggers harmony on the following locative suffix. When a dorsal consonant follows a front vowel, as in (42h-j), the dorsal surfaces as velar.

(42) Dorsal alternations in Kyrgyz, Kazakh, and Uyghur

	Gloss	Kyrgyz	Kazakh	Uyghur
a.	head	bɑʃ	bɑs	bɑʃ
b.	head-LOC	bɑʃ-tɑ	bɑs-tɑ	bɑʃ-tɑ
c.	head-DAT	bɑʃ-qɑ	bɑs-qɑ	bɑʃ-qɑ
d.	honey-DAT	bɑl-ʁɑ		bɑl-ʁɑ
e.	honey-PL-DAT	bɑl-dɑr-ʁɑ	bɑl-dɑr-ʁɑ	bɑl-lɑr-ʁɑ
f.	five	beʃ	bies	bæʃ
g.	five-LOC	beʃ-te	bies-tie	bæʃ-tæ
h.	five-DAT	beʃ-ke	bies-kie	bæʃ-kæ
i.	lower back-DAT	bel-ge	biel-gie	bæl-gæ
j.	lower back-PL-DAT	bel-der-ge	biel-dier-gie	bæl-lær-gæ

The dorsals are allophonic in Kyrgyz, but contrastive in Kazakh in Uyghur. In Kazakh, the dorsals contrast only word-initially (compare /qij/ 'manure' and /kij/ 'wear.IMP'); in Uyghur, the dorsals contrast only within roots (e.g. /qæjt/ 'pen point' and /kæjn/ 'rear'; /ɑq/ 'white' and /næq/ 'cash'). In all three languages, the dorsals are allophonically distributed in suffixes.

If non-initial vowels are not phonologically specified for [back], they should not trigger these dorsal alternations. Instead, only initial-syllable vowels, which are underlying specified for [back], should trigger these alternations. However, as seen above, this is incorrect. Non-initial vowels trigger these alternations (e.g. 42c,f) for place and manner of articulation, which demonstrates that non-initial vowels must be specified for [back] in the phonological component of the grammar. The fact that /ɡ/ and /ʁ/ alternate for harmony further suggests that this is a phonological alternation. There is no obvious phonetic reason for a slight difference in tongue body backness to result in a difference in constriction. In fact, this alternation is constrained by the inventory of these languages. In Lexical Phonology terms, this alternation is structure-preserving, one hallmark of phonological alternations. As a result, the underspecification account's requirement that harmony fall out from phonetic rather than phonological forces is simply incorrect.

Additionally, these three analyses also make different predictions regarding the economy of non-initial vowel production. The underspecification and hypoarticulation analyses both predict that non-initial vowels are produced in such a way as to minimize effort. The underspecification analysis posits effort minimization in the transition from an initial-syllable [+back] vowel to the edge of the utterance in the study, specified as [-back]. The hypoarticulation analysis posits reduction of effortful [+back] articulations in predictable (i.e. non-initial) contexts. Furthermore, the hypoarticulation analysis contends that all non-initial vowel production is constrained by economy, not just [+back] vowel realization. Thus, the absence of any significant positional shift among the front vowels must be due to their relative ease of articulation. In other words, this predicts that the F2 values characterizing the canonical front vowels in Uyghur and Kazakh are a proxy for the least effortful articulatory configurations in those languages.

However, this prediction is not consistent with high vowel fronting in Uyghur. One would expect high vowel fronting to produce vowels with the same F2 value as /i/, but this process results in vowel far more anterior than /i/. If F2 of typical /i/ (more similar to [ɪ]) in Uyghur is the least effortful articulatory configuration, then the exceedingly peripheral /i/ produced by this neutralization process must result in more effort. In short, word-final high vowel fronting is not economizing, and all non-initial vowel production should be under the hypoarticulation account. Moreover, I can envision no straightforward phonetic explanation that circumvents this, like increased perceptibility or some other plausible phonetic reason for this non-economizing pattern in Uyghur.

Defining what is more effortful is also a problem for the hypoarticulation account. If the account makes reference to [+back] completely independent of the articulatory or acoustic properties of the individual segments that comprise that class of vowels, the analysis depends on a phonological rather than phonetic class of sounds. If one tries to define a phonetic class of sounds corresponding to the phonologically [+back] vowels in either articulatory or acoustic terms, no single class is able to group /ɑ ɯ u/ together. In articulatory terms, as Honda (1996) and Esling (2005) point out, the high back vowels are produced with different muscular coordination than the low back vowels. The low back vowels are produced with significant activation of the hyoglossus, which triggers retraction of the tongue, in contrast to the styloglossus, which is activated for the production of the high back vowels, resulting in tongue body raising. Thus, in a strict articulatory sense, the back vowels differ in their articulatory configuration. There is thus an articulatory difference between /ɑ/ and /ɯ u/, which receives no straightforward explanation under this analysis. A similar issue arises if the back vowels are classified in terms of their acoustics. Acoustically, the problem vowel is /ɯ/. Almost all previous descriptions conflate /ɯ/ with /i/, missing the fact that these vowels exhibit a degree of contrast that isn't predictable if they are simply allophones of /i/ (see Chapter 2). Acoustically, /ɯ/ exhibits much higher F2 than the other back vowels, so much so that it is often misperceived as a front vowel. I can attest to my own trouble differentiating the two high unrounded vowels. Moreover, this vowel is more acoustically similar to /i ø y/ than /u o ɑ/, so defining a natural class based on acoustics that groups /ɯ/ with /u ɑ/ to the exclusion of some front vowels is an issue. The essence of the problem is this – the vowels that participate in gradient fronting are definable in terms of a relatively abstract phonological class that isn't entirely reducible to articulatory or acoustic phonetics. Again, if one argues that backness should be more generally construed, generalizing over the articulatory differences that characterize these vowels, the analysis becomes phonological. To the degree that abstracted generalizations define the analysis, the analysis become more phonological and less phonetic.

Thus, the patterns of fronting in Uyghur and Kazakh are best analyzed as gradient phonology and not some form of phonetic reduction. These alternating vowels trigger structure-preserving alternations on the dorsal obstruents (e.g. [ɡ]~[ʁ]), consistent with a morphophonological pattern. In addition, at a more conceptual level, the gradient phonology account straightforwardly captures the fact that the [+back] vowels undergo fronting due to the petering out of harmony. If, as in hypoarticulation account, harmony assimilates each vowel fully in the phonology only to be masked by phonetic reduction the analysis involves two steps resembling a Duke of York derivation (Pullum 1976). Harmony operates on a class of vowels to make them [+back]. Then phonetics manipulates that very same set of vowels to make them less backed. If phonology may be gradient, as is already hinted at in works as early as Kiparsky (1985), then the gradient phonology analysis accounts for the data without recourse to multiple mechanisms. From a diachronic point of view, if phonology can access and manipulate gradient variables (e.g. Kiparsky 1985; Hayes et al. 2004; Lionnet 2017), there is nothing to prevent an evolutionary account whereby gradient phonology emerges from hypoarticulation. If phonology is limited to categorical variables only, the hypoarticulation may persist, but if gradience is permissible in both modules of the grammar, the phonological account offers a simpler explanation than hypoarticulation.

4.5 Discussion

I will discuss some of the implications of this claim and its relevance for phonological analysis in Chapters 6 and 7. However, one important question is worth addressing at this point – am I claiming that all of phonology is gradient? In essence, should one extrapolate from the finding that morphophonology may be gradient to contend that all of phonology is, too?

Whether all phonology is gradient or not is a significant question, and my answers are distributed across the following chapters. I will preview my arguments at this point. First, although a number of writers have claimed that all phonology is gradient (e.g. Pierrehumbert et al. 2000; Silverman 2006; Tucker & Warner 2010), I am not. It is a more conservative step to claim that morphophonology may be gradient than contending that it all is gradient. Until we have more empirical evidence, it seems prudent to make the smaller claim. Second, the languages examined in this chapter exhibit both categorical and gradient patterns, suggesting that both types coexist in the phonologies of these speakers. In Uyghur, backness harmony is gradient but word-final high vowel neutralization is categorical. In Kyrgyz, on the other hand, backness harmony is categorical, but word-final high vowel neutralization is gradient, only partially neutralizing the contrast between the back and front vowels. The Kyrgyz pattern conforms to the predictions of Lexical Phonology, with a categorical morphophonological pattern, harmony, paired with a gradient post-lexical pattern, word-final high vowel neutralization. The Uyghur data, though, suggest that morphophonology may be gradient while post-lexical phonology is categorical. The claim that all phonology is gradient thus appears to gloss over what may be very important distinctions between patterns within and across languages.

Even though I maintain a distinction between the categorical and the gradient, the two can be unified formally. In Chapter 6, I introduce gradient variables into Harmonic Grammar to derive gradient as well as categorical harmony patterns. Constraint weighting in a gradient Harmonic Grammar is sufficient to account for both, allowing a unified representational system. There is therefore no need to represent Kyrgyz with a different set of representations from Kazakh and Uyghur. In this way, the categorical is only a special subcase of the gradient.

4.6 Summary

In this chapter I've presented acoustic evidence from Kyrgyz, Uzbek, Kazakh, and Uyghur addressing the categoricality of phonological patterns. In both Kyrgyz and Uzbek, vowels undergo gradient centralization while in Kazakh and Uyghur, [+back] vowels undergo asymmetric fronting. I discussed these results in the context of several competing phonetic and phonological analyses, contending that backness harmony is categorical in Kyrgyz, with gradient phonetic reduction. For Kazakh and Uyghur, though, backness harmony is gradient and phonological, which I have supported with the acoustic results reported along with independent evidence from the phonologies of these two languages. I have further presented evidence that gradient harmonies are probably present in a number of other language families, strengthening the claim that morphophonology, not just Turkic morphophonology, may be gradient.

If gradience is not the sole diagnostic for phonological versus phonetic patterns, then phonological theory needs a different set of criteria to differentiate the two. The next chapter and Chapter 7 focus on this topic. In the next chapter, I present results from several perception studies conducted during fieldwork to assess the perceptibility of the patterns described above, and in Chapter 7 I discuss a range of useful diagnostics for determining what is phonological and what is phonetic.

Chapter 5: The perceptibility of backness shifts in Kyrgyz, Kazakh, and Uyghur

The analysis presented in the previous chapter depends, at least in part, on the perceptibility of the shifts found in each language. In essence, if a subphonemic effect is phonological it should be perceivable. If a given effect is not perceivable, then the likelihood that speakers manipulate this information in their phonological grammar is seriously diminished. Generally speaking, phonological contrasts that are more distinct should be more easily maintained synchronically and more persistent diachronically. If a pattern or contrast is not perceivable, how do children learn it? If the child cannot attend to a slight acoustic or articulatory distinction, then what allows them to attend to and acquire that distinction in their language? This chapter presents evidence supporting the perceptibility of the acoustic effects reported in the previous chapter, and by extension, the phonological status of gradient harmony in Uyghur and Kazakh.

The chapter is structured as follows. In Section 5.1, I discuss previous work on "just noticeable differences" and the perceptibility of F2 shifts. In Section 5.2, I describe a perceptual experiment investigating perceptual boundaries in the three languages studied that exhibit relatively robust harmony systems, Kyrgyz, Kazakh, and Uyghur. In Section 5.3, I discuss how these perceptual findings relate to the claim that backness harmony is gradient in Kazakh and Uyghur but categorical in Kyrgyz, and the proposed requirement that phonological gradience be perceivable.

5.1 Just noticeable differences

The first and most general way to determine if the shifts in vowel backness reported in the previous chapter are perceivable is to compare the magnitude of these shifts with known thresholds in vowel perception. These have been called a just noticeable difference (sometimes JND) or difference limen (Flanagan 1955). In Flanagan's (1955) study, he finds that for differences in F2 (the primary acoustic correlate of tongue backness), the JND is approximately 3-5% of the original stimulus item. To concretize this generalization, a vowel with F2 of 1500 Hz should be distinguishable from a second vowel with F2 of 1575 Hz (75 = 5% of 1500). A vowel with F2 of 1800 Hz is therefore predicted to be distinguishable from a second vowel with F2 of 1890 (90 = 5% of 1800). Mermelstein (1978) finds that listeners are better able to discriminate between vowels in isolation than in consonantal contexts, with a 31% decrease in performance in CVC syllables. More specifically, Mermelstein (1978) reports a mean JND of 142 Hz for the vowel /ɪ/ in isolation, 174 Hz in /bɪb/ syllables, and 199 Hz in /gɪg/ syllables. In Flanagan's terms, these difference limens occur at roughly 6.5-9.5% of original /ɪ/ F2. In contrast, Kewley-Port (2001) reports an average difference limen of 1.5% in optimal listening conditions (i.e. in isolation), and 5% in longer phonetic contexts with greater uncertainty. Throughout the chapter, I use the terms perceptually salient and perceivable interchangeably. If, according to the diagnostics used, a positional shift is perceivable, then it is perceptually salient.

5.1.1 Predictions

In the previous chapter, I claimed that asymmetric fronting of back vowels in Kazakh and Uyghur is phonological, while positional variation among the front vowels is phonetic. While it is plausible to consider JNDs for both increases and decreases in F2, below I consider decreasing F2 for the front vowels. In essence, I assume that initial-syllable vowels are the most categorically front or back, with potential reduction (phonological or phonetic) in non-initial syllables. Thus, if a given back vowel shift involves a decrease in F2 or a front vowel shift involves an increase in F2, these will be assumed to be phonetic, and the shift will not be compared to the predicted JND.

To determine how positional shifts in Uyghur, Kazakh, and Kyrgyz compare to these JND findings, it is not possible to use z-score normalized acoustic data. Since z-score normalized data produces values that span 0, comparing percentages is problematic. For instance, a 5% JND for a vowel with a mean F2 (z) of 0.0 is 0. To meaningfully assess JNDs for these data, it is preferable to compare differences in raw Hz for each language with predicted JNDs. Below I compare mean group positional shifts in Hz with predicted JNDs from previous studies. If phonological distinctions must be perceivable,[7] then the strongest claim is that for a given vowel, the difference in mean F2 (Hz) for each syllable should be greater than an expected JND.[8] Throughout we will assume a relatively conservative JND of 5%. For a vowel with an initial-syllable mean F2 of 1500 Hz, a second-syllable mean should be at least 1575 Hz (1500 *1.05 = 1575) to be perceived as distinctly more anterior than the initial syllable. Similarly, a third-syllable mean should be at least 1654 Hz (1575 * 1.05 =1653.8) to be perceived as distinctly more anterior than the second syllable, and a fourth-syllable mean should be at least 1737 Hz (1654 * 1.05 = 1736.7) to be perceived as distinctly more anterior than the third syllable. In total, this would involve a shift of 237 Hz across four syllables, or a shift of approximately 79 Hz per syllable. A weaker claim is that some positional shifts must exceed a JND, but that not all must. Finally, if all vowel shifts in Kyrgyz are phonetic in nature, then one might predict that all such shifts should not exceed reported JNDs. This is probably too strong of a prediction. Phonetic shifts may be perceivable, so perceptibility is not the differentiating factor between phonological and phonetic patterns. Instead, if phonological patterns must be perceptually salient and a given pattern is not, then we can infer that that pattern is not phonological in nature. In sum, if the positional shifts reported for Uyghur and Kazakh in Chapter 4 are perceptually salient, then at least some of these must exceed known difference limens. Conversely, if the positional shifts reported for Kyrgyz in Chapter 4 are not phonologically meaningful, then those positional shifts should be less likely to exceed known difference limens. To evaluate these predictions, mean F2 (Hz) as well as beta values from a regression model will be examined. The

[7] Near-mergers and abstract contrasts are thus problematic for the prediction (see McCollum and Essegbey 2019 for a harmony-related example).
[8] This is distinct from another possible claim, that each vowel should exceed an expected JND when compared to the initial-syllable vowel.

regression model is the same mixed-effects model described in the previous chapter, except the dependent variable is in raw Hz rather than normalized z-scores.

5.1.2 Predicted JND with production results

5.1.2.1 Uyghur

Nine Uyghur speakers Chunja (five females, mean age: 44.4 years, range: 19-63 years) living in Chunja, Kazakhstan produced 5,927 vowels in one- to five-syllable words. Mean F2 (in Hz) by vowel and syllable number with standard deviations is reported below in Table 5.1. Observe that F2 of back vowels generally increases by syllable while F2 of front vowels does not systematically vary by syllable. Across some syllables, back vowel F2 actually decreases, e.g. /ɑ/ from syllable 4 to 5. Such variation is not predicted to be significant, since it does not follow the direction of the general pattern discussed in the previous chapter. Positional variation in Table 5.1 is also presented in Table 5.2, which compares positional shifts with estimated JNDs, assuming a threshold of 5%.

Table 5.1: By-syllable F2 in Hz of alternating vowels in Uyghur. Parentheses show standard deviations.

	Syllable 1		Syllable 2		Syllable 3		Syllable 4		Syllable 5	
	F2	n	F2	n	F2	n	F2	n	F2	n
ɑ	1213 (150)	592	1382 (175)	523	1467 (181)	469	1524 (181)	71	1485 (168)	59
ɯ			1481 (251)	114	1582 (210)	130	1572 (183)	85	1612 (133)	28
u	1010 (331)	290	1174 (275)	135	1475 (254)	21	1363 (160)	12	1516 (131)	7
æ	1891 (344)	311	1882 (302)	465	1867 (261)	310	1833 (241)	65	1947 (246)	36
i			1820 (196)	123	1799 (173)	92	1780 (179)	52	1792 (200)	12
y	1815 (289)	328	1817 (239)	256	1749 (181)	55	1759 (138)	6	1763 (154)	4
raised ɯ			1303 (208)	290	1447 (159)	113				
raised u			1261 (254)	123	1416 (196)	35				
raised i			1992 (262)	81	1900 (187)	45				
raised y			1831 (234)	105	1768 (164)	25				
Total		1,521		2,215		1,295		291		146

In Table 5.2, observe the difference in shifts that exceed the estimated JND by vowel backness. Of the 26 positional shifts presented below, 8 of 13 shifts for the [+back] vowels exceed estimated JNDs while none of 13 shifts for the [-back] vowels exceed their estimated JNDs. Observe also that perceptible shifts are largely localized to the first three syllables. All seven shifts in the first three syllables are predicted to be perceptible for the [+back] vowels below, but only one of six is predicted to be perceptible in syllables four and five. Additionally, although positional differences for /u/ in syllables 3 and 4 and /æ/ in syllables 4 and 5 are greater than predicted JNDs, they are critically in the wrong direction, since the front vowel is fronted and the back vowel is backed.

Table 5.2: By-syllable shifts in F2 in Hz compared with estimated just noticeable differences (JND; 5% threshold) in Uyghur. Shifts which exceed JNDs but in the wrong direction are marked with W.D.

		σ1- σ2			σ2- σ3			σ3- σ4			σ4- σ5		
		ΔF2	JND	ΔF2 > JND	ΔF2	JND	ΔF2 > JND	ΔF2	JND	ΔF2 > JND	ΔF2	JND	ΔF2 > JND
	ɑ	169	61	✓	85	63	✓	57	66		-39	70	
	ɯ				101	74	✓	-10	78		30	82	
[+bk]	ɯ$_{Raised}$				144	65	✓						
	u	164	51	✓	301	53	✓	-112	56	W.D.	153	59	✓
	u$_{Raised}$				155	63	✓						
	æ	-9	95		-15	94		-34	93		114	92	W.D.
	i				-21	91		-19	90		12	89	
[-bk]	i$_{Raised}$				-92	100							
	y	2	91		-68	91		10	87		4	88	
	y$_{Raised}$				-63	92							

When the statistical model from the previous chapter is used to analyze F2 (Hz), the model predicts that front vowels exhibit a non-significant shift in F2 by position [β= -0.2, $\chi^2(1)=0$, p = 1]. On the other hand, the back vowels exhibit an increase of approximately 152 Hz by syllable [β= 152.3, $\chi^2(1)$= 35.13, p < .001]. The beta value output by the statistical model, 152.1 Hz (152.3-0.2=152.1), far exceeds any of the estimated JNDs in Table 5.2. Thus, the statistical model predicts that [+back] vowel positional shifts are, in general, perceptible in Uyghur, but [-back] vowel shifts are not. Although the model predicts that the general trend toward fronting of back vowels, the model does not differentiate between the larger, earlier shifts and later, smaller shifts. In Table 5.2, the vast majority of shifts that exceed proposed JNDs are in the first three syllables, with only two shifts in the final two syllables exceeding the proposed JND.

5.1.2.2 Kazakh

Nine Kazakh speakers (seven females, mean age: 33.4 years, range: 19-49 years) living in Taldykorgan, Kazakhstan produced 5,342 vowels in one- to four-syllable words. Mean F2 (in Hz) by vowel and syllable number with standard deviations is reported in Table 5.3. Observe that F2 of the back vowels generally increases in non-initial syllables while F2 of front vowels does not systematically shift by position. Among the front vowels, though, observe that F2 increases in fourth syllables. Also, non-initial /ʏ/ shows monotonic increases in F2, similar to the back vowels. The shifts in Table 5.3 are also demonstrated in Table 5.4, which compares positional shifts with estimated JNDs, again assuming a threshold of 5%.

Table 5.3: By-syllable F2 in Hz of alternating vowels in Kazakh. Parentheses show standard deviations.

	Syllable 1		Syllable 2		Syllable 3		Syllable 4	
	F2	n	F2	n	F2	n	F2	n
ɑ	1228 (200)	277	1487 (195)	519	1565 (180)	345	1610 (144)	64
ɯ	1231 (200)	264	1293 (227)	231	1511 (264)	167	1670 (128)	35
ʊ	986 (133)	264	1182 (234)	236	1499 (279)	61	1696 (88)	6
ie	2313 (227)	258	2199 (252)	724	2189 (279)	464	2270 (240)	83
ɪ	1879 (193)	212	1847 (189)	293	1828 (199)	209	1846 (222)	53
ʏ	1625 (282)	234	1712 (173)	258	1726 (229)	76	1846 (102)	9
Total		1,509		2,261		1,322		250

The predicted perceptibility of vowel shifts in Kazakh differs largely by vowel backness. As in Uyghur, [+back] vowels exhibit more perceptible shifts than [-back] vowels. Of the nine positional shifts among [+back] vowels in Table 5.4, eight of them are predicted to be perceptible. In contrast, only one of the nine [-back] vowel shifts is predicted to be perceptible. Similar to Uyghur, the perceptibility of [+back] vowel shifts is larger in the first three syllables, but unlike Uyghur, shifts among both high vowels are predicted to exceed thresholds of perception in fourth syllables, too.

Table 5.4: By-syllable shifts in F2 in Hz compared with estimated just noticeable differences (JND; 5% threshold) in Kazakh. Shifts which exceed JNDs but in the wrong direction are marked with W.D.

		σ1- σ2			σ2- σ3			σ3- σ4		
		ΔF2	JND	ΔF2 > JND	ΔF2	JND	ΔF2 > JND	ΔF2	JND	ΔF2 > JND
[+bk]	ɑ	259	61	✓	78	74	✓	45	78	
	ɯ	62	62	✓	218	65	✓	159	76	✓
	ʊ	196	49	✓	317	59	✓	197	75	✓
[-bk]	ie	-144	116	✓	-10	110		81	109	
	ɪ	-32	94		-19	92		18	91	
	ʏ	87	81	W.D.	14	86		120	86	W.D.

As a second avenue for comparing positional shifts with estimated thresholds of perception, the statistical model from the previous chapter predicts insignificant reductions in front vowel F2 by position [β= -24.1, $\chi^2(1)$= 0.71, p = .4], but highly significant increases in back vowel F2 by position [β= 129.7, $\chi^2(1)$= 7.14, p < .001]. Thus, the statistical model predicts that back vowels will exhibit an average increase in F2 of around 106 (129.7-24.1=105.6) Hz by syllable. Such a shift would exceed all predicted JNDs in Table 5.4, suggesting that back vowel shifts are perceptible in Kazakh.

5.1.2.3 Kyrgyz

Thirteen Kyrgyz speakers (eleven females, mean age: 35.0 years, range: 18-57 years) living in Bishkek, Kyrgyzstan produced 9,374 vowels in one- to four-syllable words. Mean F2 (in Hz) by vowel and syllable number with standard deviations is reported in Table 5.5. Observe that, in contrast to the Uyghur and Kazakh data above, positional variation is much smaller in Kyrgyz. Generally speaking, back vowels show increasing F2 in non-initial syllables while front vowel show relatively symmetrical decreases in F2 in non-initial syllables. The most acoustically central vowels /ɯ ø/ do not entirely conform to these generalizations, though. For instance, F2 of /ø/ increases monotonically while F2 of /ɯ/ decreases from syllable one to two below. The shifts in Table 5.5 are also demonstrated in Table 5.6, which compares positional shifts with estimated JNDs, assuming as in Uyghur and Kazakh a threshold of 5%.

Table 5.5: By-syllable F2 in Hz of alternating vowels in Kyrgyz. Parentheses show standard deviations.

	Syllable 1		Syllable 2		Syllable 3		Syllable 4	
	F2	n	F2	n	F2	n	F2	n
ɑ	1296 (173)	514	1520 (168)	644	1539 (151)	420	1573 (125)	62
o	1054 (177)	438	1152 (139)	327	1159 (122)	156	1173 (130)	27
ɯ	1702 (243)	587	1608 (259)	496	1797 (330)	231	1886 (260)	38
u	922 (204)	276	1073 (193)	239	1139 (208)	106	1133 (166)	13
e	2246 (223)	550	2125 (226)	535	2069 (227)	341	2026 (275)	64
ø	1553 (238)	423	1588 (147)	406	1611 (140)	254	1650 (134)	65
i	2428 (227)	423	2312 (241)	385	2345 (292)	224	2326 (244)	38
y	1790 (354)	425	1768 (176)	393	1691 (274)	228	1817 (141)	46
Total		3,636		3,425		1,960		353

Comparing the magnitude of the positional shifts in Kyrgyz, which are shown in Table 5.6, with the shifts in Kazakh and Uyghur, the shifts in Kyrgyz are much smaller. Of the twelve potentially perceptible shifts for each class of [back] vowels in the language, only five [+back] vowel shifts exceed their associated

JND, and only one [-back] vowel shift exceeds its associated JND. In Kazakh, eight of nine, and in Uyghur eight of thirteen [+back] vowel shifts exceed their estimated JND, but in Kyrgyz only five of twelve [+back] vowel shifts exceed their estimated perceptual threshold. In Kyrgyz, these differences largely occur from the first to the second syllable, consistent with the analysis in McCollum (2019), which proposes that positional shifts in Kyrgyz reflect strengthening of the initial syllable.

Table 5.6: By-syllable shifts in F2 in Hz compared with estimated just noticeable differences (JND; 5% threshold) in Kyrgyz

		σ1- σ2			σ2- σ3			σ3- σ4		
		ΔF2	JND	ΔF2 > JND	ΔF2	JND	ΔF2 > JND	ΔF2	JND	ΔF2 > JND
[+bk]	ɑ	224	65	✓	19	76		34	77	
	o	98	53	✓	7	58		14	58	
	ɯ	-94	85	wrong direction	189	80	✓	89	90	
	u	151	46	✓	66	54	✓	-6	57	
[-bk]	e	-121	112	✓	-56	106		-43	103	
	ø	35	78		23	79		39	81	
	i	-116	121		33	116		-19	117	
	y	-22	90		-77	88		126	85	wrong direction

Using the predictions from the linear regression as a tool to generally compare positional shifts with estimated JNDs, front vowels exhibit a statistically significant reduction in F2 by position [β= -37.5, $\chi^2(1)$ = 4.50, p = .03]. In addition, the back vowels exhibit a significant increase in F2 by position [β= 57.5, $\chi^2(1)$ = 5.41, p = .02]. Thus, front vowels exhibit a by-syllable decrease in F2 of approximately 38 Hz, and back vowels exhibit a by-syllable increase in F2 of approximately 20 Hz (57.5 – 37.5 = 20). When these values are compared to the estimated JNDs in Table 6 above, though, these predicted positional differences do not exceed or even approximate estimated JNDs. Thus, the statistical model, in tandem with a relatively conservative estimate of a 5% JND, predicts that attested shifts in Kyrgyz are not perceptible.

5.1.3 Summary

Mean differences in F2 as well as predictions from the mixed effects regression discussed in Chapter 4 are used above to assess the incremental perceptual salience of positional shifts in Uyghur, Kazakh, and Kyrgyz. Generally speaking, [+back] vowel shifts are predicted to be more perceptually salient in Uyghur and Kazakh, while not reaching the threshold of perceptibility in Kyrgyz. Moreover, [-back] vowel shifts, generally, are not predicted to be perceivable in any of these three languages. Results from both types of comparisons above are summarized in Table 5.7. In the topmost portion of the table, the individual back vowels vowels whose positional shifts exceed predicted JNDs are counted and listed. For instance, in Uyghur, F2 of /ɑ/ and /u/ increases from syllable one to two exceed expected JNDs. Similarly, the individual front vowels whose positional shifts exceed predicted JNDs are counted and listed in the bottom portion of the table.

Table 5.7: Summary of positional shifts compared to just noticeable differences (JNDs; using a 5% threshold)

<table>
<tr><td></td><td>Perceptibility</td><td colspan="4">Uyghur</td><td colspan="3">Kazakh</td><td colspan="3">Kyrgyz</td></tr>
<tr><td rowspan="4">[+back]</td><td rowspan="3">Mean positional shift > JND?</td><td colspan="4">8/13</td><td colspan="3">8/9</td><td colspan="3">4/12</td></tr>
<tr><td>σ_{1-2}</td><td>σ_{2-3}</td><td>σ_{3-4}</td><td>σ_{4-5}</td><td>σ_{1-2}</td><td>σ_{2-3}</td><td>σ_{3-4}</td><td>σ_{1-2}</td><td>σ_{2-3}</td><td>σ_{3-4}</td></tr>
<tr><td>ɑ, u</td><td>ɑ, ɯ, u, ɯ$_{R}$, u$_{R}$</td><td></td><td>u</td><td>ɑ, ɯ, ʊ</td><td>ɑ, ɯ, ʊ</td><td>ɯ, ʊ</td><td>ɑ, o, u</td><td>ɯ, u</td><td></td></tr>
<tr><td>Predicted shifts from regression model > JND?</td><td colspan="4">✓</td><td colspan="3">✓</td><td colspan="3">✗</td></tr>
<tr><td rowspan="4">[-back]</td><td rowspan="3">Mean positional shift > JND?</td><td colspan="4">0/13</td><td colspan="3">1/9</td><td colspan="3">1/12</td></tr>
<tr><td>σ_{1-2}</td><td>σ_{2-3}</td><td>σ_{3-4}</td><td>σ_{4-5}</td><td>σ_{1-2}</td><td>σ_{2-3}</td><td>σ_{3-4}</td><td>σ_{1-2}</td><td>σ_{2-3}</td><td>σ_{3-4}</td></tr>
<tr><td></td><td></td><td></td><td></td><td>ie</td><td></td><td></td><td>e</td><td></td><td></td></tr>
<tr><td>Predicted shifts from regression model > JND?</td><td colspan="4">✗</td><td colspan="3">✗</td><td colspan="3">✗</td></tr>
</table>

The findings above support a general distinction between Uyghur and Kazakh on the one hand, and Kyrgyz on the other. Attested incremental positional shifts in Uyghur and Kazakh are, generally speaking, more likely to be perceptually salient than those in Kyrgyz. Returning to the predictions from the beginning of this section, only Kazakh approaches the strongest requirement, namely that all positional shifts should exceed predicted JNDs. Recall that the weaker requirement demands only that some positional shifts must exceed their associated JNDs to be considered perceivable, and thus potentially phonological in nature. All three languages, to varying degrees, meet this weaker requirement. In sum, these findings provide general support for a distinction between these languages, but these estimates cannot in and of themselves really answer the question at hand, are these positional shifts perceivable. In part, these results cannot definitively establish the perceptibility of these shifts because JND estimates are based on perception research on other languages, mainly English, and may not translate to these particular languages. Manuel (1990) argues that variation in production is predictable based on the size and density of a language's vowel space. In languages with more vowel contrasts coarticulatory effects are smaller. One plausible prediction from Manuel's findings is that vowel perception is similarly modulated by the number of vowel contrasts in a given language. In the three languages examined above, JNDs for the non-low vowels may thus be smaller, given the nature of each language's vowel space. Therefore, it would be much more informative to assess the perceptual capacities of native speakers. The next section does just that, investigating perceptual thresholds in each of these three languages.

5.2 Perceptual thresholds

This section describes a perception experiment conducted during fieldwork to determine general perceptual boundaries between alternating vowels in Uyghur, Kazakh, and Kyrgyz. Section 5.2.1 lays out a set of predictions, Section 5.2.2 describes the methods used, and Section 5.2.3 discusses results from each language. Moving on from there, Section 5.2.4 discusses results and how they relate to the larger issue of perceptual salience and phonological phenomena.

5.2.1 Predictions

It has been argued for some time that consonant perception is, in some ways, categorical (Eimas et al. 1971; Eimas & Corbit 1973). Eimas and colleagues find that the perception of stop voicing, for instance, is relatively categorical along the voice onset time (VOT) continuum. An adapted plot from Eimas & Corbit (1973:104) is shown below in Figure 5.1. In the plot, all stimuli with VOT less than 30 milliseconds are perceived as /b/ while all stimuli with VOT greater than or equal to 50 milliseconds are perceived as /p/ with a steep curve spanning the relatively small perceptual boundary between the two.

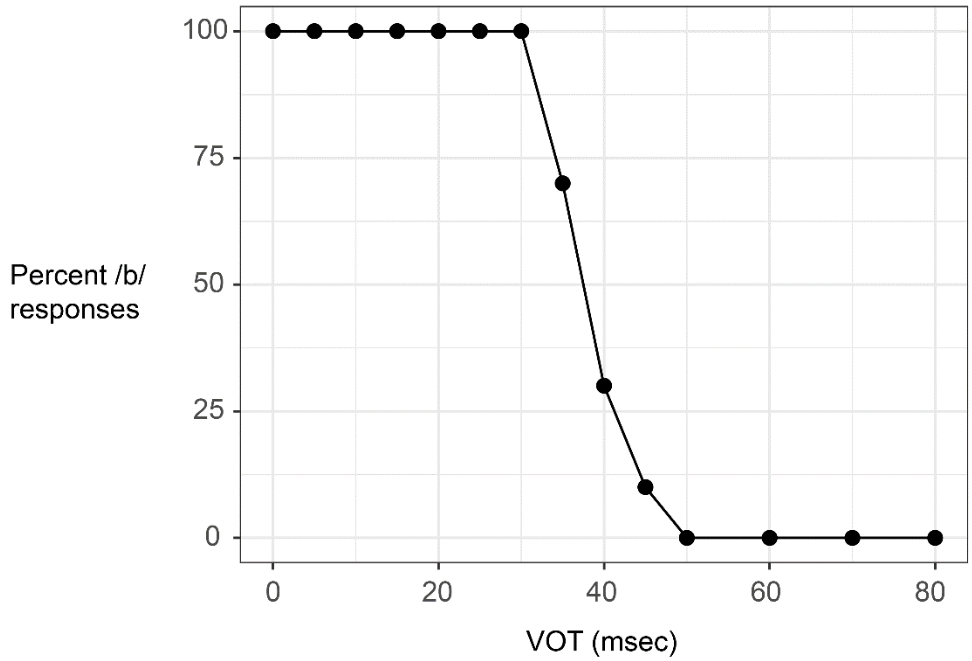

Figure 5.1: Stop voicing perception in terms of voice onset time (VOT) for the English /b-p/ contrast (adapted from Eimas & Corbit 1973:104)

In contrast to consonant perception, vowel perception is far more gradient. The perception of vowel qualities exhibits much more indeterminacy between vowel categories than perception of consonant voicing (Fry et al. 1962; Stevens et al. 1969). For instance, in Fry et al.'s (1962) seminal study on vowel perception, they assessed native English speakers' perception of synthesized vowel stimuli. Vowels differed in both F1 and F2, ranging from /ɪ/ to /ɛ/ to /æ/. Each vowel phoneme exhibits a peak of responses surrounding its perceptual center with, crucially, gradually decrease rates of perception as distance from that center increases. This is evident in Figure 5.2 below, recreated from data in Fry et al. (1962). As F1 increases and F2 decreases along the continuum described in Fry et al. (1962), the rate of /ɪ/ responses gradually diminishes to zero. At the other end of the continuum, rates of /æ/ responses gradually increase to 100%, and in the middle, /ɛ/ responses dominate.

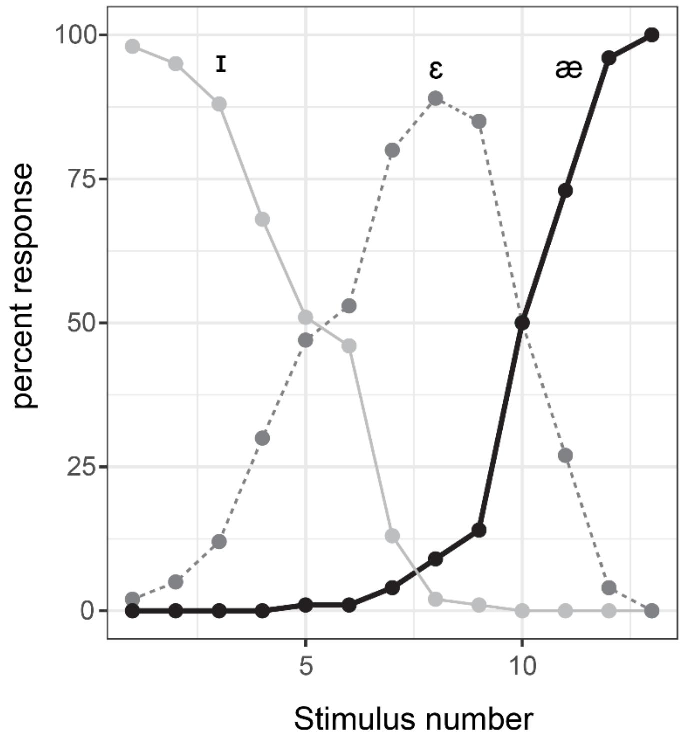

Figure 5.2: Percent vowel response by stimulus number (continuum varying between /ɪ/ and /æ/)

Vowel perception plots can often be divided into three regions for a two-way vowel contrast (e.g. /ɪ/-/ɛ/), which is schematized below in Figure 5.3. In Figure 5.3, a sigmoid curve approximating the perception of two neighboring vowel categories, A and B, is shown. I suggest there are three regions of such a plot. Two regions represent relatively consistent percepts; those areas are shaded in grey. The third area is the intermediate portion of the acoustic space wherein more uncertainty exists. In this intermediate space, both inter- and intra-speaker variability is often more pronounced, as speakers attempt to categorize stimuli that do not fall neatly into Category A or B.

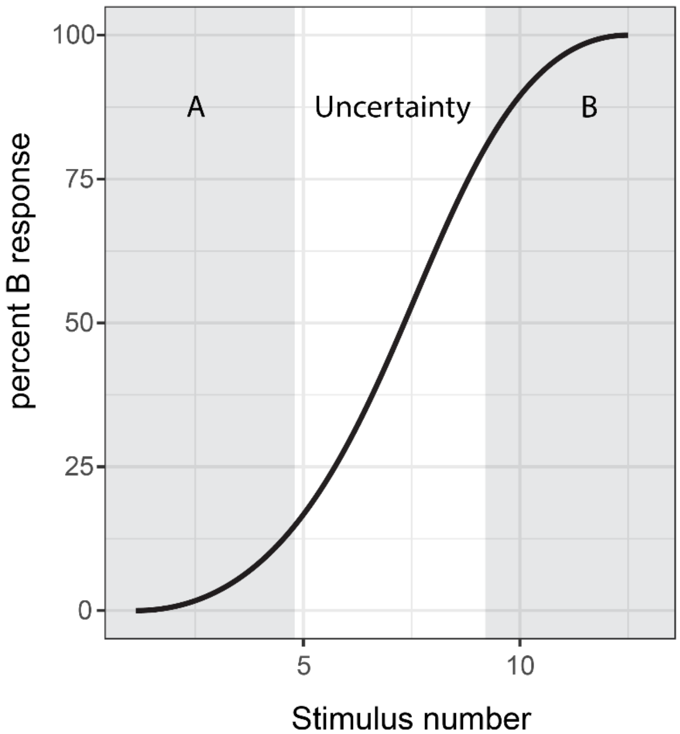

Figure 5.3: Schematized perception of two vowel categories, /A/ and /B/

If the perceptibility of the positional shifts examined is conceived of in these terms, one strong prediction is that all perceptible shifts should cross a boundary between a relatively constant percept and the zone of uncertainty. In terms of Figure 5.3, if vowel /A/ shifts toward vowel /B/ in some positions, and if Stimuli 1 and 2 represent categorical tokens of vowel category /A/, then vowel /A/'s positional shifts are only perceptually salient if they result in acoustic qualities that approximate or exceed Stimulus 5, which serves as an approximate boundary between the zone of uncertainty and the domain of consistent /A/ perceptual judgments. Instead, if /B/ exhibits positional shifts toward /A/, and if Stimuli 14 and 15 represent categorical tokens of /B/, then category /B/'s shifts are perceptually salient if and only if they approximate or cross over the acoustic qualities of Stimulus 9, the approximate boundary between the zone of consistent /B/ percepts and the zone of uncertainty. Uncertainty is a key diagnostic for the perceptibility of a given shift. If native speakers are not able to consistently judge a given vowel token as either /A/ or /B/, but exhibit variation both within- and across-participants, this suggests that the surface quality of that vowel is not categorically /A/ or /B/, but something between the two. I will treat positional shifts that produce vowels that cross over into this zone of uncertainty as perceivable. Note that this complements the incremental assessment of positional shifts above. The question here does not relate to syllable-to-syllable differences so much, but rather to the relationship between each member of an alternating pair. Under the criterion that phonological patterns must be perceivable, shifts that span from one zone of reliability into the zone of uncertainty are candidates for being phonological. In contrast, vowel shifts that do not extend into the zone of uncertainty will be regarded as phonetic, since they are not obviously perceivable. This sort of variation is construed as normal within-phoneme variation.

5.2.2 Methods

5.2.2.1 Stimuli

Stimuli were generated for alternating [back] pairing in each language. Four pairs of stimuli were created for Kyrgyz, /ɑ/-/e/,/o/-/ø/, /ɯ/-/i/, and /u/-/y/, three pairs were created for Kazakh, /ɑ-ie/,/ɯ-ɪ/,and /ʊ-ʏ/, and two pairs for Uyghur, /ɑ-æ/, and /u-y/. Stimuli were generated from a single female Kazakh speaker residing in San Diego, CA. Given the similarity of the vowel spaces in these languages and lack of access to speakers of other languages, modifying her tokens to more accurately reflect the phonetic properties of vowels in the other two languages was preferable to synthesizing vowels. The resultant sounds were more natural sounding than computer-synthesized vowel, and only minor modifications were necessary to approximate Uyghur and Kyrgyz vowel qualities.

Vowels were synthesized using a script from Matt Winn (Winn & Litovsky 2015; see also mattwinn.com/praat.html). The script takes two input sound files and synthesizes a specified number of sounds on a continuum between the two inputs. For this study, a seven-step continuum was created by manipulating F1 and F2, with each endpoint representing an input sound file. All other acoustic parameters, e.g. higher formants, duration, were held constant across each continuum. Between each harmonic pairing five roughly equidistant intermediate qualities were created. Three acoustic adjustments were made to better reflect the phonetic characteristics of Kyrgyz and Uyghur. First, F1 of Kazakh /ʊ/ was reduced to more closely approximate /u/ in the other two languages. Second, F1 of Kazakh /ɪ/ was reduced prior to continua construction via the formant modulation portion of Winn's script. Similarly, F2 was increased to more closely reflect /i/ in Kyrgyz; third, the initial portion of Kazakh /ie/ was excised, leaving a more monophthongal [e]-like sound, which better reflects Kyrgyz /e/. Finally, Kazakh /æ/ does not regularly alternate with /ɑ/ for harmony, but was used to create Uyghur /æ/.

All sounds synthesized were embedded within an actual (C)VC word from each language producecd by a Kazakh speaker. Two words were created for each member of each alternating pair. A list of sample stimuli for each language is presented below in Table 5.8. Additionally, since all pairings are constrained by phonological height (or potentially length in Kazakh), vowel duration does not substantially differ within each pairing. For consistency, all vowel durations within each pairing were identical, and the intensity for all stimuli was scaled to 75 dB and overlaid with an identical, gradually falling f0 track. F1-F2 plots for the Kazakh /ɪ-ɯ/ and /ʏ-ʊ/ continua are shown in Figure 5.4. Observe that /ɪ/ and /ɯ/ exhibit differences in both F1 and F2, and so the continua between the two forms a relatively linear interpolation between the two prototype endpoints. In contrast, observe that no such F1 differences exist between Kazakh /ʏ/ and /ʊ/, so the continuum points produced by the script center around 570 Hz throughout the continuum. For the /ɑ-ie/ contrast in Kazakh, the continuum involved differences in both F2 at vowel midpoint and the slope of F2. F2 decreases throughout /ie/ but is constant across /ɑ/; intermediate stimuli exhibited higher F2 at vowel midpoint and more negative slope as they approached /ie/.

