

UNIVERSITY OF CALIFORNIA

Los Angeles

On the Relation between Phonotactic Learning and Alternation Learning

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

by

Junxiang Adam Chong

2017

© Copyright by
Junxiang Adam Chong
2017

ABSTRACT OF THE DISSERTATION

On the Relation between Phonotactic Learning and Alternation Learning

by

Junxiang Adam Chong

Doctor of Philosophy in Linguistics

University of California, Los Angeles, 2017

Professor Megha Sundara, Co-chair

Professor Kie Ross Zuraw, Co-chair

This dissertation examines the question of how phonological alternations are learnt. In constraint-based models of phonological learning, it is hypothesized that prior learning of phonotactics from the lexicon facilitates the learning of alternations. While this is an influential assumption, the empirical evidence for it is equivocal. In this dissertation, I investigate this link by examining the learning outcomes in cases where phonotactics and alternations mismatch, particularly in cases of derived environment effects. For example, in Korean, /t/ palatalizes to [c] before [i] across a morpheme boundary, yet [ti] sequences are attested within stems. Derived-environment effects have proven theoretically challenging to account for precisely because of the mismatch in generalizations within stems and across morpheme boundaries.

Using an artificial grammar learning paradigm, I first show that alternation learning is facilitated when the phonotactics in the lexicon match the alternation. Participants who were trained on a language with vowel harmony within stems and across morpheme boundaries successfully learnt the vowel harmony alternation. Conversely, those who were trained on a language without vowel harmony in stems (derived-environment effect language) failed to learn the alternation across the morpheme boundary, despite being trained on this pattern. This supports the hypothesis of current models of phonological learning. Yet, when participants were only trained on a static phonotactic generalization, they did not readily extend this to

unseen alternations. This suggests that learners are conservative, and need experience with alternations before extending a learnt phonotactic generalization.

What does this mean for phonological patterns with a mismatch between phonotactics and alternations? I present corpus analyses as well as computational learning simulations of two cases which show these derived-environment patterns - Korean palatalization and Turkish velar deletion. I show that in both cases the reported mismatches between phonotactics and alternations are superficial. Korean palatalization is an active alternation supported by a gradient phonotactic constraint in the lexicon. Turkish velar deletion, however, is a morphologically circumscribed alternation with no accompanying phonotactic generalization in the lexicon. I also briefly discuss how this statistical pattern in Turkish is similarly the case with the well-known Finnish assibilation pattern. This undermines assumptions in previous analyses of these patterns. Instead, I argue that there is a bias to maintain similar generalizations, captured by more general constraints, in phonotactics as well as alternations. This further supports the claim that learning phonotactics aids in learning alternations, and further suggests that derived-environment patterns are typologically dispreferred.

I conclude by exploring how a bias for more general constraints might be implemented using a Maximum Entropy learner, making specific use of the model's prior. Specifically, I propose a bias for the learner to prefer generalizations to be explained by more general constraints which are blind to morphological structure, over constraints which reference morphological structure. I show that, under such a model, a canonical derived-environment pattern with a mismatch in phonotactics and alternations is never accurately learnt.

The dissertation of Junxiang Adam Chong is approved.

Sharon Peperkamp

Bruce P. Hayes

Robert Daland

Kie Ross Zuraw, Committee Co-chair

Megha Sundara, Committee Co-chair

University of California, Los Angeles

2017

TABLE OF CONTENTS

1	Introduction	1
1.1	Outline of dissertation	2
1.2	What guides phonological learning?	4
1.2.1	Phonotactics and static regularities in the lexicon	4
1.2.2	Alternation learning	5
1.2.3	Linking phonotactic and alternation learning	7
2	Alternations and phonotactics: Matches & Mismatches	10
2.1	Matches in static generalizations and dynamic alternations	10
2.2	Mismatches in static and active generalizations	13
2.2.1	Derived-Environment Blocking: Static phonotactic generalization but no alternation	14
2.2.2	Non-isomorphism between static generalization and active process	15
2.2.3	Derived-environment effects: active alternation with no phonotactic support	16
2.3	Learning mismatches: exploring possible outcomes	22
2.3.1	What does the initial stages of alternation learning look like?	23
2.3.2	The learning data	24
2.3.3	OT: Constraint demotion	25
2.3.4	Maximum entropy grammars	31
2.3.5	Summary: Preliminary modeling	38
3	Learning derived-environment patterns: experimental investigations	40
3.1	Introduction	40

3.1.1	Previous experimental investigations: phonotactics and alternations . . .	42
3.2	Experiment 1: Full language	46
3.2.1	Methods	47
3.2.2	Blick test: phonotactic generalizations	52
3.2.3	Wug test: phonological alternations	54
3.2.4	Correlation between phonotactics and alternations	56
3.2.5	Discussion	57
3.3	Experiment 2: Distributional learning of full language	58
3.3.1	Methods	59
3.3.2	Results - Blick test: phonotactic generalizations	60
3.3.3	Wug test: phonological alternations	61
3.3.4	Discussion	62
3.4	Experiment 3: Stem-only training	62
3.4.1	Methods	63
3.4.2	Results	64
3.4.3	Discussion	65
3.5	Experiment 4: Trisyllabic stems	67
3.5.1	Methods	67
3.5.2	Results	69
3.5.3	Discussion	71
3.6	General Discussion	72
4	Derived environments and the lexicon	85
4.1	Introduction	85
4.2	Morphological derived environment effects and analytic assumptions	89

4.2.1	The Derived environment condition	89
4.2.2	Phonotactic ‘productivity’	91
4.3	Korean palatalization	92
4.3.1	Historical origins and further background	92
4.3.2	Corpus study	94
4.3.3	NAKL	95
4.3.4	Child-directed speech	103
4.3.5	Corpus Summary	104
4.3.6	Learning a phonotactic grammar of Korean	105
4.3.7	Modeling Summary	108
4.3.8	A new analysis	108
4.4	Turkish velar deletion	114
4.4.1	TELL Corpus	117
4.4.2	Turkish CDS	119
4.4.3	Modeling a Turkish grammar	120
4.4.4	Island of reliability: polysyllabic nouns?	122
4.4.5	Summary	123
4.5	Another brief example: Finnish Assibilation	123
4.6	General Discussion	125
4.6.1	Derived environment effects as a unified phenomemon?	126
4.6.2	Derived environment effects and the relationship between static and dynamic generalizations	128
4.6.3	Conclusion	130
5	Simulating learning in derived-environment effects	134

5.1 Favoring generality: Implementing with a prior	135
5.2 Across-the-board language	136
5.3 Toy Korean	139
5.4 A "true" derived-environment effect	141
5.5 No phonotactic generalization and free variation	144
5.6 Summary	146
6 General Discussion & Conclusion	151
6.1 Summary of dissertation	151
6.2 Implications for derived-environment effects	154
6.3 Implications for phonological learning: phonotactics and alternations	156
Bibliography	159

LIST OF FIGURES

3.1	Example pair of visual stimuli	51
3.2	Experiment 1: Rate of choosing harmonic words in (a) blick test and (b) wug test. Each black dot represents a single participant with the large red-dot indicating mean rates for each language group. The red lines indicate 95% confidence intervals.	53
3.3	Experiment 1: Correlation of between performance in blick and wug tests	56
3.4	Experiment 2: Rate of choosing Semi-Harmonic harmonic words in (a) blick test and (b) wug test. Each black dot represents a single participant with the large red-dot indicating mean rates for each language group. The red lines indicate 95% confidence intervals.	60
3.5	Experiment 3: Rate of choosing harmonic words in (a) blick test and (b) wug test. Each black dot represents a single participant with the large red-dot indicating mean rates for each language group. The red lines indicate 95% confidence intervals.	64
3.6	Experiment 3: (a) Rate of choosing [-mu] allomorph in Wug test. Each black dot represents a single participant with the large red-dot indicating mean rates for each language group. The red lines indicate 95% confidence intervals. (b) Correlation between phonotactic and alternation learning in Harmonic language.	66
3.7	Experiment 4: Rate of choosing harmonic words in (a) blick test and (b) wug test. Each black dot represents a single participant with the large red-dot indicating mean rates for each language group. The red lines indicate 95% confidence intervals.	69

3.8	Experiment 4: (a) Rate of choosing [-mu] allomorph in Wug test. Each black dot represents a single participant with the large red-dot indicating mean rates for each language group. The red lines indicate 95% confidence intervals. (b) Correlation between phonotactic and alternation learning in Harmonic language.	71
5.1	Predicted probability of palatalization	138
5.2	Model outputs: Constraint weights	138
5.3	Model outputs: Predicted probability of palatalization	140
5.4	Model outputs: Constraint weights	141
5.5	Model outputs: Predicted probability of palatalization	143
5.6	Model outputs: Constraint weights	143
5.7	Model outputs: Predicted probability of palatalization	145
5.8	Model outputs: Constraint weights	146

LIST OF TABLES

2.1	Navajo sibilants by the feature [anterior]. Table reproduced from Martin (2011)	11
2.2	Summary of interactions between phonotactics in the lexicon and alternations	13
2.3	Output forms given language type. Column headers are inputs assuming ROTB.	25
2.4	Across-the-board input	28
2.5	Derived-environment effect input	29
2.6	Derived-environment blocking input	29
2.7	Across-the-board language	34
2.8	Weights learnt after phonotactic learning for Across-the-board language	34
2.9	Predicted alternations for structure-blind learner: Across-the-board language	35
2.10	DEE language	35
2.11	Weights learnt after phonotactic learning for DEE language	36
2.12	Predicted alternations for structure-blind learner: DEE language	36
2.13	DEB language	37
2.14	Weights learnt after phonotactic learning for DEB language	37
2.15	Predicted alternations for structure-blind learner: DEB language	37
3.1	Summary of Artificial Languages	49
3.2	Occurrence of VV combinations: by V1 type (front [-back] vs. back [+back]) and by V2 type (front [-back] vs. back [+back]) . Expected counts are in parentheses. Percentages in bold: row percentages; Percentages in italics: column percentages. The cells in gray indicate disharmonic sequences (i.e. [-back][+back] and [+back][-back])	77
3.3	Experiment 1 training lexicon: Harmonic language	78
3.4	Experiment 1 training lexicon: Semi-Harmonic language	79