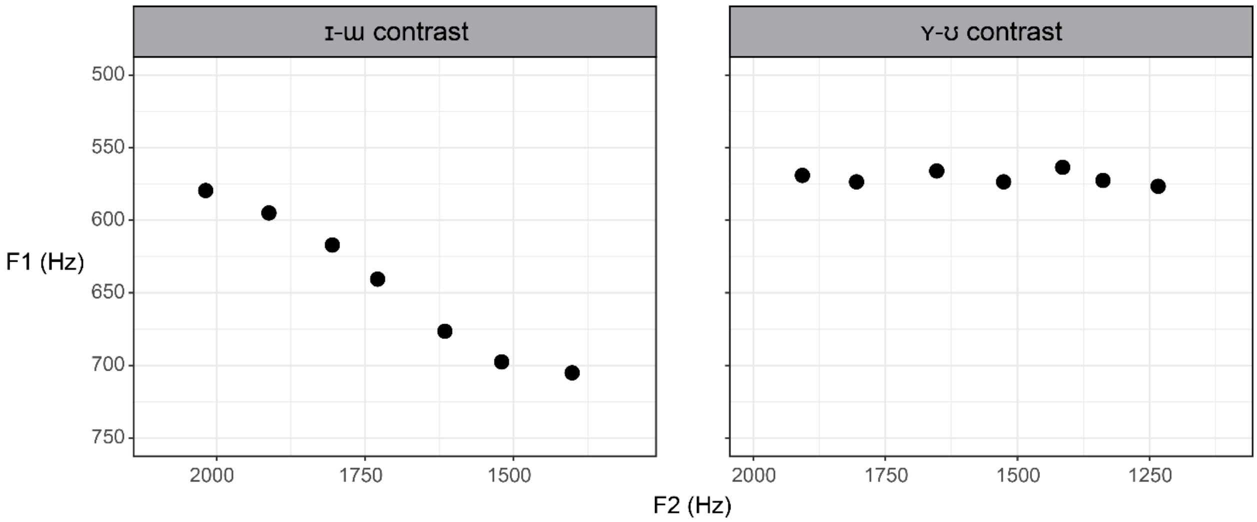

Figure 5.4: F1-F2 plots in Hz for continua between Kazakh /ɪ-ɯ/ and /ʏ-ʊ/

Table 5.8: Sample list of target words for perception study

Language	[-back]	Gloss	[+back]	Gloss
Kyrgyz	sez	feel	sɑz	swamp
	øt	gall bladder	ot	fire
	siz	2S (formal)	sɯz	damp place
	tyz	direct	tuz	salt
Kazakh	biet	face	bɑt	dive
	tɪs	tooth	tɯs	outside
	tʏs	color	tʊs	side
Uyghur	bæʃ	five	bɑʃ	head
	yʧ	three	uʧ	fly

5.2.2.2 Task

The perception task was deployed using Praat (Boersma & Weenink 2015). Participants were first instructed to follow the directions presented on the screen of the researcher's laptop computer. Directions were displayed in the target language (using the Uyghur variant of Cyrillic, not Arabic for Uyghur participants). A single variant of a stimulus word was presented auditorily to each participant. Then the participant was asked to select which of two orthographic forms was most similar to the stimulus item. The Uyghur /ɯ/-/i/ contrast was not tested because neither the modified Cyrillic Uyghur orthography used in Kazakhstan nor the modified Arabic script used in China represent this contrast (Nadzhip 1971; Hahn 1991; Engesaeth et al. 2009).

Participants first engaged in a practice round, in which they were forced to choose between two non-alternating pairs, like /ɑ/-/i/. This allowed participants to familiarize themselves with the two-alternation forced choice task using easier pairs. Moreover, by introducing the participants to the task this way allowed each speaker to hear the full range of the Kazakh speaker's (modified) vowel space, which ensured that listeners could normalize formant values for that speaker, thus allowing for a more direct comparison with the z-score normalized production results. In each language, the number of vowels each participant was exposed to during practice mirrored their own language. Thus, Uyghur speakers were exposed only to vowels in their language, and Kyrgyz speakers were exposed only to vowels present in their language. So, Kyrgyz speakers were not exposed to Uyghur /æ/, nor were Uyghur speakers exposed to Kazakh /ie/.

After completing the practice portion of the task, speakers heard seven stimulus items per word pair. Since there were two word pairs for each harmonic alternation, Kyrgyz speakers were exposed to 56 target items (4 vowel pairs * 7 continuum points * 2 lexemes), Kazakh speakers were exposed to 42 target items (3 vowel pairs * 7 continuum points * 2 lexemes), and Uyghur speakers were exposed to 28 target items (2 vowel pairs * 7 continuum points * 2 lexemes). Filler items consisted of pairs with mismatch consonants, e.g. Kyrgyz /tuz/ 'salt' vs. /tus/ 'side' or by non-relevant vowel contrasts, e.g. Kazakh /tɑs/

‘stone’ vs. /tɪs/ ‘tooth.’ Approximately 30 percent of all stimulus items were fillers, producing a total of 72 stimuli for Kyrgyz, 54 stimuli for Kazakh, and 36 stimuli for Uyghur. All stimuli were randomly ordered. For each continuum, fifty percent of [+back] response boxes were on the left. Participants used a mouse to make their responses. A sample image with instructions and response options is shown in Figure 5.5.

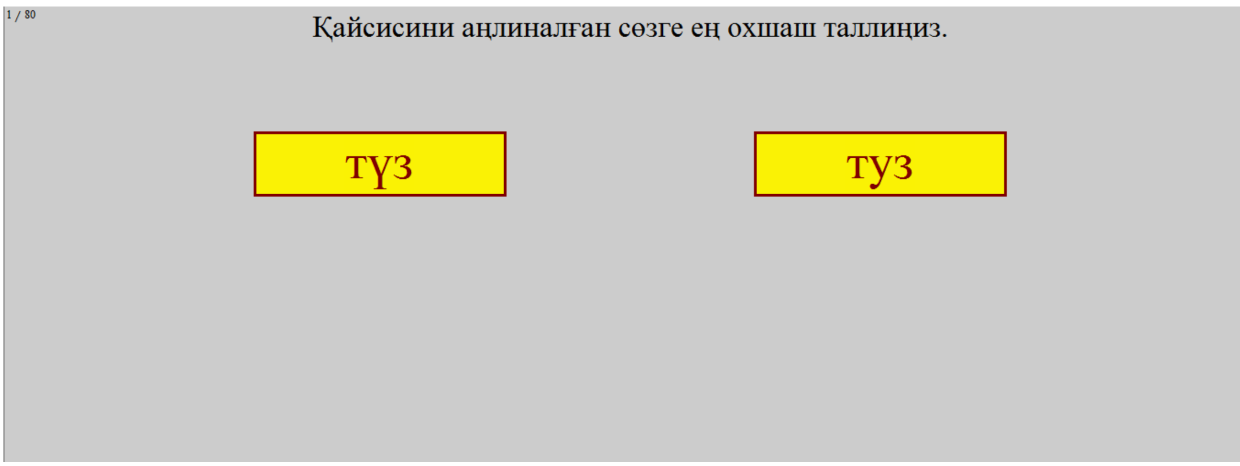

Figure 5.5: Sample image with instructions and response options from the Uyghur perception test. The prompt reads Қайсисини аңлиналған сөзге ең охшаш таллиңиз, which translates as “Select which of the two is most similar to word you heard.” In this example, the possible responses are түз “straight” and туз “salt”.

5.2.2.3 Participants

After completing the production task detailed in the previous chapter, speakers participated in the perception task. However, one Uyghur participant and one Kazakh participant declined to participate in the task. As a result, eight Uyghur speakers, eight Kazakh speakers, and thirteen Kyrgyz speakers participated in the study. For simplicity, production data from all participants is used below for production-perception comparisons.

5.2.3 Results

Responses are presented below for each harmonic pairing for each language. Responses are first presented on their own, and then they are compared with production data to determine whether or not vowel production, specifically the production of [+back] vowels in Uyghur and Kazakh results in vowel qualities that occupy the zone of uncertainty between reliable [-back] and [+back] percepts.

5.2.3.1 Kyrgyz

Results are plotted below in terms of percent [-back] response according to each harmonic pairing. The first pairing is /ɑ/-/e/. As can be seen in Figure 5.6, the two rightmost points on the continuum, indicated by dots along the plot line, were perceived as /ɑ/. The leftmost three points were perceived in all instances as /e/, and the two intervening points occupy the zone of uncertainty. Most of these two points on the continuum were perceived as /e/, but not always. This example helps to illustrate how to define the zone of uncertainty. In this plot, the zone of uncertainty spans from just about 0.0z to 0.8z (1583 Hz to 1951 Hz). Zones of reliable percepts extend right and left from these points, corresponding to the reliable percepts of /ɑ/ and /e/ respectively. As a reminder, if the positional shifts in /ɑ/ and /e/ in Kyrgyz are perceptually salient, they should cross into the zone of uncertainty.

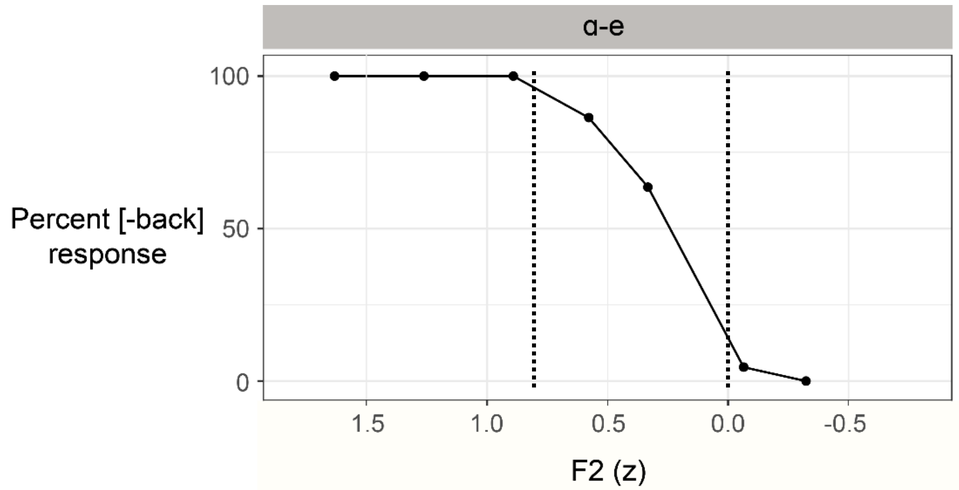

Figure 5.6: Percent [-back] responses to stimuli on Kyrgyz /ɑ/-/e/ continuum with zone of uncertainty (marked by dotted vertical lines)

With these preliminaries in order, these perception results can be compared to production results from Chapter 4. This comparison for the /ɑ/-/e/ contrast is shown below in Figure 5.7. For the back vowel, no mean values cross into the zone of uncertainty. In other words, all mean values of /ɑ/, even in the fourth syllable, should be reliably perceived as a back vowel. For the front vowel, the fourth-syllable mean of /e/ just crosses into the zone of uncertainty, suggesting that in fourth syllables, /e/ may not always be perceived as a front vowel. However, considering that the nearest stimulus in the zone of uncertainty was perceived as [-back] roughly 87% of the time, and that its F2 (z) value is noticeably lower than the fourth-syllable mean of /e/ the perceptual salience of /e/'s positional shift is not clear cut. Also, the particular thresholds posited here depend on the number of continuum steps generated. If, for instance, 21 steps were generated, these thresholds might differ slightly, and the fourth-syllable mean of /e/ might not extend in the zone of uncertainty. Though fourth-syllable /e/ crosses the posited threshold into the zone of uncertainty, the salience of this shift is still questionable.

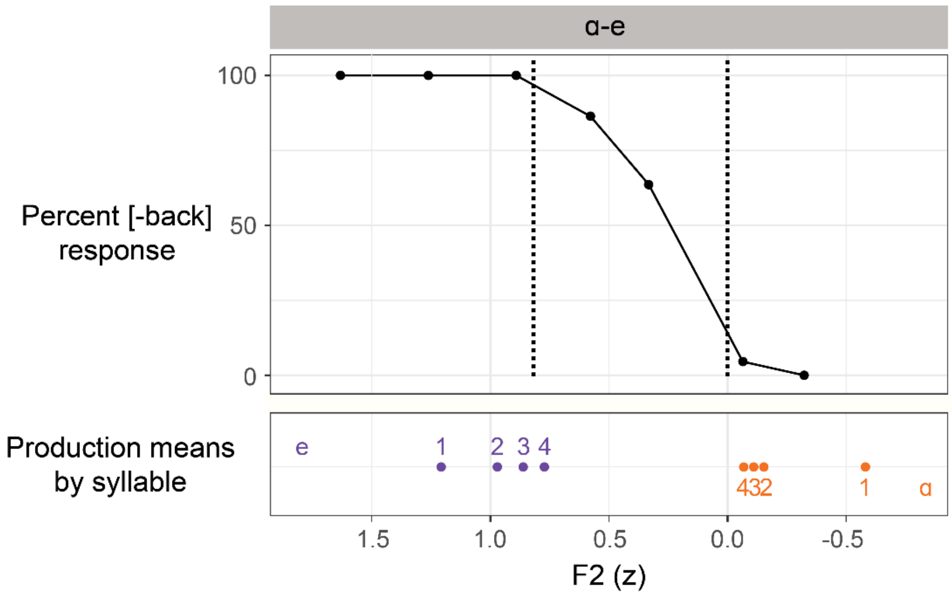

Figure 5.7: Percent [-back] responses to stimuli on Kyrgyz /ɑ/-/e/ continuum compared to production means (zone of uncertainty is marked by dotted vertical lines)

One thing to note here is that the plot in Figure 5.7 compares F2 only, ignoring F1 differences. For a contrast like /ɑ/-/e/, the F1 dimension probably contributes significantly to the distinction between the two. The /ɑ/-/ie/ contrast also faces this same limitation, although the other contrasts considered involve front and back vowel pairs that are roughly equivalent in F1, eliminating this potential confound.

Moving on to the second alternation under study, /o/-/ø/, perception results are shown in Figure 5.8. The rightmost three points on the continuum of stimuli were reliably perceived as /o/, while the leftmost two points were most reliably perceived as /ø/. Note that no stimuli were always perceived as a front vowel. Even the most anterior stimulu item was only perceived as front vowel 91% of the time. I assume this derives from either the stimulus creation process, or simply from noise in the response data. Regardless, the zone of uncertainty in Figure 5.8 spans from around -0.25z to around 0.35z (1468 Hz to 1744 Hz).

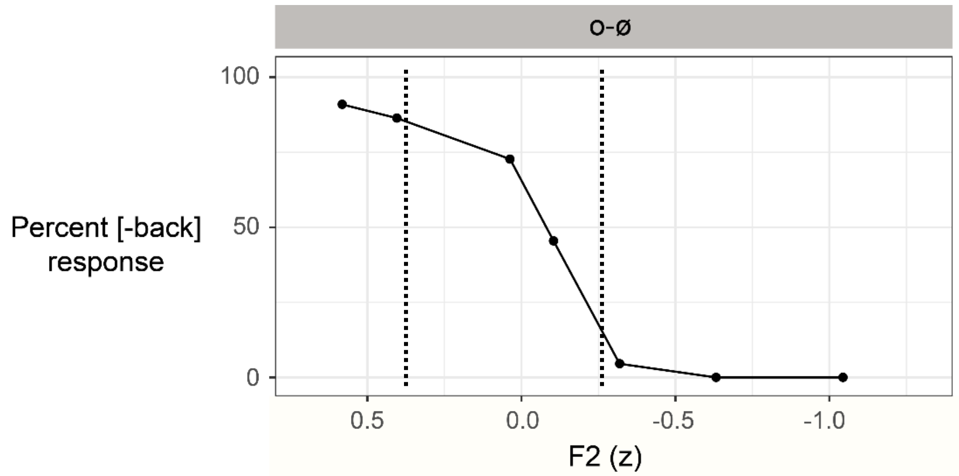

Figure 5.8: Percent [-back] responses to stimuli on Kyrgyz /o/-/ø/ continuum with zone of uncertainty (marked by dotted vertical lines)

Perception results are compared to production data for /o/ and /ø/ below. In Figure 5.9, production values for /o/ never approach the zone of uncertainty. They should thus always be perceived as back vowels, regardless of position in the word. However, all production values for /ø/ correspond to the center of the zone of uncertainty in perception. There is an interesting mismatch to be noted here. Two stimulus items for the perception study had much higher F2 and were still perceived as [+back] vowels by some participants. This might suggest that the stimuli did not have sufficiently high F2 to trigger a more reliable [-back] percept. However, all production values occur at points in which perceptual judgment were least reliable.

Two more points should be noted before examining the next alternation. First, the trajectory of /ø/'s positional shifts is the opposite of other shifts obvserved thus far. Both /ɑ/ and /e/, as well as /o/ shift from more reliable perceptual regions of the vowel space toward regions of more uncertainty in non-initial syllables. However, in Figure 8, initial-syllable /ø/ is predicted to be least reliably perceived as [-back], and fourth-syllable /ø/ is predicted to be most reliably perceived as [-back]. This is an interesting difference between /ø/ and the other vowels investigated thus far. Second, it is possible to compare the size of the zone of uncertainty with the difference in F2 between positional variants of /o/ and /ø/. In other words, the size of the normalized region between /o/ and /ø/ may persist across both production and perception. The predicted perceptual zone of uncertainty spans 0.6z. Mean F2 of fourth-syllable /o/ has the highest F2, -0.81z, for that vowel category and mean F2 of initial-syllable /ø/ has the lowest F2 for that vowel category, -0.08z. The difference between the two is 0.73z, slightly larger than the size of the zone of uncertainty from perception results. This provides some evidence that despite differences in the perception and production of /o/ and /ø/, perceptual categories are maintained in production.

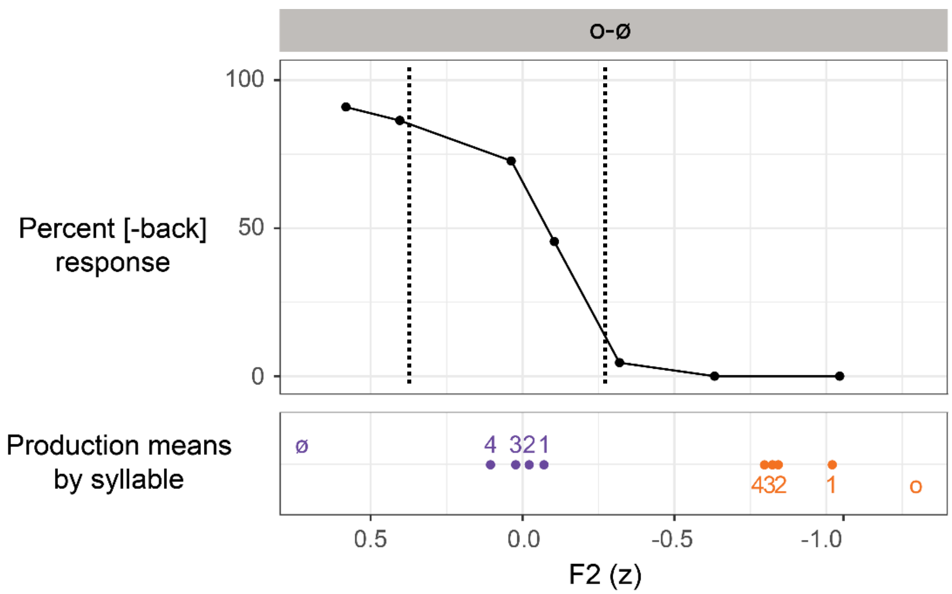

Figure 5.9: Percent [-back] responses to stimuli on Kyrgyz /o/-/ø/ continuum compared to production means (zone of uncertainty is marked by dotted vertical lines)

The first high vowel alternation to consider is /ɯ/-/i/. In Figure 5.10 below, perceptual judgments on stimuli with higher F2 resemble those concerning /ø/ above. No stimuli elicited 100% perception of a front vowel, but the leftmost two stimuli triggered fairly reliable [-back] responses, and are considered within the zone of reliable [-back] percepts. The rightmost four stimuli always prompted [+back] responses, and will be considered the zone of reliable [+back] percepts. The zone of uncertainty extends from 0.35z to 0.7z (1744 Hz to 1905 Hz). Only a single stimulus items elicited variable judgments, and was perceived as [+back] around 40% of the time.

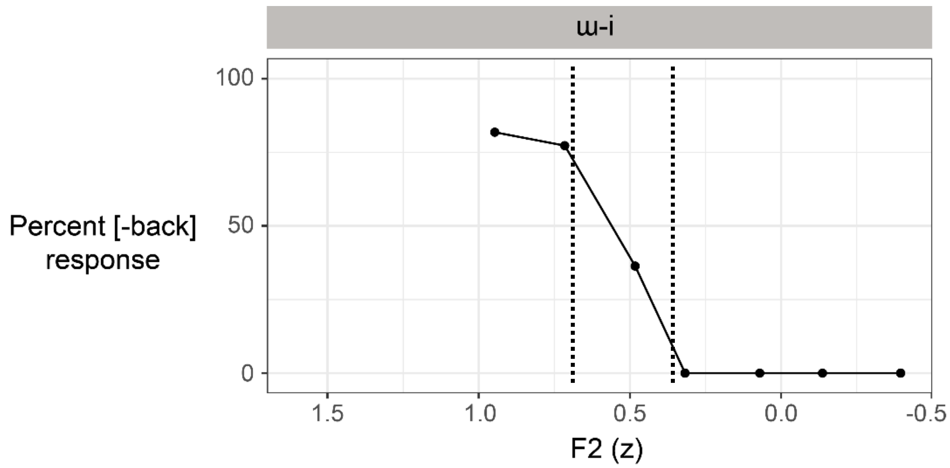

Figure 5.10: Percent [-back] responses to stimuli on Kyrgyz /ɯ/-/i/ continuum with zone of uncertainty (marked by dotted vertical lines)

These results can now be compared with production means from the previous chapter. In Figure 5.11 it is obvious that the stimuli created for /i/ did not adequately approximate the frontness of Kyrgyz /i/. These stimuli, modified from Kazakh /ɪ/, exhibited far lower F2 than actual Kyrgyz production values. For the first time, a [+back] vowel's positional shifts stretches into the predicted zone of uncertainty. Mean F2 of third-syllable /ɯ/ almost reaches the threshold near 0.35z and fourth-syllable /ɯ/ exhibits F2 of 0.5z, well beyond the predicted threshold of uncertainty. This is the best evidence yet for a positional shift in Kyrgyz exhibiting true perceptual salience.

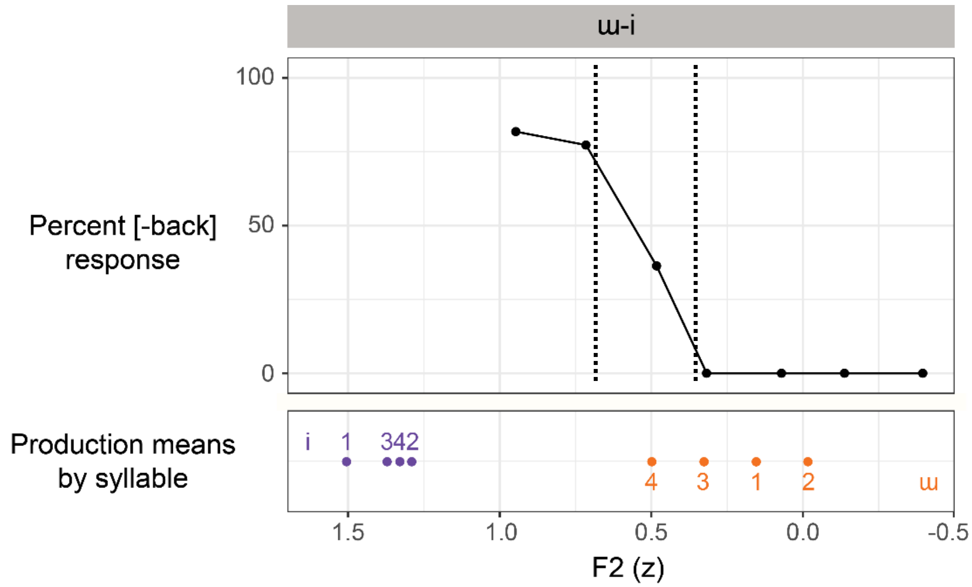

Figure 5.11: Percent [-back] responses to stimuli on Kyrgyz /ɯ/-/i/ continuum compared to production means (zone of uncertainty is marked by dotted vertical lines)

However, word-final /ɯ/, as discussed in the previous chapter, is subject to a gradient fronting process similar to the categorical fronting of word-final high vowels in Uyghur, also discussed in the previous chapter. In Figure 5.12, final-syllable /ɯ/ in open syllables exhibits a mean F2 of 0.5z, while final-syllable /ɯ/ when followed by a consonant exhibits F2 below 0, -0.06z. Figure 5.12 demonstrates that F2 of /ɯ/ is consistent across final and non-final closed syllables. In open syllables, /ɯ/ is more fronted, with additional fronting in final open syllables.

The high F2 of both third- and fourth-syllable /ɯ/ is explainable as word-final fronting. Third-syllable /ɯ/ occurred in the word-final syllable in 216 of 230 tokens recorded. Thus both third- and fourth-syllable means are affected by a separate pattern that produces fronting, which undermines the potential perceptibility of backness harmony-related shifts in vowel quality. If positional shifts in the third- and fourth-syllables are perceptually significant, it is likely due to word-final high vowel fronting and not backness harmony.

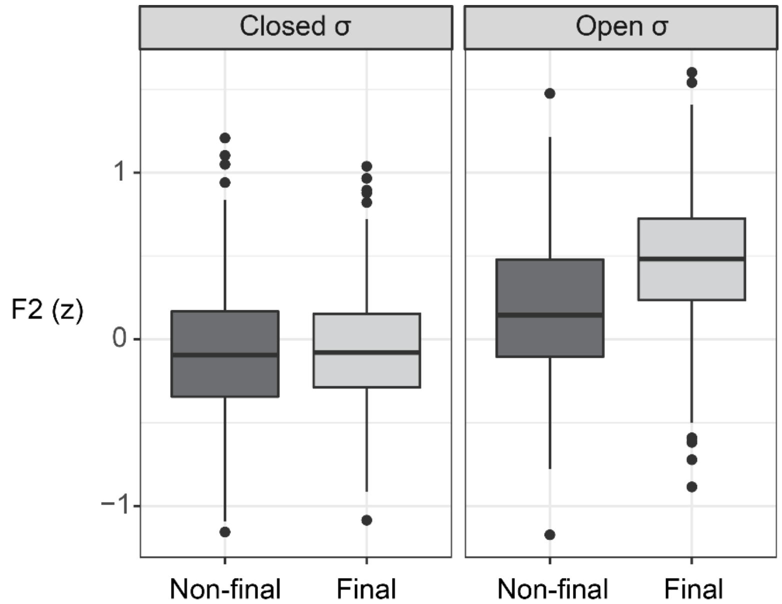

Figure 5.12: F2 (z) of /ɯ/ by position and syllable type in Kyrgyz

The final vowel alternation left to discuss is /u/-/y/. Results from the perceptual study are presented in Figure 5.13. In this figure the rightmost two points on the seven-point continuum were always perceived as [+back] vowels while the leftmost three points were always perceived as [-back] vowels. In between these two zones of reliability, two points exist in the zone of uncertainty, which span from around -0.48z to 0.05z (1362 Hz to 1606 Hz).

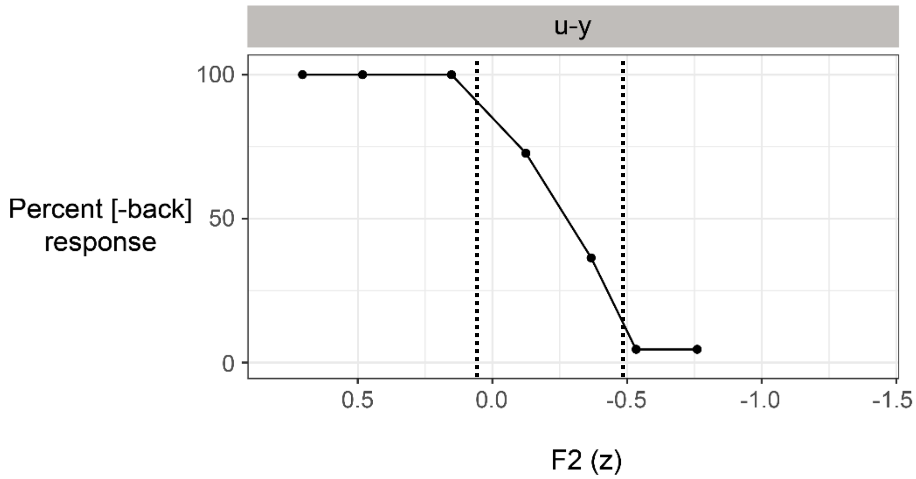

Figure 5.13: Percent [-back] responses to stimuli on Kyrgyz /u/-/y/ continuum with zone of uncertainty (marked by dotted vertical lines)

When production means are compared to perceptual results for the /u/-/y/ alternation it is clear that neither positional shifts for /u/ or /y/ cross into the zone of uncertainty. In Figure 5.14, positional shifts for /u/ never approximate F2 of stimulus items ever perceived as [-back], and shifts for /y/ never diverge from stimuli which were always perceived as [-back] during perception testing.

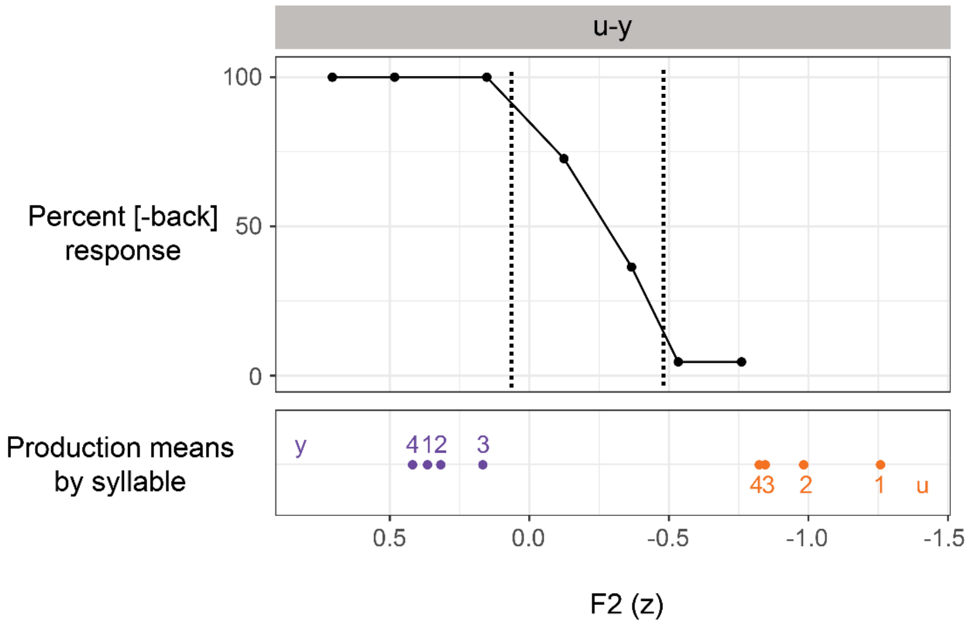

Figure 5.14: Percent [-back] responses to stimuli on Kyrgyz /u/-/y/ continuum compared to production means (zone of uncertainty is marked by dotted vertical lines)

In sum, no positional shifts in vowel quality clearly produced vowels crossed a perceptual threshold. Evidence from /e/ came closest to crossing such a boundary. Also, the interpretation of the /o-ø/ pairing is tentative, given the production-perception mismatches above. Returning to the prediction above, if phonological patterns must, at minimum be perceivable, the patterns of positional variation in Kyrgyz do not generally meet this requirement, and as a result, must be regarded as phonetic in nature.

5.2.3.2 Kazakh

Unlike Kyrgyz, the non-high vowels are not targets for rounding harmony in Kazakh. Therefore, only three alternations are attested in non-initial syllables, /ɑ/-/ie/, /ɪ/-/ɯ/, and /ʊ/-/ʏ/. Perceptual judgments to stimuli created on a continuum between /ɑ/ and /ie/ are shown in Figure 5.15. The leftmost three stimuli were reliably perceived as the [-back] vowel while the rightmost two stimuli were always perceived as the [+back] vowel. In between these two zones of reliability, two stimulus items were variably perceived as /ie/ or /ɑ/, within the zone of uncertainty. The edges of the zone of uncertainty are shown below, ranging from 0.0z to approximately 0.85z. (1583 Hz to 1974 Hz).

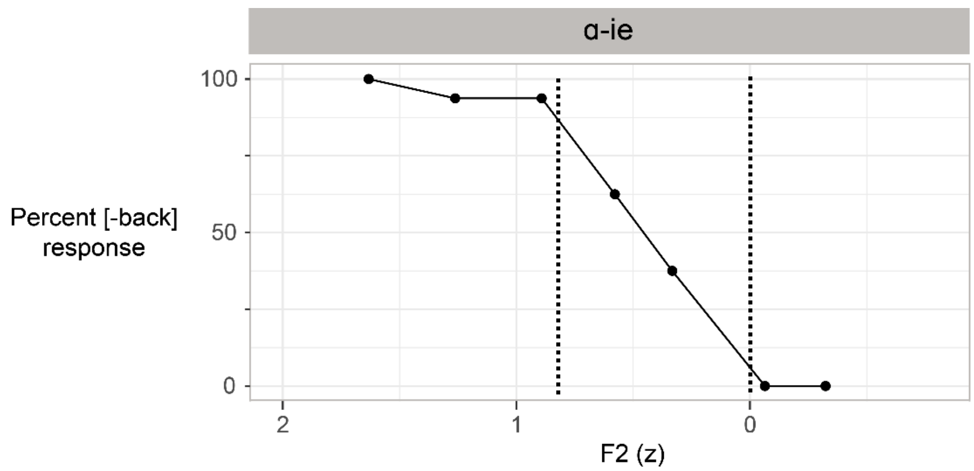

Figure 5.15: Percent [-back] responses to stimuli on Kazakh /ɑ/-/ie/ continuum with zone of uncertainty (marked by dotted vertical lines)

Perceptual judgments are compared against production means in Figure 5.16. Observe that mean F2 for the front vowel is firmly positioned within the zone of reliable [-back] percepts. On the other hand, positional variants of /ɑ/ extend into the zone of uncertainty. Mean F2 of second-syllable /ɑ/ nears the threshold, and mean F2 values for third- and fourth-syllable /ɑ/ are situated outside the zone of reliable [+back] percepts. I thus predict that normal third- and fourth-syllable variants of /ɑ/ should be perceived as a front vowel in some instances. To make my claims clear, in the previous chapter I did not claim that fronted back vowels in production must approximate a front vowel for gradient harmony to produce phonological results. Likewise, I am not claiming here that a fronted vowel must be categorically perceived as front for the effect to be phonological. Instead, I contend that the shift much simply be large enough so that the output vowel is not longer perceptually equivalent to categorical /ɑ/. It is not that gradience must produce the opposite feature value in production or perception, but simply that it must produce something substantially different than the expected output if harmony were categorical.

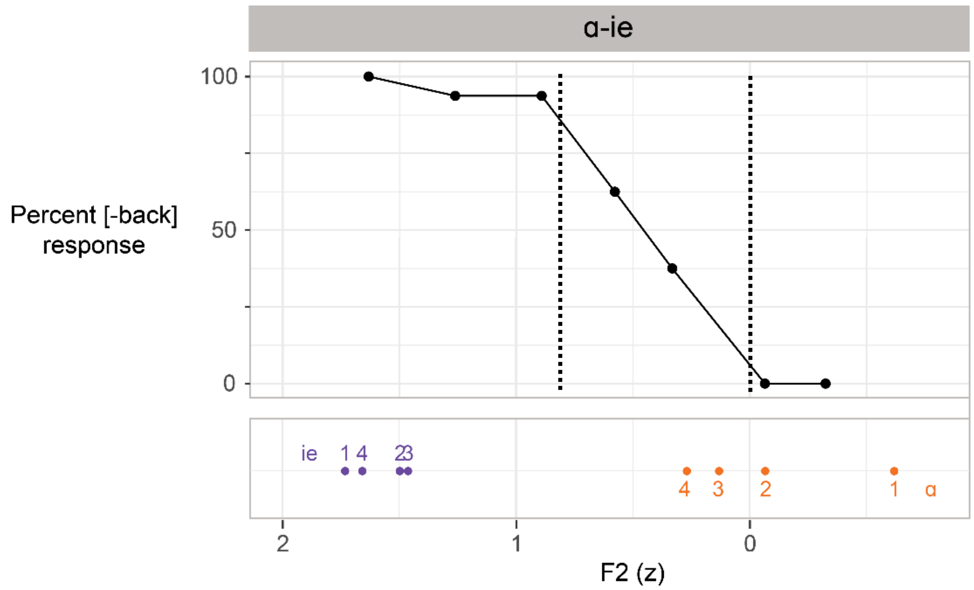

Figure 5.16: Percent [-back] responses to stimuli on Kazakh /ɑ/-/ie/ continuum compared to production means (zone of uncertainty is marked by dotted vertical lines)

In addition to the low back vowel, the high back vowels exhibit similar fronting patterns, as discussed in the previous chapter. Above, the perception results suggest that the low vowel shifts produce third- and fourth-syllable vowels that are perceptibly different from prototypical /ɑ/, satisfying my claim that phonological effects must be perceivable. To examine the perception of the unrounded vowel /ɯ/ and its front counterpart, /ɪ/, consider Figure 5.17 below. The leftmost two stimuli were typically perceived as [-back]. It is unclear why the stimulus with the highest F2 was perceived as a front vowel less often that next stimulus item. Nonetheless, these two vowels were each perceived as front vowels over 87% of the time, and the 100% perception of the second-to-leftmost stimulus item suggests that the dip in [-back] percepts is not meaningful. The rightmost three stimuli were always judged to be [+back], and like above, two stimuli occupied the zone of uncertainty, with variable [-back] and [+back] judgments during the study. The zone of uncertainty extends from roughly 0.15z to 0.62z (1652 Hz to 1868 Hz).

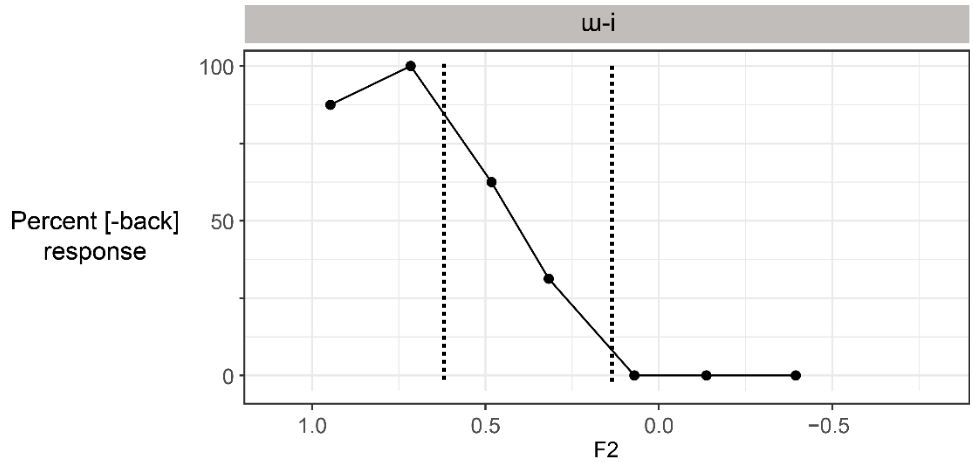

Figure 5.17: Percent [-back] responses to stimuli on Kazakh /ɯ/-/ɪ/ continuum with zone of uncertainty (marked by dotted vertical lines)

Mean F2 of /ɪ/ shows little variation in Figure 5.18, suggesting that it is reliably perceived as a front vowel in all syllables. In contrast, /ɯ/ exhibits much higher F2 in non-initial syllables, with mean F2 of third-syllable /ɯ/ approaching, and mean F2 of fourth-syllable /ɯ/ exceeding the threshold between reliable [+back] judgments and the zone of uncertainty. In other words, a normal token of fourth-syllable /ɯ/ is predicted to be perceived as a front vowel in some instances.

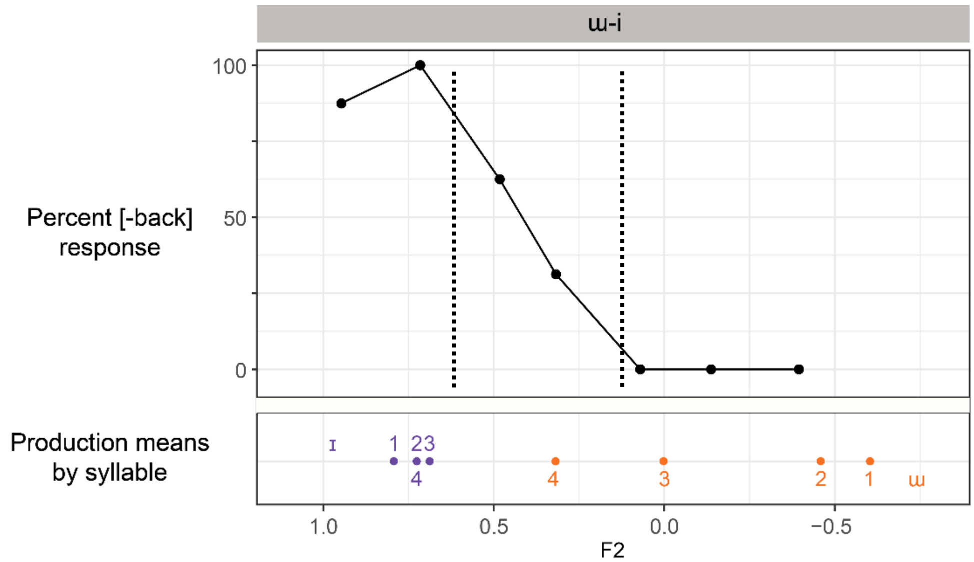

Figure 5.18: Percent [-back] responses to stimuli on Kazakh /ɯ/-/ɪ/ continuum compared to production means (zone of uncertainty is marked by dotted vertical lines)

As for the final backness alternation in Kazakh, Figure 5.19 presents results from the perception task for /ʊ/-/ʏ/. The rightmost three stimuli below were always perceived as /ʊ/ while the leftmost stimuli were almost always perceived as /ʏ/. The leftmost stimulus was always judged to be /ʏ/ while the next two stimuli were perceived as [-back] 88% of the time. The judgments for these three stimuli are considered the zone of reliable [-back] percepts, and the judgments for the rightmost three stimuli are considered to be the zone of reliable [+back] percepts. This leaves the middle stimulus on the seven-vowel continuum as the only vowel within the zone of uncertainty. Below, the zone of uncertainty extends from -0.30z to 0.10z (1445 Hz to 1629 Hz).

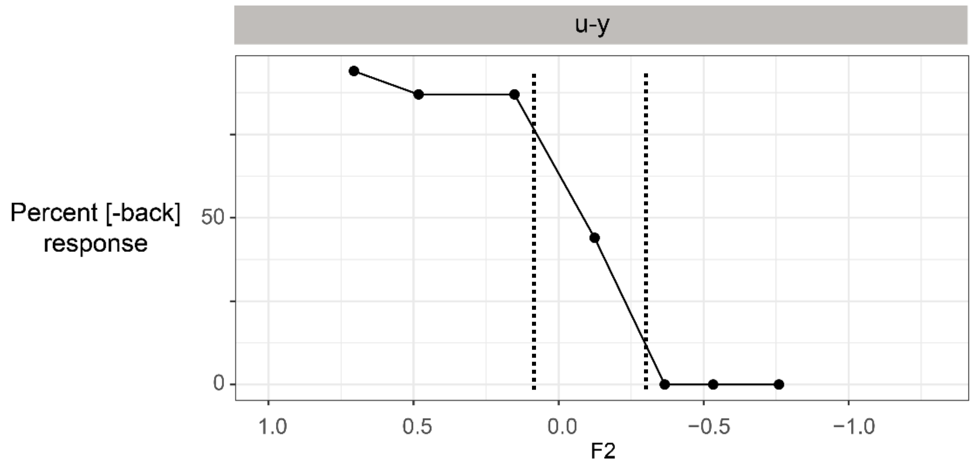

Figure 5.19: Percent [-back] responses to stimuli on Kazakh /ʊ/-/ʏ/ continuum with zone of uncertainty (marked by dotted vertical lines)

The comparison between production and perception of these alternating vowels in Figure 5.20 is illustrative, and somewhat distinct from the other vowels considered thus far. Mean F2 of fourth-syllable /ɑ/ and mean F2 of third- and fourth-syllable /ɯ/ occupy the proposed zone of uncertainty for each harmonic pairing. In this pairing, though, third- and fourth-syllable /ʊ/ not only extend into the zone of uncertainty, but fourth-syllable /ʊ/ is situated firmly within the zone of reliable [-back] percepts. Mean F2 of fourth-syllable /ʊ/ is greater than mean F2 of initial-syllable /ʏ/. Observe that mean F2 of the front vowel also increases by syllable, so fourth-syllable /ʊ/ is not identical to fourth-syllable /ɪ/, but the difference between the two means diminishes noticeably across syllables. In initial syllables, the difference between these two counterparts is 0.86z. In second syllables, the difference is 0.91z. In third and fourth syllables, though, the difference shrinks to 0.52z and 0.53z, respectively. Thus, even though the front vowel shifts forward by position, the distinction between /ʏ/ and /ʊ/ is reduced due to back vowel fronting.

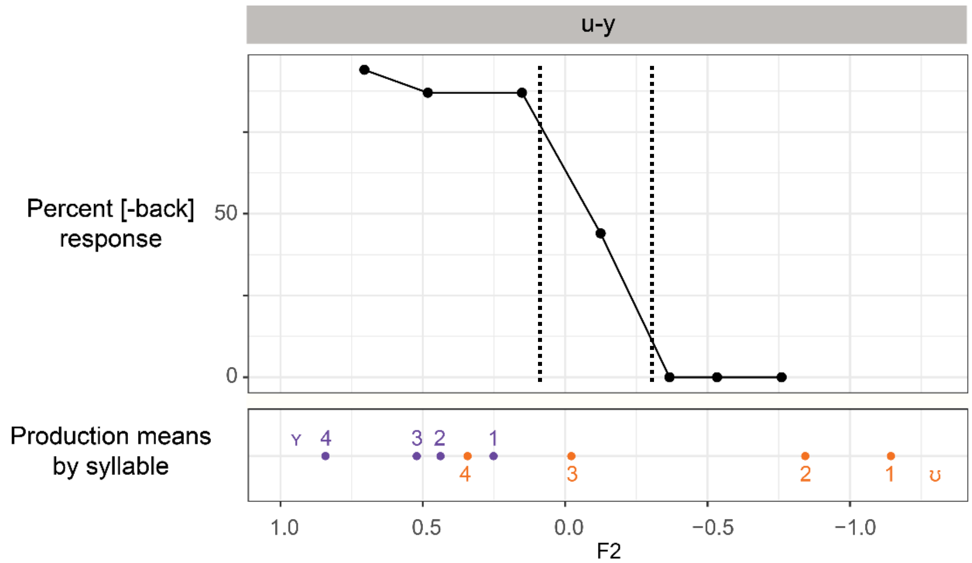

Figure 5.20: Percent [-back] responses to stimuli on Kazakh /ʊ/-/ʏ/ continuum compared to production means (zone of uncertainty is marked by dotted vertical lines)

5.2.3.3 Uyghur

Only two vowel alternations were investigated in Uyghur, /ɑ/-/æ/, and /u/-/y/. Despite the productivity of the /ɯ/-/i/ alternation in non-initial syllables, the conflation of these two in both the Cyrillic and Arabic orthographies used by Uyghurs in Kazakhstan and China prevented any meaningful examination of this third backness alternation.

Results for the /ɑ/-/æ/ alternation are presented below in Figure 5.21. Observe that the rightmost three points on the continuum were always perceived as [+back]. The leftmost two points were always perceived as [-back], and two intermediate stimuli were variably perceived, occupying what I've termed the zone of uncertainty, which extends from approximately 0.08z to 0.48z (1620 Hz to 1804 Hz).

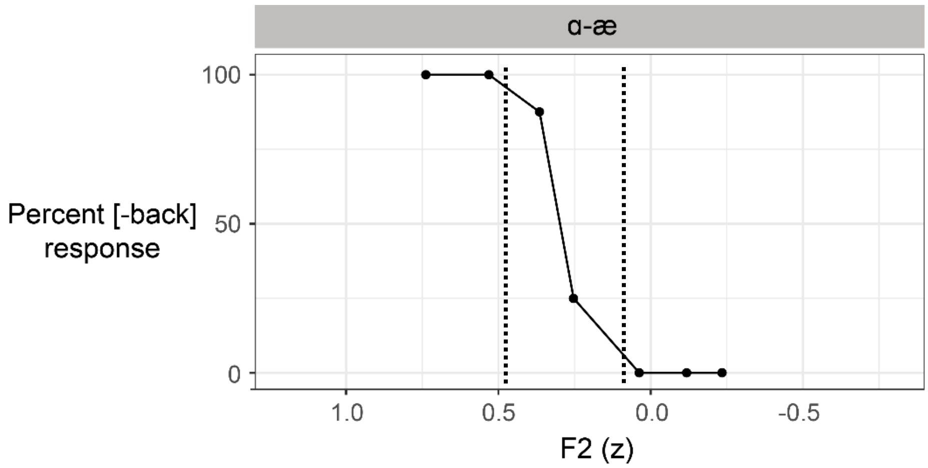

Figure 5.21: Percent [-back] responses to stimuli on Uyghur /ɑ/-/æ/ continuum with zone of uncertainty (marked by dotted vertical lines)

Production means are compared to perceptual judgments in Figure 5.22. Production means exploit a much larger range of F2 values than stimulus items. F2 of the front vowel /æ/ never approaches the zone of uncertainty, although F2 of fourth-syllable /ɑ/ does approximate the threshold between uncertainty and reliable [+back] percepts. According to the prediction above, if [+back] spreads gradiently throughout the word and if this gradience is phonological, then the outputs of this pattern should extend from the zone of reliable [+back] percepts into the zone of uncertainty. One mean value for /ɑ/ approaches this perceptual threshold, but does not convincingly cross it. Therefore, results from the Uyghur /ɑ/-/æ/ do not conclusively suggest that gradience in production is perceptually salient.

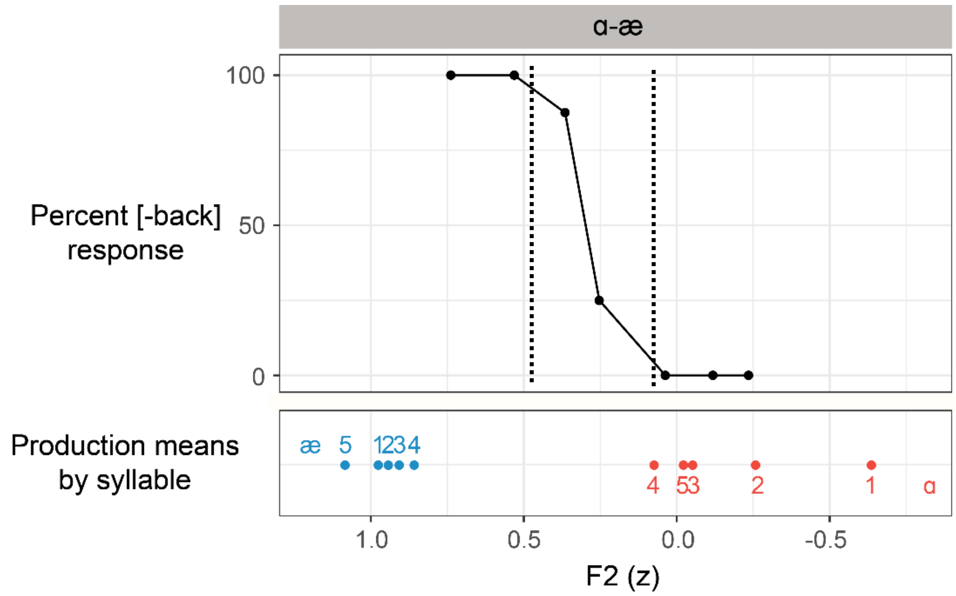

Figure 5.22: Percent [-back] responses to stimuli on Uyghur /ɑ/-/æ/ continuum compared to production means (zone of uncertainty is marked by dotted vertical lines)

Next, perceptual judgment for the /u/-/y/ contrast are presented in Figure 22. Observe that the leftmost three points were all judged as [-back] 75% of the time, and no tokens were rated as [-back] more than 75% of the time. In Figure 5.23, it is clear that these stimuli approximate F2 of Uyghur /y/, so this differs somewhat from the Kyrgyz /o-ø/ stimuli. I suggest that this sort of noise in the data may have been introduced by either the task setting, the participants, or perhaps the forced choice task itself. Putting aside this apparent noise in the data, I suggest that the three leftmost points represent the zone of reliable [-back] percepts. On the right side of the continuum, the rightmost two points were always judged to be [+back], suggesting that they occur within the zone of reliable [+back] percepts. In Figure 5.23, the zone of uncertainty reaches from approximately -0.45z to 0.15z (1376 Hz to 1652 Hz).

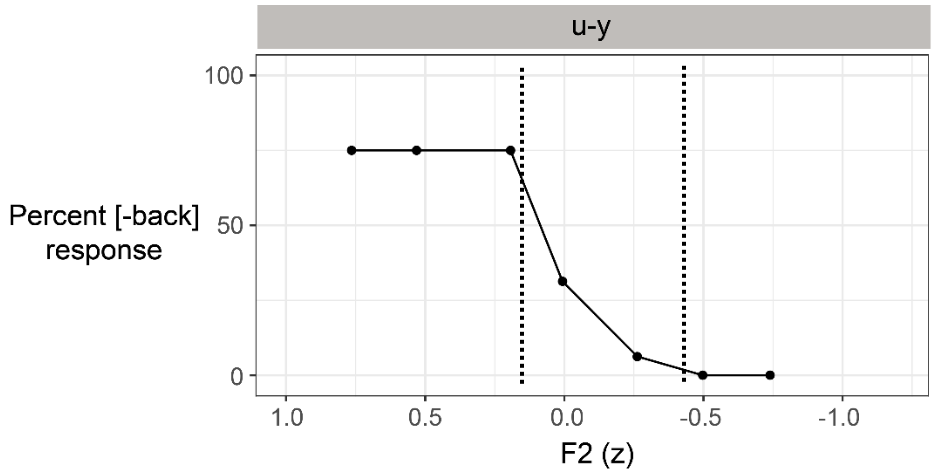

Figure 5.23: Percent [-back] responses to stimuli on Uyghur /u/-/y/ continuum with zone of uncertainty (marked by dotted vertical lines)

The comparison between production and perception data for this alternation, as shown in Figure 5.24, is very different from the Kyrgyz alternations or the /ɑ/-/æ/ alternation in Uyghur. In this alternation, production means for /u/ far exceed the zone of reliable [+back] percepts, extending throughout the entire zone of uncertainty. In fact, the mean F2 of fifth-syllable /u/ approaches the threshold of reliable [-back] percepts. In sum, positional shifts for this alternation very clearly exhibit the kind of perceptual salience predicted above. Mean F2 for third-, fourth-, and fifth-syllable vowels is consistent with variably perceived stimuli from the perception study. In contrast, production means for /y/ are all at the most categorical edge of the zone of reliable [-back] percepts.

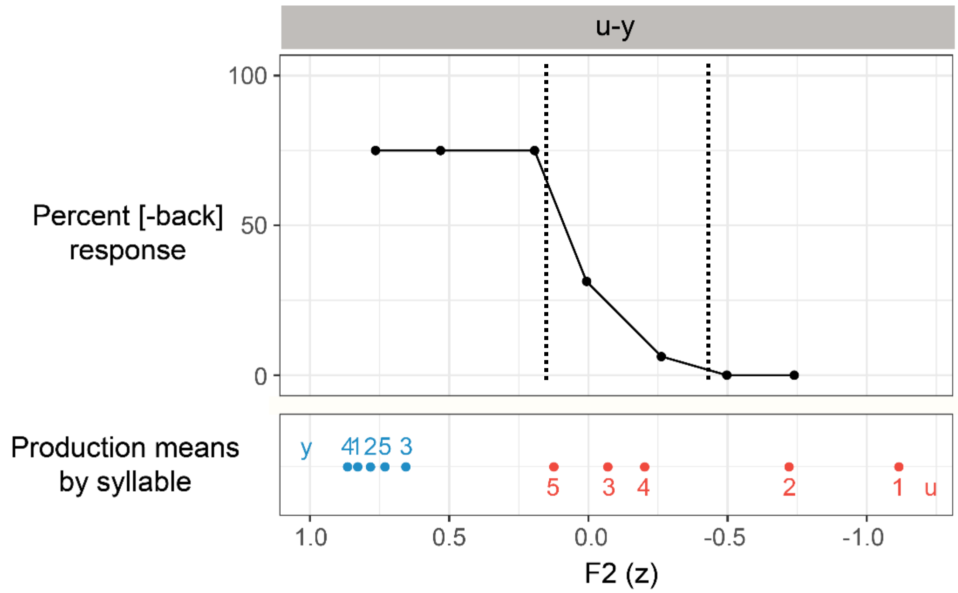

Figure 5.24: Percent [-back] responses to stimuli on Uyghur /u/-/y/ continuum compared to production means (zone of uncertainty is marked by dotted vertical lines)

5.2.4 Summary

Results from this section corroborate the essential difference in magnitude between positional shifts in Kyrgyz versus those in Kazakh and Uyghur. The positional shifts in Kyrgyz do not generally cross the proposed perceptual threshold between reliable [±back] percepts and the zone of uncertainty. In contrast, shifts in Kazakh and Uyghur typically do cross these thresholds. In fact, positional variation for the high round vowels in each language appears to not only cross this threshold, but to also approach the zone of reliable [-back] percepts in these two languages. These findings are summarized in Table 5.9. In Kyrgyz, two of four back vowel shifts approach and one shift crosses into the zone of uncertainty , although the shift for /ɯ/ (indicated by an asterisk) is due to word-final fronting and not harmony. Positional variation for one Kyrgyz front vowel also approaches and (barely) crosses into the zone of uncertainty. In Kazakh, all three back vowel shifts cross into the zone of uncertainty, and /ʊ/'s positional shift approaches and crosses into the zone of reliable /ʏ/ percepts. No front vowel shifts even approach the zone of uncertainty in Kazakh. Of the two vowel alternations examined in Uyghur, both back vowels approach the zone of uncertainty, and /u/ even approaches the threshold for reliable /y/ percepts. In contrast, no front vowel shifts approach even approach the zone of uncertainty in Uyghur.

Table 5.9: Summary of experimental results (* indicates that result derives from word-final fronting, not harmony)

Language	Perceptibility	ɑ-{e, ie, æ}		{u,ʊ}-{y,ʏ}		ɯ-{i,ɪ}		o-ø	
		[+bk]	[-bk]	[+bk]	[-bk]	[+bk]	[-bk]	[+bk]	[-bk]
	Approach [α back] threshold	✓	✓			✓*			
Kyrgyz	Cross into zone of uncertainty		✓			✓*			
	Approach [-α back] threshold								
	Approach [α back] threshold	✓		✓		✓			
Kazakh	Cross into zone of uncertainty	✓		✓		✓			
	Approach [-α back] threshold			✓					
	Approach [α back] threshold	✓		✓					
Uyghur	Cross into zone of uncertainty			✓					
	Approach [-α back] threshold			✓					

5.3 Discussion

In Section 5.1, just noticeable differences were predicted for each harmonic pairing and then compared against mean changes in F2 by syllable as well as predicted by-syllable shifts from the statistical model. In Section 5.2, a study investigated the perception of harmonic pairs in the context of monosyllabic roots, comparing the perception of alternating vowels with their by-syllable production means. Both sections address whether or not the asymmetric shifts in vowel backness in Uyghur and Kazakh, and the relatively symmetric shift in Kyrgyz are perceptually salient. I predicted that at minimum, phonological patterns should be perceptible, be they gradient or categorical. As a consequence, this predicts that patterns of positional variation in Uyghur and Kazakh are perceptually salient in a way that positional variation in Kyrgyz is not. This brief section summarizes results from the previous two sections as a final assessment of the prediction that phonological patterns must be perceptually salient.

Results from Sections 5.1-5.2 are summarized in Tables 5.10 and 5.11 below. When mean F2 for positional variants of each alternating vowel were compared against predicted just noticeable differences (JNDs) using a 5% threshold, 62% of back vowel shifts in Uyghur and 89% of back vowel shifts in Kazakh were predicted to be salient. Meanwhile, only 42% of back vowel shifts in Kyrgyz exceeded predicted JNDs. For the front vowels, only 8% of Uyghur vowels, 33% of Kazakh vowels, and 17% of Kyrgyz vowels exceeded predicted JNDs. The difference between Kyrgyz and the other two languages was not as distinct as one might predict, since all three languages exhibited back vowel shifts with at least 42% exceeding their predicted JND. Additionally, the predicted distinction between back and front vowel shifts was not as great as one might expect, since at least some front vowel shifts were perceptible.

Results from the statistical model were distinct from predicted JNDs, predicting that only back vowel shifts in Uyghur and Kazakh generally exceed their relevant JNDs. In contrast, all front vowel shifts, and all shifts in Kyrgyz did not exceed their predicted JNDs. The model thus predicts that these shifts are not perceptually salient, and thus not qualified to be phonological. Results from Section 5.2 refine the perspective offered by both the statistical model and the by-vowel positional means discussed in Section 5.1. Results from the continua-based perception experiment and their comparison with production means suggest that all back vowel shifts are salient in Kazakh since they all corresponded to the zone of reliable [+back] percepts in earlier syllables but crossed into the zone of uncertainty in later syllables. In fact, the production mean for fourth-syllable /u/ is firmly situated within the zone of reliable [-back] percepts, suggesting that fourth-syllable /u/ would be regularly perceived as a front vowel. For Uyghur, results suggest that positional variation in F2 for /u/ is significant, while results for /ɑ/ only approached the threshold for the zone of uncertainty. For Kyrgyz, results for the /ɑ/-/e/ alternation were similar to Uyghur results, with positional variants approaching the threshold for the zone of uncertainty. Additionally, all positional variants of /ø/ corresponded to the zone of uncertainty in the /o/-/ø/ continuum, making the interpretation of these results somewhat challenging. Furthermore, fourth-syllable /ɯ/ in Kyrgyz crossed the threshold from reliable [+back] percepts into the zone of uncertainty, but this result was analyzed as a product of word-final high vowel fronting, not backness harmony.

Table 5.10: Summary of production means compared to just noticeable differences (JNDs; using a 5% threshold)

<table>
<tr><td></td><td>Perceptibility</td><td colspan="4">Uyghur</td><td colspan="3">Kazakh</td><td colspan="3">Kyrgyz</td></tr>
<tr><td rowspan="4">[+back]</td><td rowspan="3">Mean positional shift > JND?</td><td colspan="4">8/13</td><td colspan="3">8/9</td><td colspan="3">4/12</td></tr>
<tr><td>σ_{1-2}</td><td>σ_{2-3}</td><td>σ_{3-4}</td><td>σ_{4-5}</td><td>σ_{1-2}</td><td>σ_{2-3}</td><td>σ_{3-4}</td><td>σ_{1-2}</td><td>σ_{2-3}</td><td>σ_{3-4}</td></tr>
<tr><td>ɑ, u</td><td>ɑ, ɯ, u, ɯ$_{R}$, u$_{R}$</td><td></td><td>u</td><td>ɑ, ɯ, ʊ</td><td>ɑ, ɯ, ʊ</td><td>ɯ, ʊ</td><td>ɑ, o, u</td><td>ɯ, u</td><td></td></tr>
<tr><td>Predicted shifts from regression model > JND?</td><td colspan="4">✓</td><td colspan="3">✓</td><td colspan="3">✗</td></tr>
<tr><td rowspan="4">[-back]</td><td rowspan="3">Mean positional shift > JND?</td><td colspan="4">0/13</td><td colspan="3">1/9</td><td colspan="3">1/12</td></tr>
<tr><td>σ_{1-2}</td><td>σ_{2-3}</td><td>σ_{3-4}</td><td>σ_{4-5}</td><td>σ_{1-2}</td><td>σ_{2-3}</td><td>σ_{3-4}</td><td>σ_{1-2}</td><td>σ_{2-3}</td><td>σ_{3-4}</td></tr>
<tr><td></td><td></td><td></td><td></td><td>ie</td><td></td><td></td><td>e</td><td></td><td></td></tr>
<tr><td>Predicted shifts from regression model > JND?</td><td colspan="4">✗</td><td colspan="3">✗</td><td colspan="3">✗</td></tr>
</table>

Table 5.11: Summary of experimental results (* indicates that result derives from word-final fronting, not harmony)

Language	Perceptibility	ɑ-{e, ie, æ}		{u,ʊ}-{y,ʏ}		ɯ-{i,ɪ}		o-ø	
		[+bk]	[-bk]	[+bk]	[-bk]	[+bk]	[-bk]	[+bk]	[-bk]
Kyrgyz	Approach [α back] threshold	✓	✓			✓*			
	Cross into zone of uncertainty		✓			✓*			
	Approach [-α back] threshold								
Kazakh	Approach [α back] threshold	✓		✓		✓			
	Cross into zone of uncertainty	✓		✓		✓			
	Approach [-α back] threshold			✓					
Uyghur	Approach [α back] threshold	✓		✓					
	Cross into zone of uncertainty			✓					
	Approach [-α back] threshold			✓					

At a very general level, this chapter has argued that perceptilibity is a prerequisite for phonological patterns. Essentially, if a pattern is not perceivable then children are unlikely to learn it. I can think of two challenges to this proposed prerequisite, abstract contrast and laryngeal transparency. By abstract contrast, I refer to segments whose phonological behavior is inconsistent with their phonetic realization. For instance, Tutrugbu possesses two vowels that exhibit a abstract level of contrast, transcribed as /ɔ-ʊ/ (or /ɔ-$ɔ^H$/) in McCollum & Essegbey (2019) and McCollum et al. (2019). In Tutrugbu, /ɔ/ triggers progressive rounding harmony on prefixes while /ʊ/ does not. In addition, /ɔ/ and /ʊ/ behave differently in hiatus. Thus, these two sounds exhibit differential phonological behavior. However, acoustic results in McCollum & Essegbey (2019) suggest that these two sounds are acoustically indistinguishable, since they exhibit no significant differences in F1, F2 or duration. I have claimed that for these two sounds to be contrastive, their should be perceivable differences between them. If we localize perception of the segment, then this is a problem for my claim that perceptibility is necessary for phonological patterns and contrasts. If however, we consider that these vowels are perceived in context in natural speech, then their phonological behavior serves as a perceptual cue to their differential phonological status. The fact that round vowels may follow /ɔ/ but not /ʊ/ provides perceptual evidence for this distinction that is sufficient to support the contrast in the language. Thus, perception can be maintained as a precondition for phonology if differences in context are taken into account.

A second challenge for my claim comes from patterns of laryngeal transparency, particularly well-noted in nasal harmony. In languages with nasal harmony, the laryngeal consonants (e.g. /h/ and /ʔ/) often behave as transparent. Walker (1999) finds that Guaraní glottal stop does not differ acoustically in oral and nasal contexts. Phonetically, this is not unsurprising, since a constriction at the glottis prevents air from escaping through the nasal cavity despite the fact that the velum appears to be lowered throughout the nasal (see Cohn 1990, 1993; Walker 1998, 2003; Walker & Pullum 1999; Borroff 2007 for more discussion). Yet, if this allophonic distinction between [ʔ] and [ʔ̰] is phonological, it should be perceivable. My claim is the same as above, that perception in context is sufficient to maintain this allophonic distinction, as argued by Walker & Pullum (1999). The realization of nasality on immediately adjacent segments is what learners are exposed to, and I assume that they use this contextual information to bootstrap perception of oral and nasalized laryngeals. Perception of a subtle contrast likely involves segment-external coarticulatory and phonological effects across languages. In Tutrugbu, /b/ and /bʰ/ are contrastive, but the difference between the two is not, as far as I tell, realized on the consonant itself, but on the following vowel. Vowels after /bʰ/ are breathy while vowels immediately following /b/ are produced with modal voice (see Seyfarth & Garellek 2018 for more discussion of vowel phonation and aspirated consonants). In effect, if perception is contingent upon not only segment-internal phonetic cues, but also coarticulatory and flanking segmental cues, then the perception of subtle distinctions like Tutrugbu /ɔ-ʊ/ and allophonic nasalization of laryngeals in nasal harmony.

This chapter has used two different evaluation metrics to assess the phonological significance of the gradient shifts in vowel backness discussed in Chapter 4. Just noticeable differences have been evaluated almost entirely for English, and this study provides a first step toward evaluating these differences for lower-resource languages. Future work should continue this path of investigation to determine more accurately what differences in vowel quality can be discriminated by speakers of these languages. In addition, this chapter has compared perception of vowels intermediate between prototypical [±back] vowels with gradient production results to evaluate whether or not vowel perception supports the claims advanced in Chapter 4. Future work can extend this method for relatively direct comparison of production and perception to better understand the nature of phonological patterns, phonological contrast, as well as the relationship between perception and production. Moving forward, future work would benefit from using stimuli generated from speakers of each target language. Using stimuli from a Kazakh speaker for these three languages is better than some options, but still falls short of the more realistic stimuli necessary to more confidently assert the perceptual significance of the backness shifts in question. Additionally, this study has assumed that vowel perception is constant across positions, comparing initial-syllable vowel perception with production from non-initial syllables. There are likely differences in perception in these positions, which deserve a fuller investigation. Also, future work should compare results from several different experimental methods beyond the particular task employed here. Some possibilities include the use of an AXB paradigm, likert-scale testing, an oddball paradigm, or eye-tracking.

Despite the challenges, particularly to the interpretation of the Kyrgyz results, findings from Sections 5.1 and 5.2 suggest two distinctions that are relevant for discussion. First, these results further support the distinction between Kyrgyz and the other two languages. In Uyghur and Kazakh, [+back] vowel positional shifts were more likely to be salient under every evaluation metric used, while results from Kyrgyz were always regarded as less likely to be perceivable. Additionally, these results further establish a difference between [+back] and [-back] vowel positional shifts. Back vowel shifts were always more likely to be salient under these predictive methods, especially for Uyghur and Kazakh. In sum, these results generally suggest that the positional increases in vowel F2 for back vowels in Uyghur and Kazakh are perceivable, which accords with the earlier prediction that phonological patterns must be perceivable. Moreover, the general insignificance of positional variation in Kyrgyz further suggests the fundamentally phonetic nature of backness variation in that language.

The distinction between phonological variation in Uyghur and Kazakh, in contrast to phonetic variation in Kyrgyz will factor into the next two chapters. In the next chapter, an analysis of all three languages is presented within a gradient Harmonic Grammar. In Chapter 7, the distinction between phonological and phonetic patterns is discussed, connecting the perception results from this chapter with a number of other diagnostics to differentiate each type of sound pattern.

Chapter 6: Incorporating gradience into the formal analysis

In the previous chapter I demonstrated that backness harmony in Uyghur and Kazakh is gradient, in contrast to categorical harmony in Kyrgyz. More importantly, I argued that morphophonological patterns may be gradient in a manner not derivable from phonetic implementation. I thus proposed that phonology is not by definition categorical, but may exhibit gradience as well as categoricality. This chapter develops a formal account of backness harmony in Uyghur, Kazakh, and Kyrgyz by introducing gradient representations into Harmonic Grammar (HG; Legendre et al. 1990).

This chapter is organized as follows. In Section 6.1, I discuss previous formal accounts of gradience. In Section 6.2, I use the normalization method described in Lionnet (2017) to compare the relative backness of vowels across all harmonic pairings in each languages. In Section 6.3, I go on to develop an analysis in HG, demonstrating that the gradient HG analysis is superior to a categorical analysis. Finally, in Section 6.4, I summarize findings from the chapter and look forward to the larger discussion of gradience in Chapter 7.

6.1 Previous work

Gradience plays, as far as I can tell, almost no role whatsoever in derivational formalisms. This likely falls out from several different factors. First, derivational models following Chomsky & Halle (1968) were often willing to incorporate a high degree of abstraction into the formalism, effectively removing the locus of explanation away from the sort of phonetic detail that comes hand-in-hand with subphonemic gradience (e.g. Vago 1973; Hyman 1970; Harms 1973; Selkirk & Vergnaud 1973; Kiparsky 2012). Second, in a formalism based on rewrite rules, there is no direct way to encode gradience. If phonology is populated by binary variables only, this is obvious, but even if one introduces gradient representations into a phonological grammar, the issue is far from settled. For instance, how does one encode the fact that in the second syllable a [+back] vowel is only 80% back relative to its initial-syllable counterpart? Direct reference to such a generalization, e.g. V $\rightarrow$ [0.8back] / # C_0V[+back]C_0___ is inelegant, since each subsequent syllable would need some slightly different formulation of the above rule to account for a gradient harmony pattern. Within a rule-based formalism, it would also be possible to overlay a relatively simple rule, e.g. V $\rightarrow$ [+back] / V[+back]C_0 ____ with a function that outputs the realization of [back] for a given vowel in syllable *x* based on the realization of [back] in a vowel in syllable *x-1*. Although this could drastically reduce the number of rules necessary to account for a given gradient pattern, this, as far as I know, has never been attempted in a derivational model. A third reason for the absence of gradient analyses in serial models was an inability to conduct detailed phonetic analysis sufficient to address the issue of gradience before the 1980s, particularly due to the lack of usable technology. The technology available made detailed phonetic work challenging. Using wax cylinders does not make fine-grained acoustic analysis easy, and after they were phased out of linguistic research, computer-based analysis was not very affordable or portable.

In constraint-based models, though, there is a small but growing body of work implementing gradient insights into the phonological formalism. This work has followed two paths: one, gradience as subphonemic, phonetic effects and two, gradience as exceptionality. These two bodies of work are discussed in turn.

6.1.1 Phonetic gradience in phonology

Gradience has played a significant role in the development of phonological theory over the past three or four decades, particularly within Optimality Theory. First, and in contrast to derivational formalisms, Kirchner (1997) argues that Optimality theory can account for both categorical and gradient patterns by introducing gradient, non-contrastive representations into the grammar. To understand Kirchner's analysis, one must first understand the basic working assumptions he addresses. First, as he notes, many phonologists have assumed that feature specifications that are never contrastive in any language must be excluded from the phonology, and relegated to phonetic implementation (e.g. Keating 1984; Lombardi 2018). Furthermore, the features that are non-contrastive in a particular languages are typically assumed to be absent from underlying representations and earlier (i.e. Lexical) portions of phonological computation (e.g. Kiparsky 1985; Archangeli & Pulleyblank 1994). Non-contrastive features along with gradience are introduced into the grammar in post-lexical and phonetic modules of the grammar (e.g. Keating 1988, 1990; Cohn 1993). Kirchner then demonstrates that constraint ranking can produce patterns of categoricality and gradience without stipulating the nature of underlying, surface, and phonetic representations if we assume that contrastive properties have relevant faithfulness constraints, while non-contrastive properties do not. When a constraint-based grammar is given access to detailed phonetic representations, Kirchner shows that both categorical and gradient patterns can be produced, effectively conflating the phonological and phonetic modules of the grammar. He writes:

> [T]he contrastiveness of a particular feature depends entirely on whether there is a corresponding faithfulness constraint which is satisfied under some set of mappings, which in turn depends on the position of the constraint within the constraint hierarchy. This result extends even to properties which are not contrastive in any language, if we simply assume that such properties lack a corresponding PRES constraint [PRES = the set of relevant Input-Output faithfulness constraints]. Finally, the distinction between categorical and gradient properties, standardly assumed to characterize the difference between phonological and phonetic representation, proves to be a special case of the previous result. Consequently, we may capture the categorical and contrastive *behavior* of particular phonetic properties (and the predictable or gradient behavior of the remainder) in terms of constraint interaction, while using *representations* which in principle may contain complete phonetic detail, including gradient properties (1997:107, emphasis in the original).

In essence, OT can account for both phonological and phonetic patterns so long as faithfulness constraints are stipulated to be related to contrastive patterns only.

A very similar stance is taken in work like Flemming (2001) and Flemming & Cho (2017). These authors derive gradient patterns of segmental and tonal coarticulation from the interaction of gradient constraints. These works differ from Kirchner (1997) in that they assume no stipulative difference between the categorical and the gradient, deriving both from quantitative optimization. For instance, Flemming (2001) describes a pattern of fronting in Cantonese whereby flanking coronal consonants cause the back vowel /u/ to be fronted to [y]. In other contexts, /u/ and /y/ may both occur, but in between two coronals the contrast is neutralized to [y]. Flemming argues that this pattern can be account for if the grammar optimizes effort over acoustic targets, including targets for vowels as well as consonant F2. In Flemming's model, some constraints penalize divergence from target consonant and vowel F2 while an additional constraint penalizes articulatory effort. This effort reduction constraint, in effect, produces patterns of undershoot, where neither consonantal nor vocalic targets are reached. This model differs from OT by using quantitative rather than discrete mathematical optimization. In OT, conflict is resolved by strict domination, whereas in Flemming's model summative effects are possible, similar to 'gang effects' in HG. In the undershoot case, gradient patterns are produced by constraint resolution, but Flemming contends that in some contexts the benefits are outweighed by the effort required to maintain a given contrast, yielding categorical neutralization from this quantitative approach to optimization.

The common theme between these two approaches is that they effectively disregard any distinction between phonology and phonetics. These unified models are thus more explicit formalizations of Pierrehumbert's proposal that there is no interface between phonology and phonetics because they both exist on a single continuum (e.g. Pierrehumbert 1994; Pierrehumbert et al. 2000; Cohn 2006; Tucker & Warner 2010). Throughout I assume there is no discernible formal difference between Flemming's unified model and Pierrehumbert's continuum proposal. To my knowledge, no formal analyses have implemented Pierrehumbert's claim. In such a model, how does one differentiate the phonological from the phonetic? One presumably cannot simply order one before the other. If one differentiates the two based on their representations, how does one prevent gradient phonetics from interacting with the categorical phonology to produce gradient phonology? If gradient representations exist in the formalism, these become the lowest common denominator, and as a consequence, the formalism becomes identical to the unified model.

Although the unified view of phonology and phonetics (including views that completely unify the two and views that remove the categorical distinction between them) has been influential, more particularly in the laboratory phonology community, phonological analyses incorporating gradience have also developed under a modular view of phonology and phonetics. One strand of this work encodes phonological contrasts as scalar variables. Flemming (2013) defines the vowel space in terms of a discretized grid of vowel height and backness, turning the binary features, [high], [low], [front], and [back] into scales, as shown below in (43). For instance, the feature [back] has up to five values, as opposed to the binary distinction employed in a featural representation. While it seems that Flemming bases his scale on and IPA chart, there is nothing to prevent backness from being scaled across a much larger number of possible values, as noted and advocated for in McCollum (2018).

(43) Flemming's (2013:30) representation of vowel height and backness

F2

5	4	3	2	1		
i	y	ɨ	ɯ	u	1	
ɪ	ʏ			ʊ	2	F1
e	ø	ə	ɤ	o	3	
	ɛ	œ	ʌ	ɔ	4	
		a	ɑ		5	

In addition to Flemming, other work has argued for non-binary distinctions among vowels, particularly for vowel height (Parkinson 1996; Gnanadesikan 1997). These approaches are similar to older proposals like Schane's particle phonology, in which the representation of vowel height is represented in terms of an aperture particle, a. In Schane (1984), the high vowels have no aperture particle, the mid vowels have one, and the low vowels have two, yielding representations like /i/ = i, /e/ =

ai, /æ/ = aai. The aperture value in Schane's theory produces ternary distinctions from a unary primitive, indirectly encoding a scalar notion of height. If phonological distinctions may be scalar and n-ary rather than strictly binary, this introduces the possibility for gradience. While Schane's formalism is constrained according to traditional conceptions of high, mid, and low, it is only an assumption of categoricality than prevents representing /æ/ as aaaaaai and /e/ as aai, or aaii, since these longer representational strings would introduce the possibility of more representational categories. By increasing representational categories, Kirchner (1997) demonstrates that the formal analysis shifts toward gradience. Scalar variables have been used widely in HG in more recent work, but in the vast majority of cases, constraint weights are scaled in relation to some factor like lexical stratum rather than some phonological feature (e.g. Hsu & Jesney 2016, 2017).

Although work like Flemming (2001) conflates phonology and phonetics, deriving gradient patterns from a unified model, a different line of work incorporates gradience into phonological computation via phonetic gradience. This work claims that phonology has access to fine-grained phonetic detail, which shapes human sound patterns. The amount of detail varies by analysis, but the overarching thrust of this work is that the typology of phonological patterns is explainable in terms of phonetic content (e.g. Hayes et al. 2004). For instance, Zhang (2004) argues that the cross-linguistic tendency for contour tones to occur in word-final syllables is derivable from the relatively low-level details of phonetic lengthening in final syllables, encoding facts about phonetic duration directly into the formalism (see also R. Kirchner 1998). As another example, Kaun (2004) analyzes the typology of rounding harmony as a consequence of variation in degree of phonetic lip rounding by vowel height and backness. Kaun and Zhang's work thus encodes phonetic detail indirectly, formulating markedness constraints that refer to relatively abstracted phonetic generalizations.

On the other end of the spectrum, some work has used far more detailed phonetic information to motivate phonological analysis. McCollum (2018) incorporates mean phonetic values for F1-F2 into Kaun's larger analysis, showing that the language-specific distribution of [round] harmonic pairings in a normalized F1-F2 vowel space predicts the activity or inertness of vowels in rounding harmony. McCollum (2018) uses prototype-based representations that encode mean phonetic information for each vowel category, which is directly referenced in the phonology. In my analysis, violations of the constraint driving harmony are scaled by the phonetic distance between the triggering [+round] vowel and its [-round] counterpart. Whereas Kaun's (2004) analysis encodes a cross-linguistic generalization, McCollum (2018) builds a competing analysis from language-specific detail, contending that the real source of generalization is observable from fine-grained, language-specific detail. Similarly, Lionnet (2017) analyzes rounding harmony in Laal as the consequence of detailed phonetic knowledge employed by speakers. In Laal, some vowels undergo regressive rounding harmony if they agree in backness and height with the following vowel, and crucially, if they are adjacent to a labial consonant. If the target vowel is not flanked by a labial consonant, harmony fails to apply. Lionnet argues that doubly-triggered harmony is dependent on phonetic coarticulation, since coarticulatory lip rounding from the adjacent consonant serves to define a natural class of undergoers. In his analysis, coarticulatory lip rounding is part of speaker knowledge, which is exploited at the phonology-phonetics interface by continuous subfeatures (e.g. a vowel may be $[\![x \text{ round}]\!]$, where $0 \leq x \leq 1$), a hierarchically organized level of representation linking abstract distinctive features with low-level phonetic patterns. In both Lionnet (2017) and McCollum (2018), phonetic detail is available to the phonological grammar, but crucially, in each analysis, the phonological grammar still manipulates only abstract phonological variables (see also Stanton 2019).

In contrast to the preceding work, Braver (2013, 2019) accounts for various cases of incomplete neutralization by weighted phonetic constraints. In his work, phonology operates categorically over abstract variables, like the vast majority of generative work. His work differs, though, in that it formalizes the phonetic patterns he claims produce gradient outputs. His two-level model uses the same quantitative optimization employed in Flemming (2001) but without conflating phonetic and phonological modules. Under Braver's analysis, incomplete neutralization is a phonetic fact rather than a phonological one. Smolensky et al. (2014) suggest the potential for a phonological analysis of incomplete neutralization, and perhaps for other types of gradient sound patterns, Braver argues that gradience is phonetic. The various analyses just discussed are summarized in Table 6.1 below. For each analysis, the source of gradience, the role of gradience in phonology, and the nature of phonological and phonetic optimization are listed with an example citation. Note that blank spaces in the table represent areas where the analysis makes no explicit claims. These four views are compared and contrasted with the view developed here, that phonology may be gradient. If phonology may be gradient, then phonetics is not the only source of gradience in the grammar. Furthermore, if phonology may be gradient, discrete optimization becomes problematic, as discussed below.

Table 6.1: Summary of previous work on gradience in phonology

<table>
<tr><th></th><th>Unified model</th><th>Phonetically-based model</th><th>Direct reference model</th><th>Two-level model</th><th>Gradient phonology</th></tr>
<tr><td>Source of gradience</td><td></td><td>abstract phonetic generalizations</td><td>detailed phonetic information</td><td>phonetics</td><td>phonology</td></tr>
<tr><td>Gradience in phonology</td><td>yes</td><td>as indirect support for discrete constraints</td><td>as direct support for discrete constraints</td><td>no</td><td>yes</td></tr>
<tr><td>Phonological optimization</td><td rowspan="2">unified; quantitative</td><td>discrete</td><td>discrete</td><td>discrete</td><td>quantitative</td></tr>
<tr><td>Phonetic optimization</td><td></td><td></td><td>quantitative</td><td></td></tr>
<tr><td>Example citation</td><td>Flemming (2001)</td><td>Zhang (2004)</td><td>Lionnet (2017)</td><td>Braver (2019)</td><td></td></tr>
</table>

6.1.2 Gradience and exceptionality

One recent line of work has been introducing gradience into phonological analysis, not to account for subphonemic gradience, but rather to address exceptionality. This body of work has centered around Gradient Symbolic Computation (GSC; Smolensky et al. 2014; Smolensky & Goldrick 2016). As an example, Smolensky & Goldrick (2016) analyze liaison in French using gradient representations. The basics of the French liaison pattern are illustrated in (44). In (44a,b), the final <t> of <petit> is pronounced before <ami>, but not before <copain>. In (44c), the second <t> of <petite> is pronounced before <copine>, in contrast to (1b), where orthographic <t> is not pronounced. The real question is why [t] (among several other consonants) is pronounced at the ends of some words but not others. Finally, in (1d), no [t] is pronounced between <joli> and <ami>. There is no orthographic evidence to support an underling /t/ in either <joli> or <ami>. This form thus serves to show that [t] is not simply epenthetic in (1a), repairing hiatus between final /i/ in <petit> and initial /a/ in <ami>. In turn, this supports /t/ being present underlyingly in these word pairs. If /t/ is present underlyingly in at least some of these forms, the question is why it surfaces in some forms (44a,c) and not others (44b).

(44) French liaison

	Orthography	Surface form	Gloss
a.	petit ami	pø.ti.ta.mi	‘little.MASC friend.MASC’
b.	petit copain	pø.ti.ko.pɛ̃	‘little.MASC friend.MASC’
c.	petite copine	pø.tit.ko.pin	‘little.FEM friend.FEM’
d.	joli ami	ʒo.li. a.mi	‘nice.MASC friend.MASC’

As discussed in Smolensky & Goldrick (2016), two general types of analyses have been posited for French. The first claims that the liaison consonant is part of the lexical entry of the first word, making <petit> underlyingly /pøtit/. Crucially, final [t] only surfaces in some contexts, and is deleted in others. Under the second type of analysis posited in the literature, the liaison consonant is actually affiliated with the second word, whose segmental form is determined by the preceding word. Thus, the word ‘friend.MASC’ has the allomorphs, /tami/, /ami/, and several others, which are determined based on preceding context. The key question is thus, where does the [t] come from? Based on (44d), liaison [t] does not appear to be epenthetic, suggesting that it is underlyingly part of at least one of the words in each phrase in (44). While phonologists have argued that the liaison consonant is part of either the first or second word, Smolensky & Goldrick (2016) instead argue that the liaison consonant is actually present in both words. Therefore, the lexical entries for <petit> and <ami> from (44a) are /pøtit/ and /tami/. However, the final /t/ in /pøtit/ and the initial /t/ in /tami/ are different from other consonants in French. They claim that these are defective, and encode this by assigning them gradient values, like 0.5 /t/ and 0.3 /t/. Gradient representations are present underlyingly across the lexicon for liaison consonants at the left edge of vowel-initial words. Similarly, the gradient value associated for the orthographically represented consonant is underlyingly present in word-final position. Under their analysis, it is the combined effect of multiple defective (i.e. gradient) consonants that yields a pronounced liaison consonant.

A similar line of reasoning is exploited in Zimmermann (2018 a,b,c,d, 2019) to derive exceptional tonal and segmental patterns across a range of languages. Zimmermann uses gradient input representations to predict what she calls "ghost segments." In many cases, she examines categorical surface representations from gradient inputs. This is largely in line with Smolensky & Goldrick (2016), using gradient representations as a representational solution for exceptionality within a categorical formalism. Yet, in some cases she considers the possibility of gradient surface representations, arguing that markedness constraints must be able to evaluate gradient outputs. For instance, in Zimmermann (2018d), she accounts for variable consonantal realization in Nuu-chah-nulth with gradient inputs. Some suffixes in the language occur with their underlying onset consonants post-vocalically. Post-consonantally, though, these suffixes are realized without their initial consonants. These unstable segments differ from regular suffixes, whose initial consonants are realized both post-vocalically and post-consonantally. Zimmermann (2018d) contends that these unstable consonants are underlyingly specified as gradient. Moreover, these gradient inputs may surface as gradient outputs, too. She argues that the input /q/ in /ʔatła-qumł/ 'two-round' is specified as 0.5 /q/, and surfaces as 0.5 [q] in the output. Problematically, since she does not tie gradience to any verifiable acoustic or articulatory properties, it is unclear what 0.5 [q] would sound like. More generally, Zimmermann does not connect her representational assumptions with actual phonological production. Without an articulated view of how these gradient forms are actually realized in speech production, it is challenging to evaluate Zimmerman's claim. This abstraction is not the only problem to be found in this use of gradience. Additionally, it is not at all clear what aspects of [q] are specified as 0.5. Is the [consonant] feature gradient? Or is the [dorsal] feature? This treatment of gradience parallels that in Smolensky & Goldrick (2016), which appears to treat the segment as the fundamental building block, ignoring the key role that subsegmental features play in phonology. Without an understanding or real commitment to the actual surface properties of gradient segments, gradience is simply a representational device to account for exceptionality. In this way, gradient representations look very similar to diacritic features in Chomsky & Halle (1968) and other representational strategies to account for exceptional forms. Additionally, since Zimmermann (2018a,b,c,d, 2019) and Smolensky & Goldrick (2016) tie gradience to segments rather than features, these analyses make very few testable phonetic predictions.

In this line of work, gradience is tied to exceptionality and not subphonemic effects. Despite this overarching generalization, Smolensky et al. (2014:1126-27) suggests the possibility for linking gradience to surface phonetic properties in a more principled way. "Gradient Symbol Processing allows us to go further, and account for the empirically documented incompleteness of German final voicing neutralization. As in the empirical data, these simulations show that the final t's output for /rad/ → [rat] 'wheel' and for /rat/ → [rat] 'advice' differ slightly. The former is slightly closer than the latter to /d/, showing a gradient trace of the underlying lexical form." Beyond this comment, there has been little work in GSC connecting gradient representations with subphonemic effects and incomplete neutralization.

One of the most significant challenges for encoding gradience in the formalism is translating between phonetics and phonology. The mapping between something like F2 and the phonological feature [back] is not always straightforward. To account for gradient backness across a number of alternations, in the next section I detail a normalization procedure that allows for comparisons both across-speakers and across-alternations. By doing so, this allows for a level of generalization not possible under raw Hz or z-score normalized values for F2.