3.5	Experiment 1 training lexicon: Non-Harmonic language	80
3.6	Experiment 4 training lexicon: Harmonic language	81
3.7	Experiment 4 training lexicon: Non-Harmonic language	82
3.8	Blick test stimuli	83
3.9	Wug test stimuli	84
4.1	No. of words that contain [ti], [t ^h i], [t* ⁱ], [tj], [t ^h j] and [t*j] in NAKL corpus (and by lexical strata).	96
4.2	Occurrence of CV combinations: by consonant type (T vs. other Cs) and vowel type (i, j vs. other Vs). Expected counts are in parentheses. Percentages in bold: row percentages; Percentages in italics: column percentages.	97
4.3	Occurrence of CV combinations: by consonant type (CH vs. other Cs) and vowel type (i, j vs. other Vs) in the entire NAKL. Expected counts are in parentheses. Percentages in bold: row percentages; Percentages in italics: column percentages.	98
4.4	Occurrence of CV combinations: by consonant type (T vs. other Cs) and vowel type (i, j vs. other Vs) - Native lexicon. Percentages in bold: row percentages; Percentages in italics: column percentages.	100
4.5	Occurrence of CV combinations: by consonant type (T vs. other Cs) and vowel type (i, j vs. other Vs) - Sino-Korean lexicon. Percentages in bold: row percent- ages; Percentages in italics: column percentages.	100
4.6	Occurrence of CV combinations: by consonant type (T vs. other Cs) and vowel type (i, j vs. other Vs) - Loanwords. Percentages in bold: row percentages; Percentages in italics: column percentages.	101
4.7	Observed/Expected counts of English CV sequences (loaned/not loaned against TI/CV). Percentages in bold: row percentages; Percentages in italics: column percentages.	102

4.8	Occurrence of CV combinations: by consonant type (T vs. other Cs) and vowel type (i, j vs. other Vs) - CDS corpus. Percentages in bold: row percentages; Percentages in italics: column percentages.	104
4.9	Top weighted constraints learned from NAKL corpus. Grey cells indicate constraints that also fall into the top six in the simulation with Child-Directed Speech (Table 4.10)	106
4.10	Top weighted constraints learned from CDS corpus. Grey cells indicate constraints that also fall into the top six in the simulation with the NAKL corpus (Table 4.9)	107
4.11	Occurrence of /k,g/ compared to other stops/affricates in V_V vs. other contexts. Expected counts are in parentheses. Percentages in bold: row percentages; Percentages in italics: column percentages.	118
4.12	Occurrence of /k, g/ compared to other stops/affricates in V_V vs. other contexts in Turkish nouns. Expected counts are in parentheses. Percentages in bold: row percentages; Percentages in italics: column percentages.	118
4.13	Occurrence of /k, g/ compared to other stops/affricates in V_V vs. other contexts in Turkish verbs. Expected counts are in parentheses. Percentages in bold: row percentages; Percentages in italics: column percentages.	119
4.14	Occurrence of /k, g/ compared to other stops/affricates in V_V vs. other contexts in Turkish CDS. Expected counts are in parentheses. Percentages in bold: row percentages; Percentages in italics: column percentages.	120
4.15	Occurrence of /k, g/ compared to other stops/affricates in V_V vs. other contexts in Turkish polysyllabic nouns. Expected counts are in parentheses. Percentages in bold: row percentages; Percentages in italics: column percentages.	122
4.16	Occurrence of /t/ compared to other stops before /i/ vs. other contexts in Finnish. Expected counts are in parentheses. Percentages in bold: row percentages; Percentages in italics: column percentages.	125
4.17	Static phonotactics vs. alternations	129

4.18 Korean Features	132
4.19 Turkish Features	133
5.1 Input for across-the-board language	137
5.2 Novel test items	137
5.3 Input for Toy Korean	140
5.4 Input for a “true” derived-environment language	142
5.5 Input for a free variation language	145
5.6 Summary: “Best” σ^2 values for *T+I for each language.	147

ACKNOWLEDGMENTS

This dissertation would not have been possible were it not for the help, guidance and feedback from many people. I would first like to sincerely thank my advisors, Megha Sundara and Kie Zuraw. I could not have asked for a better team for advice on navigating the dissertation, the academic job market and also a career beyond graduate school. Megha has been a stalwart in my research career in the graduate program, and I am so very grateful for her reading of numerous drafts and her mentorship over the years. I have always felt welcome in her office to bat around research ideas and discuss my research projects at any time, something which I am very thankful for. Kie has also been a fantastic advisor these past few years. I remember the excitement at listening to the voicemail message from Kie informing me of my admission into the UCLA graduate program in Linguistics. Since then, I have benefited from Kie's fantastic teaching in the classroom and mentoring outside of it. She always offered invaluable candid advice on a whole range of academic and non-academic topics, and I believe our advisor-advisee relationship has developed into one between colleagues, one which I feel very privileged to have.