6.2 Normalization

For the analysis developed later in the chapter, I use a relative measure of backness that generalizes across all harmonic alternations in a language. One might wonder if the raw Hertz or z-score normalized phonetic representations from the previous chapter are appropriate. If they aren't, what is the character of those representations that precludes them from adequately serving as a phonological representation? To answer this question, consider the realization of the Uyghur vowels /ɑ/ and /u/. In initial syllables, their mean F2(z) is -0.66 and -1.12, respectively. If z-values are used, how can one generalize across particular harmonic pairings? Should -0.66z represent a categorically [+back] vowel? Or should -1.12z represent a categorically [+back] vowel? It should be obvious that neither value provides a workable point of reference for generalizing over all [+back] vowels. Of course, one key difference between /ɑ/ and /u/ is lip rounding, and it is possible to generate a single z-value for [back] and [round] by taking the beta values from the linear regression developed in the previous chapter. There are at least two problems with such an approach. First, there is no guarantee that every regression model fits the data, and without an additional stipulation regarding model fit, it is unclear how to ensure that the point of reference output by the model is actually appropriate. Second, the model relativizes over the different fixed effects (i.e. features) used to predict vowel acoustics, and so can only really output individual points of reference for each feature combination used in the fixed effect structure. Stated differently, such a model for Uyghur could output points of reference for [+back,-round], [+back, +round], [-back, -round], and [-back, +round] vowels, but not for [+back] and [-back] without reference to rounding. Since backness and rounding both primarily affect F2, it can be difficult to separate effects of rounding and backness from one another. For these reasons it is more advantageous to use a normalization method that generalizes across harmonic pairings in such a way that the feature [back] is completely separated from [round], something which raw Hertz or z-score normalized F2 values cannot easily do.

The method used to normalize across harmonic pairing here is based on the method used in Lionnet (2017) to account for gradient rounding harmony in Laal. Lionnet represents the roundness of a high back vowel in Laal as its relative linear position between the two relevant phonemic categories, /ɨ/ and /u/, or /ə/ and /o/. So, a high vowel produced with an F2 value halfway between the two endpoints, /ɨ/ and /u/ would receive a value of 0.5 round, which Lionnet formalizes as a subfeatural representation, indicated by ⟦round⟧, e.g. ⟦0.5 round⟧. This type of representation can be stated more explicitly, as in (45) and (46), which provide the basis for the gradient representation of [back] throughout.

(45) For F2 of a particular vowel token *x* for an alternation between vowel categories Y and Z, with mean initial-syllable values, *y* and *z*, the ⟦back⟧ representation of that particular token is equivalent to the ratio of the distance between $F2_x$ and $F2_y$, and $F2_y$ and $F2_z$, $\Delta F2_{y\text{-}x}$ and $\Delta F2_{y\text{-}z}$, where *y* is the unmarked value of the feature [back].

(46) $[\![\text{back}]\!]_{\text{x}} = \frac{|\text{F2y-F2x}|}{|\text{F2y-F2z}|}$ or $\Delta\text{F2}_{y-x} \,/\, \Delta\text{F2}_{y-z}$

To exemplify how this produces a gradient representation, consider vowel *x* somewhere between Uyghur /æ/ and /ɑ/, with F2 of 1500 Hz. If F2 of initial-syllable /æ/ and /ɑ/ are 1850 and 1300 Hz, respectively, then the ⟦back⟧ value for this particular vowel would be represented in (47).

(47) $[\![\text{back}]\!]_{\text{x}} = \frac{|1850\text{-}1500|}{|1850\text{-}1300|} = \frac{350}{550} = 0.64$

A different normalization equation was used for each alternating pair, so a [+high, +round] vowel with ⟦back⟧ = 0.5 has a different F2 value than a [-high, -round] vowel with the same gradient backness value, ⟦back⟧ = 0.5, because backness values are relativized to each harmonic pair. Furthermore, to account for interspeaker variation, each ⟦back⟧ is relativized for each speaker. So, backness values take pair- as well as speaker-specific information to produce a generalized backness value that can be compared across harmonic pairings and speakers. Generally speaking, a value of 0 is equivalent to a categorical [-back] vowel, while a value of 1 is equivalent to a categorical [+back] vowel. Note that values below 0 and above 1 are possible, and occur in cases where the F2 value of a given token is hyperarticulated relative to its mean initial-syllable value. I assume that all values below 0 are treated like 0, and all values above 1 are treated like 1. These values reflect phonetic hyperarticulation, which is not the concern of this analysis, and so all hyperarticulated vowels are treated as equivalent to categorical [±back] vowels in the phonologies of these languages. Finally, it is also worth noting that the normalized method employed is linear. This reflects the nature of vowel perception, which is far more linear than the perception of consonant voicing (Fry et al. 1962; Lisker & Abramson 1970).

6.2.1 Uyghur

When applied to the Uyghur data, the normalization procedures just described resulted in the plots in Figure 6.1. The goal of the phonological analysis is to generate an approximation of the attested patterns for all word-types collected. Since the normalization method employed allows us to generalize across the particular backness alternations in the language, the remaining task is to account for words of varying lengths. To that end, fourteen possible word-types require an analysis, varying in length from one to five syllables, and differing in length of the morphological root from one to two syllables.

The same generalizations from Chapter 4 re-emerge in Figure 6.1 below. First, the front vowels (light grey) show no systematic shifts by-position, although there is a tendency for second-syllable vowels in words formed from disyllabic roots to be more fronted than their suffixal counterparts. Second, among the back vowels (dark grey), as syllable number increases, vowels tend to be produced with decreasing backness. Third and finally, second-syllable, and to a lesser degree, third-syllable vowels in words formed from disyllabic roots tend to be realized with more categorically ⟦back⟧ values. This suggests that backness is more categorical within roots than across morpheme boundaries.

Mean values from Figure 6.1 are treated as data points to which to fit the HG analysis throughout this chapter. In effect, instead of analyzing each token, this allows for a reasonable abstraction over production patterns in the data. A great deal of phonetic variation is removed, and we are left with a more general depiction of harmony.

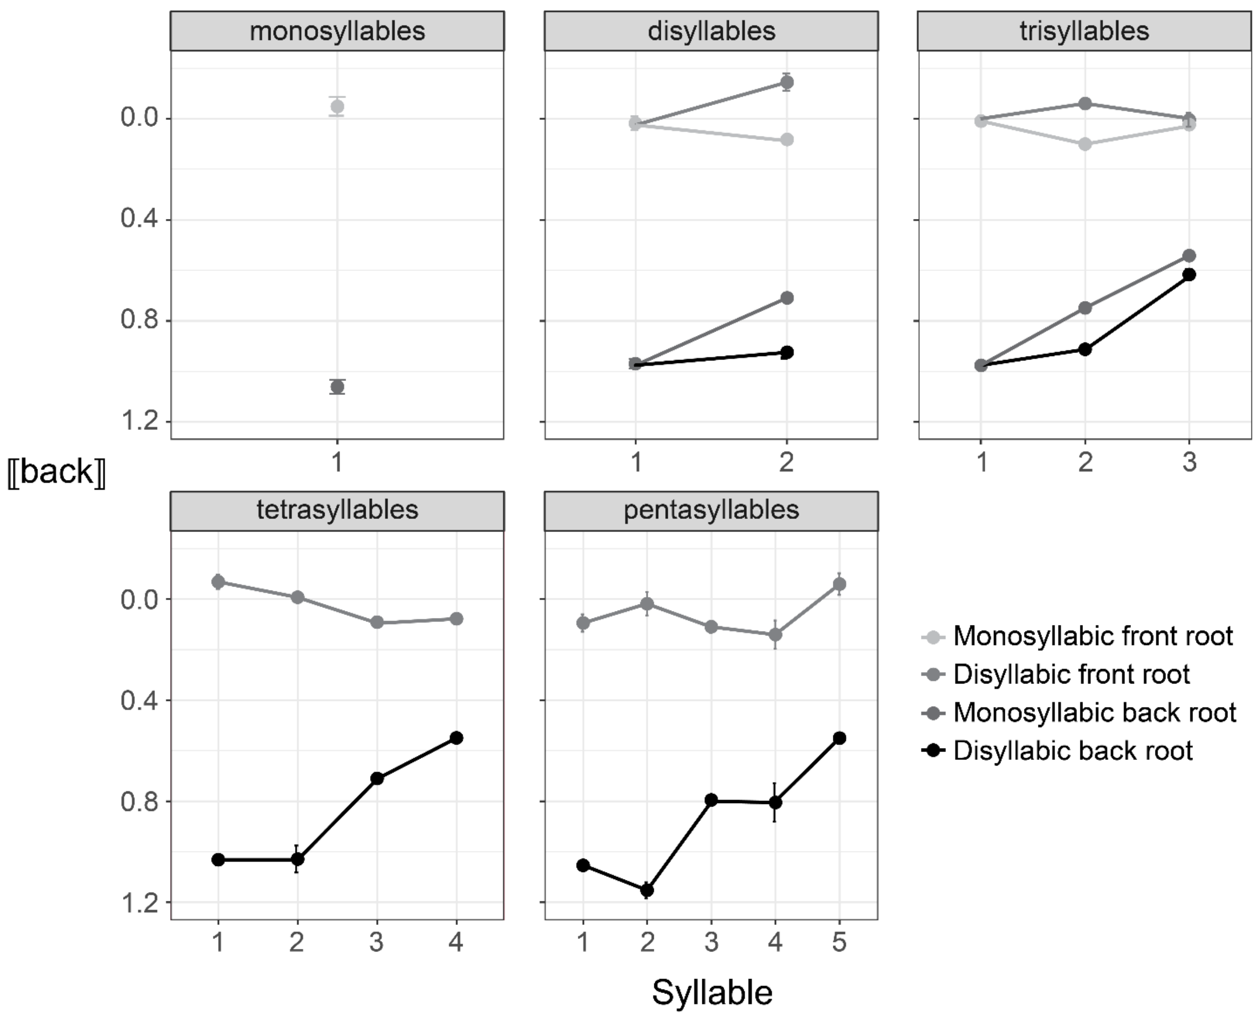

Figure 6.1: Normalized backness of Uyghur words by backness, length and root type (with SE bars). Lighter grey indicates [-back], darker grey indicates [+back]; darker colors indicate disyllabic roots.

6.2.2 Kazakh

Kazakh vowel quality after normalization generally follows the same pattern as in Uyghur, which was discussed in Chapter 4. Front vowels exhibit no systematic shifts by position while the back vowels undergo gradient fronting in non-initial syllables, shown in Figure 6.2 below. The only position that shows significant centralization for the front vowels is in the third syllable of four-syllable words. I suggest that consonantal context plays a significant role here. In roots with [-high] vowels, like /kørpie/ 'sleeping mat', the third syllable of a four-syllable word was always /-lier/ 'PL.' In this context, /r/ triggers anticipatory lowering of F2, producing a pattern of phonetic centralization. In roots with [+high] vowels, like /ılım/ 'science', the third syllable of four-syllable was always the root-final consonant plus /-ım/ 'POSS.1S.' The bilabial nasal following the third-syllable vowel, as one would expect, triggers slight depression of F2 due to anticipatory lip rounding. Like F2 lowering before /r/, this effect is orthogonal to the general harmony pattern, so I do not attempt to analyze this effect. Before moving on, note that these same suffixes were elicited in other positions, but in those positions a number of other suffixes, e.g. /-die/ 'LOC', /-dien/ 'ABL', /-nı/ 'ACC' and /-ı/ 'POSS.3S', were also elicited, nullifying the low-level effects of the suffix-final consonants in PL and POSS.1S.

In contrast to Uyghur, morphology appears to exert no significant effect on Kazakh vowel quality. In words formed from disyllabic roots, there is no general trend toward more peripheral vowel qualities, which suggests a slightly different harmony pattern in Kazakh. In Uyghur, harmony is defined by both morphology and a general gradient fronting of back vowels. In Kazakh, though, only the gradient fronting pattern is evident.

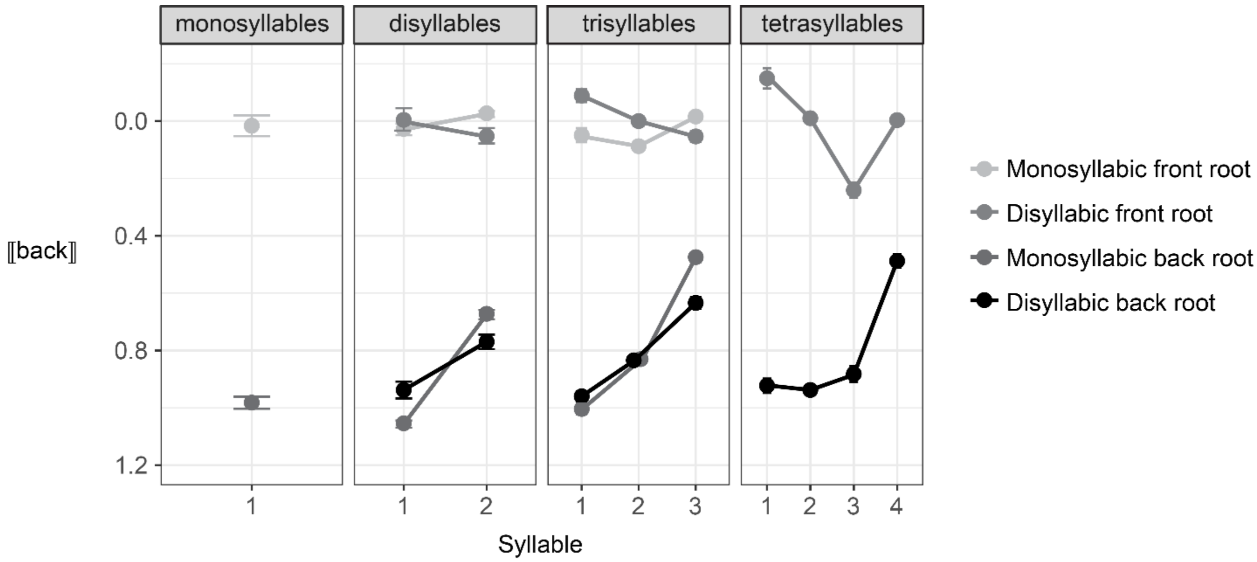

Figure 6.2: Normalized backness of Kazakh words by backness, length and root type (with SE bars). Lighter grey indicates [-back], darker grey indicates [+back]; darker colors indicate disyllabic roots.

6.2.3 Kyrgyz

The Kyrgyz data after normalization shows the same pattern observed in Chapter 4, gradient, symmetrical centralization, which is shown in Figure 6.3.[9] The ⟦back⟧ values for both front and back vowels shift toward a more central value in non-initial syllables. Like Uyghur, morphology affects the realization of vowel backness, with words formed from disyllabic roots being produced with more peripheral vowel qualities. Similar to Kazakh, the ⟦back⟧ value of third-syllable front vowels in four-syllable words is surprisingly high. The same explanation offered for Kazakh holds for Kyrgyz, too, as vowels in these positions were always followed by the /r/ of the PL suffix, or /m/ of the first-person possessive suffix.

[9] In this plot, words formed from monosyllabic [-back] roots don't show centralization although words formed from disyllabic [-back] roots do. Perhaps disyllabic roots are hypearticulated, and this phonetic effect diminishes throughout the word. If words from [-back] roots don't generally undergo centralization, then perhaps Kyrgyz exhibits gradience, too. If Kyrgyz harmony is construed as gradient, though, assimilation is still more complete than in Kazakh or Uyghur.

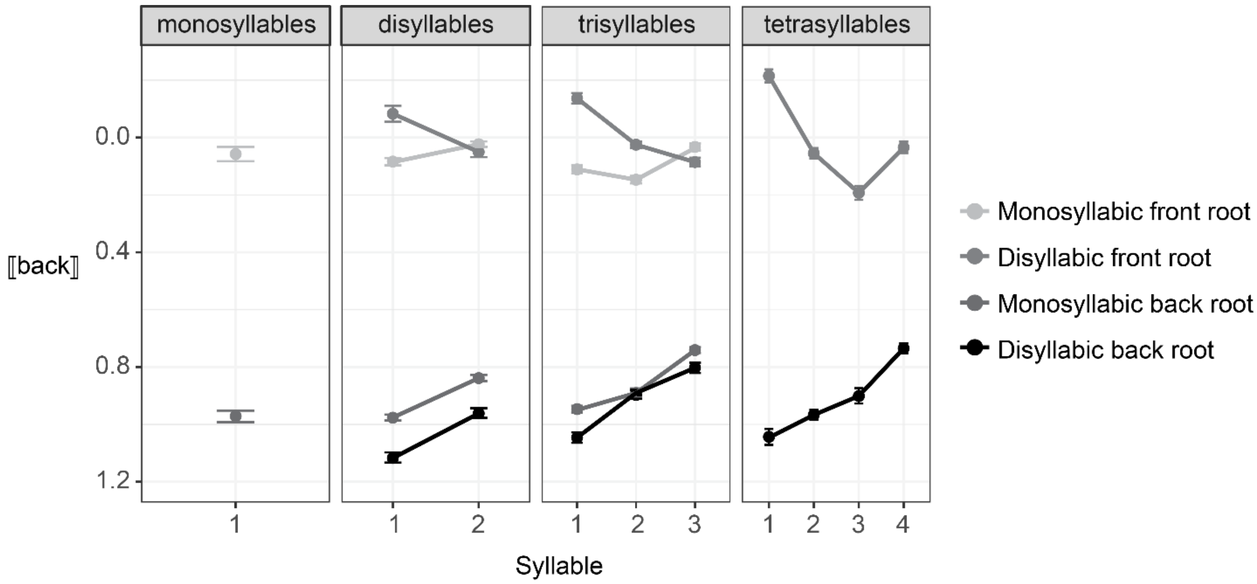

Figure 6.3: Normalized backness of Kyrgyz words by backness, length and root type (with SE bars). Lighter grey indicates [-back], darker grey indicates [+back]; darker colors indicate disyllabic roots.

6.3 Analysis

In this section I demonstrate how to incorporate gradient representations into HG, building an analysis of gradient backness harmony in Uyghur and Kazakh, which is compared to the analysis of categorical harmony in Kyrgyz.

6.3.1 Uyghur

6.3.1.1 Categorical model

Thus far I have argued that the typical conception of phonology, which manipulates categorical variables only, cannot adequately account for the patterns attested in languages like Uyghur and Kazakh. In Chapter 4, I argued that known phonetic forces of reduction and underspecification-dependent interpolation are inconsistent with the attested shifts in vowel quality in these two languages. If this is the case, the phonologist cannot assume that all gradient backness is phonetic in nature, and if not phonetic, these gradient values require a phonological account.

To lay the groundwork for the central analysis developed in this chapter, I provide a brief HG analysis of backness harmony in Turkic. Since the harmony is for vowel backness, we need markedness and faithfulness constraints that refer to the feature [back]. The markedness constraint in (48) penalizes every instance of [+back] on the surface, and the faithfulness constraint in (49) penalizes all input-output mismatches for [back].

(48)	*BK	assign a violation to every vowel that bears the feature [+back]
(49)	IDENT-IO[BK]	assign a violation to every input-output pair that disagrees for [back]

In addition to these two constraints, we need a constraint to motivate harmony. The AGREE constraint in (50) prohibits vowel bigrams that disagree in backness. Agree does not impose directionality on assimilation (Lombardi 1999; Baković 2000), so to enforce that fact that the initial-syllable controls harmony in Turkic, I use a positional faithfulness constraint to protect the initial-syllable vowel from assimilation for [back] (51).

(50)	AGREE[BK]	assign a violation to every syllable-adjacent vowel pair that disagrees for the feature [+back]
(51)	IDENT-IO[BK]$_{\sigma 1}$	assign a violation to every input-output pair in the initial syllable that disagrees for [back]

The four constraints introduced in (48-51) and their weights are exemplified in (52). The most highly weighted constraint is the positional faithfulness constraint, enforcing the fact the initial syllables never alternate for harmony. Since AGREE[BK] is weighted much more highly than both *BK and ID-IO[BK] in (52), categorical harmony obtains in the language. Candidate (52a) below assimilates the first-syllable vowel to the front vowels elsewhere in the word, incurring a violation of both the positional faithfulness and general faithfulness constraints. Since the weights of these are 10 and 1, respectively, and all violations are standardly assigned negative values, this output candidate receives a Harmony score of -11. Candidate (52b), the faithful candidate, fails to assimilate any of the non-initial front vowels, incurring a single violation of AGREE[BK] and the markedness constraint *BK. As a consequence, candidate (52b) receives a Harmony score of -9 ((-8 * 1) + (-1 * 1) = -9). Candidate (52c) assimilates the second-syllable vowel to [+back], but harmony on the third- and fourth-syllable vowels fails, incurring a single violation of AGREE, two violations of *BK, and one violation of ID-IO[BK]. Since these two constraints were assigned weights of 8, 1, and 1, respectively, this output candidate is given a Harmony score of -11. Candidate (52d) achieves harmony on the second- and third-syllable vowels, with only the fourth-syllable vowel failing to surface as [+back], incurring one more violation of *BK and ID-IO[BK] than candidate (52c). These additional violations result in a Harmony score of -13. Candidate (52e), the optimal candidate, does not violate the positional faithfulness constraint or AGREE[BK], but does incur four violations of *BK and three violations of ID-IO[BK]. The low weighting of these constraints results in full harmony in (52), as candidate (52e) receives the highest (least negative) Harmony score, -7. Although the categorical analysis is framed in terms of constraint weights rather than strict domination, the predicted outcomes are identical due to the representational assumptions that underlie both formalisms. Generally, both assume that phonological representations are categorical, and as a result, harmony either obtains or fails; there is no middle ground of partial assimilation.

(52)

	Cand.	/pɑltæ-lær-dæ/	10 ID-IO[BK]$_{\sigma1}$	8 AGREE[BK]	1 *BK	1 ID-IO[BK]	HARMONY
	a.	pæltæ-lær-dæ	-1			-1	-11
	b.	pɑltæ-lær-dæ		-1	-1		-9
	c.	pɑltɑ-lær-dæ		-1	-2	-1	-11
	d.	pɑltɑ-lɑr-dæ		-1	-3	-2	-13
☞	e.	pɑltɑ-lɑr-dɑ			-4	-3	-7

To further demonstrate the predictions of the typical phonological model, consider how model outputs compare to the attested backness values in Uyghur. This comparison is shown in Figure 6.4, with predictions from this categorical model being plotted with hollow circles and dotted lines.

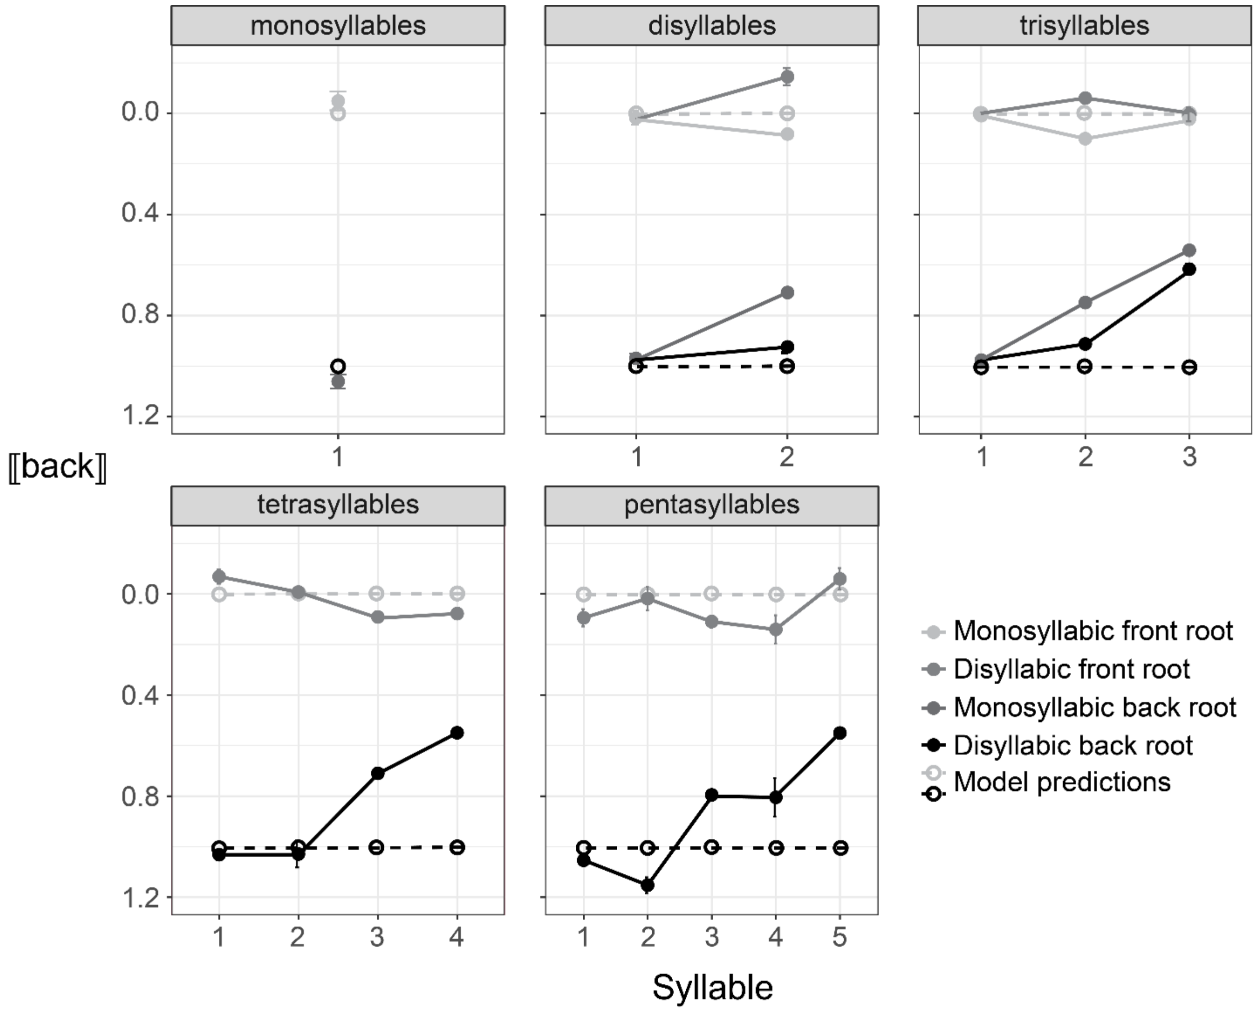

Figure 6.4: Categorical model predictions compared to attested outputs for Uyghur by backness, length and root type (with SE bars).

As can be seen in Figure 6.4, the categorical model fits the attested data relatively well for first- and second-syllable vowels, but then back vowels in the third, fourth, and fifth syllables diverge significantly from the predictions of the categorical model. In total, the categorical model accounts for roughly 84% of the attested variance in the data. I will return to this point in Section 6.3.3, where I discuss the metric and method used to calculate model fit. By recognizing [+back] and [-back] in the phonology of Uyghur, the categorical model can account for most of attested variation in the data. One might wonder if adding or modifying existing constraints could improve model fit, but crucially, since the model is constrained to produce categorical values of [back], no modification can improve the model.

6.3.1.2 Gradient representations

In order to improve model fit, and more meaningfully account for the variation attested in Uyghur, continuous variables must be introduced into the formalism. To do that, only one modification needs to be made to HG, GEN must be modified to produce gradient representations. Thus, in the gradient HG developed here, GEN may produce continuous values for a feature [F]. I will represent categorical [-F] as 0 and [+F] as 1, with continuous values in between representing gradient values for [F]. All values output from the normalization described above that result in numbers greater than 1 or less than 0 are treated in the grammar as 1 or 0. In other words, outputs that are more categorically front or back than their initial-syllable counterparts are treated equivalently to those counterparts.

If GEN can generate candidates with gradient representations, the rest of the architecture is already equipped to deal with these representations. First, consider CON. For the constraints in CON to be able to evaluate gradient representations in the candidate set, CON must contain constraints capable of gradient evaluation. Yet, at first glance CON might not seem equipped to do this. To see that CON, in its current form, is able to evaluate gradient representations, consider the evaluation of a markedness constraint, like *[BK], which penalizes output [+back] vowels. In a categorical framework, possible values to evaluate are {0,1}, while in a gradient model, they are continuous between 0 and 1, producing potentially gradient constraint violations. I propose that any amount of backness incurs a violation of *[BK]. Thus, an output candidate with vowel *V* bearing a [back] value of *n* incurs exactly *n* violations of *[BK] (53).

(53) For any output form containing vowel *V,* whose value for ⟦back⟧ = *n*, the total violations to *[BK] for *V* = *n*.

Allowing a representation with continuous values for [F], e.g. [back], does not change the essential nature of the constraint. Thus, a vowel with output ⟦back⟧ = 0.3 incurs exactly 0.3 violations of *[BK]. In this way, the expansion of representations from {0,1} to {0…1} results in no fundamental changes in how markedness constraints are evaluated.[10] There are other ways one could define a constraint like *[BK]. One could penalize all values v, $0.5 < v < 1$ ⟦back⟧, ignoring all output forms with backness $0 < v < 0.5$ ⟦back⟧. In effect, this operationalization of *[BK] would, in tandem with an equivalent *[FRONT] constraint, result in a preference for central vowels over either front or back vowels. If *[FRONT] assigns violations to all vowels with values w, $0 < v < 0.5$ ⟦back⟧, then the least marked vowel will have 0.5 ⟦back⟧, since this is the only vowel not penalized by either constraint. The definition of *[BK] in (53) instead favors categorically front vowels, since only vowels with zero backness avoid violating the constraint.

Now, consider the evaluation of a faithfulness constraint, which is a straightforward extension of the above observation for the evaluation of gradient markedness constraints. For instance, a constraint like IDENT-IO[BK] evaluates correspondent input-output pairs for mismatches in [back]. If an input vowel is specified with ⟦back⟧ = 0, and its output correspondent is specified with ⟦back⟧ = 0.7, the difference between these two is 0.7. In its most general sense, a faithfulness constraint evaluates this difference. In binary terms, the possible differences are {0,1}, but in scalar or gradient terms, these differences increase. As an example, consider Clements' (1991) analysis of height assimilation. He argues that vowel height can be treated as a single phonetic distinction, and represented as a single scalar feature, with the height of /a/ = 4,/ ɛ ɔ/ = 3, /e o/ = 2, and /i u/ =1. If height is defined in these terms, an output form involving a shift input /a/ to [i] could incur 3 violations of IDENT-IO[HEIGHT], since the

[10] I have limited discussion to simple markedness constraints, but co-occurrence constraints, like *[+round, -high] in Kaun (2004) behave differently, since a vowel with ⟦round⟧ = 0.5, and ⟦high⟧ = 0.5 could produce a non-linear interaction between ⟦F⟧ and constraint violations. For an output with a vowel bearing ⟦round⟧ = 0.5, and ⟦high⟧ = 0.5, perhaps the form would result in 0.25 (0.5 * 0.5) violations of *[+round, -high]. For more on this topic, see Mai et al. (2018).

difference in vowel height between /a/ and [i] is 3. Returning to the original example, if an input vowel is specified as ⟦back⟧ = 0, and its output correspondent is specified as ⟦back⟧ = 0.7, the difference between these two is 0.7, and this difference equals the number of constraint violations incurred. In essence, categorical changes to backness are evaluated exactly the same as gradient ones, as shown in (54). The categorical is only a special case of the regular evaluation of a faithfulness constraint.

(54) For any pair of input-output correspondents, *x* and *y*, the total violations to IDENT-IO-[BK] are equal to $|\text{Input[back]}_x - \text{Output[back]}_y|$ or $|\Delta_{x\text{-}y}|$.

Finally, the introduction of continuous variables into HG requires no modifications to candidate evaluation. Given a set of candidates whose summed product of constraint violations multiplied by their respective constraint weights equals their individual Harmony scores, the existing evaluation of candidates involves selecting the candidate with the least negative Harmony score as optimal. HG already outputs numerical values, since constraints are assigned numerical weights, so the introduction of continuous values into GEN, along with the concomitant introduction of continuously-valued constraint violations has no effect on the evaluation of Harmony, since Harmony scores are not constrained by the representational assumptions that typically hold over features. By allowing numerical weights in HG, the evaluation function must already deal with continuous values, so the introduction of continuous values at the level of GEN does not affect evaluation. With these preliminaries laid out, we can proceed to the details of the analysis.

6.3.1.3 Constraint set

The analysis proposed here utilizes four types of constraints to account for gradient harmony in Uyghur: markedness, faithfulness, constraints against gradient representations, and a harmony-driving constraint. The first and second, the markedness and faithfulness constraints from above, *[BK] and IDENT-IO[BK] are explicitly recast in gradient terms below in (55-56).

(55)	*⟦BK⟧	Assign n violations to every output vowel with ⟦BACK⟧ = n.
(56)	IDENT-IO⟦BK⟧	For any pair of input-output correspondents, p and q, assign r violations, where $r = \lvert \text{Input[back]}_p - \text{Output[back]}_q \rvert$ or $\lvert \Delta_{p\text{-}q} \rvert$.

A second faithfulness constraint is necessary to prevent changes to the initial-syllable vowel, which is always part of the morphological root (Beckman 1997, 2013). Since harmony is controlled by the initial syllable, this constraint prevents changes to the initial syllable to escape the harmonic impetus (57).

(57)	IDENT-IO⟦BK⟧$_{\sigma 1}$	For any pair of input-output correspondents, p and q, in the initial syllable, assign r violations, where $r = \lvert \text{Input[back]}_p - \text{Output[back]}_q \rvert$ or $\lvert \Delta_{p\text{-}q} \rvert$.

In addition to markedness, faithfulness, and positional faithfulness, the analysis uses constraints on gradience, as in (58). The larger claim is that gradience is, in some sense, phonologically marked. At this point I conjecture that categorical phonological patterns are easier to learn than gradient patterns, lending a learnability benefit to categoricality.

(58)	BECAT⟦BK⟧	Assign s violations to every output form with ⟦BACK⟧ = t, where s is the minimum of $\lvert t-1 \rvert$ or $\lvert t-0 \rvert$.

To see how BECAT⟦BK⟧ works, consider an output with a single vowel with ⟦BACK⟧ = 0.3. This constraint assesses the featural distance from t to both 0 and 1, and in this case assigns 0.3 (|0.3-0)| violations to this hypothetical form. In a form with a single vowel with ⟦BACK⟧ = 0.8, this constraint would assign 0.2 violations since 0.2 (i.e. |0.8-1|) is less than 0.8 (i.e. |0-0.8|). Intuitively, BECAT determines the shortest path to a categorical feature value, either 0 or 1, and returns that distance as a number of violations. These constraints are very similar to the constraints used in Lionnet (2017) and Zimmerman (2018) banning gradient outputs. While no constraints actually penalize gradient outputs in the grammatical framework, Gradient Symbolic Computation (Smolensky & Legendre 2006; Smolensky et al. 2014; Smolensky & Goldrick 2016), their model incorporates a quantization over gradient representations, producing categorical outputs. Thus, the quantization parameter in their framework limits gradience far more strictly than the constraint-type proposed above.

Lastly, the constraint set needs a constraint to motivate harmony. The constraint used below in (59) is a gradient version of Mahanta's (2008) sequential markedness constraints, which are equivalent to the harmony-drivers in Pulleyblank (2002). She recasts AGREE (Lombardi 1999; Baković 2000) as asymmetrical, prohibiting sequences where the active feature value is immediately followed (or preceded) by the opposite feature value.

(59) *[+BK][-BK] For vowels in syllables v and $v+1$ in the output, assign w violations if ⟦BK⟧ of vowel$_{v+1}$ < ⟦BK⟧$_v$, where $w = (⟦BK⟧_{v+1} - ⟦BK⟧_v)^2$.

This formulation of AGREE essentially parallels quantitative formulations in Kirchner (1997), Flemming (2001), Flemming & Cho (2017), and Braver (2013, 2019). Two properties of this constraint are noteworthy. First, since the constraint squares the distance between syllable-adjacent vowels, it produces a quadratic function, which is useful for optimization since it has a minimum, in contrast to e.g. a linear function (see discussion in Braver 2013, 2019). Second, since differences in ⟦BK⟧ are squared, it prefers gradual changes in ⟦BK⟧ over drastic changes. If this constraint were defined in linear terms, similar to a faithfulness constraint, it would not be able to differentiate between words like those in (60). If harmony is evaluated linearly and backness shifts monotonically throughout the word, the total number of violations will equal the difference between the backness of the most categorically back vowel and the most fronted vowel in the word, regardless of its trajectory. To see this, consider (61), where the linear constraint assigns 1 violation to both words while the quadratic constraint assigns 1 violation to the first and 0.5 violations to the second.

(60) V – V – V V – V – V

⟦1⟧-⟦1⟧-⟦0⟧ ⟦1⟧-⟦0.5⟧-⟦0⟧

(61)

	⟦1⟧-⟦1⟧-⟦0⟧			⟦1⟧-⟦0.5⟧-⟦0⟧		
	Violations			Violations		
	V1-V2	V2-V3	Sum	V1-V2	V2-V3	Sum
*[+BK][-BK] (quadratic)	$(1-1)^2 = 0$	$(1-0)^2 = 1$	1	$(1-0.5)^2 = 0.25$	$(0.5-0)^2 = 0.25$	0.5
*[+BK][-BK] (linear)	1-1 = 0	1-0 = 1	1	1-0.5 = 0.5	0.5-0 = 0.5	1

One thing to note is that the notion of gradience assumed throughout runs contrary to autosegmental formulations of harmony. In an autosegmental analysis, a single feature is associated with a number of segments, which is problematic given the petering out of assimilation discussed in Chapter 4. If the effect of harmony peters out of the course of the word, and if harmony is modeled as the association of a single [back] feature to multiple vowels, gradience must emerge from somewhere else. In an AGREE-based analysis, though, agreement is enforced between multiple instances of a given feature, which allows for a more straightforward formalization of diminishing assimilatory effects in non-initial syllables.

6.3.1.4 Gradient model

The basic harmony pattern is evident below in (62). The weight of the positional faithfulness constraint prevents changes to the initial syllable, and the weight of the harmony-driver is greater than the weight of the markedness and faithfulness constraints, consistent with the categorical analysis sketched in

(52) above. For ease of presentation two simplifications are made throughout. First, all gradient values are rounded to the nearest tenth, reducing the number of potential candidates to present in a tableau. Second, in a disyllabic word, 121 possible forms with varying ⟦back⟧ values are generated (ignoring the infinite other possible candidates). In the tableaux below I compare the plausible contenders with the optimal candidate while also showing something of the spectrum of violation profiles generated. To manage the size of the tableaux, I omit some candidates whose Harmony scores are much lower than the optimal candidate.

(62)

	Cand.	1 - 0	30	20	3	2	
		/bɑl-dæ/	ID-IO⟦BK⟧$_{\sigma 1}$	*[+BK][-BK]	*⟦BK⟧	ID-IO⟦BK⟧	HARMONY
	a.	0 - 0 bæl-dæ	-1			-1	-32
	b.	0.1 - 0 bæl-dæ	-0.9	-0.01	-0.1	-0.9	-29.5
	c.	0.2 - 0 bæl-dæ	-0.8	-0.04	-0.2	-0.8	-27
		…					
	d.	1 – 0.6 bɑl-dɑ		-0.16	-1.6	-0.6	-9.2
	e.	1 – 0.7 bɑl-dɑ		-0.09	-1.7	-0.7	-8.3
	f.	1 – 0.8 bɑl-dɑ		-0.04	-1.8	-0.8	-7.8
☞	g.	1 – 0.9 bɑl-dɑ		-0.01	-1.9	-0.9	-7.7
	h.	1 - 1 bɑl-dɑ			-2	-1	-8

In the above example, the suffix vowel is underlyingly specified as ⟦back⟧= 0, in conformity with the unmarked status of front vowels in the language. However, per OT's insistence that constraint ranking (or in the case of HG, weighting) drives the analysis and not language-specific representational solutions (Prince & Smolensky 1993:209), the tableau in (63) demonstrates how the HG analysis differs from (62) when the locative suffix is underlyingly specified with ⟦back⟧= 1. In this case, the analysis predicts that harmony categorically assimilates the suffix vowel. Why do the analyses differ, though? Their constraint set and weights are the same, but they predict different outputs. The answer is the general faithfulness constraint, ID-IO⟦BK⟧. The difference in prediction lies in the second syllable, so the positional faithfulness constraint indexed to the initial syllable is not the cause. Moreover, neither the markedness constraint, *⟦BK⟧, nor the harmony-driving constraint, *[+BK][-BK], are the cause, since the difference emerged from a change to the underlying representation of the suffix. This leaves only the general faithfulness constraint, which makes sense because the faithfulness constraint evaluates the transduction from input to output, and only this constraint evaluates the representational change made to the second-syllable vowel in (63).

(63)

	Cand.	1 - 1 /bɑl-dɑ/	30 ID-IO⟦BK⟧$_{\sigma 1}$	20 *[+BK][-BK]	3 *⟦BK⟧	2 ID-IO⟦BK⟧	HARMONY
	a.	0 - 0 bæl-dæ	-1			-2	-34
	b.	0.1 - 0 bæl-dæ	-0.9	-0.01	-0.1	-1.9	-31.3
	c.	0.2 - 0 bæl-dæ	-0.8	-0.04	-0.2	-1.8	-29
		…					
	d.	1 – 0.6 bɑl-dɑ		-0.16	-1.6	-0.4	-8.8
	e.	1 – 0.7 bɑl-dɑ		-0.09	-1.7	-0.3	-7.5
	f.	1 – 0.8 bɑl-dɑ		-0.04	-1.8	-0.2	-6.6
	g.	1 – 0.9 bɑl-dɑ		-0.01	-1.9	-0.1	-6.1
☞	h.	1 - 1 bɑl-dɑ			-2		-6

These differences, albeit small, suggest a fundamental problem with faithfulness in the analysis. The problem derives from the introduction of gradient representations, and the resultant grammar's ability to compromise. OT and HG cannot produce compromise, since both operate over discrete variables, and both strict domination and cumulativity produce categorical outcomes. By introducing gradient representations into the computation, the grammar is now able to produce outputs that aren't preferred by any constraint. In essence, *[+BK][-BK] compels categorical harmony while *⟦BK⟧ compels wholesale vowel fronting. The third force at work, faithfulness, simply seeks to avoid change, regardless of input specification. When underlying ⟦back⟧ = 0, faithfulness teams up with *⟦BK⟧, resulting in more gradient assimilation of the suffix. However, when underlying ⟦back⟧ = 1 or 0.9, faithfulness teams up with harmony to produce more categorical assimilation. As a result, I exclude the general faithfulness constraint from the following tableaux. This is equivalent to weighting faithfulness so lowly that its effect is negligible. With ID-IO[BK] excised from the tableaux below, the input specification of the suffix vowel no longer makes divergent predictions. Instead, in both tableaux in (64) and (65), the prediction is the same.

(64)

	Cand.	1 - 0 /bɑl-dæ/	30 ID-IO⟦BK⟧$_{\sigma1}$	20 *[+BK][-BK]	3 *⟦BK⟧	HARMONY
	a.	0 - 0 bæl-dæ	-1			-30
	b.	0.1 - 0 bæl-dæ	-0.9	-0.01	-0.1	-27.5
	c.	0.2 - 0 bæl-dæ	-0.8	-0.04	-0.2	-25.4
		…				
	d.	1 – 0.6 bɑl-dɑ		-0.16	-1.6	-8
	e.	1 – 0.7 bɑl-dɑ		-0.09	-1.7	-6.9
	f.	1 – 0.8 bɑl-dɑ		-0.04	-1.8	-6.2
☞	g.	1 – 0.9 bɑl-dɑ		-0.01	-1.9	-5.9
	h.	1 - 1 bɑl-dɑ			-2	-6

(65)

	Cand.	1 - 1 /bɑl-dɑ/	30 ID-IO⟦BK⟧$_{\sigma 1}$	20 *[+BK][-BK]	3 *⟦BK⟧	HARMONY
	a.	0 - 0 bæl-dæ	-1			-30
	b.	0.1 - 0 bæl-dæ	-0.9	-0.01	-0.1	-27.5
	c.	0.2 - 0 bæl-dæ	-0.8	-0.04	-0.2	-25.4
		…				
	d.	1 – 0.6 bɑl-dɑ		-0.16	-1.6	-8
	e.	1 – 0.7 bɑl-dɑ		-0.09	-1.7	-6.9
	f.	1 – 0.8 bɑl-dɑ		-0.04	-1.8	-6.2
☞	g.	1 – 0.9 bɑl-dɑ		-0.01	-1.9	-5.9
	h.	1 - 1 bɑl-dɑ			-2	-6

By doing so, the constraint weightings in (64) and (65) derive the predicted outputs patterns independent of the input feature specification of the suffix, in conformity with Richness of the Base.

Note that the input specification of the initial syllables above was always 1. One might wonder what the HG analysis predicts when the initial syllable is assigned a value like ⟦back⟧ = 0.6. In (66), the predicted output crucially maintains the input specification for ⟦back⟧ in the initial syllable. This is not unexpected, since the highest weighted constraint bans changes to the initial-syllable input. However, this is problematic because the actual data from all three languages with harmony strongly suggests that initial-syllable vowels are not produced with gradient degrees of backness. Before addressing this issue, (66) demonstrates one way in which the analysis developed here produces a restricted typology of expected patterns. Note in (66) that all candidates with second-syllable vowels that are more ⟦back⟧ than the initial-syllable vowel are harmonically bounded. The predicted output, with ⟦back⟧ values equaling ⟦0.6⟧⟦0.5⟧ harmonically bounds all candidates with backness values like ⟦0.6⟧⟦0.7⟧, ⟦0.6⟧⟦0.8⟧, and ⟦0.6⟧⟦0.9⟧. The only constraint that distinguishes between these candidates is *⟦BK⟧ and since the number of violations to *⟦BK⟧ for a candidate with increasing backness will always be greater than the number of violations to *⟦BK⟧ for a candidate with non-increasing backness, all candidates with increasing backness are guaranteed to be suboptimal.