I would also like to thank the other members of my dissertation committee: Bruce Hayes, Robert Daland and Sharon Peperkamp. Bruce has been a de-facto third advisor to me and my meetings with him always pushed me to go deeper into a given research question. I deeply appreciate all his input on my work both in terms of the dissertation but also other research projects. Thank you to Robert Daland especially for all those informal chats on coffee runs to Northern Lights over the years. Robert has always been supportive in my research endeavors and always helped me see the novelty of the question that I was asking, even if I sometimes lost sight of that myself. I am also very grateful for the advice of Sharon Peperkamp, who so graciously agreed to be the external member of my dissertation committee. Despite the long distance and time zone difference, Sharon was always willing to discuss various aspects of the

dissertation and always asked insightful questions that forced me to reconsider the results from a different angle.

In addition to those at UCLA, I am indebted to Sharon Inkelas, Karen Jesney and Stephanie Shih for discussion about various aspects of this work. I would also like to thank audiences at AMP 2015 (UBC), 2016 (USC), LSA 2016, 2017, SCaMP 2016 (UCSD), as well as audience members at Phorum at UC Berkeley, PhLunch at USC, University of Melbourne Phonetics Lab, various UCLA Phonology Seminars, University of Hawai'i at Mānoa, Queen Mary - University of London and NYU. A modified version of Chapter 4 of this work is currently under review for publication, and I thank three anonymous reviewers and an Associate editor for their helpful comments. I would like to thank Eleanor Glewwe for helping record the stimuli used in this experiment and Audrey Kim for helping run a pilot version of the experiments. I am also greatly indebted to Daniel Szeredi and Michael Becker for their help with Experigen.

I am truly grateful for having the opportunity to undertake my graduate training at UCLA's Department of Linguistics. It has been an exceptional learning and research environment, and I am extremely thankful for all the excellent teachers and graduate student colleagues I have had over the last six years. It has been a particular privilege being part of the UCLA Phonetics Lab and learning under Pat Keating and Sun-Ah Jun. Under their guidance, I have honed my own skills as a phonetician and am fortunate enough to have both of them not only as mentors but as research collaborators. I feel very lucky to have such a supportive network of friends and colleagues, too many to enumerate, here at UCLA, and beyond. I would also like to thank my undergraduate mentors at the University of Melbourne, in particular Brett Baker and Janet Fletcher, for first introducing me to phonology and phonetics and developing my interests in these areas. Finally, I would like to thank Henry Tehrani for his technical help at various points during my time in the UCLA Phonetics Lab.

Outside of UCLA and Linguistics, I am so thankful for having been part of the Choir of All Saints' Episcopal Church in Beverly Hills. You have all provided me with a musical and spiritual home outside of academia, one that has kept me grounded throughout these past six years. I will miss singing with you all dearly!

Finally, it would be remiss of me to not thank my family. Pa and Ma, your unconditional support and love from the other side of the world, has always strengthened and nourished me through life's ups and downs. I apologize for all those weekly phone calls where I was a little grumpy! Thank you for being so patient with me. I want you to know how much this means to me and that I would not be here without the both of you being there along the way. To my brother, Nicholas, who has been with me on this crazy rollercoaster ride that is the academic job market: thank you for being there and sharing the ups and downs of academic life, and congratulations on the new job! Lastly, to my dear Chad, thank you for being my rock day in and day out these past five years, and for putting up with all my neuroses. Thank you for your discussion of various aspects of my research both dissertation-related and beyond. Your support got me across the finish line, and I could not have done this without you!

This dissertation was funded by a UCLA Dissertation Year Fellowship.

VITA

- 2009 Dean's List, Faculty of Arts
University of Melbourne
Melbourne, Victoria, Australia
- 2009 Australian-Asian Association of Victoria Prize (Beginners Arabic)
University of Melbourne
Melbourne, Victoria, Australia
- 2010 Summer Research Scholarship
Australian National University
Canberra, ACT, Australia
- 2011 B.A. (Hons.) with First-Class Honours (Linguistics)
University of Melbourne
Melbourne, Victoria, Australia
- 2015 M.A. (Linguistics)
University of California, Los Angeles
Los Angeles, California, USA
- 2015-2016 UCLA Dissertation Year Fellowship
University of California, Los Angeles
Los Angeles, California, USA
- 2012-2017 Research Assistant, Teaching Assistant/Associate/Fellow
Department of Linguistics
University of California, Los Angeles
Los Angeles, California, USA

PUBLICATIONS

Chong, A.J. & German, J.S. (2017). The accentual phrase in Singapore English. *Phonetica*, 74(2), 63-80.

Chong, A.J. (2016). Learning consequences of derived-environment effects. *Proceedings of the Linguistic Society of America*, Vol 1.

Chong, A.J. & Sundara, M. (2016). Perceptual similarity modulates online compensation for phonological variation. *Proceedings of the Annual Meeting on Phonology 2*.

Chong, A.J. & German, J.S. (2015). Prosodic phrasing and F0 in Singapore English. In The Scottish Consortium for ICPHS 2015 (Ed.), *Proceedings of the 18th International Congress of Phonetic Sciences*. Glasgow, UK: the University of Glasgow.

Bishop, J.B., **Chong, A.J.**, Jun, S. (2015). Individual differences in prosodic strategies to sentence parsing. In The Scottish Consortium for ICPHS 2015 (Ed.), *Proceedings of the 18th International Congress of Phonetic Sciences*. Glasgow, UK: the University of Glasgow.