(66)

	Cand.	0.6 - 0 /bɑl-dæ/	30 ID-IO⟦BK⟧$_{\sigma 1}$	20 *[+BK][-BK]	3 *⟦BK⟧	HARMONY
	a.	0 - 0 bæl-dæ	-0.6			-18
	b.	0.1 - 0 bæl-dæ	-0.5	-0.01	-0.1	-15.5
	c.	0.2 - 0 bæl-dæ	-0.4	-0.04	-0.2	-13.4
		…				
	d.	0.6 - 0 bɑl-dæ		-0.36	-0.6	-9
	e.	0.6 – 0.1 bɑl-dæ		-0.25	-0.7	-7.1
	f.	0.6 – 0.2 bɑl-dæ		-0.16	-0.8	-5.6
	g.	0.6 – 0.3 bɑl-dæ		-0.09	-0.9	-4.5
	h.	0.6 – 0.4 bɑl-dæ		-0.04	-1	-3.8
☞	i.	0.6 – 0.5 bɑl-dɑ		-0.01	-1.1	-3.5
	j.	0.6 – 0.6 bɑl-dɑ			-1.2	-3.6
	k.	0.6 – 0.7 bɑl-dɑ			-1.3	-3.9
	l.	0.6 – 0.8 bɑl-dɑ			-1.4	-4.2
	m.	0.6 – 0.9 bɑl-dɑ			-1.5	-4.5
	n.	0.6 – 1 bɑl-dɑ			-1.6	-4.8
		…				
	o.	1 – 0.6 bɑl-dɑ	-0.4	-0.16	-1.6	-20
	p.	1 – 0.7 bɑl-dɑ	-0.4	-0.09	-1.7	-18.9
	q.	1 – 0.8 bɑl-dɑ	-0.4	-0.04	-1.8	-18.2
	r.	1 – 0.9 bɑl-dɑ	-0.4	-0.01	-1.9	-17.9
	s.	1 - 1 bɑl-dɑ	-0.4		-2	-18

Although ID-IO⟦BK⟧$_{\sigma 1}$ can prevent changes to the initial syllable, it cannot by itself ensure that the initial syllable is produced categorically. This role is handled by a BECAT constraint indexed to the root, defined in (67).

(67) BECAT⟦BK⟧$_{ROOT}$ Assign s violations to every output form in the morphological root with ⟦BACK⟧ = t, where s is the minimum of $|t-1|$ or $|t-0|$

Introducing this constraint into the analysis and assigning it a high weight resolves the issue of initial-syllable gradience. As seen below, even if the initial syllable is underlyingly given some gradient value for ⟦back⟧, the analysis correctly predicts that the initial syllable should be output categorically.

(68)

	Cand.	0.6 - 0 /bɑl-dæ/	40 BECAT ⟦BK⟧$_{RT}$	30 ID-IO ⟦BK⟧$_{\sigma1}$	20 *[+BK][-BK]	3 *⟦BK⟧	HARMONY
	a.	0 - 0 bæl-dæ		-0.6			-18
	b.	0.1 - 0 bæl-dæ	-0.1	-0.5	-0.01	-0.1	-19.5
	c.	0.2 - 0 bæl-dæ	-0.2	-0.4	-0.04	-0.2	-26.8
		…					
	d.	0.6 - 0 bɑl-dæ	-0.4		-0.36	-0.6	-25
	e.	0.6 – 0.1 bɑl-dæ	-0.4		-0.25	-0.7	-23.8
	f.	0.6 – 0.2 bɑl-dæ	-0.4		-0.16	-0.8	-21.6
	g.	0.6 – 0.3 bɑl-dæ	-0.4		-0.09	-0.9	-20.5
	h.	0.6 – 0.4 bɑl-dæ	-0.4		-0.04	-1	-19.8
	i.	0.6 – 0.5 bɑl-dɑ	-0.4		-0.01	-1.1	-19.5
	j.	0.6 – 0.6 bɑl-dɑ	-0.4			-1.2	-19.6
	k.	0.6 – 0.7 bɑl-dɑ	-0.4			-1.3	-19.9
	l.	0.6 – 0.8 bɑl-dɑ	-0.4			-1.4	-20.2
	m.	0.6 – 0.9 bɑl-dɑ	-0.4			-1.5	-20.5
	n.	0.6 – 1 bɑl-dɑ	-0.4			-1.6	-20.8
		…					
	o.	1 – 0.6 bɑl-dɑ		-0.4	-0.16	-1.6	-20
	p.	1 – 0.7 bɑl-dɑ		-0.4	-0.09	-1.7	-18.9
	q.	1 – 0.8 bɑl-dɑ		-0.4	-0.04	-1.8	-18.2
☞	r.	1 – 0.9 bɑl-dɑ		-0.4	-0.01	-1.9	-17.9
	s.	1 - 1 bɑl-dɑ		-0.4		-2	-18

In addition to resolving the issue of initial-syllable gradience, the particular BECAT constraint employed above serves as second purpose in the analysis. Recall from Chapter 4 and from Figure 6.1 above that second-syllable vowels in Uyghur surface more categorically when they are part of the morphological root than when they are suffixal. More generally, root-internal vowels are produced with more peripheral vowel qualities. I see two straightforward ways to account for this, BECAT⟦BK⟧$_{\text{ROOT}}$ or a separate harmony-driving constraint indexed to root-internal vowels (i.e. *[+BK][-BK]$_{\text{ROOT}}$). Using a second harmony-driver adds cumulative pressure for root-internal vowels to agree in backness with the preceding vowel, since these vowels would be evaluated by both a general harmony-driving constraint and the positional harmony-driving constraint. This same effect, though, is derivable from BECAT⟦BK⟧$_{\text{ROOT}}$ since the attested outputs are categorically ⟦back⟧, regardless of whether this is encoded as harmony, or as a general preference for roots to be produced with more peripheral vowel qualities. The constraint banning gradient roots serves a second purpose, though, as discussed above, forcing initial-syllable vowels to surface as either 0 or 1. This second use of BECAT⟦BK⟧$_{\text{ROOT}}$ is not replicable with another constraint motivating harmony, since that constraint is just as inert as the general harmony-driver to enforce categoricality on the trigger itself. The constraint driving harmony can prevent a second-syllable vowel from being less ⟦back⟧ than the initial-syllable vowel, but it has nothing to say concerning the actual gradience or categoricality of that triggering vowel. Thus, I adopt BECAT⟦BK⟧$_{\text{ROOT}}$ to account both for the categoricality of all root-internal vowels in Uyghur.

In (69), BECAT⟦BK⟧$_{\text{ROOT}}$ successfully prevents initial-syllable vowels, regardless of their underlying feature specification, from surfacing as gradient. In the tableau below, this same constraint ensures that root-internal vowels are categorically assimilated to the backness of the initial syllable. The weight of BECAT⟦BK⟧$_{\text{ROOT}}$ compels both root-internal vowels to categorical values for ⟦back⟧, and the fact that *[+BK][-BK] is more highly weighted than *⟦BK⟧ produces categorically back vowels within a disyllabic root.

(69)

	Cand.	1 - 0 /pɑltæ/	40 BECAT⟦BK⟧$_{ROOT}$	30 ID-IO⟦BK⟧$_{\sigma 1}$	20 *[+BK][-BK]	3 *⟦BK⟧	HARMONY
	a.	0 - 0 pæltæ		-1			-30
	b.	0.1 - 0 pæltæ	-0.1	-0.9	-0.01	-0.1	-31.5
	c.	0.2 - 0 pæltæ	-0.2	-0.8	-0.04	-0.2	-33.4
		…					
	d.	1 – 0.6 pɑltɑ	-0.4		-0.16	-1.6	-24
	e.	1 – 0.7 pɑltɑ	-0.3		-0.09	-1.7	-18.9
	f.	1 – 0.8 pɑltɑ	-0.2		-0.04	-1.8	-14.2
	g.	1 – 0.9 pɑltɑ	-0.1		-0.01	-1.9	-9.9
☞	h.	1 - 1 pɑltɑ				-2	-6

While the previous tableaux have focused on back vowels, the following tableau provides an analysis of a disyllabic root with front vowels. Despite the second-syllable vowel's underlying specification, the analysis correctly predicts that both vowels in this word will surface with categorically front vowels. As one might be able to glean from this single tableau, no constraints penalize surface front vowels, and because all gradient vowels, irrespective of where they fall on the continuum, are assigned some value for ⟦back⟧, and as a result, they incur violations of *⟦BK⟧. Stated differently, the least marked vowel in the system is a categorically front vowel, so the HG analysis pushes all front vowels to be categorically front

(70)

	Cand.	0 - 1 /sællɑ/	40 BECAT ⟦BK⟧$_{\text{ROOT}}$	30 ID-IO ⟦BK⟧$_{\sigma 1}$	20 *[+BK][-BK]	3 *⟦BK⟧	HARMONY
☞	a.	0 - 0 sællæ					0
	b.	0.1 - 0 sællæ	-0.1	-0.1	-0.01	-0.1	-7.5
	c.	0.2 - 0 sællæ	-0.2	-0.2	-0.04	-0.2	-15.4
		…					
	d.	0 – 0.1 sællæ	-0.1			-0.1	-4.3
	e.	0 – 0.2 sællæ	-0.2			-0.2	-8.6
	f.	0 – 0.3 sællæ	-0.3			-0.3	-12.9
	g.	0 – 0.4 sællæ	-0.4			-0.4	-17.2
	h.	0 – 0.5 sællɑ	-0.5			-0.5	-21.5
	i.	0 – 0.6 sællɑ	-0.4			-0.6	-17.8
	j.	0 – 0.7 sællɑ	-0.3			-0.7	-14.1
	k.	0 – 0.8 sællɑ	-0.2			-0.8	-10.4
	l.	0 – 0.9 sællɑ	-0.1			-0.9	-6.7
	m.	0 – 1 sællɑ				-1	-3
		…					
	n.	1 – 0.9 sɑllɑ	-0.1		-0.01	-1.9	-9.9
	o.	1 - 1 sɑllɑ		-1		-2	-36

To reinforce this point, consider a three-syllable word derived from a monosyllabic root (71). Since BECAT⟦BK⟧$_{\text{ROOT}}$ only compels categoricality within roots, it should not affect the realization of the second- and third-syllable vowels. Moreover, the surface quality of all three vowels in this example is not a by-product of their underlying specification, since they're all assigned gradient values of ⟦back⟧ underlyingly.

(71)

	Cand.	0.4-0.8-0.2 /bæl-lɑr-dæ/	40 BECAT ⟦BK⟧$_{ROOT}$	30 ID-IO⟦BK⟧$_{\sigma 1}$	20 *[+BK][-BK]	3 *⟦BK⟧	HARMONY
☞	a.	0 - 0 - 0 bæl-lær-dæ		-0.4			-12
	b.	0.1 – 0 - 0 bæl-lær-dæ	-0.1	-0.3	-0.01	-0.1	-13.5
	c.	0.2 – 0 - 0 bæl-lær-dæ	-0.2	-0.2	-0.04	-0.2	-15.4
		…					
	d.	0 – 0.1 - 0 bæl-lær-dæ		-0.4	-0.01	-0.1	-12.5
	e.	0 – 0.2 - 0 bæl-lær-dæ		-0.4	-0.04	-0.2	-13.4
	f.	0 – 0.3 - 0 bæl-lær-dæ		-0.4	-0.09	-0.3	-14.7
	g.	0 – 0.4 - 0 bæl-lær-dæ		-0.4	-0.16	-0.4	-16.4
	h.	0 – 0.5 - 0 bæl-lær-dæ		-0.4	-0.25	-0.5	-18.5
	i.	0 – 0.6 - 0 bæl-lær-dæ		-0.4	-0.36	-0.6	-21
	j.	0 – 0.7 - 0 bæl-lær-dæ		-0.4	-0.49	-0.7	-23.9
	k.	0 – 0.8 - 0 bæl-lær-dæ		-0.4	-0.64	-0.8	-27.2
	l.	0 – 0.8 – 0.1 bæl-lær-dæ		-0.4	-0.49	-0.9	-24.5
	m.	0 – 0.8 – 0.2 bæl-lær-dæ		-0.4	-0.36	-1	-22.2
		…					
	n.	0.4 – 0.8 – 0.2 bæl-lær-dæ	-0.4		-0.36	-1.4	-27.4

Thus, words with front vowels surface categorically because categorical front vowels are the least marked in the system. Compared to words with back vowels, the analysis of front vowel words is simple and straightforward.

6.3.1.5 Finding the weights

Thus far I have shown a constraint set with particular weights that output the generalizations from the Uyghur data. However, I have not discussed how those weights were found. Determining the constraint weights for HG typically involves setting up a system of inequalities wherein the attested output is guaranteed to receive a higher (less negative) Harmony score than all losing candidates (see Potts et al. 2010 for an overview). Using the data in Figure 6.1, reproduced below in Figure 6.5, to build an HG analysis proved problematic because there is no single set of constraint weights that produce the exact means for each word type in the figure.

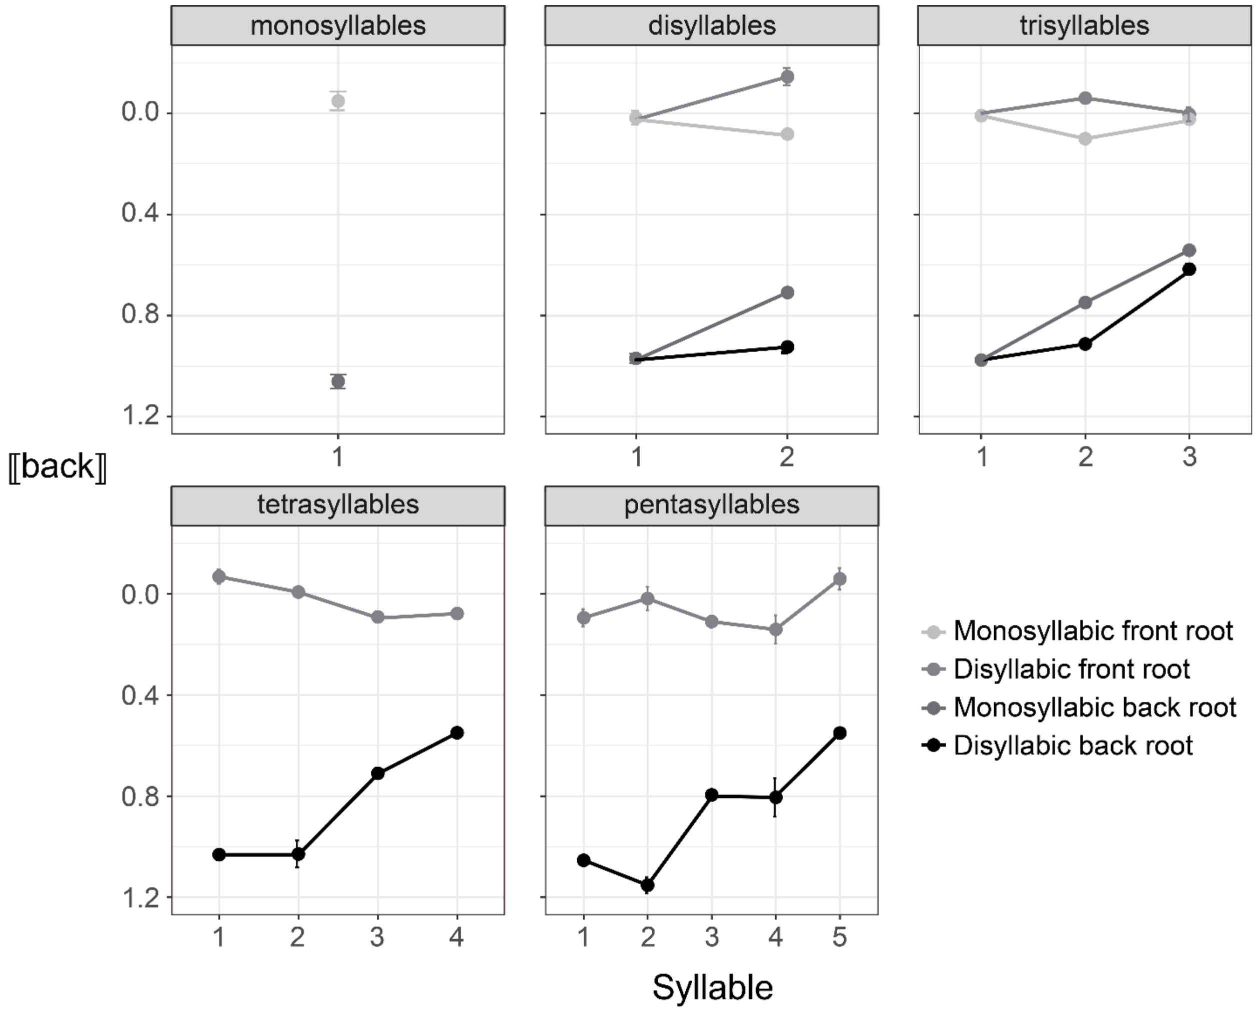

Figure 6.5: Normalized backness of Uyghur words by backness, length and root type (with SE bars). Lighter grey indicates [-back], darker grey indicates [+back]; darker colors indicate disyllabic roots.

The linear programming algorithm in Potts et al. (2010) fails for several reasons. First, some data in Figure 6.5 is greater than 1 or less than 0, and so outside the bounds imposed on the phonological grammar above. To correct this, one could simply transform all negative values to 0 and all values greater than 1 to 1. This would not resolve the second problem, random variation. Phonetic variation within each syllable for each word type has been eliminated by taking mean values, but random variation across syllables is still present. Consider the realization of second-syllable vowels in four- and five-syllable words. At the phonological level, is there any truly meaningful reason for why second-syllable vowels in five-syllable words are more categorical than those in four-syllable words? I think not, especially since there is no evidence for something like secondary stress or iterative footing that might result in the attested differences. Both four- and five-syllable words were generated from the exact same lexical items, and differ only in the concatenation of an additional suffix in the five-syllable forms. Thus, I see no phonological, or for that matter, phonetic, reason for this variability. Another instructive example emerges from five-syllable words. Elsewhere there is a consistently linear and monotonic trend toward more fronted vowels, but in five-syllable words, the mean realization of fourth-syllable vowels is slightly more backed than that of third-syllable vowels. Is there some significant phonological (or phonetic) reason for this mildly exceptional behavior in five-syllable words? I cannot see any meaningful phonological reason, as they derive from the same lexical roots, and involve the same sets of suffixes. Additionally, the metrical structure imposed on these forms doesn't provide a reason for this minor change. In terms of stress, all words, in conformity with the regular stress pattern in Uyghur, were stressed on the final syllable (Nadzhip 1971; R. F. Hahn 1998). Without some obvious conditioning factor, I attribute this slight variation in the general pattern to randomness in the data.

While the analyst may see such variance and generalize over it, the algorithm cannot. To circumnavigate this difficulty, one could either transform the recalcitrant data points to some value more consistent with the general pattern, or employ some other method for constraint weight estimation. The first method, to manually transform the data involves a level of subjectivity that is problematic. Just how much can one adjust a data point? Which data points can one adjust? Manually adjusting the data, at best introduces the possibility for "cooking the data," and at worst guarantees it. I thus employed a different method to estimate constraint weights, namely r^2, which assesses the amount of variance accounted for by a given model. Since the data are continuous, it was straightforward to find constraint weights that maximize r^2. To do this, I calculated the total sum of squares () as well as the sum of squares due to error (72). To calculate the sum of squares (SS), for the i^{th} data point, y_i, square its divergence from the mean,

y_M. Then multiply the squared difference by a weight, w, which was (and is typically) set to 1. Repeat this procedure for the entire data set, and then sum over all data points.

(72) Sum of Squares (SS) = $\sum_{i=1}^{n} w_i (y_i - y_M)^2$

To calculate the sum of squares due to error (SSE), for the i^{th} data point, y_i, square its divergence from the predicted value from the model, p_i. Then multiply the squared difference by a weight, w, which was again set to 1. Repeat this procedure for the entire data set, and then sum over all data points.

(73) Sum of Squares due to Error (SSE) = $\sum_{i=1}^{n} w_i (y_i - p_i)^2$

With SS and SSE in hand, determining r^2 is straightforward (74) – subtract the quotient of the SSE divided by SS from 1.

(74) $r^2 = 1 - \frac{SSE}{SS}$

To find constraint weights that maximized r^2, I set up tableaux in Excel for each word type, including gradient representations for all logically possible output candidates. To manage the size of the candidate space, I constrained all output representations to one decimal place (e.g. 0, 0.1, 0.2, 0.3, 0.4 … 1). I then assessed constraint violations and provided initial weights for each of the four constraints under consideration, which in turn produced Harmony scores for each candidate in each tableau. For each tableau, the output ⟦back⟧ representation (e.g. ⟦1⟧-⟦0.7⟧-⟦0.6⟧) for the candidate with the highest Harmony score (the predicted optimal candidate) was compared to the attested mean, providing y_i and p_i from (28). I then used the Evolutionary algorithm in Excel's Solver add-in (Fylstra et al. 1998) to find constraint weights to maximize r^2. The algorithm was constrained so that all constraint weights were positive integers less than 50. The highest r^2 discovered by Solver was .956, meaning that the gradient HG analysis accounts for roughly 96% of the variance in the dataset (mean ⟦back⟧ values by syllable number and word type). The predicted outputs are compared against attested outputs in Figure 6.6.

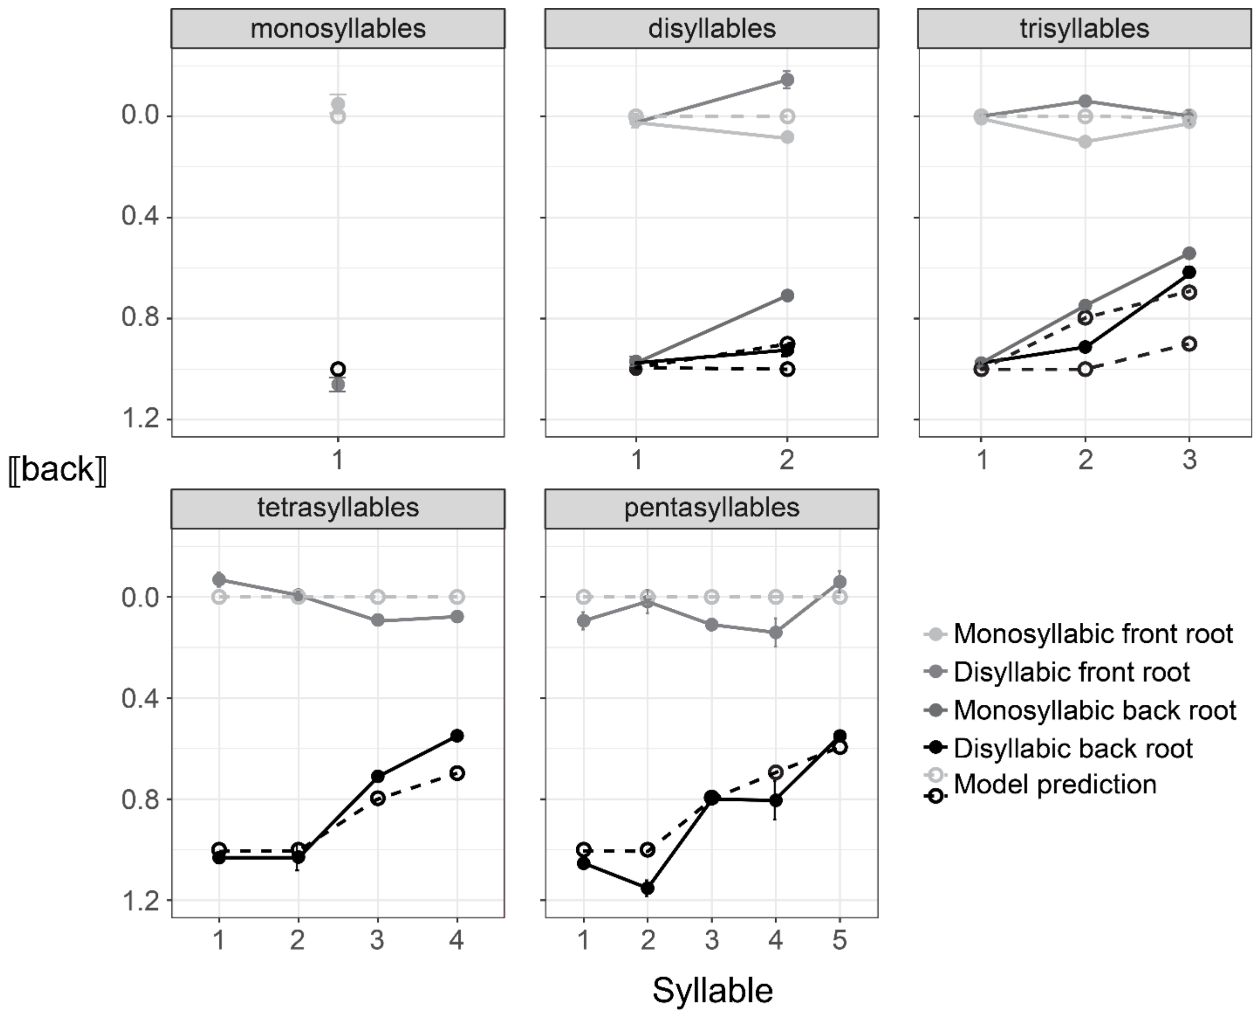

Figure 6.6: Gradient HG model predictions compared to attested outputs for Uyghur by backness, length, and root type (with SE bars)

By incorporating gradience into the analysis, model fit shows marked improvement, accounting for 12% more variance in the Uyghur data than the categorical analysis above.

As reported in Chapter 4, morphology plays a key role in the realization of back vowels in Uyghur. Within roots, back vowels undergo categorical harmony but are realized with gradient backness outside of roots. To help assess the formal role that morphology plays in the analysis, I developed another BECAT⟦BK⟧ constraint, BECAT$⟦\text{BK}⟧_{\sigma 1}$ (75).

(75) BECAT$⟦\text{BK}⟧_{\sigma 1}$ Assign *s* violations to every output form in the initial syllable with ⟦BACK⟧ = *t*, where *s* is the minimum of |t-1| or |t-0|

Recall that BECAT$⟦\text{BK}⟧_{\text{ROOT}}$ forces root-internal vowels to be produced with categorical backness. If this constraint is truly significant in the analysis, as suggested by the statistical analysis in Chapter 4, then switching BECAT$⟦\text{BK}⟧_{\sigma 1}$ for BECAT$⟦\text{BK}⟧_{\text{ROOT}}$ should affect model fit. Specifically, using BECAT$⟦\text{BK}⟧_{\sigma 1}$ in the analysis nullifies any effect of morphology, which should decrease the goodness of fit compared to the actual model. After using this variant of the BECAT⟦BK⟧ constraint, I found the constraint weights to maximize r^2 in Excel. When BECAT$⟦\text{BK}⟧_{\sigma 1}$ was used in the analysis, the maximum r^2 discovered was .923. Thus, the model with BECAT$⟦\text{BK}⟧_{\text{ROOT}}$ accounts for 3% more variance in the data than the model with only BECAT$⟦\text{BK}⟧_{\sigma 1}$. This suggests that encoding a preference for categorical harmony within roots does in fact provide better empirical coverage, encoding the generalization that root-internally, harmony is categorical.

6.3.2 Kazakh

The analysis of Kazakh largely follows the analysis of Uyghur. However, there is one notable difference between the two, the role of morphology. In Chapter 4, morphology plays an insignificant role in the realization of F2 in Kazakh, but a highly significant role in Uyghur. To account for this, the Uyghur analysis implemented a constraint against gradience within roots, BECAT$⟦\text{BK}⟧_{\text{ROOT}}$. This constraint enforces categorical harmony within roots, which accounts for the differences found in Uyghur words like pɑltɑ ‘axe’ and bɑl-dɑ ‘honey-LOC.’ The second-syllable vowel in pɑltɑ is produced with lower F2 than the second-syllable vowel in bɑl-dɑ, and I attribute this difference to morphology. Within roots, harmony is categorical in Uyghur, but across morpheme boundaries harmony is gradient. However, in Kazakh, harmony is gradient both within roots and across morpheme boundaries. Thus, the only position in which vowels are produced with categorical backness is the initial syllable. For this reason, I use the constraint introduced in (75) above, BECAT⟦BK⟧$_{\sigma 1}$. The four constraints, BECAT⟦BK⟧$_{\sigma 1}$, ID-IO⟦BK⟧$_{\sigma 1}$, *[+BK][-BK], and *⟦BK⟧, were used to fit the gradient HG analysis of Kazakh. Like in the Uyghur analysis, Excel’s Solver add-in was used to find constraint weights that maximize r^2. The optimal constraint weights discovered by the solving algorithm are presented in (76).

(76) Optimal constraint weights for Kazakh analysis

a.	BECAT⟦BK⟧$_{\sigma 1}$	30
b.	ID-IO⟦BK⟧$_{\sigma 1}$	10
c.	*[+BK][-BK]	5
d.	*⟦BK⟧	1

These four constraints and the weights in (76) fit the attested data well, accounting approximately 95% of the variance in the data, as shown in Figure 6.7.

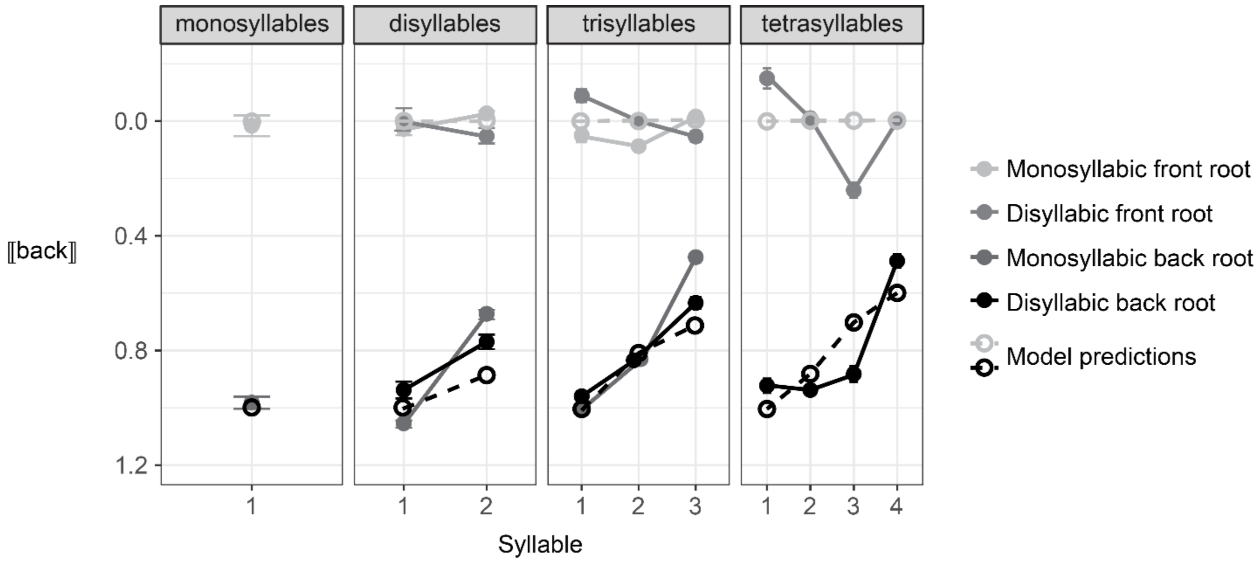

Figure 6.7: Gradient HG model predictions compared to attested outputs for Kazakh by backness, length, and root type (with SE bars)

To see these constraints and weights at work, consider the realization of a three-syllable word like [bɑltɑ-dɑ] ‘axe-LOC’ in (77). The mean ⟦back⟧ values for a three-syllable word like [bɑltɑ-dɑ] with an initial-syllable back vowel are ⟦0.96⟧-⟦0.83⟧-⟦0.63⟧ and the predicted backness values from the gradient HG analysis are ⟦1⟧-⟦0.8⟧-⟦0.7⟧.

(77)

	Cand.	0.8-0.8-0.2 /bɑltɑ-dæ/	30 BECAT ⟦BK⟧$_{\sigma1}$	10 ID-IO ⟦BK⟧$_{\sigma1}$	5 *[+BK][-BK]	1 *⟦BK⟧	HARMONY
	a.	0 - 0 - 0 bɑltæ-dæ		-0.8			-8
		…					
	b.	0.8 – 0 - 0 bɑltæ-dæ	-0.2	-0.2	-0.64	-0.8	-12
	c.	0.8 – 0.1 - 0 bɑltæ-dæ	-0.2	-0.2	-0.50	-0.9	-11.4
		…					
	d.	1 – 0.8 – 0.5 bɑltɑ-dɑ		-0.2	-0.09	-2.3	-4.75
	e.	1 – 0.8 – 0.6 bɑltɑ-dɑ		-0.2	-0.04	-2.4	-4.6
☞	f.	1 – 0.8 – 0.7 bɑltɑ-dɑ		-0.2	-0.01	-2.5	-4.55
	g.	1 – 0.8 – 0.8 bɑltɑ-dɑ		-0.2		-2.6	-4.6
		…					
	h.	1 – 0.9 – 0.5 bɑltɑ-dɑ		-0.2	-0.17	-2.4	-5.25
	i.	1 – 0.9 – 0.6 bɑltɑ-dɑ		-0.2	-0.1	-2.5	-5
	j.	1 – 0.9 – 0.7 bɑltɑ-dɑ		-0.2	-0.05	-2.6	-4.85
	k.	1 – 0.9 – 0.8 bɑltɑ-dɑ		-0.2	-0.02	-2.7	-4.8
	l.	1 – 0.9 – 0.9 bɑltɑ-dɑ		-0.2	-0.01	-2.8	-4.85
		…					
	m.	1 – 1 – 0.5 bɑltɑ-dɑ		-0.2	-0.25	-2.5	-5.75
	n.	1 – 1 – 0.6 bɑltɑ-dɑ		-0.2	-0.16	-2.6	-5.4
	o.	1 – 1 – 0.7 bɑltɑ-dɑ		-0.2	-0.09	-2.7	-5.15
	p.	1 – 1 – 0.8 bɑltɑ-dɑ		-0.2	-0.04	-2.8	-5
	q.	1 – 1 – 0.9 bɑltɑ-dɑ		-0.2	-0.01	-2.9	-4.95
	r.	1 – 1 – 1 bɑltɑ-dɑ		-0.2		-3	-5

Two points are worth mentioning at the conclusion of this section. First, a categorical model was compared with the gradient model. The categorical HG model accounted for 82% of the variance in the data, performing more poorly than the gradient HG model, which accounted for 95% of the variance in the data. Second, I compared the performance of these four constraints against the four constraints from the Uyghur analysis. In particular, comparing the performance of BECAT⟦BK⟧$_{\text{ROOT}}$ and BECAT⟦BK⟧$_{\sigma 1}$ was my chief concern. If morphology plays no role in the Kazakh data, as suggested by the statistical analysis in Chapter 4, then the model using BECAT⟦BK⟧$_{\text{ROOT}}$ should not perform better than the model using the more restricted BECAT⟦BK⟧$_{\sigma 1}$. Consistent with this prediction, the model with BECAT⟦BK⟧$_{\text{ROOT}}$ accounted for slightly less variance in the data, roughly 94%. Whereas BECAT⟦BK⟧$_{\text{ROOT}}$ yielded a 3% in explained variance in Uyghur, its introduction to the Kazakh analysis yielded a 1% decrease in explain variance, corroborating the insight from the statistical analysis- morphology affects non-initial vowel realization in Uyghur but not in Kazakh.

6.3.3 Kyrgyz

In Uyghur and Kazakh, gradient representations provide greater empirical coverage than discrete, categorical representations. To account for categorical harmony, one could either constrain the representations in a language like Kyrgyz to exclude gradience, but this would run into two types of problems. First, this would use language-specific representations rather than constraint weightings to derive the input-output mapping, counter to Richness of the Base. This first issue is formal, but the second is empirical. If gradient representations are forbidden in Kyrgyz, one cannot account for the gradient word-final high vowel neutralization discussed in Chapter 4. Recall that word-final high vowel fronting is gradient in Kyrgyz and Kazakh, but categorical in Uyghur. Thus, the analysis must be able to account for both gradient and categorical patterns with one set of representations. I do not account for gradient word-final high vowel fronting here for the sake of space, and since the general gradience of the backness harmony pattern is paramount. However, in a fuller analysis, both patterns deserve an analysis.

I have already shown that gradient representations outperform categorical [back] feature values for Uyghur, and I now demonstrate that gradient representations can account for categorical harmony, in addition to the gradient patterns analyzed for Uyghur and Kazakh. The difference between a language with gradient harmony and categorical harmony is only a matter of constraint weightings. The Kyrgyz analysis employs the same set of constraints and the same representations, but derives different output patterns from different constraint weights. To illustrate this, consider the three tableaux below. In each, the weight of the harmony-motivating constraint increases while the weight of the markedness constraint is held constant. In (78) the constraint driving harmony and the general markedness constraint receive the same weight, and the optimal output is (78f), ⟦1⟧-⟦0.5⟧.

(78)

	Cand.	1 - 0 /bɑl-de/	1 *[+BK][-BK]	1 *⟦BK⟧	HARMONY
	a.	1 - 0 bɑl-de	-1	-1	-1
	b.	1 – 0.1 bɑl-de	-0.81	-1.1	-1.92
	c.	1 – 0.2 bɑl-de	-0.64	-1.2	-1.84
	d.	1 – 0.3 bɑl-de	-0.49	-1.3	-1.79
	e.	1 – 0.4 bɑl-de	-0.36	-1.4	-1.76
☞	f.	1 – 0.5 bɑl-de	-0.25	-1.5	-1.75
	g.	1 – 0.6 bɑl-dɑ	-0.16	-1.6	-1.76
	h.	1 – 0.7 bɑl-dɑ	-0.09	-1.7	-1.79
	i.	1 – 0.8 bɑl-dɑ	-0.04	-1.8	-1.84
	j.	1 – 0.9 bɑl-dɑ	-0.01	-1.9	-1.91
	k.	1 - 1 bɑl-dɑ		-2	-2

When the weight of harmony is increased, unsurprisingly, the degree of assimilation increases. In (79), the weight of harmony is increased to three times that of general markedness, and as a result, the optimal output is (79i), ⟦1⟧-⟦0.8⟧.

(79)

	Cand.	1 - 0 /bɑl-de/	3 *[+BK][-BK]	1 *⟦BK⟧	HARMONY
	a.	1 - 0 bɑl-de	-1	-1	-4
	b.	1 – 0.1 bɑl-de	-0.81	-1.1	-3.53
	c.	1 – 0.2 bɑl-de	-0.64	-1.2	-3.12
	d.	1 – 0.3 bɑl-de	-0.49	-1.3	-2.77
	e.	1 – 0.4 bɑl-de	-0.36	-1.4	-2.48
	f.	1 – 0.5 bɑl-de	-0.25	-1.5	-2.25
	g.	1 – 0.6 bɑl-dɑ	-0.16	-1.6	-2.08
	h.	1 – 0.7 bɑl-dɑ	-0.09	-1.7	-1.97
☞	i.	1 – 0.8 bɑl-dɑ	-0.04	-1.8	-1.92
	j.	1 – 0.9 bɑl-dɑ	-0.01	-1.9	-1.93
	k.	1 - 1 bɑl-dɑ		-2	-2

Finally, when the weight of harmony is sufficiently high, categorical assimilation is optimal. In (80), the weight of harmony is increased to twelve times that of general markedness, and as a direct result, this particular constraint weighting predicts full assimilation of the second-syllable vowel, ⟦1⟧-⟦1⟧. For categorical assimilation to occur, the weight of *[+BK][-BK] must be greater than ten times the weight of *⟦BK⟧. For weights less than ten, gradient harmony is optimal. With a weight of ten, output candidates ⟦1⟧-⟦0.9⟧ and ⟦1⟧-⟦1⟧ tie, because the product of 10 and -0.01, the number of *[+BK][-BK] violations incurred by (80j) would equal 1. When this is added to -1.9, the total number of *⟦BK⟧ violations, the total Harmony score would equal -2, the same as the Harmony score for the categorical harmony candidate.

(80)

	Cand.	1 - 0 /bɑl-de/	12 *[+BK][-BK]	1 *⟦BK⟧	HARMONY
	a.	1 - 0 bɑl-de	-1	-1	-13
	b.	1 – 0.1 bɑl-de	-0.81	-1.1	-10.82
	c.	1 – 0.2 bɑl-de	-0.64	-1.2	-8.88
	d.	1 – 0.3 bɑl-de	-0.49	-1.3	-7.18
	e.	1 – 0.4 bɑl-de	-0.36	-1.4	-5.72
	f.	1 – 0.5 bɑl-de	-0.25	-1.5	-4.5
	g.	1 – 0.6 bɑl-dɑ	-0.16	-1.6	-3.52
	h.	1 – 0.7 bɑl-dɑ	-0.09	-1.7	-2.78
	i.	1 – 0.8 bɑl-dɑ	-0.04	-1.8	-2.28
	j.	1 – 0.9 bɑl-dɑ	-0.01	-1.9	-2.02
☞	k.	1 - 1 bɑl-dɑ		-2	-2

Thus, constraint weighting is sufficient, in and of itself, to differentiate categorical from gradient harmony. As a result, both types of harmony patterns can be accounted for under a single representational system.

In the examples above, assigning a weight of 12 to the harmony-driver produces categorical assimilation of a second-syllable vowel. One interesting property of HG is that this does not necessarily predict the same categorical assimilation in a three-syllable word. Under this weighting, a three-syllable word will exhibit gradient harmony despite the categorical nature of harmony in a two-syllable word. In (81), output candidate (81f) ⟦1⟧-⟦0.9⟧-⟦0.9⟧ receives the highest Harmony score, not the categorical output candidate (81m), ⟦1⟧-⟦1⟧-⟦1⟧.

To prevent candidates like (81f) from winning in this tableau, the weight of the harmony-driving constraint must be greater than 20. The logic is the same as above, which is explained in greater detail here. Candidate (81f) incurs 0.01 violations of *[+BK][-BK] and 2.8 violations of *⟦BK⟧. Therefore, its Harmony score is expressible in terms of the following equation, H(f)= -0.01x - 2.8y, where x is the weight of *[+BK][-BK] and y is the weight of *⟦BK⟧. Meanwhile, the attested candidate, (37m) incurs no violations of *[+BK][-BK] and 3 violations of *⟦BK⟧. As a result, its Harmony score is expressed as: H(m)= 0x - 3y, and since 0x = 0, H(m) = -3y. To make H(m) greater than H(f), the inequality must hold: -3y > -0.01x - 2.8y. To solve the inequality, add 2.8y to both sides of the equation, -0.2y > -0.01x. Then, divide both sides by 0.01, which produces, -20y > -1x, which can be restated as 20y < x. In other words, the weight of *[+BK][-BK], x, must be greater than twenty times that of *⟦BK⟧, y (see Potts et al. 2010 for more on finding weights in HG).

(81)

	Cand.	1 - 0 - 0 /bɑl-der-de/	12 *[+BK][-BK]	1 *⟦BK⟧	HARMONY
	a.	1 – 0 - 0 bɑl-der-de	-1	-1	-13
		…			
	b.	1 – 0.9-0.5 bɑl-dɑr-de	-0.17	-2.4	-4.44
	c.	1 – 0.9 - 0.6 bɑl-dɑr-de	-0.1	-2.5	-3.7
	d.	1 – 0.9 - 0.7 bɑl-dɑr-de	-0.05	-2.6	-3.2
	e.	1 – 0.9 - 0.8 bɑl-dɑr-de	-0.02	-2.7	-2.94
☞	f.	1 – 0.9 – 0.9 bɑl-dɑr-de	-0.01	-2.8	-2.92
	g.	1 – 0.9 - 1 bɑl-dɑr-de	-0.01	-2.9	-3.02
		…			
	h.	1 – 1 – 0.5 bɑl-dɑr-dɑ	-0.25	-2.5	-5.5
	i.	1 – 1 - 0.6 bɑl-dɑr-dɑ	-0.16	-2.6	-4.52
	j.	1 – 1 - 0.7 bɑl-dɑr-dɑ	-0.09	-2.7	-3.78
	k.	1 – 1 - 0.8 bɑl-dɑr-dɑ	-0.04	-2.8	-3.28
	l.	1 – 1 – 0.9 bɑl-dɑr-dɑ	-0.01	-2.9	-3.02
	m.	1 – 1 - 1 bɑl-dɑr-dɑ		-3	-3

Since BECAT⟦BK⟧$_{\text{ROOT}}$ is necessary to make sure root-internal vowel qualities are not gradient, and since no more than two suffixes were elicited, the weight of *[+BK][-BK] must be greater than 20 times the weight of *⟦BK⟧ throughout the analysis of Kyrgyz. A full tableau is shown below in (82) for a four-syllable word. Note that the input ⟦back⟧ values in the root are gradient in (82), and the analysis predicts that even from gradient root inputs, harmony is categorical. The weight of BECAT⟦BK⟧$_{\text{ROOT}}$ compels gradient roots to be produced with categorical backness. In tandem, ID-IO⟦BK⟧$_{\sigma 1}$ pushes gradient initial-syllable values for ⟦back⟧ towards the closest categorical value, 0 or 1. Due to the high weight associated with the harmony driver, even in forms like (82), categorical backness harmony is output from gradient inputs.

(82)

	Cand.	0.6 - 0.4 - 0 - 0 /balte-ler-de/	21 *[+BK][-BK]	20 ID-IO ⟦BK⟧$_{\sigma 1}$	15 BECAT ⟦BK⟧$_{ROOT}$	1 *⟦BK⟧	HARMONY
	a.	0 – 0 – 0 – 0 belte-ler-de		-0.6		-1	-13
		…					
	b.	0.6 – 0.4 - 0 - 0 balte-ler-de	-0.04		-0.8	-1	-13.84
		…					
	c.	0.6 – 0.6 – 0.6 – 0.6 balte-ler-de			-0.8	-2.4	-14.4
		…					
	d.	1 – 1 – 0.9 – 0.8 balte-ler-de	-0.02	-0.4		-3.7	-12.12
	e.	1 – 1 – 0.9 – 0.9 balte-ler-de	-0.01	-0.4		-3.8	-12.01
		…					
	f.	1 – 1 – 1 – 0.8 balte-ler-de	-0.04	-0.4		-3.8	-12.64
	g.	1 – 1 – 1 – 0.9 balte-ler-de	-0.01	-0.4		-3.9	-12.11
☞	h.	1 – 1 – 1 – 1 balte-ler-de		-0.4		-4	-12

The full model's predictions are compared against the attested ⟦back⟧ values in Kyrgyz in Figure 6.7. Using r^2 to assess model performance, approximately 93% of variance in the data is accounted for. Recall from earlier that 84% of the Uyghur data and 81% of the Kazakh data was accounted for under the categorical account. In contrast, the categorical analysis of Kyrgyz fits the data much better than the categorical analyses of those two languages. Despite this general fact, it is fairly clear in Figure 6.7 that a gradient analysis would probably fit the data better than the categorical analysis, especially for three- and four-syllable words. I did not consider a gradient analysis for the reasons articulated in Chapters 4 and 5, namely that positional variation in Kyrgyz was not statistically significant, in contrast to Uyghur and Kazakh. Second, results from Chapter 5 suggest that positional variation in Kyrgyz is not likely perceivable, thus not meeting a baseline criteria for phonological distinctions. Thus, while modeling Kyrgyz as gradient could in principle fit the data better it runs counter for the argumentation developed in the previous chapters.