Chong, A.J. & Sundara, M. (2015). 18-month-olds compensate for a phonological alternation. In Grillo, E., & Jepson, K. (Eds.), *Proceedings of the 39th Annual Boston University Conference on Language Development*, 113-126.

Chong, A.J. (2011). Lenition in Gaalpu: An Optimality Theoretic analysis. *Australian Journal of Linguistics*, 31(4), 473-490.

CHAPTER 1

Introduction

Phonology is concerned with characterizing speakers' tacit knowledge of the sound patterns of their language. One component is knowledge about phonological alternations: when a phoneme is pronounced in one way in a particular context, but in a different way in another. For example, in English, the past tense suffix '-ed' (/ -d/) is produced as [-d] when it occurs following a word with a final voiced obstruent (e.g. *lag*: /læg-d/ → [lægd]), and when / -d/ is suffixed to a word with a final voiceless obstruent, the suffix surfaces as [-t] (e.g. *lack*: /læk-d/ → [lækt]). There is, therefore, an alternation between [t] and [d]. A fundamental question in phonological research concerns how an alternation like this is learnt.

One factor that has been hypothesized to aide alternation learning is the prior learning of another component of phonological knowledge: phonotactics (Hayes 2004, Prince & Tesar 2004, Pater & Tessier 2005, Jarosz 2006, Tesar & Prince 2007, Hayes & Wilson 2008, Jarosz 2011). In addition to knowledge about phonological alternations, speakers are also tacitly aware of the legal and illegal sound sequences of words in their mental lexicon (i.e. phonotactics). For example, in English, [kt] sequences are allowed in codas (e.g. *act* [ækt]) but [kd] is not (e.g. *akd* [ækd] is not a possible word). It has further been observed that phonological alternations often enforce the phonotactic constraints of the language (Chomsky & Halle 1968, Kenstowicz & Kisseberth 1977). In the example above, the alternation from / -d/ to [-t] ensures that a morphologically complex word is not produced with the phonotactically illegal sequence [kd] (e.g. *lack*: /læk-d/ → *[lækd]). In fact, constraint-based models of phonology such as Optimality Theory (Prince & Smolensky 1993/2004) encode this close relationship by capturing generalizations about both alternations and phonotactics using a similar mechanism (i.e. the

same markedness constraint): a constraint like *_{KD} captures both the fact that words like *akd* do not occur in English (phonotactics) and that /d/ alternates to [t] in words like *lacked*. The intuition is that learning first that [kd] sequences are illegal in one's language, a phonotactic generalization, bootstraps the later learning of the alternation that avoids this sequence.

This hypothesis is plausible given that knowledge about native language phonotactics seems to emerge before knowledge about phonological alternations. Moreover, a number of computational models have showed support for the assumption that first learning phonotactic knowledge aids in later learning of alternations (Jarosz 2006, 2011, Tesar & Prince 2007). Yet while this is an influential hypothesis in models of phonological learning, it is one that lacks strong experimental support. In this dissertation, we examine this question by investigating what happens to learning when phonotactics mismatch alternations. What if [kd] were a legal sequence in English? How might one learn to change /d/ to [t] in a word like *lacked*? The overarching research questions are as follows:

1. Does phonotactic learning facilitate alternation learning?
2. How might a learning perspective on phonological mismatches shed light on how to theoretically account for these kinds of phonological patterns?

1.1 Outline of dissertation

This dissertation is structured as follows. In the rest of this chapter, I discuss the background on phonological learning in the domain of both phonotactics and alternations.

In **Chapter 2**, I provide a global overview on the general typology of mismatches between phonotactics and alternations, focusing on reviewing the literature on derived-environment effects. I then present the results of toy computational simulations that illustrate how alternation learning in these cases might be more difficult.

In **Chapter 3**, I present the results of a series of artificial grammar learning experiments that test directly the question of whether phonotactic matches aid alternation learning. The

results indicate that, indeed, learning of an alternation is facilitated by phonotactic learning. In fact, learners in the mismatch language failed to learn the alternation.

In **Chapter 4**, I examine two well-known examples of mismatches between phonotactics and alternations (derived-environment effects): Korean palatalization and Turkish velar deletion. I show that, in Korean, coronal stop + high front vowel sequences [TI], which are repaired across a suffix boundary, are actually heavily under-represented in the lexicon. I then show that a phonotactic learner easily infers a constraint against such a constraint. This indicates that stem-internal [TI] sequences are not completely well-formed, as is largely assumed in previous analyses of these patterns. Contrastively, our investigation of Turkish reveals that, although Turkish velar deletion is structurally similar on the surface as Korean palatalization, the lexical statistics are quite different. In this case, there is no strong phonotactic generalization against intervocalic velars in the lexicon. Concomitantly, the alternation is also much more constrained in applicability. I also briefly consider another example, Finnish assibilation, and show that the same statistical pattern with Turkish velar deletion is found there as well. I argue, therefore, that derived-environment effects are by no means a unitary phenomenon, and moreover, the empirical data do not support existing analytical assumptions of these patterns.

In **Chapter 5**, I present toy simulations of how the kinds of derived-environment patterns we investigated in Chapter 4 could be learnt, further showing that a ‘true’ derived-environment pattern is not easily learnable. I show how a bias for more general constraints could in principle be implemented computationally by adjusting the value of the free parameters in a Maximum Entropy grammar in favor of assigning weight to a less complex, more general, structure-blind constraint.

Finally, I end in **Chapter 6** by discussing the general implications of the results of this dissertation and suggest fruitful avenues for future work.

1.2 What guides phonological learning?

1.2.1 Phonotactics and static regularities in the lexicon

Speakers of a given language possess knowledge about what the legal and illegal sound sequences of their target language (i.e. phonotactics). Research on phonotactic acquisition in infancy has shown that by 9 months of age, infants are sensitive to the phonotactic regularities in their input and show preferences for words that conform to native language phonotactics (Jusczyk, Friederici, Wessels, Svenkerud, & Jusczyk 1993, Friederici & Wessels 1993). They are moreover able to use these distributional cues for the task of word segmentation (Jusczyk, Hohne, & Bauman 1999, Mattys & Jusczyk 2001, Mattys, Jusczyk, Luce, & Morgan 1999, see also Adriaans & Kager 2010). Infants also show the ability to distinguish words of differing phonotactic probabilities (Jusczyk, Luce, & Charles-Luce 1994), thus indicating that their sensitivity to phonotactic probabilities is fine-grained. Infants by 9-months of age then are capable of learning phonotactic regularities from short exposure to novel stimuli (Saffran & Thiessen 2003, Chambers, Onishi, & Fisher 2003, K. S. White, Peperkamp, Kirk, & Morgan 2008).

Further, previous studies have shown that adult speakers not only encode the static generalizations in their language input, but can project this knowledge onto novel unseen forms. Evidence for this comes from the domain of loanword adaptation whereby non-native words are often repaired to conform to native language phonotactics (e.g. Paradis & Lacharité 1997, Kang 2010) as well as nonce-word rating experiments (e.g. Frisch & Zawaydeh 2001, Kager & Pater 2012, Hayes & White 2013, Moore-Cantwell & Sanders 2014). Moreover, artificial language learning paradigms have shown that experimental participants with adult participants are able to extract phonotactic regularities from a short exposure to a novel language and also apply them to novel forms (Onishi, Chambers, & Fisher 2002, Linzen & Gallagher 2014). Furthermore, it seems as though the encoding of these static generalizations is sensitive to the frequency of *word-types* versus *word-tokens* (Richtsmeier 2011, Hamann, Apoussidou, & Boersma 2012, Pierrehumbert 2003, Edwards, Beckman, & Munson 2004), with word-type frequency typically being implicated in phonotactic learning. In fact, speakers' well-formedness judgments often

match the frequency with which that pattern occurs in the lexicon (Moore-Cantwell & Sanders 2014, Hayes & Londe 2006, Hayes, Siptár, Zuraw, & Londe 2009). Thus, both infant and adult learners are able to induce phonotactic regularities in natural and artificial languages. Furthermore, these studies also show that learners do not internalize these regularities in a purely categorical way. Learners seem to be sensitive to the strength of a particular regularity (as shown by their sensitivity to frequency of occurrence), suggesting that speakers' knowledge of static generalizations over their lexicons is gradient and dependent on the strength of the trends in the lexicon.

1.2.2 Alternation learning

In addition to learning phonotactics, learners must also acquire the phonological alternations in their target language. Phonological alternations occur when a particular morpheme varies in pronunciation depending on context. For example, as noted above, the English past tense /-d/ is pronounced as [t] following voiceless obstruents, but as [d] elsewhere. Part of a child's learning task must be to learn where these alternations occur, and what sounds are involved in the alternation. How does this occur?

Evidence has been accumulating that alternation learning is guided by learning biases. Biased learning of phonological alternations occurs when learners fail to pick up on phonological patterns available in the input, or make certain assumptions in the face of ambiguous input. Two common types of biases are often discussed in the literature: substantive and complexity biases (see Moreton & Pater 2012a,b for an overview). Substantive biased learning (Wilson 2006) attributes a role to phonetics in how alternations are learnt (see also phonetically-based phonology: Hayes, Kirchner, & Steriade 2004). Simply put, learners are predisposed to learning phonological patterns that are motivated by perceptual and articulatory considerations. For example, J. White (2014) and J. White & Sundara (2014) showed that adult and infant learners trained on an artificial language generalize a learnt alternation asymmetrically from more phonetically distant pairs to less distant pairs. When learners were trained on an alternation [p] → [v], they generalized the alternation to [b] → [v]. When they were trained on [b] → [v],

however, they did not generalize to [p] → [v]. In this case, learners seem predisposed to posit alternations between more phonetically similar segments than more distant ones, evidence of a learning bias based on the phonetic considerations. Not all studies that examined substantive biases, however, show evidence in support of this. For example, studies that have compared phonetically motivated vowel harmony and less phonetically motivated disharmony have failed to find a learning difference (Pycha, Nowak, Shin, & Shosted 2003, Skoruppa & Peperkamp 2011).