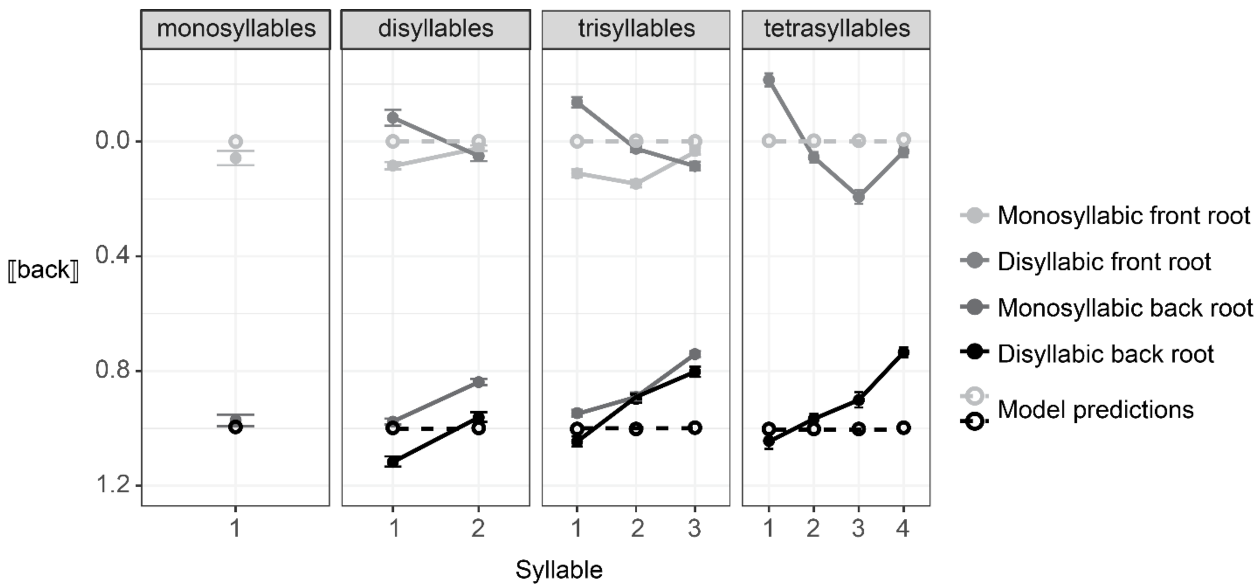

Figure 6.7: Gradient HG model predictions compared to attested outputs for Kyrgyz by backness, length, and root type (with SE bars)

Since harmony is categorical, we cannot assess the role of morphology independent of harmony, since the general pattern is categorical to begin with. One question at this point is why are non-initial vowels in words derived from disyllabic roots more peripheral than non-initial vowels in words derived from monosyllabic roots? I can envision two possible answers.

One, those vowels are more peripheral because root-internal vowels are less susceptible to centralization. In an information-theoretic sense, it is more advantageous to categorically articulate a root than an affix. To efficiently communicate a message, ensuring that the morphological root is understood outweighs the accurate perception of number, case, or possession-related suffixes. Also, affixes are more predictable than roots, if for no other reason, there are only a few dozen of them while there are myriad roots in each language, so in addition to the semantic primacy of roots, there is a much larger space of possibilities for roots, since there are tens of thousands of them, in contrast to the handful of affixes in any of these languages. If affixes are more predictable, and if more predictable morphological and syntactic units are more likely to be reduced (Seyfarth 2014), then it would make sense for roots to be produced more categorically, completely independent of the harmony system. Since predictability and informativity play a key role in speech production, it may be the case that in both Kyrgyz and Uyghur, the effect of root length is a phonetic fact and not a phonological one. This is a valid possibility, and to test this out, it would require collecting words derived from disyllabic roots that differ in their frequency across the lexicon. The prediction from this account would be that less frequent roots would be produced more categorically while more frequent roots would be produced less categorically. Since harmony in these languages does not appear to enhance contrast, the relative categoricality of these words would then affect subsequent suffixes, producing different word-level patterns of harmony.

A second possibility is that harmony is more strictly obeyed within roots. A number of languages exhibits patterns like this. In Kazakh, rounding harmony is more frequently observed within roots than across morpheme boundaries (McCollum 2018). Rounding harmony in Yeyi similarly applies more consistently within roots (Seidel 2008). Archangeli & Pulleyblank (2007) argue that ATR harmony in Ngbaka is root-delimited. Moreover, Archangeli & Pulleyblank (2002) demonstrate that leftward ATR harmony in Kinande is dictated by morphological constituency. Thus, morphology is known to play a significant role in a number of vowel harmony patterns, and it may be the case that Uyghur and Kyrgyz simply exhibit patterns similar to those immediately above.

6.4 Summary

In the previous section I demonstrated that a gradient HG analysis successfully accounts for the vast majority of the data in Uyghur and Kazakh, while a categorical HG analysis of Kyrgyz accounts well for the Kyrgyz data. Model fit for the various gradient and categorical models discussed in this chapter are presented in Table 6.2 below.

Table 6.2: Comparing model fit of gradient and categorical analyses

Language	Model fit (in r^2)	
	Gradient analysis	Categorical analysis
Uyghur	.956	.841
Kazakh	.947	.812
Kyrgyz	N/A	.932

Crucially, the gradient analyses of Uyghur and Kazakh backness harmony developed in this chapter are able to account for much of the data than their categorical counterparts. This falls out from one simple yet profound change- adopting gradient representations. By doing so, the analysis is able to make sense of asymmetric vowel fronting in these two languages, a pattern that is unaccounted for under the typical assumption that phonological computation is categorical. While previous work introducing gradience into the formalism, e.g. Lionnet 2017, assumes that phonetic detail is the source, I have argued above the phonology is the source of gradience in Uyghur and Kazakh. Furthermore, although previous work has allowed gradience to factor directly into phonological computation, by and large these works still assume that phonology is categorical, even if it has access to phonetic gradience (e.g. Hayes et al. 2004; McCollum 2018). I have argued that phonological computation itself is gradient.

The implications of the claim that phonology is gradient are further explored in the next chapter. Chapter 7 discusses evidence for gradience elsewhere in human grammar, and details a number of proposed distinctions between phonology and phonetics. I suggest that the distinctions discussed in the next chapter supplement the notion of gradience to provide a fuller understanding of the phonology-phonetics interface. From there I lay out potential avenues for future work and the continued investigation of phonological gradience.

Chapter 7: Discussion

The three previous chapters focused exclusively on the analysis of the Kyrgyz, Kazakh, and Uyghur from acoustic, perceptual, and formal points of view. This chapter, however, zooms out from Turkic, discussing the implications of gradient phonology more generally. First, however, this section presents typological data on gradient phonology, reporting on a cross-linguistic survey of impressionistic and phonetic reports of gradient harmony patterns, including vowel harmony, consonant harmony, nasal harmony, and emphasis harmony. After discussing these results, the chapter moves on to discuss a larger set of diagnostics for distinguishing phonological from phonetic patterns. Finally, the chapter discusses future directions for investigating and analyzing gradience in phonological theory.

7.1 Reports of gradient harmony

Recall from the previous chapter that the analysis developed predicts monotonic shifts from marked to unmarked. In essence, if [+back] is the active feature value in Turkic, then no target vowel will be more posterior than the trigger vowel. In both Uyghur and Kazakh, backness diminishes monotonically across positions in the word, which falls out nicely from the constraint-based account developed in the previous chapter. It is not immediately obvious, though, that this prediction is a good one. When harmony is gradient, is it always the marked feature value that is reduced? Also, is it always the case that gradient harmony results in contrast reduction, where alternating pairs further from the trigger vowel are realized with less distinctive acoustic and articulatory configurations? Based on my asearch, the answer is yes to both questions. Extant descriptions of gradient harmony always suggest that the active feature is less associated with segments further from the original trigger. Also, in every case that I've found, contrast is reduced in positions further from the original trigger.

7.1.1 Vowel harmony

At present, in addition to Uyghur and Kazakh backness harmony, I'm aware of at least thirteen other vowel harmony patterns that exhibit gradience. Among these, gradient ATR, rounding, and backness harmony patterns are all attested. ATR harmony is by far the most common of the gradient patterns reported in the literature. One case comes from Rangi (Bantu F.33). Stegen (2002) reports that Rangi has seven short vowel contrasts, /i ɪ ɛ a ɔ ʊ u/; [+hi, +ATR] vowels triggering harmony on preceding non-low vowels, creating mid-vowel allophones [e o]. Stegen (2002:137) writes, "[t]his process seems to be gradient, i.e., having diminished effect with increasing distance from the [+ATR] spreading vowel."

Gradient phrasal harmony is attested in Kinande, Nawuri, and Akan. Kinande (Bantu JD.42) has an inventory of seven underlying vowels, /i ɪ ɛ a ɔ ʊ u/, and ten surface vowels, which include three [+ATR] allophones, /ə e o/ (Mutaka 1995; Gick et al. 2006). The feature [+ATR] spreads bidirectionally, although only leftward harmony may cross word boundaries (Mutaka 1995). Regressive phrasal harmony is dependent on vowel height, occurring only when the final vowel of the target word is [+high], although harmony may affect both mid and high vowels. Mutaka's (1995:52) observations parallel Stegen's description of Rangi,

> this ATR spreading is 'gradient' and it depends upon the rate of speech. In a very careful speech, the first word may be entirely [-ATR]. Following Pulleyblank's explanation (personal communication, July 19, 1994), gradient spreading, when applied to Kinande, should be understood in two senses. The first sense is that the left edge of the harmonic domain may optionally extend one, two, three, etc. vowels away from the source. *The second sense is in the degree of advancement; that is the source of phrasal harmony becomes fully advanced, with vowels getting gradually less and less advanced, in a non discrete fashion, as a vowel gets further and further from the source of harmony* [emphasis mine].

Nawuri (Kwa) has a nine-vowel inventory, /i ɪ e ɛ a ɔ o ʊ u/ with bidirectional [+ATR] harmony. [+ATR] may spread across word boundaries in both directions, but rightward harmony is non-iterative, while leftward phrasal harmony may affect a number of vowels (Casali 2002). Regarding this leftward phrasal harmony, Casali (2002:25) comments:

> The extent of leftward postlexical [+ATR] spreading is variable and dependent on rate or style of speech. *My impression is that assimilation is in many such cases only partial, resulting in a vowel of intermediate quality*. In casual speech, however, cases of [+ATR] spreading over several syllables, as in the above examples are not uncommon. As there is no indication that leftward [+ATR] spreading is limited to any specifiable number of syllables, I assume that the process is unbounded, *although its effect may tend to diminish with distance from the triggering [+ATR] morpheme*, as has been claimed for Akan (Clements 1981) [emphasis mine].

As Casali references, Akan is also reported to exhibit gradient [+ATR] harmony across word boundaries. Akan (Kwa) has an inventory of nine oral vowels /i ɪ e ɛ a ɔ o ʊ u/ with a tenth surface vowel, a [+ATR] allophone of the low vowel. Bidirectional harmony within the word is commonly reported (cf. Casali 2012), but in the leftward direction, harmony may cross word boundaries (Clements 1981; Dolphyne 1988; Kügler 2015; Berry 1957). Describing this leftward phrasal harmony, (Clements 1981:157) reports: “As first observed in (Berry 1957), Vowel Raising [phrasal ATR harmony] is not local to the syllable immediately preceding the conditioning syllable *but influences the articulation of preceding syllables as well, causing them to acquire increasingly raised variants in a gradual “crescendo” as the conditioning syllable is approached* [emphasis mine].” More nuanced results are reported in subsequent studies. Hess (1992) does not find evidence for gradient assimilation of /a/, but does report a gradual, gradient effect in mid vowels. Casali (2012) shows evidence for a gradient effect on /ɪ/, but Kügler (2015) reports that only the word-final vowel is assimilated to [+ATR].

Gradience is demonstrated within a different domain in Somali (Cushitic). According to (Saeed 1999), the language has an inventory of ten vowels, /i ɪ e ɛ æ ɑ ɵ ɔ ʉ u/. These vowels divide into two sets, called *culus* “heavy” (~[+ATR]) and *fudud* “light” (~[-ATR]). In Nilsson & Downing (2018), harmony is reported to operate within the phonological word, with a stronger effect in the leftward direction. Generally, Nilsson & Downing (2018) notes that “the effects of the harmony are somewhat gradient, so that there is more of an effect close to the triggering culus morpheme." Bwisi (Bantu JE.102) is another African language reported to exhibit gradient ATR harmony. Much like Stegen's description of Rangi, Tabb (2001:11) writes, "[t]he farther the usually afftected vowel is from the +ATR vowel, the less it is affected." It appears that speakers are even aware of harmony petering out. Tabb notes that speakers consider at least some high vowels in a [+ATR] spreading domain to surface as [-ATR]. Based on his discussion, this does not appear to be an issue of exceptional non-undergoers, but rather the result of harmony failing to spread throughout a relatively large domain. Gradience is also observed in Ikoma (Bantu E.45) ATR harmony, subject to vowel backness. While the effect of vowel backness is well attested in Bantu height harmony, (e.g. Hyman 1999), Ikoma shows a similar pattern for ATR harmony. Ikoma has a seven-vowel inventory, /i e ɛ a ɔ o u/, with gradient rightward harmony. Suffixes with underlying /ɛ/ undergo categorical assimilation to [e] after a [+ATR] morpheme while suffixes with underlying /ɔ/ only undergo gradient assimilation to a vowel in between [o] and [ɔ] (Higgins 2011:§6.4.2). Higgins goes on to show statistically significant differences between F1 of derived and underlying back mid vowels, although she finds no statistically significant differences in F1 for derived and underlying front mid vowels, further suggesting that harmony is categorical among the front mid vowels and gradient among the back mid vowels.

French exhibits a restricted type of vowel harmony whereby a final-syllable vowel triggers assimilation of the preceding vowel. The harmonizing feature has been variously described as tense~lax, height, or ATR. Nguyen & Fagyal (2008) investigates harmony in dialects of continental French, finding gradient assimilation of penultimate vowel targets. They argue that harmony in French is not simply phonetic assimilation due to across-speaker and across-dialectal differences, in tandem with the particular trajectories of F1 shifts observed, which are incongruent with expected patterns of coarticulation. Related work on Canadian French has observed similarly gradient assimilation of penultimate vowels. Jeffrey Lamontagne (personal communication) suggests that harmony may operate categorically for the mid vowels but gradiently for the high vowels in Canadian French (see also Dalton 2012; Lamontagne 2018).

In addition to the ATR patterns described above, gradience is attested among a number of rounding harmony patterns. Although the focus of the acoustic study in this chapter was backness harmony, gradient effects of rounding harmony are also discernible in Kazakh and Uyghur. In particular, observe the F2 differences between [round] harmonic pairs like ɯ-ʊ in Kazakh and ɯ-u in Uyghur across positions in Figures 7.1 and 7.2 below. In Uyghur, note that the F2 difference between /ɯ/ and /u/ is relatively large earlier in the word but is substantially reduced by the third syllable. Impressionistically, it is challenging to hear a rounding later in the word, although there is a slight visible difference in lip posture.

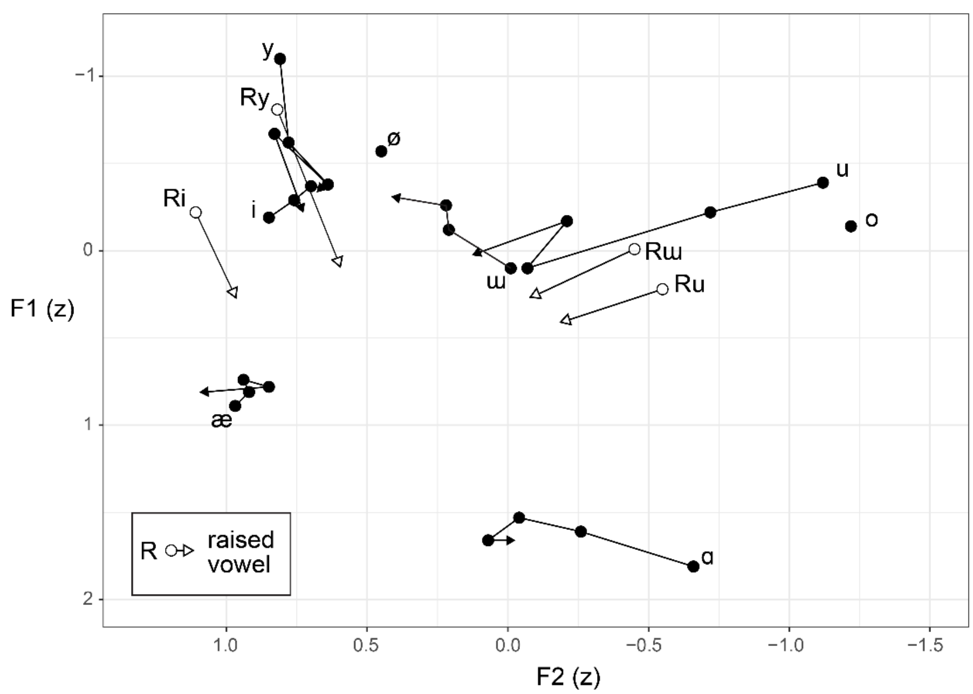

Figure 7.1: Mean F1-F2 (z) by position in Uyghur

The petering out of rounding harmony is even more noticeable in Kazakh, seen in Figure 7.2. In particular, observe that by syllable 3 there is effectively no difference between /ɯ/ and /u/. Remember that these codes represent the realization if harmony is categorical, so /u/ represents a high vowel in a context which could undergo rounding harmony (e.g. /qurum-ɪm-dɪ/ → [qurumumdu] ~ [qurumumdɯ] ~ [qurumɯmdɯ] ‘soot-POSS.1S-ACC’).

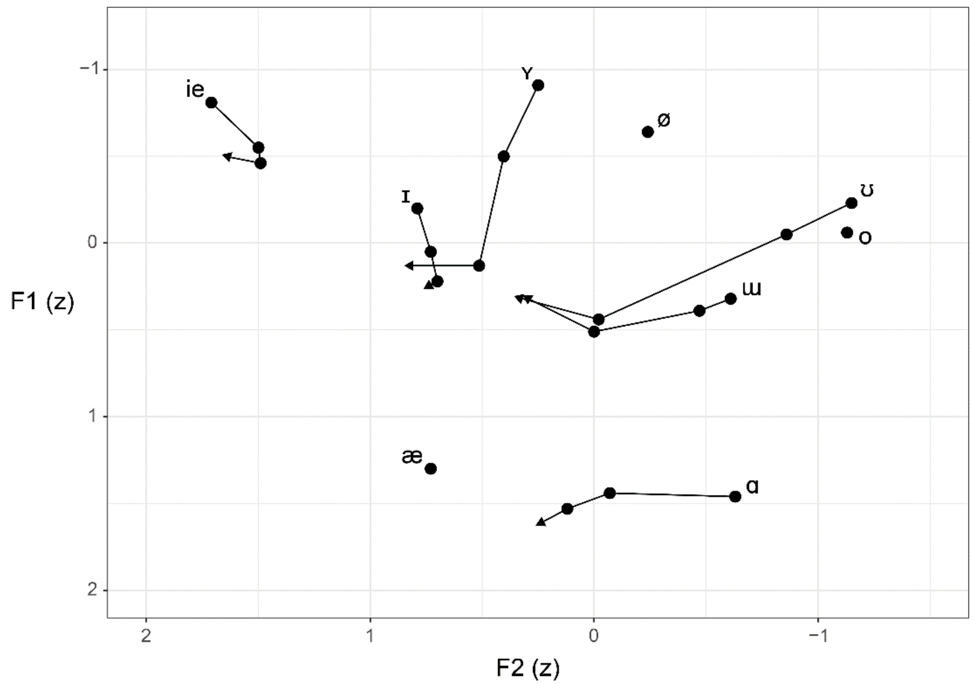

Figure 7.2: Mean F1-F2 (z) by position in Kazakh

The finding that rounding harmony is gradient in Uyghur is new. Standard Uyghur is reported to exhibit categorical rounding harmony that is lexically restricted, affecting only certain suffixes (Vaux 2000; Lindblad 1990; R. Hahn 1991). In other dialects of the language, harmony is more pervasive, affecting a larger subset of suffixes in the Turfan dialect (Yakup 2005), and affecting a larger subset of the vowel inventory in the Lopnor dialect (Abdurehim 2014).

In contrast, the finding that rounding harmony is gradient in Kazakh is relatively unsurprising. (Balakayev 1962:101) reports that the third syllable is only partially assimilated, and Kirchner (1998b:320-321) very clearly describes the petering out of rounding harmony in the language. He writes, “the strength of rounding decreases with the distance from the first syllable. both low and high vowels are affected, but the range of rounding is larger with high vowels.” As I have remarked elsewhere

(McCollum 2015, 2018), harmony exerts little effect on the mid vowels in the speech of most Kazakhs today, but Kirchner's observations for the high vowels corroborate claims by Balakayev and the data in Figure 7.2.

Gradient rounding harmony is also attested in Yeyi (Bantu R.40). Yeyi has a five-vowel inventory, /i e a o u/, with iterative regressive rounding of underlying /i/ before the two round vowels, /u o/ or the labiovelar glide, /w/ (Seidel 2008:47). According to Seidel, "the resulting vowel quality can vary between a centralized [ʉ] and a full fledged [u]." Stated differently, in many cases underlying /i/ is only partially assimilated to [u], resulting in output [ʉ].

In addition to the patterns of partial assimilation found in a number of ATR and rounding harmonies, gradient backness harmony is also reported. In particular, Hungarian exhibits the same type of gradient [+back] spreading found in Uyghur and Kazakh. Hungarian has seven short and long vowel phonemes, /i y e ø ɑ o u/, with the backness of the initial syllable dictating the backness of subsequent vowels. Booij (1984) contends that [+back] is the active feature value in Hungarian, and Szeredi (2012) finds asymmetric contraction of the vowel space in non-initial syllables, with the back vowels /ɑ o/ being realized with higher F2 in non-initial positions while the front vowel /e/ shows no comparable lowering of F2. Szeredi interprets his result in terms of stress, since stress falls on the initial syllable in Hungarian, concluding that stressed syllables show more peripheral vowel qualities than unstressed and secondarily stressed syllables. However, stress-related centralization is typically more symmetrical than what Szeredi reports, making an interpretation based on gradient harmony imminently plausible for Hungarian (see Benus & Gafos 2007 for incomplete neutralization for Hungarian front unrounded vowels).

7.1.2 Consonant harmony and vowel-consonant harmony

Gradient harmony is not restricted to vowel harmonies only, but is reported for both consonant harmony and vowel-consonant harmonies. During my search I found eight languages (including three

Arabic dialects) that exhibit gradient consonant harmony or vowel-consonant harmony. Below I describe reports for consonant harmony in Ventureño Chumash, as well as nasal harmony in Guarani, Kaiwa, and Umbundu, emphasis harmony in Arabic, and uvularization in Papantla Totonac.

In Ventureño Chumash (Chumashan), as in the related Inseneño Chumash, which figures largely in the consonant harmony literature (e.g. Rose and Walker 2004; Bennett 2015; Hansson 2010), sibilants trigger regressive assimilation for anteriority. For our purposes, it is most significant that neutralization is incomplete. Harrington (1974:5)writes:

> [t]he intention is perfect assimilation. But in actual practice the raising or lowering is largely only partial and frequently does not occur at all. Intermediate sounds between [s] and [ʃ], here written as [ṣ], arise by such imperfect assimilation or by a lowering of [s] sounds before [t l n]…The assimilation is moreover less through with some speakers than with others. Especially in slow speech and when detached words are furnished it is apt to be absent.

Russell (1993:147) and Bird (1995:106) take Harrington's description to mean that harmony in Ventureño Chumash is phonetic and not phonological. However, Poser (2004) responds with several indications that harmony is in fact phonological in the language. First, Poser argues that the harmony rule must be crucially ordered with respect to a palatalizing rule, that changes alveolar sibilants into palatal sibilants before coronal non-sibilants (i.e. /t l n/). This rule is indirectly referenced in Harrington's quote above. Second, Poser marshals evidence from Inseneño Chumash, where harmony must be crucially ordered with respect to consonant germination, which in turn interacts with a morphological reduplication rule, further suggesting the phonological character of harmony in Chumash (Applegate 1972:122, 130-137).

Partial assimilation is described for several nasal harmonies, notably Guarani, Kaiwa, and Umbundu. In Guarani (Tupian), nasal harmony is triggered by stressed nasal vowels, and may spread nasality iteratively leftward, and to a more limited degree rightward in the word (Demolin 2015; Lunt 1973). Lunt (1973:134) notes that in a nasal span, "[t]he velum may be fully raised at the outset and be gradually lowered toward the peak syllable, which contains the strongest nasalization." Lunt goes on to formalize nasal harmony as two rules, one categorical, which assigned contrastive nasality, and a second, gradient rule, which assigns n-ary degrees of nasality, in conformity with his description of incremental velum opening. Demolin (2015) also acknowledges numerous degrees of nasality in the language, suggesting that their differences may be due to phonetic enhancement, resulting in more perceptually salient distinctions between nasal and oral segments (see also Walker 1999 on the transparency of voiceless obstruents).

Harrison & Taylor (1971) describe a very similar pattern of nasal harmony in related Kaiwa. Most pertinent is their description of nasal spreading preceding a stressed nasal vowel. They note that the domain of leftward nasal spreading “may be completely nasal, completely oral, or partly oral and nasal (16-17). They also describe modulated effects due to glottal stops and speech rate, which suggests that this process is not an optional, categorical process, but rather a gradient process. Thus, both Guarani and Kaiwa spread gradient degrees of nasality leftward within the word. Gradient nasal harmony occurs in at least one African language as well, Umbundu (Bantu R11). During Schadeberg's (1982) description of leftward nasal spreading in the language, he notes discrepancies between the extent of nasalization in related words. Moreover, he describes differences in his own and native speaker impressions regarding the extent of nasal spreading. He claims that distinct degrees of nasality are clearly audible throughout the domain of nasal harmony (116-117). Thus, the same apparent diminution of nasal effect described in Guarani and Kaiwa is reported in Umbundu, too.

Arabic (Semitic) possesses a set of consonants produced with secondary pharyngeal or uvular constriction, often called emphatics. These consonants are known to trigger spreading of a pharyngeal constriction gesture, although the directionality and domain of spreading differs by dialect (Watson 2002, 1999; Davis 1995). In a number of experimental studies the degree of pharyngealization varies gradiently, affecting segments to lesser degrees as distance from the triggering segment increases (Ali & Daniloff 1972; Hassan & Esling 2011; Zawaydeh & de Jong 2011; Alarifi & Tucker 2016). In these studies, speakers of Iraqi, Jordanian, and Najdi dialects are shown to produce decreasing pharyngealization across the domain of emphasis harmony. While a number of phonetic, coarticulatory explanations are offered in these accounts, Ali & Daniloff (1972) argue against a purely phonetic interpretation of gradient spreading. They write, "[t]he source of the coarticulation observed in this study can not be attributed solely to mechano-intertial factors or to phonetic context…This leads to the conclusion that the neural command(s) responsible for the backing gesture of emphatic sounds are centrally programmed at a very high level" (101). Thus, emphasis harmony may show some of the same petering out effects attested in the vowel and consonants harmonies surveyed above. Shahin (2003) also discusses a number of faucal assimilation patterns from Palestinian Arabic and St'át'imcets (Salishan). However, Shahin's data suggests temporal changes throughout each vowel that appear more like Cohn's (1993) interpolation account of English nasalization. In both, a gradient cline marks the articulation of the vowel, in contrast to the plateaus predicted of phonological effects in Cohn (1993), and seen in Turkic (see Bessell 1998 for more on faucal harmony in Salishan).

The final case of reported gradience comes from Papantla Totonac (Totonacan), which exhibits a pattern of vowel lowering around uvulars that is similar to Arabic. The phonetic details appear to differ, since harmony in Arabic involves pharyngealization rather than uvularization, but there are many parallels between the two. Papantla Totonac has a five-vowel inventory, /i e a o u/, with vowel lowering triggered by uvulars (Levy 1987; Puente 1996). In this language a uvular triggers lowering of an adjacent vowel, e.g. /u/ → [ɔ], but the lowering effect extends gradiently across several syllables, resulting in vowel sequences like [ɔ…o…u] derived from underlying /u…u…u/ (see Levy 1987:28,70; Puente 1996:187 for examples and discrepancies in transcription).

7.1.3 Summary

I have presented a number of reports suggesting that other languages exhibit gradient harmony, in addition to the two cases, Uyghur and Kazakh backness harmony, described in Chapter 4. These patterns occur on almost every continent, among numerous language families, and involve a variety of harmonic features. These facts indicate that gradient harmony is not genetic, areally, or featurally restricted, but may occur in a variety of linguistic contexts. Furthermore, I have adduced examples of progressive (e.g. Ikoma ATR), regressive (e.g. Rangi ATR), and bidirectional (e.g. Arabic emphasis) harmonies, which suggest that gradience is not directionally constrained. Moreover, the domains of harmony discussed above include post-lexical phrasal harmonies along with metrically-defined (e.g. Guarani) and word-level (e.g. Hungarian) patterns. Thus, the domain of gradience is not constrained to some smaller span of phonetic coarticulation, or to putatively post-lexical phrasal harmonies. For all these reasons I conclude that gradient harmony is widely present, even if not always reported. It seems likely that fieldwork-based grammars underreport gradient harmony due to a variety of factors. One, the fieldworker often has many topics to address in their grammar, making in-depth analysis of one particular portion of the phonology of a language more difficult. Two, the fieldworker may not notice the pattern, since in my experience, these patterns are very subtle. Three, the fieldworker may not have the appropriate recording equipment to properly assess the gradience or categoricality of a given pattern. These factors, in conjunction with the general challenges of fieldwork, may have prevented earlier studies from reporting these sorts of patterns.

One significant generalization that emerges from the descriptions just summarized, as well as my own findings, is that gradient harmony appears to always reduce the amount of acoustic or articulatory distinctiveness between alternating vowels. This falls out naturally if gradience is tied to the marked feature value. If the marked value is reduced across positions in the word, segments should be produced with less marked feature values, producing the kind of contrast reduction described above, and exemplified in Kazakh and Uyghur. However, this is not the only possibility; if gradience is not directly linked to the marked feature value a different set of predictions emerges. Pearce (2008) argues that vowel harmony blocks phonetic reduction in the form of undershoot. From this line of reasoning it's possible to construct an analysis in which all gradience in harmony derives from phonetic reduction, similar to the adaptive hypoarticulation analysis discussed in Chapter 4. However, if the spreading of some feature value, [+F], blocks reduction, then the only feature value that may reduce is the unmarked value, [-F]. If phonetic reduction may preferentially target the unmarked feature value, there is no reason to expect gradience in harmony to produce positional reduction of contrast. To illustrate this, consider a language with nine contrastive vowels, /i ɪ e ɛ a ɔ o ʊ u/, and ATR harmony. As a growing body of work has shown, 9-vowel languages with [ATR] contrasts among the high vowels almost exclusively exhibit [+ATR] dominance (Casali 2003, 2008, 2016). In such a language, the [-ATR] vowels are unmarked, and under an analysis whereby phonetic reduction targets the unmarked value of the feature, this class of vowels could undergo centralization. If they did, this would result in contrast enhancement, as seen in Figure 7.3. The acoustic contrast between harmonic pairs, i-ɪ, e-ɛ, o-ɔ, and u-ʊ would actually increase since the [-ATR] vowels are centralized and the [+ATR] vowels maintain their vowel qualities regardless of position. The acoustic distance between a pair like /i/ and /ɪ/ would actually be smaller in triggering position, and potentially much larger in target position. The net result of such a pattern is increased contrast in syllables farther from the original trigger. If the language exhibits leftward harmony to prefixes, then this would produce initial syllables with relatively large acoustic contrasts for [ATR], but with much smaller contrasts on roots. To-date, this is unattested, but is predicted to be possible if harmony can block phonetic reduction of the marked/spreading value. If gradience is tied directly to the marked feature value, then such a pattern is not predicted.

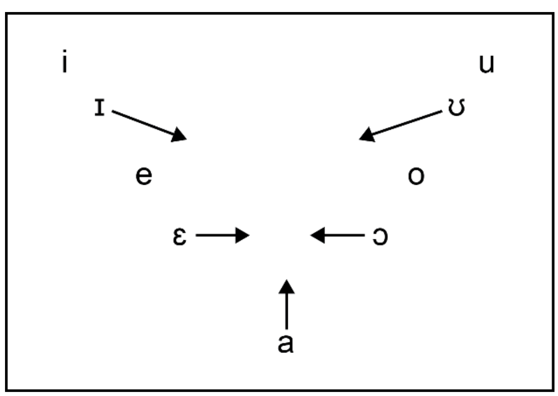

Figure 7.3: Hypothetical pattern of unmarked vowel reduction in a language with [+ATR] harmony

While we still know very little about gradience in harmony, the descriptions above overwhelmingly suggest that gradience produces contrast reduction rather than enhancement. In every case, gradient harmony produces partial assimilation, and as a consequence, a diminished contrast between two segments. If this is generally the case, then, linking all gradience to the marked feature value and distinguishing between phonological and phonetic gradience restricts the typology of expected patterns, excluding the pattern sketched above. If phonetic patterns of reduction target all segments, then we do not expect to find phonological markedness blocking reduction. As a result, only two types of patterns are expected, symmetrical reduction of all relevant segments, like backness harmony in Kyrgyz, and reduction of the marked feature value, producing contrast-reducing asymmetrical gradience, as in Kazakh and Uyghur, which conforms to the descriptions surveyed above.

7.2 Diagnostics for phonological and phonetic effects

Phonological gradience is certainly connected to phonetic gradience, in much the same way that phonology generally is intimately connected to phonetics (e.g. Kingston 2007). Due to the interconnectedness of phonetics and phonology, in all likelihood, some aspects of the harmony patterns described earlier in this section may be better analyzed as phonetic and not phonological. Particularly for patterns where the domain is relatively small (e.g. less than a word) or the pattern is described as post-lexical, it is typically possible to construct a phonetic account of the facts (e.g. Liberman & Pierrehumbert (1984) claim that all post-lexical patterns are phonetic). In general, the question arises how to differentiate phonological from phonetic effects. In this section, I discuss previous diagnostics, and then propose a set of characteristics that can distinguish phonological and phonetic forces.

7.2.1 Previous proposals

One early division of labor, as proposed in Chomsky and Halle (1968:293) differentiates phonology and phonetics in terms of language-specific and universal properties. Specifically, phonological patterns are proposed to be language-specific, while phonetic patterns are universal, physically-constrained patterns that are outside the realm of grammatical competence. This lines up with the early Generative distinction between competence and performance, with phonological competence being divorced from phonetic performance. This proposal was rejected, though, as it became clear that phonetic implementation is not universal, but differs across languages. Phonetic categories are both distinct across languages, and among other things, phonetic coarticulation differs strikingly between different languages (Keating 1985b; Keating 1985a; Kingston & Diehl 1994). As an example of these differences, consider anticipatory nasalization in English. Cohn (1993) demonstrates a cline of gradually increasing nasal airflow during the vowel immediately preceding a nasal consonant. In contrast, French exhibits minimal anticipatory nasalization, which Cohn interprets to be due to phonological factors in French. French contrasts oral with nasal vowels, thus crucially referring to [nasal] for vowel segments. English does not use a [nasal] contrast for vowels, allowing phonetic patterns of gradient interpolation across vowels unspecified for [nasal]. For nasal coarticulation patterns in non-European languages, consult e.g. Butcher 1999; Stoakes et al. 2019 for nasal coarticulation in languages of Australia. In short, Chomsky & Halle's (1968) proposal cannot account for language-specific phonetic patterns, like nasal coarticulation, suggesting a different division of labor is necessary to adequately model a modular phonology and phonetics.

A second proposal, which still holds significant sway in the contemporary literature (Cohn (2006:27) calls this the "consensus view"), is that phonology manipulates categories while phonetics manipulates gradient variables. This view conforms with Lexical Phonology (Mohanan 1982; Kiparsky 1985), wherein the core of phonological alternations are structure-preserving, while peripheral, post-lexical patterns may produce sounds not contained in the underlying inventory of the language. In essence, phonological knowledge relates to abstract, symbolic categories while phonetic knowledge relates to their translation into continuous space and time. Phonological representations are sparse while phonetic representations are significantly enriched. The proponents of this view include (Kingston 2007; Cohn 2006, 1993, 1998, 2003; Keating 1988, 1990; Zsiga 1997, 1995) among many others. Throughout these works, temporal and spectral gradience, as well as subphonemic phonetic outputs are diagnostics of phonetic effects, in contrast to categorical phonological effects. Zsiga (1995) and Zsiga (1997) use this criterion of gradience to distinguish post-lexical from lexical palatalization in English, and phonetic assimilation from harmony in Igbo. Chomsky & Halle's (1968) conception of phonetics was discarded based on a growing tide of experimental methods applied to human sound systems. In a very similar way, this "consensus" view is now under more scrutiny due to continued experimentation and richer phonological and phonetic data.

First, the prediction that phonological alternations are categorical by default expels gradient alternations from the phonology. Problematically, many putatively phonological alternations have been shown to be gradient, in that they produce sounds that are not identical to their non-alternating counterparts. For example, Warner et al. (2004) finds that final devoicing in Dutch does not produce categorically voiceless obstruents, as seen by increased vowel duration immediately preceding the target sound. Voiced obstruents trigger phonetic lengthening, and underlyingly voiced obstruents trigger a small degree of similar lengthening even though they are largely devoiced on the surface. This same result has been found in German (Port et al. 1981; Port & O'Dell 1985; Dinnsen 1985) and Catalan (Charles-Luce & Dinnsen 1987; Dinnsen & Charles-Luce 1984), among other cases of final devoicing. Some other cases of incomplete neutralization include unstressed vowel reduction in Russian (Padgett & Tabain 2005) and moraic lengthening in Japanese (Braver 2013). In many cases, researchers have utilized a variety of experimental methods to demonstrate that these effects are not artefactual, and that the incompleteness of neutralization bears on issues of linguistic and phonological competence. One common trait among the patterns examine in this literature is their status in Lexical Phonology. In

Lexical Phonology, final devoicing, moraic lengthening, and vowel reduction can be analyzed as post-lexical. Patterns like these are not cyclic, do not interact with morphology, and are not structure-preserving. Given that attested patterns of gradient, incomplete neutralization can be analyzed as post-lexical, one might think the domain of gradience is co-extensive with post-lexical phonology. Kiparsky (1985:86) argues that post-lexical rules may be gradient versions of lexical rules,

> some postlexical processes are truly phonological, feature-changing rules. I further suggest that even gradient application might not suffice to ban a process from the phonology, on the grounds that such gradient postlexical processes are in a number of interesting cases the same as rules which apply categorically in the lexical phonology of the same language, and that their gradience might be as predictable as a general property of the postlexical application of certain types of rules.

Crucially, no known cases of gradience have been reported in earlier, lexical phonology. This makes it possible to, like Kiparsky, allow gradience to enter the edges of phonology, but not the core. Cohn (2006:36) states the issue thusly: "Morphophonemic [lexical phonological] alternations are at the very core of what most phonologists think of as phonology. Most alternations are understood to be quite categorical in nature…if these sorts of cases are shown to involve gradience, this would strike at the core of our understanding of the phonology, since these are the least disputable cases for 'being phonology.'"

The findings from Chapter 4 bear directly on Cohn's claim: morphophonological harmony in Uyghur and Kazakh is gradient. These patterns are not congruent with post-lexical patterns, and yet still produce subphonemic effects that are not derivable from known phonetic forces of centralization, underspecification-based interpolation, or reduction. In tandem with the numerous cases of incomplete neutralization, the finding herein supports a revised view of phonology, allowing both categorical and gradient patterns to be modeled as part of the phonology of a language. These two types of patterns are seen in all three of the languages under study. In Kyrgyz, harmony is categorical, but high vowel fronting is gradient. In Uyghur, harmony is gradient but high vowel fronting is categorical, while in Kazakh harmony is gradient and high vowel fronting is gradient. These patterns suggest that phonological alternations, even morphophonological alternations, may be gradient or categorical.

As a result, gradience is not the sole litmus test for phonological and phonetic patterns. This claim isn't new, having been made by a number of authors for the past twenty years. Pierrehumbert et al. (2000:32-33) claims that knowledge of sound structures cannot be divided into categorical and gradient, but all fall on a cline from granular (i.e. more phonological) to more detailed (e.g. more phonetic). The role of gradient representations is paramount in work like Steriade (2000), Kirchner (1997), and Flemming (1995, 2001), which use access to gradient phonetic information to drive phonological patterning. This same approach is found in more recent work, too (e.g. Lionnet 2017; McCollum 2018). In Lionnet (2017) and McCollum (2018), phonology "sees" gradient phonetic detail, but still produces categorical effects. In Chapter 4 and 6 I argued that phonology does not just "see" gradient information, it outputs it. The findings from Chapter 4 strongly suggest that while gradience may play a role in defining phonological from phonetic, it is not the sole distinguishing factor between the two.

Tucker & Warner (2010:318) discuss the issue of gradience, noting a number of the diagnostics proposed in the literature and their challenges.

> Throughout this paper, we have been considering some factors that cause sound patterns to be phonological and others to be phonetic. However, we have not defined a single set of criteria for which is which. The literature includes a variety of criteria or definitions: a phenomenon may be considered phonetic if it is gradient and variable, if it is caused by biological or aerodynamic necessity, if it is too subtle to detect by ear without instrumental measurement, if it results in sounds that are not lexically distinctive in the language or in any language, if it is triggered or influenced by other factors that are considered phonetic, etc. . . A phenomenon is considered phonological if it is the opposite of these things: if it is consistent and categorical, involves only complete sounds that are distinctive and is perceptually obvious. A phenomenon may also be considered phonological if it interacts with highly abstract, phonological phenomena, such as morphophonology or foot structure. Furthermore, if an alternation seems arbitrary, so that it is not phonetically natural, it may be considered phonological.

Of these criteria, we have already discussed gradience as well as the Structuralist notion of biological necessity in Chomsky & Halle (1968). Chapter 6 argues that perceptibility is a prerequisite for phonological patterns, since they must be perceived if they are to be manipulated in the minds of speakers and listeners. Tucker & Warner (2010) also note that some have argued that phonological patterns are structure preserving in the Lexical Phonology sense, producing only sounds that are present in the underlying inventory of the language. A more relaxed view is also possible, as noted above, requiring that phonological patterns produce sounds that are contrastive (i.e. lexically distinctive) in some language. If there exists an alphabet of all exploited phonological contrasts (e.g. the IPA), then a pattern yielding one of these sounds could be regarded as phonological since Tucker & Warner (2010) suggests that phonetic patterns do not produce sounds that are ever contrastive in a language. The issue of structure preservation and contrast in some language are tied intimately with gradience. Structure preservation is a more restrictive requirement than contrast, since it imposes (typically categorical) contrast over a restricted set of phonetic categories. Contrast is also an issue of categories, which is simply at loggerheads with gradience generally because contrast is typically framed in absolute terms (though see Hall 2009 for a gradient notion of contrast).

In addition, Tucker & Warner (2010) notes that phonetic patterns are thought to interact with other known phonetic patterns while phonological patterns are thought to interact with more abstract, morphological, metrical, or syntactic patterns. As they note, none of these are foolproof, with many sound patterns exhibiting traits that resemble both phonetic and phonological patterns. In the following two subsections I discuss one of the diagnostics Tucker & Warner (2010) mention, naturalness, and then discuss another dimension to the phonology-phonetics interface that has not figured into previous work, computational complexity.

7.2.2 Naturalness

Naturalness is a very loaded term, and has factored largely into the development of phonological theory throughout the last half of the twentieth century and the first two decades of the twenty-first century (see especially Natural Phonology, e.g. Donegan & Stampe 1979; Natural Generative Phonology, e.g. Hooper 1976). Before discussing naturalness, we must first venture a definition of the term. Blevins (2008:126) proposes the following definition, "Natural sound patterns are sound patterns grounded in articulatory or perceptual properties of speech." This same dependence on phonetic patterns, both articulatory and perceptual, to determine the range and scope of phonological patterns, is proposed in earlier work, like Donegan & Stampe (1979), and Ohala (1979, 1990). Moreover, the role of naturalness is evident in markedness constraints in OT, and factors largely into the significant body of phonetically-based and functionalist work within OT (e.g. Hayes et al. 2004). Whether viewed in synchronic terms, as in OT, or in diachronic terms, like in work by Ohala and Blevins, the basis for naturalness in phonology is naturalness in phonetics. Phonetics is constrained by both articulatory, mechanical factors as well as perceptual and auditory factors. Significantly, in the naturalness debates that have sprung up during the last fifty or so years, no one has argued that phonetics is anything but natural. The logic goes thusly, if phonetics is natural, and if phonetics and phonology are intimately connected, then phonological patterns should be natural in the sense that they are explainable in terms of articulation and/or perception.

One recalcitrant problems for natural theories of phonology is what Bach & Harms (1972) call "crazy rules." In essence, phonological patterns are not always natural, and in some cases exhibit alternations that cannot easily be derived from any phonetic motivation, articulatory or perceptual. For many, the existence of these patterns suggests that phonology must be viewed from a diachronic perspective. For instance, Yu (2004) argues that the unnatural pattern of final obstruent voicing in Lezgian derives from several historical changes (see also Beguš 2018; Beguš and Nazarov 2019 for more on unnatural patterns). Likewise, Buckley (2000) argues that phonological patterns, like one in Kashaya, in which /i/ is realized as [u] after [d], are not ontologically constrained to be natural. Buckley contends that, in actuality, the computational nature of phonology does not privilege the putatively natural over the unnatural, but instead produces both types without distinction (see also Hale & Reiss 2000, 2008). In a similar vein, Blevins (2004) argues that during the evolution of a sound pattern it passes from a phonetic stage, in which it is constrained by naturalness to a phonological stage, in which naturalness no longer holds sway. While the debate over the naturalness of phonology and the role of diachronic factors in explaining typology is worthwhile, a simple generalization emerges for the present discussion: phonetic patterns are natural. This is perhaps so obvious that it has not been leveraged significantly to help address the phonological or phonetic status of a given sound pattern. If the large body of work from the Neogrammarians onward is correct, which it seems almost inevitably so, then one very good diagnostic for determining what is phonological and what is phonetic is naturalness. If a pattern is unnatural, then it is definitely not phonetic. As is probably obvious here, this is a unidirectional test. If a pattern is natural, this simple diagnostic cannot decide if the pattern is phonological or phonetic.

To further demonstrate the relevance of naturalness for demarcating phonological from phonetic, consider a few patterns that have been advanced in the literature as instances of unnatural phonology. First, as mentioned above, Lezgian obstruents undergo voicing in word-final position (Yu 2004). Given the preponderance of word-final devoicing patterns, and the general phonetic tendency to turn off vocal fold vibration at the end of an utterance, there is no obvious phonetic motivation for word-final voicing, and thus constitutes a prime case of unnatural phonology (cf. Ozburn and Kochetov 2018 for a re-analysis of Lezgian). Second, in Kashaya the high front vowel /i/ is backed and rounded to [u] after [d], although coronals are known to trigger vowel fronting cross-linguistically (Buckley 2000). Buckley outlines a number of reconstructed suffix variants, as well as the historical changes necessary to produce this particular rule. In short, the rule appears to derive from the historical durative suffix, /-adu/, whose final vowel was, at some point, reinterpreted as epenthetic. This reanalysis of /u/ as epenthetic was then generalized such that the typical epenthetic [i] was replaced by [u] immediately following [d]. As a result, the rule i → u / d ___ is very productive in Kashaya, and Pomoan in general, and according to Buckley, treated equivalently with other, more natural phonological patterns in the language. Just to make it clear, the reason why high vowel backing/rounding in Kashaya is unnatural is because coronals tend to front adjacent vowels in other languages. For instance, Flemming (2001) shows that coronals trigger fronting of /u/ to [y] in Cantonese, and given that there is no known retroflexion or articulatory property of /d/ in Kashaya that might explain backing and rounding of /i/ to [u], we must conclude that the context does not provide the phonetic motivation for the rule (see Hume 1996; Flemming 2003 for more on coronal consonants and front vowels). In short, the change affected is inconsistent with the trigger context, making the rule unnatural by the definition above.

As a final example, palatalization in Ojibwa represents a kind of unnatural pattern derivable from independent diachronic changes. As detailed in (Piggott 1980), dialects of Ojibwa exhibited a historical pattern of palatalization before the high front vowels, *i, *iː, and *y. Specifically, *t and *θ alternated with *ʧ and *ʃ in the context of the high front vowels. However, *t and *θ merged with *l and *n, producing the contemporary alternation between /n/ and /ʃ/. There is no straightforward phonetic explanation available to the child learning Ojibwa why the coronal nasal is realized as [ʃ] before a high front vowel, making the alternation unnatural. Moreover, since word-final short vowels are deleted in the contemporary language, there are instances where /n/ is realized as [ʃ] even without the conditioning high vowel, furthering the unnaturalness of the pattern. Finally, when *l merged with *n, the reflexes of *n did not begin to undergo the pattern, so [ni] sequences are attested so long as they derive from *n, but are illicit if they derive from *t, *θ, or *l. Thus, both the triggering context and the targeted segments do not conform to any reasonable definition of natural, and in turn, represent the capacity for phonological patterns to be unnatural.

If naturalness is thus used as a diagnostic, it behaves much like gradience for distinguishing between the phonological and phonetic. If a pattern is categorical, it is likely phonological, but if it is gradient, one cannot be sure, as I have demonstrated from backness harmony in Turkic. To present one example of gradience and categoricality in this regard, consider palatalization in English. Examples are presented in (83) below. Within words, pairs like *press*-*pressure*, and *confess*-*confession* abound. When examining the phonetic consequences of this pattern, Zsiga (1995) finds that [ʃ] in words like *pressure* and *confession* are acoustically and articulatory indistinguishable from underlying /ʃ/, as in *mesh* or *trash*. In this regard, word-internal paltalization in English is categorical, and as a result, phonological.

(83) Word-internal palatalization in English

	Underived	Derived
a.	press [pɹ̥ɛs]	pressure [pɹ̥ɛʃʊɹ]
b.	confess [kəɱfɛs]	confession [kəɱfɛʃn̩]

Zsiga (1995) also investigates phrasal palatalization in English, finding that across word boundaries, palatalization does not yield categorical assimilation to the voiceless alveo-palatal fricative. Instead, in phrases like *press your luck* and *confess your crimes*, the final consonants of the phrase-initial words are produced with articulatory characteristics in between canonical [s] and [ʃ] in English. Thus, Zsiga (1995) argues that phrasal palatalization is gradient while word-internal palatalization is categorical. As a result, word-internal palatalization is phonological, while the status of phrasal palatalization is not immediately discernable from this diagnostic. In sum, if a pattern is either categorical or unnatural, then it is likely phonological. If a pattern is neither, then a different diagnostic is necessary to adequately assess its status in the grammar of a given language.

7.2.3 Iterativity and unboundedness

In addition to gradience, naturalness, and the other potential criteria discussed above, one other defining feature of phonetic patterns is their computational complexity. In recent work on computational phonology, the regular region of the Chomsky Hierarchy has been divided into a number of nested regions, each exhibiting different computational properties (e.g. Heinz 2018; Chandlee & Heinz 2018). The innermost region of the subregular region of the Chomsky hierarchy is input strictly local. Input strictly local patterns can be modeled as a transduction from an input string to an output string in either direction with very limited memory. So, for instance, an input strictly local mapping can account for word-final devoicing if reading the input tape from right-to-left or left-to-right. From right to-left, the input tape is /#d/ and the output tape is [#t]. The transducer encounters a word boundary symbol, and after so doing, the next symbol must be voiceless if it is an obstruent. If reading the opposite direction, from left-to-right, the transducer reads an input string /d#/ and just like the other transducer, outputs [t#]. The left-to-right transducer, after encountering a voiced obstruent, needs to know if the next symbol is another segment or the word boundary symbol. If the next symbol is a segment, the underlying /d/ and the following segment are output faithfully, but if the next symbol is a word boundary symbol, /d/ is mapped to [t] and the boundary symbol is output faithfully. The ability to account for input strictly local patterns from either direction, in reality, depends on a fundamental property of input strictly local patterns- they are bounded. In other words, the trigger and target for the pattern are never too far from each other.

This generalization about input strictly local patterns lines up nicely with patterns of phonetic coarticulation. For instance, Öhman (1966) finds that vowels may affect one another across an intervening consonant, and Magen (1997) demonstrates the vowel-to-vowel coarticulation in English may span an intervening schwa. Crucially, though, the extent of a vowel coarticulatory force does not extend indefinitely throughout the word. Instead, the range of coarticulation is fixed, which is consistent with input strictly local functions. Furthermore, consonantal patterns of coarticulation are the same in this regard, extending over a relatively fixed domain. A number of studies have investigated the extent of liquid coarticulation in British as well as North American English (Kelly & Local 1986; Tunley 1999; West 1999, 2000; Hall et al. 2017). These studies report that coarticulatory effects on the second and third formants extend up to two to three syllables in both directions from the liquid. Although the range of effects is relatively large, it is still bounded. For instance, one does not expect /r/ to affect the realization of a segment separated by ten or eleven syllables. For this reason the range of phonetic coarticulation is bounded, unlike iterative phonological patterns. Consider the words below in (84), and note that even though they are very long, harmony affects all vowels within the word. Even though harmony is gradient in Kazakh, the initial syllable still controls the quality of each subsequent vowel and dorsal obstruent in the word. Surely, /ɑ/ and /ɯ/ later in these two words surface with more fronted tongue positions than their initial-syllable counterparts, but this is distinct from the iterativity of this pattern. No studies have examined the realization of vowels in words this long in Kazakh, but I predict that final-syllable vowels will still exhibit some degree of backness linking them with the initial syllable. The initial syllable iteratively spreads its feature for [back] to the rest of the syllables in the word, unlike coarticulation, wherein a segment may non-iteratively affect segments within a relatively fixed domain.

(84) Word / Gloss

a. qɑnɑʁɑt-tɑn-dɯr-ɯl-mɑ-ʁɑn-dɯq-tɑr-ɯŋɯz-dɑn
'satisfaction-VRB-CAUS-PASS-NEG-PFV-NMZLR-PL-2P-ABL'
"because you (pl) were not satisfied"

b. tɑldɯqorʁɑn-dɑ-ʁɯ-lɑr-ɯmɯz-dɯŋ
'Taldykorgan-LOC-REL-PL-POSS.1P-GEN'
"ours who live in Taldykorgan's"

This distinction between iterative and non-iterative also provides a way to differentiate phonological from phonetic patterns. I suggest that phonetic patterns are all input strictly local, while phonological patterns may be more expressive. Although Chandlee & Heinz (2018) note that roughly 95% of phonological patterns in PBase (Mielke 2018) are input strictly local, a number of authors have shown that phonological patterns like vowel harmony are more expressive (Heinz & Lai 2013; Jardine 2016). The increased expressivity of patterns like vowel harmony boils down to iterativity. Thus, if a sound pattern in question is iterative, then it is phonological.

In addition to iterativity, if a pattern exhibits a dependency that is not definable within a fixed segmental window, it is also phonological. One general type of pattern exemplifying long-distance dependencies is consonant harmony (Rose & Walker 2004; Hansson 2010; McMullin 2016). For example, in a sibilant harmony like in Aari words with sibilants that disagree in anteriority are illicit, e.g. *…ʃ…s... (Hayward 1990). When a two sibilants occur within a word, the first determines the realization of the second, regardless of the distance between them, so /…ʃ…s…/ surfaces as […ʃ…ʃ…]. Even when consonant harmony patterns are not iterative, if the dependency between trigger and target is unbounded, the pattern is more computationally complex than the input strictly local functions. In some cases, a pattern may exhibit both iterativity and long-distance dependencies. In Tutrugbu, the root spreads [+ATR] iteratively to prefixes. However, when a high vowel occurs in the initial syllable and is followed by a low-vowel prefix at any distance, harmony is blocked by the low-vowel prefix (see McCollum et al. 2019; McCollum & Essegbey 2018 for more on the Tutrugbu pattern). Most important for our purposes, though, is the distinction between patterns that do not exhibit iterativity or unbounded dependencies, and those that do. If a pattern exhibits either iterativity or an unbounded dependency, it is definitely phonological. On the other hand, if it exhibits neither characteristic, it may be phonological or phonetic.

The diagnostics discussed above can be summed up thusly. A sound pattern is phonological if it is at least one of the following: categorical, unnatural, iterative, or exhibits an unbounded dependency. If the relevant pattern exhibits at least one of those traits, then we can confidently classify it as phonological. If the pattern under consideration exhibits none of those properties, then I suggest that it is phonetic. In addition to these four diagnostics, a phonological pattern should be perceivable. As argued in Chapter 5, if a listener cannot perceive a pattern of variation, then it seems unlikely that they can manipulate that in their phonological system. Perceptibility also offers a way to further develop the formal analysis of gradience, which I discuss below in Section 6.4.2. I realize that these diagnostics will not adequately demarcate all questionable patterns, but given that the consensus diagnostic has been categoricality~gradience only, the addition of these other diagnostics provides additional insight into the phonology-phonetics interface. These diagnostics bring together views from diachronic and evoluationary perspectives with insights from computational phonology to further our understanding of what is phonetic and what is phonological.

7.2.4 Predictions of a unified phonology and phonetics

As Tucker & Warner (2010:218) note, many within the laboratory phonology community have questioned the putative distinction between phonology and phonetics, suggesting instead that the two exist on a continuum(e.g. Pierrehumbert et al. 2000). Whether viewed this way, or in entirely conflationist terms (e.g. Flemming 2001), these predict that sound patterns should likely fall into a single distribution within and across languages. To spell this out more clearly, given some continuum, be it gradience or some other factor, if phonology and phonetics are not distinct, then they should form a unimodal distribution along the relevant continuum. Gradience serves as a nice example for demonstration. If, as suggested by Pierrehumbert and colleagues, phonological patterns tend to be coarse-grained to categorical, while phonetic patterns tend to be finer-grained, this suggests a single distribution of sound patterns over the spectrum of categorical-to-fine-grained information, as shown in Figure 4 below. In the leftmost plot, sounds patterns are normally distributed across the continuum, with the largest number of patterns exhibiting intermediate degrees of categoricality/gradience. In the middle plot, the distribution is highly skewed toward more gradient patterns, and in the rightmost plot, there is almost no distinction in the number of patterns by categoricality~gradience. The real question is thus: what is the empirical distribution of patterns along a continuum like categorical~gradient? It is not clear at present, although current results suggest that distributions like the left- and rightmost plots below are incorrect. If the vast majority of sound patterns are distributed as in these plots, it is unclear why phonetics and phonology would've developed separately. Also, the kind of distribution in the leftmost plot predicts that near-categorical patterns, like incomplete neutralizations should be relatively few, but laboratory results suggest that they are, in fact, quite common. These incomplete neutralizations are almost categorical, and can only be accounted for under such a unified model if sound patterns are skewed toward the categorical, like the mirror image of the middle plot. This seems equally unlikely, as the number of language-specific low-level phonetic patterns would be remarkably small under this view of the world, and that is clearly inconsistent with the large body of literature carefully describing the phonetic patterns of the world's languages.

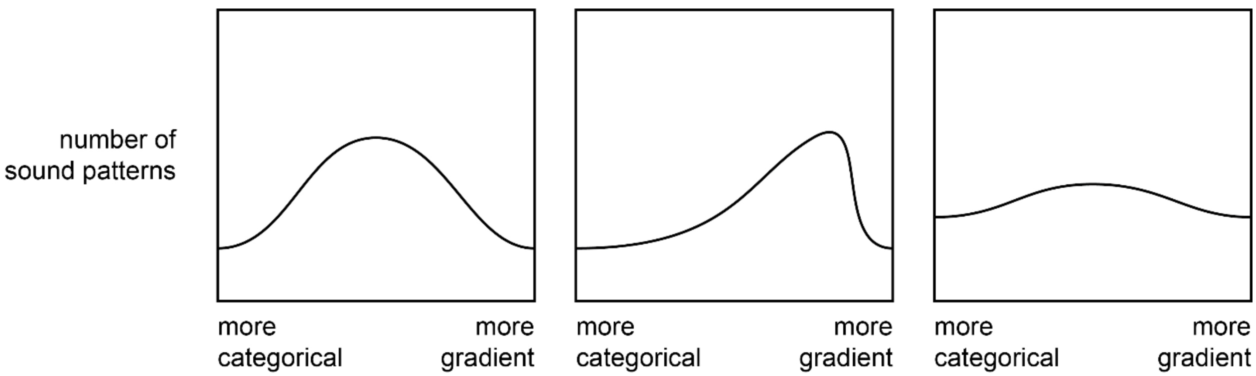

Figure 7.4: Potential distributions of sound patterns under a unified view of phonology and phonetics

In contrast to a unified view of phonology and phonetics, a modular view predicts a bimodal distribution of sound patterns, like those in Figure 7.5. In these plots, phonological and phonetic patterns are largely but not entirely distinct along the categorical~gradient continuum. Phonological patterns tend to be more categorical while phonetic patterns tend to be more fine-grained and gradient. Regardless of the particular skews found in the distributions below, in each plot there is a clear distributional distinction between sound patterns that are phonological and phonetic. Yet, in each there are cases where the evidence may be ambiguous and it may be difficult to determine whether the pattern under study derives from phonology or phonetics. As before, it should be noted that we don't yet know enough to really pick one distribution over another. That being said, the history of Structuralist and Generative research that has led us to our current understanding is predicated on such a view. More significantly in my estimation is the intuitive bimodality of the distribution of sound patterns. A unimodal distribution might predict that the fieldworker spends most of their time trying to ascertain is a given alternation is phonetic or not. In my experience, relatively few alternations fall into this category, and far more are judged to be decidedly phonetic or phonological. While one upshot of the laboratory phonology movement is a heightened distrust for impressionistic judgments, the fact that phonological patterns are constrained by human perceptual faculties (even if they have vastly different experiences informing them) suggests that aural impressions are not to be entirely discarded.

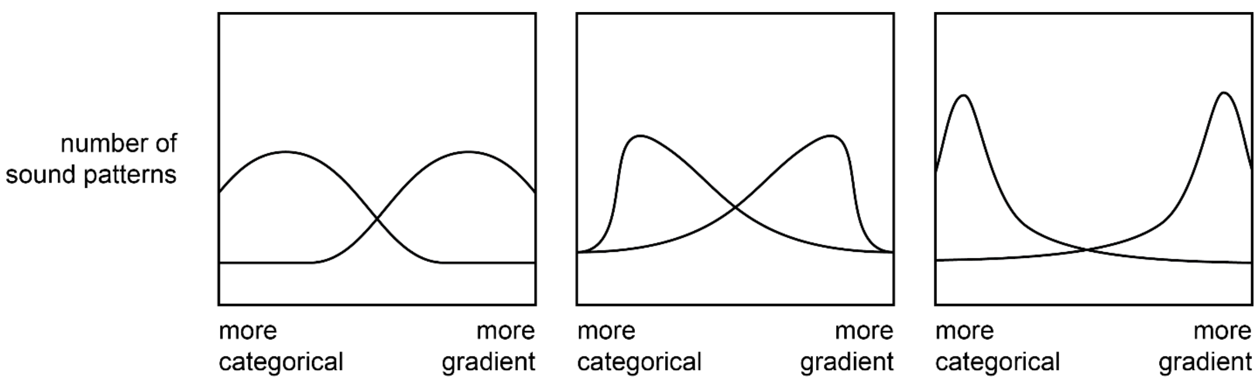

Figure 7.5: Potential distributions of sound patterns under a modular view of phonology and phonetics

A second prediction of the increasingly popular unified view of phonology and phonetics is that there exist only two levels of sound representations, underlying and surface. Under this view, only underlying representations, with whatever representational content they possess, and surface forms with all their attendant details are objects for research. This differs from the modular view, which allows for a third level of representation. Within a modular view, there is an input to phonology, the underlying representation. There is also a level for the output of phonology that is distinct from the surface phonetic outputs.

7.3 Future directions

7.3.1 Empirical directions

Looking forward, several distinct directions exist to develop this line of research, ultimately aimed at understanding the role of gradience in phonology. First, from an empirical perspective, the number and variety of languages described as gradient earlier in the chapter suggests that the same methodologies applied to Turkic can and should be applied to other languages to assess the reported gradience in these, and perhaps other, languages. For instance, all the languages examined herein exhibited progressive harmony, so investigating a language, like Rangi, with regressive harmony would be a natural next step. Although extant descriptions report gradience in both directions, it may be the case that progressive and regressive harmonies produce different patterns of gradience. One recent paper, Olejarczuk et al. (2019), examines the phonetic properties of ATR harmony in Komo, a language with bidirectional ATR harmony. They find that the acoustic properties of harmony are largely similar, with slight differences based on duration and spectral slope for leftward and rightward harmony. Directionality has been linked to differences in coarticulation (e.g. Beddor & Yavuz 1995), and further work may unearth differences at the phonological level, as well. Additionally, it would be fruitful to investigate gradient propagation of other features. Turkic exhibits backness, and to a lesser degree, rounding harmony, but a number of ATR harmony languages have been reported to exhibit gradience, as well. One recurrent issue in ATR harmony is defining the acoustic correlates of the harmonic feature (Lindau et al. 1972; Fulop et al. 1998; Guion et al. 2004; Olejarczuk, Otero et al. 2019), and it would be beneficial to investigate the realization of the harmonic feature both within triggers, and the potentially gradient realization of the tongue root contrast across target positions to better understand the nature of the feature, ATR. In addition to examining languages with differing directionality and different harmonic features, investigating gradience among consonant and vowel-consonant harmonies provides a way to further unite (or alternatively, further distinguish) these feature-spreading patterns from one another. These different harmonies have been shown to exhibit some notable differences (Hansson 2010; Rose &

Walker 2011). For instance, consonant harmonies rarely exhibit blocking effects, while blocking is common in vowel harmony. Consonant harmonies are also more likely to exhibit long-distance alternations, while vowel harmonies typically exhibit dependencies across a single syllable boundary.

This dissertation has used acoustic information as a proxy for articulatory information. It is well known that acoustic properties are not always congruent with articulatory facts, and for that reason it would be worthwhile to extend research on gradience to incorporate articulatory methods, as well. Ultrasound, for example, provides a relatively non-invasive means by which to examine the articulatory postures associated with harmony and how those postures differ by position, which may provide further insight into the Turkic patterns discussed above, as well as the other harmonies listed at the beginning of this chapter. For instance, I have assumed throughout that the harmonic feature in Turkic is [back]. There are, however, a few authors who have argued that Turkic actually spreads the feature [ATR] (Vajda 1994; Washington 2016). If both features co-vary by harmonic class, this may not matter greatly. That being said, comparing positional differences in tongue body backness and tongue root position may help to assess which feature is actually being manipulated in Turkic.

Thus, the empirical opportunities just discussed offer ways to further understand the distribution of gradient harmony patterns as well as the articulatory nature of gradient harmony. Outside of harmony, using gradience as a window to re-evaluate the analysis of incomplete neutralization, as well as in vowel excrescence and vowel epenthesis are additional opportunities to assess the role of gradience in human phonology (Braver 2013; Hall 2006).

7.3.2 Formal directions

In addition to the empirical prospects for developing a larger theory of gradience in phonology, the formal opportunities to refine and develop the analysis presented here are equally significant. The biggest challenging facing the analysis presented in the previous chapter is how to restrict the formal analysis. While the analysis is restrictive in that it excludes contrast-enhancing gradience, the formalism itself cannot determine what gradience is meaningful to the phonology of a language. For instance, if a language exhibits progressive backness harmony and produces ⟦back⟧ values like ⟦1⟧-⟦0.99⟧-⟦0.98⟧-⟦0.97⟧, the formalism cannot distinguish this, presumably phonetic gradience from true phonological gradience. In Chapter 5 I argued that phonological gradience should produce effects large enough to be perceived by native speakers. If this is the case, and if values like ⟦1⟧-⟦0.99⟧-⟦0.98⟧-⟦0.97⟧ are not distinguishable by native speakers, then this type of low-level phonetic gradience should be excluded from the formalism. One intuitive way to accomplish this is to adopt a discretized representational system, like Flemming (1995), which is shown below in (85). In Flemming's grid space, vowels are differentiated from one another by at least one horizontal (F2) or vertical (F1) grid space.

(85) Flemming's (1995:30) representation of vowel height and backness

F2

5	4	3	2	1		
i	y	ɨ	ɯ	u	1	
ɪ	ʏ			ʊ	2	F1
e	ø	ə	ɤ	o	3	
	ɛ	œ	ʌ	ɔ	4	
		a	ɑ		5	

Flemming's grid is an idealized IPA-based representation of the range of the human F1-F2 space. It provides a first approximation of what vowel categories humans use in language, but it does not provide the detail of representation necessary to account for the gradience found in Kazakh and Uyghur. To see this, if we replace Flemming's /a/ with Uyghur /æ/, there is no space in between back /ɑ/ and front /æ/, and so this representational schema cannot account for the range of low vowel F2 in third, fourth, and fifth syllables, which can't be neatly categorized as /ɑ/ or /æ/. To modify this, one could simply expand the range of F1 and F2 grid spaces to provide more space between phonemic categories, and thus enable the representational system to account for subphonemic gradience. In this sample representation (86), the amount of subphonemic variation differs across alternating pairs. Five grid spaces intervene between /y/ and /u/, four intervene between /i/ and /ɯ/, while only three spaces separate /æ/ and /ɑ/. These differences reflect the relationship between tongue backness and jaw lowering. As jaw opening increases, the range of front-back differences diminish, thus the relatively small number of distinct low vowel categories attested in human language. If this relationship extends to subphonemic variation, then one might expect differential patterns based on vowel height. There is some evidence from Chapter 4 that high vowels behave slightly differently than low vowels in Kazakh and Uyghur. In both languages, the low vowels did not shift as much in the relativized vowel space as the high vowels. High back vowels, /ɯ/ and /u/ almost approximate the spectral properties for their front vowel counterparts in later syllables, but /ɑ/ never approximates /æ/ or /ie/. In fact, this result suggests that the low vowels could be separated by a larger number of grid spaces to reflect a potential resistance to subphonemic variation.

(86) Flemming's representational system modified for the three Uyghur backness alternations. Darker cells represent higher ⟦back⟧ values and lighter cells represent lower ⟦back⟧ values.

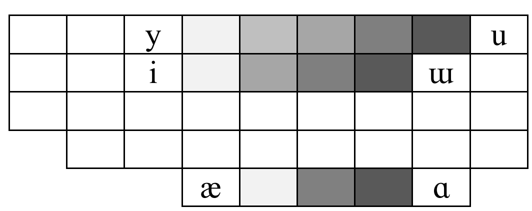

Another possible modification to Flemming's system is to eschew articulatory differences based on jaw opening, and uniformly distribute grid spaces between harmonic counterparts, as in (87). In this vowel space, which is more similar to the representational system used in Chapter 6, there is not inter-category difference across harmonic alternations. The same number of grid spaces intervene between each alternating pair.

(87) A second variant of Flemming's representational system modified for the three Uyghur backness alternations. Darker cells represent higher ⟦back⟧ values and lighter cells represent lower ⟦back⟧ values.

y						u
i						ɯ
æ						ɑ

At present, I cannot argue for one representational system over another, but there are some avenues through which to determine the relatives merits of the various possible discretized representational schemas. First and most intuitively, perception studies should inform the nature of the representational system. If a Uyghur cannot perceive five vowel qualities between /æ/ and /ɑ/, then perhaps the representation in (87) is overly expressive. Another way to address this is to determine what just-noticeable-differences exist for each phonemic category and use these to determine the distance between two alternating pairs. For instance, if the small perceivable difference between some variant of /æ/ is 100 Hz in the F2 dimension, and roughly the same threshold exists for /ɑ/, then we could construct a representation with this information. Given these just-noticeable-differences, if the difference in F2 for these two contrastive vowels is 500 Hz in initial syllables, this could support a representation with 5 subphonemic grid spaces, as in (87). These details are significant, but largely unknown at present. To better determine how to structure the phonological representation of gradience, further studies on vowel perception are necessary.

In addition to developing an adequate representational system that connects production patterns with perceptual capabilities, developing the kind of discretized representations shown in (86-87) offers a very significant formal benefit to the analysis. By discretizing the vowel space, the hypothesis space is significantly constrained. Consider the current analysis- for the feature [back], an infinite number of real-valued numbers exists between each category, defined as 0 and 1 above. If this infinitude of possibilities is restricted to 5, 6, or 7, this would drastically reduce the hypothesis space for the learner, and make the analysis more formally tractable. At present, the speaker/learner must keep up with all real-valued numbers between 0 and 1, but in a discretized vowel space, this is reduced to (in all likelihood) single digits.

Returning to one of the issues raised above, developing a discretized representational system should rule out pseudo-gradient patterns like ⟦1⟧-⟦0.99⟧-⟦0.98⟧-⟦0.97⟧. If five grid spaces intervene between two categorical production values, this entails a difference of approximately 0.14 ⟦back⟧ between each grid space. As a result, one might predict that all values between ⟦1⟧ and ⟦0.86⟧ should fall within the same grid space. In turn, a pattern like ⟦1⟧-⟦0.99⟧-⟦0.98⟧-⟦0.97⟧ would be interpreted as categorical in this representational system without any additional dependence on the analyst. In Chapters 4 and 6, I drew on phonetic and phonological literature to determine which of the three harmonies are gradient and which are categorical. If the formalism is able to remove the analyst's subjectivity from this stage of the analysis, even if only to a small degree, this is likely a benefit to the larger theory. For expedience, I only considered backness values rounded to the nearest tenth. This move was purely to reduce the size of the spreadsheets used. Further reducing the number of available backness categories for the model to predict will decrease model fit, but this is not necessarily an unfortunate outcome. Reducing model fit by reducing the representational capacity of the formalism allows for phonetics to account for variation over and above that found in the formal analysis. It is thus an issue of division of labor. In the current analysis, some phonetic facts are likely captured under the formal analysis, but with less powerful representations, the amount of phonetic detail capture in the phonological analysis would be reduced, which would be advantageous, since the goal is to account for phonological patterns of gradience and not phonetic patterns of reduction.

In addition to the representation of gradience, the constraint-based computation of gradience also deserves further attention. Specifically, the BECAT constraints in the analysis in Chapter 6 support the idea that phonological patterns are preferentially categorical. Yet, the fact that these constraints exist in a grammatical architecture in HG allows for gradience due to higher ranked constraints. If these BECAT constraints were lowly ranked and a set of faithfulness constraints were highly ranked, then some input with gradient feature values (e.g. [back] values like /⟦0.5⟧-⟦0.3⟧-⟦0.8⟧/) would be predicted to surface faithfully. This scenario seems highly unlikely in natural language, which may suggest that some other device is responsible for promoting categoricality. In Gradient Symbolic Computation, a quantization parameter forces gradient featural activations to surface as entirely or almost categorical (Smolensky et al. 2013). As discussed in Mai et al. (2018), quantization is still not well understood, but further development of this formalism could provide a more satisfactory account of preferential categorcality in phonology.

7.4 Conclusion

This chapter has marshalled evidence from a variety of languages that suggests that gradient harmony patterns may occur across a range of language families, and further, that these harmony may manipulate a number of different features. I proposed that, despite the number of descriptions included in this chapter, that gradience in harmony is probably underreported. How gradience is manifested across languages, across the different features active in harmony, and in the differing directions in which harmony applies is a question that may influence how gradience is conceived, formalized, and contextualized in phonology and phonetics.

In addition to the typology of gradient harmony, the chapter discusses a number of the proposed distinctions between phonology and phonetics. I agree generally with many of these distinctions, but advocate that they be used in concert to provide a more nuanced perspective on phonological and phonetic patterns. In addition to existing proposals, I added an additional way to distinguish phonological from phonetic, unboundedness. Phonological patterns may be unbounded, while phonetic patterns (at least segmental phonetic patterns) are bounded. As a result, phonological patterns may exhibit increased complexity, relative to phonetic patterns, which I suggest occupy the least expressive class of subregular languages.

Finally, I have laid out some possible directions for future work, including both the empirical investigation of harmony as well as its formal analysis. Combining these two provides a variety of fruitful possibilities. Formal questions may drive the empirical, and empirical results can inform theory. The claims advanced in this as well as previous chapters demands a more definitive answer, drawing on data from articulatory phonetics, psycholinguistics, typology, and formal phonology.

Chapter 8: Conclusion

I have used experimental fieldwork on Turkic to address two phonological questions. One, are phonological patterns local? And two, is phonology by definition categorical, or may it exhibit gradience? To the first question, results are mixed. Evidence from Uyghur indicates that high vowels that have been transcribed as /i/ and analyzed as transparent to backness harmony do, in fact, alternate for harmony, as discussed in Chapter 2. The upshot of this finding is that backness harmony in Uyghur is decidedly more local than previous accounts suggest. Moving from locality among vowels, Chapter 3 examines the realization of laterals and coronal fricatives to determine whether backness and rounding harmony are strictly local in Turkic, triggering subphonemic alternations on all segments, both vowels and consonants. Results indicate that although consonants typically undergo low-level alternations consistent with harmony, that their alternation is not universal, as predicted by advocates of strict locality (e.g. Gafos 1999) and much areal work (e.g. Dzhunisbekov 1980, 1991; see also Johanson 1991; Csató 1999). Instead, results indicate that the intrinsic articulatory properties of intervening consonants factor significantly in their (non-)participation in harmony. Future work will need to determine which consonantal segments alternate for harmony, which do not, and how to adequately analyze the set of (non) alternating consonants. In the preliminary analysis posited here, variation among the sibilants and lateral is attributed to phonetic implementation, and as such, these consonants are not interpreted as phonological targets for harmony.

In Chapters 4-6, the dissertation examines the second main question, is phonology categorical, or may it exhibit gradience, arguing that backness harmony in Uyghur and Kazakh gradiently peters out throughout the word. In contrast, backness harmony in Kyrgyz is categorical, with phonetic centralization producing symmetrical reduction of contrast in non-initial syllables. Chapter 5 demonstrates that the acoustic effects of gradient [+back] spreading in Uyghur and Kazakh are perceivable, proposing that perceptibility is a fundamental requirement of phonological phenomena. Chapter 6 formalizes the analysis of gradient and categorical harmony within a Harmonic Grammar by introducing gradient representations into the grammar. In so doing, the analysis is able to account for gradient harmony in Kazakh and Uyghur as well as categorical harmony in Kyrgyz with the same set of representations and constraints. Finally, Chapter 7 draws on evidence from elsewhere among the world's languages to argue that gradience is not a corner issue in the phonology of Turkic, but rather a characteristic of a variety of morphophonological patterns from languages around the world. This chapter also discusses diagnostics of phonological and phonetic patterns, introducing computational complexity as an additional distinction between the two.

Looking forward, this thesis has demonstrated a means by which to experimentally and formally evaluate the nature of phonological patterns, their locality and their putative categoricality. The methods used in this dissertation extend naturally to other questions, and the findings here suggest fundamental changes to how we view phonology. Most importantly, the finding that morphophonology may be gradient provides a question to be further investigated to more fully understand the nature of phonological computation and the role of gradience in human language.

References

Abdurehim, Esmael. 2014. *The Lopnor Dialect of Uyghur: A Descriptive Analysis*. Publications of the Institute for Asian and African Studies 17. University of Helsinki.

Abuov, Zhoumagaly. 1994. The Phonetics of Kazakh and the Theory of Synharmonism. *UCLA Working Papers in Phonetics* 88: 39–54.

Alarifi, Abdulaziz, and Benjamin V. Tucker. 2016. An Acoustic Investigation of the Categorical and Gradient Spread of Pharyngealization in Urban Najdi Arabic. presented at LabPhon 2016, Cornell University.

Ali, Latif Hassan, and Raymond G. Daniloff. 1972. A Contrastive Cinefluorographic Investigation of the Articulation of Emphatic-Non Emphatic Cognate Consonants. *Studia Linguistica* 26.2: 81–105.

Allen, Blake, Douglas Pulleyblank, and Ọládiípọ̀ Ajíbóyè. 2013. Articulatory Mapping of Yoruba Vowels: An Ultrasound Study. *Phonology* 30.2: 183–210.

Amanzholov, Sarsen. 1959. *Voprosy Dialektologii i Istorii Kazakhskogo Jazyka*. Alma-Ata: National Instructional Institute in the name of Abai.

Applegate, Richard. 1972. Ineseño Chumash Grammar. PhD thesis, University of California, Berkeley.

Archangeli, Diana Bennett. 1984. Underspecification in Yawelmani Phonology and Morphology. PhD thesis, Massachusetts Institute of Technology.

Archangeli, Diana, and Douglas Pulleyblank. 1994. *Grounded Phonology*. MIT Press.

Archangeli, Diana, and Douglas Pulleyblank. 2002. Kinande Vowel Harmony: Domains, Grounded Conditions and One-Sided Alignment. *Phonology* 19.2: 139–188.

Archangeli, Diana, and Douglas Pulleyblank. 2007. Harmony. *The Cambridge Handbook of Phonology*, ed. by Paul de Lacy, 353–378.

Avcu, Enes. 2018. Experimental Investigation of the Subregular Hypothesis. In *Proceedings of the 35th West Coast Conference on Formal Linguistics*, 77–86.

Bach, Emmon, and Robert T Harms. 1972. How Do Languages Get Crazy Rules. *Linguistic Change and Generative Theory* 1-21

Baković, Eric. 2000. Harmony, Dominance, and Control. PhD thesis, Rutgers University.

Balakayev, M.B. 1962. *Sovremennij Kazakhskij Jazyk*. Alma-Ata: Nauka.

Bates, Douglas, Martin Machler, Ben Bolker, and Steve Walker. 2015. Fitting Linear Mixed-Effects Models Using Lme4. *Journal of Statistical Software* 67.1: 1–48.

Batmanov, I.A. 1938. *Severnye Dialekty Kirgizskogo Jazyka*. Frunze: Kirgizgosizdat.

Batmanov, L.A. 1940. *Foneticheskaya Sistema Sovremennogo Kirgiskogo Jazyka*. Nauka.

Beckman, Jill N. 1997. Positional Faithfulness, Positional Neutralisation and Shona Vowel Harmony. *Phonology* 14.1: 1–46.

Beckman, Jill N. 2013. *Positional Faithfulness: An Optimality Theoretic Treatment of Phonological Asymmetries*. Routledge.

Beddor, Patrice, and Handan Yavuz. 1995. The Relationship between Vowel-to-Vowel Coarticulation and Vowel Harmony in Turkish. In *Proceedings of the 13th International Congress of Phonetic Sciences*, 2: 44–51.

Beguš, Gašper. 2018. Unnatural Phonology: A Synchrony-Diachrony Interface Approach. PhD thesis, Harvard University.

Beguš, Gašper, and Alexei Nazarov. 2019. Gradient Trends against Phonetic Naturalness: The Case of Tarma Quechua. In *Proceedings of the 48th Meeting of the Northeast Linguistics Society*.

Bennett, Ryan, Boris Harizanov, and Robert Henderson. 2018. Prosodic Smothering in Macedonian and Kaqchikel. *Linguistic Inquiry* 49.2: 195–246.

Bennett, Wm G. 2015. *The Phonology of Consonants: Harmony, Dissimilation and Correspondence*. Vol. 147. Cambridge University Press.

Benus, Stefan, and Adamantios I. Gafos. 2007. Articulatory Characteristics of Hungarian 'transparent' Vowels. *Journal of Phonetics* 35.3: 271–300.

Berry, Jack. 1957. Vowel Harmony in Twi. *Bulletin of the School of Oriental and African Studies* 19.1: 124–130.

Berta, Árpád. 1998. West Kipchak Languages. In *The Turkic Languages*, 301–317.

Bessell, Nicola J. 1998. Local and Non-Local Consonant–Vowel Interaction in Interior Salish. *Phonology* 15.1: 1–40.

Binnick, Robert I. 1991. Vowel Harmony Loss in Uralic and Altaic. In *Studies in the Historical Phonology of Asian Languages*, edited by William G. Boltz and Michael C. Shapiro, 35–52. Current Issues in Linguistic Theory 77. John Benjamins.

Bird, Steven. 1995. *Computational Phonology: A Constraint-Based Approach*. Cambridge University Press.

Blevins, Juliette. 2004. *Evolutionary Phonology: The Emergence of Sound Patterns*. Cambridge University Press.

Blevins, Juliette. 2008. Natural and Unnatural Sound Patterns A Pocket Field Guide. *Naturalness and Iconicity in Language* 7: 121.

Bloomfield, Leonard. 1933. *Language*. Holt, Rinehart and Winston.

Bodrogligeti, András J.E. 2001. *A Grammar of Chagatay*. Languages of the Word 155. Lincom Europa.

Boersma, Paul, and David Weenink. 2015. *Praat: Doing Phonetics by Computer* (version 5.4.18).

Booij, Geert E. 1984. Neutral Vowels and the Autosegmental Analysis of Hungarian Vowel Harmony. *Linguistics* 22.5: 629–642.

Borroff, Marienne. 2007. A landmark underspecification account of the patterning of glottal stop. PhD dissertation, Stony Brook University.

Boyce, Suzanne E. 1990. Coarticulatory Organization for Lip Rounding in Turkish and English. *The Journal of the Acoustical Society of America* 88.6: 2584–2595.

Braver, Aaron. 2013. Degrees of Incompleteness in Neutralization: Paradigm Uniformity in a Phonetics with Weighted Constraints. PhD thesis, Rutgers University.

Braver, Aaron. 2019. Modeling Incomplete Neutralization with Weighted Phonetic Constraints. *Phonology* 36.

Buckley, Eugene. 2000. On the Naturalness of Unnatural Rules. In *Proceedings from the Second Workshop on American Indigenous Languages. UCSB Working Papers in Linguistics* 9: 1–14.

Butcher, Andrew. 1999. What Speakers of Australian Aboriginal Languages Do with Their Velums and Why: The Phonetics of the Nasal/Oral Contrast. In *Proceedings of the International Congress of Phonetic Sciences 1999*, 479–482.

Calabrese, Andrea. 1995. A Constraint-Based Theory of Phonological Markedness and Simplification Procedures. *Linguistic Inquiry*, 373–463.

Carter, Paul, and John Local. 2007. F2 Variation in Newcastle and Leeds English Liquid Systems. *Journal of the International Phonetic Association* 37.2: 183–199.

Casali, Roderic F. 2002. Nawuri ATR Harmony in Typological Perspective. *Journal of West African Languages* 29.1: 3–43.

Casali, Roderic F. 2003. ATR] Value Asymmetries and Underlying Vowel Inventory Structure in Niger-Congo and Nilo-Saharan. *Linguistic Typology* 7.3: 307–382.

Casali, Roderic F. 2008. ATR Harmony in African Languages. *Language and Linguistics Compass* 2.3: 496–549.

Casali, Roderic F. 2012. [+ATR] Dominance in Akan. *Journal of West African Languages* 39.1: 33–59.

Casali, Roderic F. 2016. Some Inventory-Related Asymetries in the Patterning of Tongue Root Harmony Systems. *Studies in African Linguistics* 45 (1 & 2): 96–99.

Chandlee, Jane, and Jeffrey Heinz. 2018. Strict Locality and Phonological Maps. *Linguistic Inquiry* 49.1: 23–60.

Charles-Luce, Jan, and Daniel A Dinnsen. 1987. A Reanalysis of Catalan Devoicing. *Journal of Phonetics* 15.1: 187–190.

Chitoran, Ioana, and Abigail C. Cohn. 2009. Complexity in Phonetics and Phonology: Gradience, Categoiality, and Naturalness. In *Approaches to Phonological Complexity*, edited by Francois Pellegrino, Egidio Marsico, Ioana Chitoran, and Christophe Coupé, 21–46. de Gruyter.

Cho, Taehong. 2005. Prosodic Strengthening and Featural Enhancement: Evidence from Acoustic and Articulation Realizations of /ɑ, i/ in English. *Journal of the Acoustical Society of America* 117.6: 3867–3678.

Cho, Taehong, and Patricia Keating. 2009. Effects of Initial Position versus Prominence in English. *Journal of Phonetics* 37.4: 466–485.

Chomsky, Noam, and Morris Halle. 1968. *The Sound Pattern of English.* Harper & Row.

Clements, George N. 1981. Akan Vowel Harmony: A Non-Linear Analysis. In *Harvard Studies in Phonology 2*, 108–177.

Clements, George N. 1991. Vowel Height Assimilation in Bantu Languages. In *The Proceedings of the 17th Annual Meeting of the Berkeley Linguistics Society*, 17: 25–64.

Clements, George N. and Elizabeth V. Hume. 1995. The internal organization of speech sounds. *The Handbook of Phonological Theory*, 245-317.

Clements, George N., and Engin Sezer. 1982. Vowel and Consonant Disharmony in Turkish. *The Structure of Phonological Representations* 2: 213–255.

Cohn, Abigail C. 1993. Nasalisation in English: Phonology or Phonetics. *Phonology* 10.1: 43–81.

Cohn, Abigail C. 1998. The Phonetics-Phonology Interface Revisited: Where's Phonetics? In *Proceedings of the 1998 Texas Linguistics Society Conference*, 25–40.

Cohn, Abigail C. 2003. Phonological Structure and Phonetic Duration: The Role of the Mora. *Working Papers of the Cornell Phonetics Laboratory* 15: 69–100.

Cohn, Abigail C. 2006. Is There Gradient Phonology? In *Gradience in Grammar: Generative Perspectives*, edited by Gilbert Fanselow, Caroline Féry, Matthias Schlesewsky, and Ralf Vogel, 25–44. Oxford University Press.

Comrie, Bernard. 1981. *The Languages of the Soviet Union*. Cambridge University Press.

Crosswhite, Katherine. 2003. Formant Logger Script. http://phonetics.linguistics.ucla.edu/facilities/acoustic/formant_logging.txt.

Csató, Éva Ágnes. 1999. Analyzing Contact-Induced Phenomena in Karaim. In *The Proceedings of the 25th Annual Meeting of the Berkeley Linguistics Society*, 25: 54–62.

Dalton, Will. 2012. Deconstructing Canadian French Vowel Harmony. paper presented at the Montreal Ottowa Toronto. https://www.ling.upenn.edu/Events/PLC/plc36/abstracts/Dalton_PLC36.pdf.

Daniloff, Raymond, and Kenneth Moll. 1968. Coarticulation of Lip Rounding. *Journal of Speech and Hearing Research* 11.4: 707–721.

Davis, Stuart. 1995. Emphasis Spread in Arabic and Grounded Phonology. *Linguistic Inquiry* 26.3: 465–498.

Demolin, Didier. 2015. Insights from the Field. In *The Handbook of Speech Production*, 477–504.