Alongside substantive biases, learners have also exhibited complexity biases in favor of alternations involving simpler patterns over more complex ones. Typically, the notion of phonological complexity has been operationalized in terms of the number of phonological features involved in defining the class of segments that are subject to a particular phonological process, or that are required to change due to a phonological process (see Moreton & Pater 2012a for an overview). For example, participants in Pycha et al.'s (2003) study learned a vowel harmony alternation much better when the generalization could be stated using one feature (e.g. all front vowels take [-εk], all back vowels take [-Λk]) than more features (e.g. stems with final [i, ae, υ] take [-εk], stems with final [u, a, ɪ] take [-Λk]). In the latter case, the generalization is formally more complex than the former, no matter how the rule is formulated. Thus learners in Pycha et al.'s (2003) study were better able to learn an alternation which was simpler, involving fewer features, than one which was more complex, involving more features.

Thus it is clear that some kinds of alternation patterns are easier to learn than others, based on phonetic similarity between the segments in alternation or in terms of formal complexity, and further that learners are biased to assume that phonological patterns operate over a word-sized domain. In the next section, I review one other factor that has been argued to aid in alternation learning: phonotactic knowledge.

1.2.3 Linking phonotactic and alternation learning

In addition to the learning biases discussed in the previous section, another factor that has been argued to aid in alternation learning is the prior knowledge about phonotactics (Hayes 2004, Prince & Tesar 2004, Pater & Tessier 2005, Jarosz 2006, Tesar & Prince 2007, Hayes & Wilson 2008, Jarosz 2011). In fact, in constraint-based phonological models, such as Optimality Theory, both phonotactic generalizations and alternations are often assumed to be encoded using the same constraints, leading us to the prediction that they are learnt using a single mechanism. Computational models that incorporate this assumption have had success in modeling the developmental trajectory of alternations (e.g. Jarosz 2011). Further it is generally hypothesized that phonotactic learning precedes alternation learning, for which the earliest evidence from natural language learning has been shown at 12 months (Sundara, Kim, White, & Chong 2014). Infants have further been shown to be able to learn alternations in an artificial grammar paradigm at the same age (K. S. White et al. 2008, J. White & Sundara 2014). Given that the learning of phonotactics in a native language seems to emerge earlier at around 9 months (Jusczyk et al. 1994, 1993, Friederici & Wessels 1993), the ability for phonotactics to aid in alternation learning seems a possible learning scenario.

While phonotactic learning is hypothesized to aid in alternation learning in many models of phonological learning, the experimental evidence in support of this hypothesis is equivocal. Past research has focused on whether or not learners extend a static generalization or a learnt alternation to novel items that involve the same type of generalization. That is, learners are trained on alternations and tested on novel alternations, or trained on a phonotactic generalization and tested on knowledge about that learnt phonotactic generalization. Typically studies ignore *how* these types of phonological knowledge interact from a learning perspective.

In the most well-known study that examined this link, Pater & Tessier (2005), English speakers were trained on one of two possible alternations. Participants assigned to Language 1 in (1), were trained on an alternation that relied on phonotactic constraint in English that lax vowels do not occur in open syllables. This phonotactic generalization, however, does not

engender any alternations in English. In Language 1, [t]-epenthesis occurred when a surface vowel was lax and in an open syllable, such as in the singular forms of (1)(a-c). Epenthesis did not occur in the plural forms or when the vowel was tense as in (1)(d-e). Crucially epenthesis applied to ensure conformity to the existing English phonotactic constraint.

(1) **Language 1** (Pater & Tessier 2005)

	Root	Plural	Singular
a.	/bɪ/	[bɪsɔ]	[bɪt]
b.	/gɛ/	[gɛsɔ]	[gɛt]
c.	/flʌ/	[flʌsɔ]	[flʌt]
d.	/blej/	[blejso]	[blej]
e.	/glɛk/	[glɛkso]	[glɛk]

Participants assigned to Language 2 (2) were trained on a similar epenthesis alternation, except in this case the alternation was conditioned by the frontness of the vowel. Epenthesis occurred following front vowels (2)(a-c) but not back (2)(d-e). Compared to Language 1, Language 2 was of the same formal complexity. Pater & Tessier (2005) argued that if alternations were learnt by mere pattern induction without reference to phonotactic knowledge, then these two languages should be equally easy to learn (equivalent accuracy). They predicted that if learners utilized their *existing* knowledge of the phonotactics of English, then Language 1 should be easier to learn than Language 2.

(2) **Language 2** (Pater & Tessier 2005)

	Root	Plural	Singular
a.	/lij/	[lijso]	[lijt]
b.	/blej/	[blejso]	[blejt]
c.	/træ/	[træso]	[træt]
d.	/fuw/	[fuwso]	[fuw]
e.	/gluwk/	[gluwkso]	[gluwk]

Conforming to Pater & Tessier's (2005) prediction, learners learnt the alternation better in the language with phonotactic support (Language 1) than the one without (Language 2). However, the authors point out that not only is Language 2 not phonotactically motivated, it is also typologically unattested. Given that we know that unnatural patterns are at the very least dispreferred by learners (Hayes et al. 2009, Becker, Ketrez, & Nevins 2011, Hayes & White 2013), it is possible that the poorer performance in Language 2 could be explained by this alone. Due to this confound, a fairer and stronger test for a link between phonotactic and alternation learning would be to look at whether learners generalize better (or faster) equally natural generalizations where one has support from phonotactics but the other does not.