Dinnsen, Daniel A. 1985. A Re-Examination of Phonological Neutralization. *Journal of Linguistics* 21.2: 265–279.

Dinnsen, Daniel A, and Jan Charles-Luce. 1984. Phonological Neutralization, Phonetic Implementation and Individual Differences. *Journal of Phonetics* 12.1: 49–60.

Dolphyne, Florence Abena. 1988. *The Akan (Twi-Fante) Language: Its Sound Systems and Tonal Structure*. Ghana Universities Press.

Donegan, Patricia J, and David Stampe. 1979. The Study of Natural Phonology. *Current Approaches to Phonological Theory* 126-173.

Dye, Amanda. 2015. Vowel Harmony and Coarticulation in Wolof and Pulaar: An Ultrasound Study. PhD thesis, New York University.

Dzhunisbekov, Alimkhan. 1980. *Singarmonizm v Kazakhskom Jazyke*. Alma-Ata: Nauka.

Dzhunisbekov, Alimkhan. 1991. *Stroi Kazakhskogo Jazyka*. Alma-Ata: Gylym.

Eckmann, János. 2017. *Chagatay Manual*. Routledge.

Eimas, Peter D, and John D Corbit. 1973. Selective Adaptation of Linguistic Feature Detectors. *Cognitive Psychology* 4.1: 99–109.

Eimas, Peter D., Einar R. Siqueland, Peter Jusczyk, and James Vigorito. 1971. Speech Perception in Infants. *Science* 171.3968: 303–306.

Engesaeth, Tarjei, Mahire Yakup, and Arienne M Dwyer. 2009. *Greetings from the Teklimakan: A Handbook of Modern Uyghur*. University of Kansas.

Ernestus, Mirjam. 2011. Gradience and Categoricality in Phonological Theory. In *The Blackwell Companion to Phonology*, 2115–2136. Wiley-Blackwell.

Esling, John H. 2005. There Are No Back Vowels: The Laryngeal Articulator Model. *Canadian Journal of Linguistics* 50.1–4: 13–44.

Farnetani, Edda, and Daniel Recasens. 1997. Coarticulation and Connected Speech Processes. In *The Handbook of Phonetic Sciences*, edited by William J. Hardcastle, John Laver, and Fiona E. Gibbon, 1: 316–352. Wiley-Blackwell.

Finley, Sara. 2008. Formal and Cognitive Constraints on Vowel Harmony. PhD thesis, Johns Hopkins University.

Finley, Sara. 2010. Exceptions in Vowel Harmony Are Local. *Lingua* 120.6: 1549–1566.

Flanagan, James L. 1955. A Difference Limen for Vowel Formant Frequency. *Journal of the Acoustical Society of America* 27.3: 613–17.

Flemming, Edward. 1995. Auditory Representations in Phonology. PhD thesis, University of California, Los Angeles.

Flemming, Edward. 2001. Scalar and Categorical Phenomena in a Unified Model of Phonetics and Phonology. *Phonology* 18.1: 7–44.

Flemming, Edward. 2003. The Relationship between Coronal Place and Vowel Backness. *Phonology* 20.3: 335–373.

Flemming, Edward, and Hyesun Cho. 2017. The Phonetic Specification of Contour Tones: Evidence from the Mandarin Rising Tone. *Phonology* 34.1: 1–40.

Fougeron, Cecile, and Patricia A. Keating. 1997. Articulatory Strengthening at Edges of Prosodic Domains. *Journal of the Acoustical Society of America* 101.6: 3728–40.

Fourakis, Marios. 1991. Tempo, Stress, and Vowel Reduction in American English. *The Journal of the Acoustical Society of America* 90.4: 1816–1827.

Fourakis, Marios, Antonis Botinis, and Maria Katsaiti. 1999. Acoustic Characteristics of Greek Vowels. *Phonetica* 56.1–2: 28–43.

Fry, Dennis B., Arthur S. Abramson, Peter D. Eimas, and Alvin M. Liberman. 1962. The Identification and Discrimination of Synthetic Vowels. *Language and Speech* 5.4: 171–189.

Fulop, Sean A., Ethelbert Kari, and Peter Ladefoged. 1998. An Acoustic Study of the Tongue Root Contrast in Degema Vowels. *Phonetica* 55.1–2: 80–98.

Fylstra, Daniel, Leon Lasdon, John Watson, and Allan Waren. 1998. Design and Use of the Microsoft Excel Solver. *Interfaces* 28.5: 29–55.

Gafos, Adamantios I. 1999. *The Articulatory Basis of Locality in Phonology*. Garland Publishing.

Gafos, Adamantios I, and Amanda Dye. 2011. Vowel Harmony: Opaque and Transparent Vowels. *The Blackwell Companion to Phonology*, 1–26.

Ghosh, Satrajit S, Melanie L Matthies, Edwin Maas, Alexandra Hanson, Mark Tiede, Lucie Ménard, Frank H Guenther, Harlan Lane, and Joseph S Perkell. 2010. An Investigation of the Relation between Sibilant Production and Somatosensory and Auditory Acuity. *The Journal of the Acoustical Society of America* 128.5: 3079–3087.

Gick, Bryan, Douglas Pulleyblank, Fiona Campbell, and Ngessimo Mutaka. 2006. Low Vowels and Transparency in Kinande Vowel Harmony. *Phonology* 23.1: 1–20.

Gick, Bryan, Ian Wilson, Karsten Koch, and Clare Cook. 2004. Language-Specific Articulatory Settings: Evidence from Inter-Utterance Rest Position. *Phonetica* 61.4: 220–233.

Gnanadesikan, Amalia Elisabeth. 1997. Phonology with Ternary Scales. PhD thesis, University of Massachusetts, Amherst.

Goldsmith, John A. 1976. Autosegmental Phonology. PhD thesis, Massachusetts Institute of Technology.

Goldsmith, John A. 1995. Phonological Theory. In *The Handbook of Phonological Theory*, ed. by John A. Goldsmith, 1–23. Wiley Blackwell.

Goldsmith, John A., and Jason Riggle. 2012. Information Theoretic Approaches to Phonological Structure: The Caase of Finnish Vowel Harmony. *Natural Language & Linguistic Theory* 30: 859–896.

Gordon, Matthew. 1999. The "neutral" vowels in Finnish: How neutral are they? *Linguistica Uralica* 35.1:17-21.

Gouskova, Maria. 2004. Relational Hierarchies in Optimality Theory: The Case of Syllable Contact. *Phonology* 21.2: 201–250.

Guion, Susan G, Mark W Post, and Doris L Payne. 2004. Phonetic Correlates of Tongue Root Vowel Contrasts in Maa. *Journal of Phonetics* 32.4: 517–542.

Hahn, Reinhard. 1986. Modern Uighur Language Research in China: Four Recent Contributions Examined. *Central Asiatic Journal* 30.1/2: 35–54.

Hahn, Reinhard. 1991. *Spoken Uyghur*. University of Washington Press.

Hahn, Reinhard. 1998. Uyghur. *Johanson, Lars; Csató, Éva Ágnes, The Turkic Languages, Routledge*, 379–396.

Hale, Mark, and Charles Reiss. 2000. ‘Substance Abuse’ and ‘Dysfunctionalism’: Current Trends in Phonology. *Linguistic Inquiry* 31.1: 157–169.

Hale, Mark, and Charles Reiss. 2008. *The Phonological Enterprise*. Oxford University Press.

Hall, Kathleen Currie. 2009. A Probabilistic Model of Phonological Relationships from Contrast to Allophony. PhD thesis, The Ohio State University.

Hall, Kathleen Currie. 2013. A Typology of Intermediate Phonological Relationships. *The Linguistic Review* 30.2: 215–75.

Hall, Nancy. 2006. Cross-Linguistic Patterns of Vowel Intrusion. *Phonology* 23.3: 387–429.

Hall, Nancy, Nancy Vasquez, Muhammad Damanhuri, Francisco Aguirre, and Connor Tree. 2017. Long-Distance Liquid Coarticulation in American English. In *Proceedings of the Annual Meeting on Phonology*.

Halle, Morris, and Jean-Roger Vergnaud. 1981. Harmony Processes. In *Crossing the Boundaries in Linguistics*, 1–22. Springer.

Hansson, Gunnar Ólafur. 2010. *Consonant Harmony: Long-Distance Interactions in Phonology*. Vol. 145. University of California Press.

Hansson, Gunnar Ólafur. 2014. Agreement by (Non) Correspondence: Inspecting the Foundations. In *ABC Conference, University of California, Berkeley*, 3–62.

Hardcastle, William J, and Nigel Hewlett. 2006. *Coarticulation: Theory, Data and Techniques*. Cambridge University Press.

Harms, Robert T. 1973. How Abstract Is Nupe? *Language*, 439–446.

Harrington, Jonathan; Phil Hoole; Felicitas Kleber, and Ulrich Reubold. 2011. The physiological, acoustic, and perceptual basis for high back vowel fronting: Evidencce from German tense and lax vowels. *Journal of Phonetics* 39.2: 121-131.

Harrington, John P. 1974. Sibilants in Ventureño. *International Journal of American Linguistics* 40.1: 1–9.

Harrison, Carl H, and John M Taylor. 1971. Nasalization in Kaiwá. *Tupí Studies I* 29: 15–20.

Harrison, K David, Mark Dras, and Berk Kapicioglu. 2002. Agent-Based Modeling of the Evolution of Vowel Harmony. In *Proceedings of the 30th Meeting of the Northeastern Linguistics Society*, 217–236.

Harrison, K. David, and Abigail Kaun. 2000. Pattern-Responsive Lexicon Optimization. In *Proceedings of the Thirtieth Meeting of the North East Linguistics Society*.

Hasanboj, Zhamolxonov. 2009. *Özbek Tilining Nazarij Fonetikasi*. Toshkent: Fan.

Hassan, Zeki Majeed, and John H Esling. 2011. Investigating the Emphatic Feature in Iraqi Arabic. In *Instrumental Studies in Arabic Phonetics*, edited by Zeki Majeed Hassan and Barry Heselwood, 217–234.

Hayes, Bruce, Robert Kirchner, and Donca Steriade, eds. 2004. *Phonetically Based Phonology*. Cambridge University Press.

Hayes, Bruce, and Zsuzsa Cziráky Londe. 2006. Stochastic Phonological Knowledge: The Case of Hungarian Vowel Harmony. *Phonology* 23.1: 59–104.

Hayes, Bruce, and Colin Wilson. 2008. A Maximum Entropy Model of Phonotactics and Phonotactic Learning. *Linguistic Inquiry* 39.3: 379–440.

Hayward, Richard J. 1990. Notes on the Aari Language. *Omotic Language Studies*, 425–493.

Hebert, Raymond, and Nicholas Poppe. 1963. *Kirghiz Manual*. Uralic and Altaic Series 33. Indiana University Press.

Heinz, Jeffrey. 2011a. Computational Phonology–Part I: Foundations. *Language and Linguistics Compass* 5.4: 140–152.

Heinz, Jeffrey. 2011b. Computational Phonology–Part II: Grammars, Learning, and the Future. *Language and Linguistics Compass* 5.4: 153–168.

Heinz, Jeffrey. 2018. The Computational Nature of Phonological Generalizations. *Phonological Typology, Phonetics and Phonology*, 126–195.

Heinz, Jeffrey, and William Idsardi. 2011. Sentence and Word Complexity. *Science* 333 (6040): 295–297.

Heinz, Jeffrey and William Idsardi. 2013. What Complexity Differences Reveal about Domains in Language. *Topics in Cognitive Science* 5.1: 111–131.

Heinz, Jeffrey, and Regine Lai. 2013. Vowel Harmony and Subsequentiality. In *Proceedings of the 13th Meeting on the Mathematics of Language (MoL 13)*, 52–63.

Herman, Rebecca, Mary E. Beckman, and Kiyoshi Honda. 1997. Subglottal Pressure and Final Lowering in English. In *Proceedings of the 4th International Conference on Spoken Language Processing*, 145–48.

Hess, Susan. 1992. Assimilatory Effects in a Vowel Harmony System-an Acoustic Analysis of Advanced Tongue Root in Akan. *Journal of Phonetics* 20.4: 475–492.

Higgins, Holly Ann. 2011. Ikoma Vowel Harmony: Phonetics and Phonology. MA thesis, Trinity Western University.

Honda, Kiyoshi. 1996. Organization of Tongue Articulation for Vowels. *Journal of Phonetics* 24.1: 39–52.

Hooper, Joan B. 1976. *An Introduction to Natural Generative Phonology*. Academic Press.

Hsu, Brian, and Karen Jesney. 2016. Scalar Positional Markedness and Faithfulness in Harmonic Grammar. In *Proceedings of the Annual Meeting of the Chicago Linguistic Society*, 51: 241–255.

Hsu, Brian, and Karen Jesney. 2017. Loanword Adaptation in Québec French: Evidence for Weighted Scalar Constraints. In *Proceedings of the 34th West Coast Conference on Formal Linguistics*, 249–258.

Hudu, Fusheini Angulu. 2010. Dagbani Tongue-Root Harmony: A Formal Account with Ultrasound Investigation. PhD thesis, University of British Columbia.

Hulst, Harry van der, and Jereon van der Weijer. 1995. Vowel Harmony. In *The Handbook of Phonological Theory*, edited by John A. Goldsmith, 495–534.

Hume, Elizabeth. 1996. Coronal Consonant, Front Vowel Parallels in Maltese. *Natural Language & Linguistic Theory* 14.1: 163–203.

Hyman, Larry M. 1970. How Concrete Is Phonology? *Language*, 58–76.

Hyman, Larry M. 1988. Underspecification and Vowel Height Transfer in Esimbi. *Phonology* 5.2: 255–273.

Hyman, Larry M. 1999. The Historical Interpretation of Vowel Harmony in Bantu. *Bantu Historical Linguistics: Theoretical and Empirical Perspectives. Stanford: CSLI*, 235–295.

Ibrohimov, Sobirjon. 1967. *Özbek Tilining Andizhon Shevasi*. Toshkent: Fan.

Jardine, Adam. 2016. Computationally, Tone Is Different. *Phonology* 33.2: 247–283.

Johanson, Lars. 1991. On Syllabic Frontness Oppositions in Turkic. *Varia Eurasiatica*, 77–94.

Johanson, Lars. 1998. The History of Turkic. In *The Turkic Languages*, 81–125. Routledge.

Johanson, Lars, and Eva A. Csato, eds. 1998. *The Turkic Languages*. Routledge.

Johnson, Keith, and Jack Martin. 2001. Acoustic Vowel Reduction in Creek: Effects of Distinctive Length and Position in the Word. *Phonetica* 58.1–2: 81–102.

Jongman, Allard, Ratree Wayland, and Serena Wong. 2000. Acoustic Characteristics of English Fricatives. *The Journal of the Acoustical Society of America* 108.3: 1252–1263.

Junusaliev, BM. 1971. *Kyrgyz Tilinin Dialektologijasy*. Frunze: Mektep basmasy.

Kang, Yoonjung. 2011. Loanword Phonology. *The Blackwell Companion to Phonology*, 1–25.

Kaplan, Aaron. 2015. Maximal Prominence and a Theory of Possible Licensors. *Natural Language & Linguistic Theory* 33.4: 1235–1270.

Kara, David Somfai. 2003. *Kyrgyz*. Lincom Europa.

Kaun, Abigail. 2004. The Typology of Rounding Harmony. In *Phonetically Based Phonology*. Cambridge University Press.

Kaun, Abigail Rhoades. 1995. The Typology of Rounding Harmony: An Optimality Theoretic Approach. PhD thesis, University of California, Los Angeles.

Kavitskaya, Darya. 2010. *Crimean Tatar*. Lincom Europa.

Keating, Patricia A. 1984. Phonetic and Phonological Representation of Stop Consonant Voicing. *Language* 60.2, 286–319.

Keating, Patricia A. 1985a. CV Phonology, Experimental Phonetics, and Coarticulation. *UCLA Working Papers in Phonetics* 62: 1–13.

Keating, Patricia A. 1985b. The Phonology-Phonetics Interface. *UCLA Working Papers in Phonetics* 62: 14–33.

Keating, Patricia A. 1988. Underspecification in Phonetics. *Phonology* 5.2: 275–92.

Keating, Patricia A. 1990. The Window Model of Coarticulation: Articulatory Evidence. In *Papers in Laboratory Phonology I: Between the Grammar and Physics of Speech*, edited by John Kingston and Mary E. Beckman, 451–70. Cambridge University Press.

Kelly, John, and John Local. 1986. Long-Domain Resonance Patterns in English. In *International Conference on Speech Input/Output: Techniques and Applications*, 304–309.

Kenstowicz, Michael. 2005. The Phonetics and Phonology of Korean Loanword Adaptation. In *Proceedings of the First European Conference on Korean Linguistics*, 1: 17–32.

Kewley-Port, Diane. 2001. Vowel Formant Discrimination II: Effects of Stimulus Uncertainty, Consonantal Context, and Training. *The Journal of the Acoustical Society of America* 110.4: 2141–2155.

Kingston, John. 2007. The Phonetics-Phonology Interface. In *The Cambridge Handbook of Phonology*, edited by Paul de Lacy, 401–34. Cambridge University Press.

Kingston, John, and Randy L Diehl. 1994. Phonetic Knowledge. *Language* 70.3: 419–454.

Kiparsky, Paul. 1985. Some Consequences of Lexical Phonology. *Phonology* 2.1: 85–138.

Kiparsky, Paul. 2012. *Explanation in Phonology*. Vol. 4. Walter de Gruyter.

Kiparsky, Paul. 2015. Phonologization. In *The Oxford Handbook of Historical Phonology*.

Kiparsky, Paul, and Karl Pajusalu. 2003. Towards a Typology of Disharmony. *Linguistic Review* 20.2/4: 217–242.

Kirchner, Mark. 1998a. Kazakh and Karakalpak. In *The Turkic Languages*, 318–32. New York: Routledge.

Kirchner, Mark. 1998b. Kirghiz. In *The Turkic Languages*, edited by Lars Johanson and Eva A. Csato, 344–56. Routledge.

Kirchner, Robert. 1997. Contrastiveness and Faithfulness. *Phonology* 14: 83–111.

Kirchner, Robert. 1998. An Effort-Based Approach to Consonant Lenition. PhD thesis, University of California, Los Angeles.

Krakow, Rena A., Fredericka Bell-Berti, and Q. Emily Wang. 1995. Supralaryngeal Declination: Evidence from the Velum. In *Producing Speech: Contemporary Issues: For Katherine Safford Harris*, edited by Fredericka Bell-Berti and Lawrence J. Raphael, 333–53. AIP Press.

Krämer, Martin. 2003. *Vowel Harmony and Correspondence Theory*. Vol. 66. Walter de Gruyter.

Kügler, Frank. 2015. Phonological Phrasing and ATR Vowel Harmony in Akan. *Phonology* 32.1: 177–204.

Labov, William, Mark Karen, and Corey Miller. 1991. Near-Mergers and the Suspension of Phonemic Contrast. *Language Variation and Change* 3.1: 33–74.

Lai, Regine. 2015. Learnable vs. Unlearnable Harmony Patterns. *Linguistic Inquiry* 46.3: 425–451.

Lamontagne, Jeffrey. 2018. Evidence of Phonemicization: Lax High Vowels in Canadian French. poster presented at The Annual Meeting of Phonology, UC San Diego.

Lanfranca, Mark. 2012. An Acoustic Study of Underspecified Vowels in Turkish. MA thesis, University of Kansas.

Legendre, Géraldine, Yoshiro Miyata, and Paul Smolensky. 1990. *Harmonic Grammar: A Formal Multi-Level Connectionist Theory of Linguistic Well-Formedness: Theoretical Foundations*. Institute of Cognitive Science Technical Report 90-5. University of Colorado.

Levy, Paulette. 1987. Fonología Del Totonaco de Papantla, Veracruz. PhD thesis, Universidad Nacional Autónoma de México.

Liberman, Mark, and Janet Pierrehumbert. 1984. Intonational Invariance under Changes in Pitch Range and Length. In *Language Sound and Structure: Studies in Phonology Presented to Morris Halle by His Teachers and Students*, edited by Mark Aronoff, Richard Oehrle, Frances Kelley, and Bonnie Wilker Stephens, 157–233. MIT Press.

Lindau, Mona, Leon Jacobson, and Peter Ladefoged. 1972. The Feature Advanced Tongue Root. *UCLA Working Papers in Phonetics* 22: 76–94.

Lindblad, Vern M. 1990. Neutralization in Uyghur. MA thesis, University of Washington.

Lindblom, Bjorn. 1963. Spectrographic Study of Vowel Reduction. *Journal of the Acoustical Society of America* 35.11: 1773–81.

Lionnet, Florian. 2017. A Theory of Subfeatural Representations: The Case of Rounding Harmony in Laal. *Phonology* 34.3: 523–64.

Lisker, Leigh, and Arthur S Abramson. 1970. The Voicing Dimension: Some Experiments in Comparative Phonetics. In *Proceedings of the 6th International Congress of Phonetic Sciences*, 563–567. Academia Prague Prague, Czech Republic.

Lobanov, Boris M. 1971. Classification of Russian Vowels Spoken by Different Speakers. *Journal of the Acoustical Society of America* 49.4B: 606–8.

Lombardi, Linda. 1999. Positional Faithfulness and Voicing Assimilation in Optimality Theory. *Natural Language & Linguistic Theory* 17.2: 267–302.

Lombardi, Linda. 2018. *Laryngeal Features and Laryngeal Neutralization*. Routledge.

Lubowicz, Anna. 2012. *The Phonology of Contrast*. Equinox.

Lunt, Horace G. 1973. Remarks on Nasality: The Case of Guraní. In *A Festschrift for Morris Halle*, edited by Stephen R. Anderson and Paul Kiparsky, 131–139. Holt, Reinhart and Winston.

Maddieson, Ian, and Karen Emmorey. 1985. Relationship between Semivowels and Vowels: Cross-Linguistic Investigations of Acoustic Difference and Coarticulation. *Phonetica* 42.4: 163–174.

Magen, Harriet S. 1997. The Extent of Vowel-to-Vowel Coarticulation in English. *Journal of Phonetics* 25.2: 187–205.

Mahanta, Shakuntala. 2008. *Directionality and Locality in Vowel Harmony: With Special Reference to Vowel Harmony in Assamese*. PhD thesis, University of Utrecht.

Mahanta, Shakuntala. 2012. Locality in Exception and Derived Environments in Vowel Harmony. *Natural Language & Linguistic Theory* 30: 1109–46.

Mai, Anna, Eric Baković, and Matthew Goldrick. 2018. Phonological Opacity as Local Optimization in Gradient Symbolic Computation. poster presented at the Society for Computation in Linguistics. Salt Lake City, UT.

Manuel, Sharon Y. 1990. The Role of Contrast in Limiting Vowel-to-Vowel Coarticulation in Different Languages. *Journal of the Acoustical Society of America* 88.3: 1286–98.

Manuel, Sharon Y, and Rena A Krakow. 1984. Universal and Language Particular Aspects of Vowel-to-Vowel Coarticulation. *Haskins Laboratories Status Report on Speech Research* 77: 69–78.

Mayer, Connor, and Travis Major. 2018. A Challenge for Tier-Based Strict Locality from Uyghur Backness Harmony. In *Formal Grammar 2018*: 62–83. Lecture Notes in Computer Science.

Mayer, Connor, Travis Major, and Mahire Yakup. 2019. Wug-Testing Uyghur Vowel Harmony. poster presented at the 27th Manchester Phonology Meeting, University of Manchester, UK.

McCarthy, John J. 2004. Headed Spans and Autosegmental Spreading. unpublished manuscript, University of Massachusetts, Amherst.

McCarthy, John J, and Alan Prince. 1995. Faithfulness and Reduplicative Identity. *Linguistics Department Faculty Publication Series* 10.

McCollum, Adam G. 2015. Labial Harmonic Shift in Kazakh: Mapping the Pathways and Motivations for Decay. In *The Proceedings of the 41st Annual Meeting of the Berkeley Linguistics Society* 41: 329–52.

McCollum, Adam G. 2018. Vowel Dispersion and Kazakh Labial Harmony. *Phonology* 35.2: 287–326.

McCollum, Adam G. 2019. Positional Effects on Vowel Quality and Vowel Harmony in Kyrgyz. unpublished manuscript, UC San Diego.

McCollum, Adam G., Eric Baković, Anna Mai, and Eric Meinhardt. 2019. The Expressivity of Segmental Phonology and the Definition of Weak Determinism. unpublished manuscript, UC San Diego.

McCollum, Adam G., and Si Chen. 2019. Kazakh. *Journal of the International Phonetic Association*.

McCollum, Adam G, and James Essegbey. 2018. Unbounded Harmony Is Not Always Myopic: Evidence from Tutrugbu. In *Proceedings of the 35th West Coast Conference on Formal Linguistics*, 251–258.

McCollum, Adam G., and James Essegbey. 2019. Initial Prominence and Progressive Vowel Harmony in Tutrugbu. *Phonological Data and Analysis*.

McMullin, Kevin James. 2016. Tier-Based Locality in Long-Distance Phonotactics: Learnability and Typology. PhD thesis, University of British Columbia.

McPherson, Laura, and Bruce Hayes. 2016. Relating Application Frequency to Morphological Structure: The Case of Tommo So Vowel Harmony. *Phonology* 33.1: 125–167.

Mermelstein, Paul. 1978. Difference Limens for Formant Frequencies of Steady-State and Consonant-Bound Vowels. *The Journal of the Acoustical Society of America* 63.2: 572–580.

Mielke, Jeff. 2018. P-Base: A Database of Sound Patterns in 500+ Languages [Electronic Resource]. http://pbase.phon.chass.ncsu.edu/.

Mohanan, Karuvannur Puthanveettil. 1982. Lexical Phonology. PhD thesis, Massachusetts Institute of Technology.

Mutaka, Ngessimo. 1995. Vowel Harmony in Kinande. *Journal of West African Languages* 25.2: 41–55.

Nadzhip, Ėmir Nadzhipovich. 1971. *Modern Uigur*. Nauka.

Nevins, Andrew. 2010. *Locality in Vowel Harmony*. Vol. 55. Mit Press.

Nguyen, Noël, and Zsuzsanna Fagyal. 2008. Acoustic Aspects of Vowel Harmony in French. *Journal of Phonetics* 36.1: 1–27.

Ní Chiosáin, Máire, and Jaye Padgett. 2001. Markedness, Segmental Realization, and Locality in Spreading. In *Segmental Phonology in Optimality Theory: Constraints and Representations*, edited by Linda Lombardi. Cambridge University Press.

Nilsson, Morgan, and Laura J. Downing. 2018. Somali Vowel Qualities and Vowel Harmony. poster presented at 2nd BJALL, Université de Nantes.

Niyazgalieva, A.A., and G.G. Turganalieva. 2013. *Qazaq Dialektologijasy*. Oral: M.Otemisov atyndagy BQMU redaktsiyalyq baspa ortalygy.

Nord, Lennart. 1986. Acoustic Studies of Vowel Reduction in Swedish. In *Quarterly Progress and Status Report 4*, 19–36. KTH Royal Institute of Technology in Stockholm.

Odden, David. 1991. Vowel Geometry. *Phonology* 8.2: 261–289.

Odden, David. 1994. Adjacency Parameters in Phonology. *Language*, 289–330.

Ohala, John J. 1979. The Contribution of Acoustic Phonetics to Phonology. *Frontiers of Speech Communication Research*, 355–363.

Ohala, John J. 1990. There Is No Interface between Phonology and Phonetics: A Personal View. *Journal of Phonetics* 18.2: 153–172.

Ohala, John J. 1994. Towards a Universal, Phonetically-Based, Theory of Vowel Harmony. In *Third International Conference on Spoken Language Processing*.

Öhman, Sven E.G. 1966. Coarticulation in VCV Utterances: Spectrographic Measurements. *Journal of the Acoustical Society of America* 39.1: 151–68.

Olejarczuk, Paul, Manuel A Otero, and Melissa M Baese-Berk. 2019. Acoustic Correlates of Anticipatory and Progressive [ATR] Harmony Processes in Ethiopian Komo. *Journal of Phonetics* 74: 18–41.

Ozburn, Avery, and Alexei Kochetov. 2018. Ejective Harmony in Lezgian. *Phonology* 35.3: 407–440.

Padgett, Jaye, and Marija Tabain. 2005. Adaptive Dispersion Theory and Phonological Vowel Reduction in Russian. *Phonetica* 62.1: 14–54.

Parkinson, Frederick Brooke. 1996. The Representation of Vowel Height in Phonology. PhD thesis, The Ohio State University.

Pearce, Mary. 2008. Vowel Harmony Domains and Vowel Undershoot. In *UCL Working Papers in Linguistics*, 20: 115–40.

Pierrehumbert, Janet. 1994. Knowledge of Variation. In *Proceedings of the Annual Meeting of the Chicago Linguistic Society*, 30.2: 232–256.

Pierrehumbert, Janet, Mary E Beckman, and D Robert Ladd. 2000. Conceptual Foundations of Phonology as a Laboratory Science. In *Phonological Knowledge: Conceptual and Empirical Issues*, edited by Noel Burton-Roberts, Philip Carr, and Gerard Docherty, 273–304.

Piggott, Glynne. 1980. Implications of Linguistic Change for Concrete Phonology. *Canadian Journal of Linguistics/Revue Canadienne de Linguistique* 25.1: 1–19.

Plag, Ingo, Julia Homann, and Gero Kunter. 2017. Homophony and Morphology: The Acoustics of Word-Final S in English 1. *Journal of Linguistics* 53.1: 181–216.

Port, Robert F, and Michael L O'Dell. 1985. Neutralization of Syllable-Final Voicing in German. *Journal of Phonetics*.

Port, Robert, Fares Mitleb, and Michael O'Dell. 1981. Neutralization of Obstruent Voicing in German Is Incomplete. *The Journal of the Acoustical Society of America* 70.S1: S13–S13.

Poser, William J. 2004. On the Status of Chumash Sibilant Harmony. unpublished manuscript., University of Pennsylvania.

Potts, Christopher, Joe Pater, Karen Jesney, Rajesh Bhatt, and Michael Becker. 2010. Harmonic Grammar with Linear Programming: From Linear Systems to Linguistic Typology. *Phonology* 27.1: 77–117.

Prince, Alan, and Paul Smolensky. 1993. *Optimality Theory: Constraint Interaction in Generative Grammar*. Rutgers University Center for Cognitive Science Technical Report 2.

Puente, F Jabier Elorrieta. 1996. A Bipolar Model of Height in Vowels and Consonants. PhD thesis, University of Texas at Austin.

Pulleyblank, Douglas. 2002. Harmony drivers: No disagreement allowed. *Proceedings of the 28th Annual Meeting of the Berkeley Linguistics Society*, 249-267.

Pullum, Geoffrey K. 1976. The Duke of York Gambit. *Journal of Linguistics* 12.1: 83–102.

Purcell, Edward T. 1979. Formant Frequency Patterns in Russian VCV Utterances. *The Journal of the Acoustical Society of America* 66.6: 1691–1702.

R Core Team. 2017. *R: A Language and Environment for Statistical Computing* (version 3.4.2). R foundation for statistical computing.

Raun, Alo. 1969. *Basic Course in Uzbek*. Uralic and Altaic Series 59. Indiana University Press.

Razhabov, Nazar. 1996. *Özbek Shevashunoligi*. Toshkent: Oqituvchi.

Recasens, Daniel. 1985. Coarticulatory Patterns and Degrees of Coarticulatory Resistance in Catalan CV Sequences. *Language and Speech* 28.2: 97–114.

Recasens, Daniel. 1989. Long Range Coarticulation Effects for Tongue Dorsum Contact in VCVCV Sequences. *Speech Communication* 8 (4): 293–307.

Recasense, Daniel. 1999. Lingual Coarticulation. *Coarticulation: Theory, Data and Techniques* 4: 80–110.

Recasens, Daniel, and Aina Espinosa. 2005. Articulatory, Positional and Coarticulatory Characteristics for Clear/l/and Dark/l: Evidence from Two Catalan Dialects. *Journal of the International Phonetic Association* 35.1: 1–25.

Recasens, Daniel, Maria Dolors Pallarès, and Jordi Fontdevila. 1997. A Model of Lingual Coarticulation Based on Articulatory Constraints. *The Journal of the Acoustical Society of America* 102.1: 544–561.

Reshetov, V.V., and Sh. Shoabdurahmonov. 1978. *Özbek Dialektologiyasi*. Toshkent: Oqituvchi.

Ringen, Catherine O., and Miklós Kontra. 1989. Hungarian Neutral Vowels. *Lingua* 78.2–3: 181–91.

Ritchart, Amanda, and Sharon Rose. 2017. Moro Vowel Harmony: Implications for Transparency and Representations. *Phonology* 34.1: 163–200.

Rolle, Nicholas R. 2018. Grammatical Tone: Typology and Theory. PhD thesis, University of California, Berkeley.

Rose, Sharon, and Rachel Walker. 2004. A Typology of Consonant Agreement as Correspondence. *Language*, 475–531.

Rose, Sharon and Rachel Walker. 2011. Harmony Systems. In *The Handbook of Phonological Theory*, 2nd edition, edited by John A. Goldsmith, Jason Riggle, and Alan C. L. Yu, , 240–90. Blackwell Publishing.

Russell, Ronald Kevin. 1993. A Constraint-Based Approach to Phonology and Morphology. PhD thesis, University of Southern California.

Sabev, Mitko and Elinor Payne. 2018. Unstressed Vowel Reduction in Modern Standard Turkish. poster presented at BAAP 2, University of Kent.

Saeed, John. 1999. *Somali*. Vol. 10. John Benjamins Publishing.

Saltzman, Elliot L, and Kevin G Munhall. 1989. A Dynamical Approach to Gestural Patterning in Speech Production. *Ecological Psychology* 1.4: 333–382.

Saussure, Ferdinand de. 1993. *Saussure's Third Course of Lectures on General Linguistics (1910-1911): From the Notebooks of Emile Constantin*. Edited by Eisuke Komatsu. Translated by Roy Harris. Pergamon Press.

Savy, Renata, and Francesco Cutugno. 1998. Hypospeech, Vowel Reduction, Centralization: How Do They Interact in Diaphasic Variations. In *Proceedings of the XVIth International Congress of Linguists*.

Scarborough, Rebecca, and Georgia Zellou. 2013. Clarity in Communication:'Clear' Speech Authenticity and Lexical Neighborhood Density Effects in Speech Production and Perception. *The Journal of the Acoustical Society of America* 134.5: 3793–3807.

Schachter, Paul. 1969. Natural Assimilation Rules in Akan. *International Journal of American Linguistics* 35.4: 342–355.

Schadeberg, Thilo C. 1982. Nasalization in Umbundu. *Journal of African Languages and Linguistics*, 109.

Schane, Sanford A. 1984. The Fundamentals of Particle Phonology. *Phonology* 1: 129–155.

Seidel, Frank. 2008. *A Grammar of Yeyi: A Bantu Language of Southern Africa*. Vol. 33. Köppe.

Selkirk, Elisabeth O, and Jean-Roger Vergnaud. 1973. How Abstract Is French Phonology? *Foundations of Language* 10.2: 249–254.

Seyfarth, Scott. 2014. Word Informativity Influences Acoustic Duration: Effects of Contextual Predictability on Lexical Representation. *Cognition* 133.1: 140–155.

Seyfarth, Scott, and Marc Garellek. 2018. Plosive voicing acoustics and voice quality in Yerevan Armenian. *Journal of Phonetics* 71: 425-450.

Seyfarth, Scott, Marc Garellek, Gwendolyn Gillingham, Farrell Ackerman, and Robert Malouf. 2018. Acoustic Differences in Morphologically-Distinct Homophones. *Language, Cognition and Neuroscience* 33.1: 32–49.

Shahin, Kimary N. 2003. *Postvelar Harmony*. Vol. 225. John Benjamins Publishing.

Silverman, Daniel. 2006. *A Critical Introduction to Phonology: Of Sound, Mind, and Body*. Continuum Critical Introductions to Linguistics. Continuum.

Sjoberg, Andrée. 1963. *Uzbek Structural Grammar*. Uralic and Altaic Series 18. Indiana University Press.

Smith, Caitlin Michele. 2018. Harmony in Gestural Phonology. PhD thesis, University of Southern California.

Smolensky, Paul. 2006. Optimality in Phonology II: Harmonic Completeness, Local Constraint Conjunction, and Feature Domain Markedness. Edited by Paul Smolensky and Geraldine Legendre. *The Harmonic Mind: From Neural Computation to Optimality-Theoretic Grammar* 2: 27–160.

Smolensky, Paul, and Matthew Goldrick. 2016. Gradient Symbolic Representations in Grammar: The Case of French Liaison. Unpublished manuscript,. Johns Hopkins University and Northwestern University.

Smolensky, Paul, Matthew Goldrick, and Donald Mathis. 2014. Optimization and Quantization in Gradient Symbol Systems: A Framework for Integrating the Continuous and the Discrete in Cognition. *Cognitive Science* 38: 1102–38.

Smolensky, Paul, and Géraldine Legendre, eds. 2006. *The Harmonic Mind: From Neural Computation to Optimality-Theoretic Grammar*. Vol. 1–2. MIT Press.

Stallcup, Kenneth L. 1980. Noun Classes in Esimbi. *Noun Classes in the Grassfields Bantu Borderland*, 139–153.

Stanton, Juliet. 2019. Phonetic Lapse in American English–Ative. unpublished manuscript, New York University. *Glossa*.

Stegen, Oliver. 2002. Derivational Processes in Rangi. *Studies in African Linguistics* 31.1/2: 129–54.

Steriade, Donca. 2000. Paradigm Uniformity and the Phonetics-Phonology Boundary. *Papers in Laboratory Phonology V: Acquisition and the Lexicon* 3: 13–334.

Stevens, Kenneth N, Alvin M Libermann, Michael Studdert-Kennedy, and SEG Öhman. 1969. Crosslanguage Study of Vowel Perception. *Language and Speech* 12.1: 1–23.

Stewart, John Massie. 1967. Tongue Root Position in Akan Vowel Harmony. *Phonetica* 16 (4): 185–204.

Stoakes, Hywel M., Janet M. Fletcher, and Andrew R. Butcher. 2019. Nasal Coarticulation in Bininj Kunwok: An Aerodynamic Analysis. *Journal of the International Phonetic Association* 49.

Szeredi, Dániel. 2012. Possible Theoretical Relevance of Subphonemic Vowel Reduction in Hungarian. *Proceedings of ConSOLE XVII,* 299-319.

Szeredi, Daniel. 2016. Exceptionality in Vowel Harmony. PhD thesis, New York University.

Tabain, Marija. 2001. Variability in Fricative Production and Spectra: Implications for the Hyper-and Hypo-and Quantal Theories of Speech Production. *Language and Speech* 44.1: 57–93.

Tabain, Marija. 2003. Effects of Prosodic Boundary on /AC/ Sequences: Acoustic Results. *Journal of the Acoustical Society of America* 113.1: 516–31.

Tabb, Waller C. 2001. Lubwisi Phonology Survey. http:\\lubwisi.webonary.org.

Töjchiboyev, B., and B. Hasanov. 2004. *Özbek Dialektologijasi*. Toshkent: Abdulla Qodirij nomidagi halq merosi nashrjoti.

Toktonaliev, K.T., ed. 2015. *Kyrgyz Tilinin Zhazma Grammatikasy: Azyrky Kyrgyz Adabij Tili*. Bishkek: Avrasia Press.

Tucker, Benjamin V, and Natasha Warner. 2010. What It Means to Be Phonetic or Phonological: The Case of Romanian Devoiced Nasals. *Phonology* 27.2: 289–324.

Tunley, Alison. 1999. Coarticulatory Influences of Liquids on Vowels in English. PhD thesis, University of Cambridge.

Tursunov, U., A. Muxtorov, and Sh. Rahmatullayev. 1992. *Hozorgi Özbek Adabiy Tili*. Toshkent: Özbekiston.

Ultan, Russell. 1973. Some Reflections on Vowel Harmony. *Working Papers on Language Universals* 12: 37–67.

Vago, Robert M. 1973. Abstract Vowel Harmony Systems in Uralic and Altaic Languages. *Language*, 579–605.

Vago, Robert M. 1976. Theoretical Implications of Hungarian Vowel Harmony. *Linguistic Inquiry* 7.2: 243–263.

Vajda, Edward J. 1994. Kazakh Phonology. In *Opuscula Altaica: Essays Presented in Honor of Henry Schwarz*. Western Washington University.

Vaux, Bert. 2000. Disharmony and Derived Transparency in Uyghur Vowel Harmony. In *Proceedings of the North East Linguistic Society* 30: 672–698.

Vayra, Mario, and Carol A. Fowler. 1992. Declination of Supralaryngeal Gestures in Spoken Italian. *Phonetica* 49.1: 48–60.

Walker, Rachel. 1998. Nasalization, Neutral Segments and Opacity Effects. University of California, Santa Cruz.

Walker, Rachel. 1999. Guaraní Voiceless Stops in Oral versus Nasal Contexts: An Acoustical Study. *Journal of the International Phonetic Association* 29.1: 63–94.

Walker, Rachel. 2003. Reinterpreting transparency in nasal harmony. In *The Phonological Spectrum: Segmental Structure*, 37-72. John Benjamins.

Walker, Rachel. 2011. *Vowel Patterns in Language*. Vol. 130. Cambridge University Press.

Walker, Rachel, Dani Byrd, and Fidèle Mpiranya. 2008. An Articulatory View of Kinyarwanda Coronal Harmony. *Phonology* 25.3: 499–535.

Walker, Rachel, and Geoffrey K. Pullum. 1999. Possible and impossible segments. *Language* 75.4: 764-780.

Warner, Natasha, Allard Jongman, Joan Sereno, and Rachèl Kemps. 2004. Incomplete Neutralization and Other Sub-Phonemic Durational Differences in Production and Perception: Evidence from Dutch. *Journal of Phonetics* 32.2: 251–276.

Washington, Jonathan North. 2016. An Investigation of Vowel Anteriority in Three Turkic Languages Using Ultrasound Tongue Imaging. PhD thesis, Indiana University.

Watson, Janet. 1999. The Directionality of Emphasis Spread in Arabic. *Linguistic Inquiry* 30.2: 289–300.

Watson, Janet. 2002. *The Phonology and Morphology of Arabic*. Oxford University Press.

West, Paula. 1999. The Extent of Coarticulation of English Liquids: An Acoustic and Articulatory Study. In *Proc. 14th ICPhS*, 1901–1904.

West, Paula. 2000. Long-Distance Coarticulatory Effects of British English/l/and/r: An EMA, EPG and Acoustic Study. In *Proccedings of the 5th Seminar on Speech Production: Model and Data*, 105–108.

Wightman, Colin W., Stefanie Shattuck-Hufnagel, Mari Ostendorf, and Patti J. Price. 1992. Segmental Durations in the Vicinity of Prosodic Phrase Boundaries. *Journal of the Acoustical Society of America* 91.3: 1707–17.

Wilson, Ian, and Bryan Gick. 2006. Articulatory Settings of French and English Monolinguals and Bilinguals. *Journal of the Acoustical Society of America* 120.5: 3295–96.

Winn, Matthew B, and Ruth Y Litovsky. 2015. Using Speech Sounds to Test Functional Spectral Resolution in Listeners with Cochlear Implants. *The Journal of the Acoustical Society of America* 137.3: 1430–1442.

Wurm, Stefan. 1949. The (Kara-) Kirghiz Language. *Bulletin of the School of Oriental and African Studies* 13.1: 97–120.

Yakup, Abdurishid. 2005. *The Turfan Dialect of Uyghur*. Turcologica 63. Harrassowitz Verlag.

Yu, Alan C. L. 2007. Understanding near Mergers: The Case of Morphological Tone in Cantonese. *Phonology* 24.1: 187–214.

Yu, Alan C.L. 2004. Explaining Final Obstruent Voicing in Lezgian: Phonetics and History. *Language* 80.1: 73–97.

Yunusaliev, B.M. 1971. *Kyrgyz Dialektologiyasy*. Frunze: Mektep basmasy.

Zawaydeh, Bushra A, and Kenneth de Jong. 2011. The Phonetics of Localising Uvularisation in Ammani-Jordanian Arabic. In *Instrumental Studies in Arabic Phonetics*, edited by Zeki Majeed Hassan and Barry Heselwood, 319: 257–76.

Zhang, Jie. 2004. The Role of Contrast-Specific and Language-Specific Phonetics in Contour Tone Distribution. In *Phonetically Based Phonology*, 157–190. Cambridge University Press.

Zimmermann, Eva. 2018a. Being Exceptional Is Being Weak: Tonal Exception in San Miguel El Grande Mixtec. In *Supplemental Proceedings of the 2017 Annual Meeting on Phonology*, edited by Gillian Gallagher, Maria Gouskova, and Sora Heng.

Zimmermann, Eva. 2018b. Gradient Symbolic Representations and the Typology of Ghost Segments: An Argument from Gradient Markedness. paper presented at The Annual Meeting on Phonology 2018, UC San Diego.

Zimmermann, Eva. 2018c. Gradient Symbolic Representations in the Output: A Case Study from Moses Columbian Salishan Stress. In *Proceedings of the Northeastern Linguistics Society 48*.

Zimmermann, Eva. 2018d. The Gradience of Ghosts. paper presented at the 26th Manchester Phonology Meeting.

Zimmermann, Eva. 2019. Gradience in Phonology: The Argument from Exceptions. paper presented at the Princeton Phonology Forum.

Zsiga, Elizabeth. 1995. An Acoustic and Electropalatographic Study of Lexical and Postlexical Palatalization in American English. In *Phonology and Phonetic Evidence: Papers in Laboratory Phonology 4*, edited by Bruce Connell and Amalia Arvaniti, 282–302. Cambridge University Press.

Zsiga, Elizabeth. 1997. Features, Gesture, and Igbo Vowels. *Language* 73.2: 227–74.