Given the lack of clear experimental evidence in support of this assumption, this dissertation's goal, therefore, is to examine how phonotactic learning influences alternation learning. In particular, we will use artificial grammar learning to compare the success of alternation learning in a language with matching phonotactics with one in which the phonotactics mismatch (Chapter 3). The mismatch language is modeled after derived-environment effects (a.k.a. non-derived environment blocking; Kiparsky 1993; Chapter 2). In these cases, there is an active alternation that is not supported by phonotactics with the very same sequences which are subject to an alternation across a morpheme boundary nonetheless occurring in the lexicon. These types of patterns thus provide a nice test case with which to examine the relationship between these two types of phonological generalizations. Moreover, I hope to also move the theoretical discussion about these mismatch patterns forward by bringing a novel learning-theoretic perspective on a phonological pattern that has proven particularly thorny to account for (Chapter 4).

This dissertation is therefore as much about the relationship between phonotactic and alternation learning as it is about derived-environment effects themselves. The goal then is to use derived-environment effects as a means of probing the relationship between phonotactic and alternation learning, then, in turn, using the results obtained to inform our theoretical accounts of such patterns.

CHAPTER 2

Alternations and phonotactics: Matches & Mismatches

2.1 Matches in static generalizations and dynamic alternations

It has often been observed that same phonological generalizations seem to hold both within a morpheme and across a morpheme boundary. In English, for example, the past tense suffix /-d/ alternates to [t] following a word with a final voiceless consonant e.g. *lack*: /læk-d/ → [lækt]) to avoid the consonant cluster [kd] that is unattested in the lexicon. The qualities of these generalizations are different, however. In the latter case, these phonological generalizations give rise to alternations, an active process that changes the form of a morpheme from one context to another. In the former, these are static phonotactic generalizations that describe the possible sequences of sounds within morphemes in the lexicon. In earlier rule-based theories (Chomsky & Halle 1968), these generalizations were stated separately in the grammar as phonological rules and morpheme structure constraints. Yet this ignored the fact that phonological rules often enforced static phonotactic restrictions in the lexicon, therefore a particular generalization under a rule-based account was stated twice in the grammar (Duplication Problem: Kenstowicz & Kisseberth 1977).

An example of the isomorphism between static phonotactic generalizations and active alternations can be seen in Navajo sibilant harmony. In Navajo, (Sapir & Hoijer 1967, Kari 1976, McDonough 1991, 2003, Fountain 1998) all sibilants in a root must agree in their [anterior] feature. A list of Navajo sibilants (grouped by their value for the feature [anterior]) is given in Table 2.1. The harmony restrictions tautomorphemically mean that hypothetical disharmonic

Navajo roots in (2) are not attested, and only harmonic roots like those in (1) exist. All data cited here are reproduced from Martin (2011), cited originally from Fountain (1998).

[+anterior]	[-anterior]
s	ʃ
z	ʒ
ts ^h	tʃ ^h
ts	tʃ
ts'	tʃ'

Table 2.1: Navajo sibilants by the feature [anterior]. Table reproduced from Martin (2011)

(1) Harmonic roots

- a. /tʃ'óʒ/ 'worm'
- b. /ts'óʒí/ 'slender'

(2) Hypothetical disharmonic roots

- a. *soʃ
- b. *tʃiz

This co-occurrence restriction also holds across morpheme boundaries, as in (3), with the prefix sibilant typically harmonizing with the [anterior] feature of the root sibilant. So in (3a), the prefix /-s-/ harmonizes to the root sibilant [ʒ], surfacing as [ʃ]. In (3b), on the other hand, both the prefix and root sibilants agree in anteriority and thus the prefix sibilant does not alternate.

(3) Navajo prefix+root

- a. /ji-s-lééʒ/ → [ji-ʃ-lééʒ] 'it was painted'
- b. /ji-s-tiz/ → [ji-s-tiz] 'it was spun'

In a rule-based (SPE) account (Chomsky & Halle 1968), the generalizations that hold over individual morphemes in isolation were accounted for via context-free Morpheme Structure Constraints (MSCs; Halle 1959, Stanley 1967, Kenstowicz & Kisseberth 1977). Generalizations across morpheme boundaries, on the other hand, were captured by a separate mechanism, phonological rules. Crucially, it was assumed that MSCs applied directly in the lexicon prior to the cyclic application of any phonological rules. Due to these differences, Kenstowicz & Kisseberth (1977) argued MSCs were thus totally *distinct* from phonological rules proper. This redundancy of explanation of phonological phenomenon, where generalizations were stated twice in the grammar, was questioned since this treatment suggested that “what appears to be a single phenomenon in some sense must be treated as two unrelated phenomena” (Kenstowicz & Kisseberth 1977:136). Patterns like this in which both tautomorphemic and heteromorphemic phonological patterns can be accounted for by the same generalization led to what is known as the “Duplication Problem.”

The dispreference against this redundancy in the grammar led to the formulation of Optimality Theory (Prince & Smolensky 1993/2004:henceforth OT). OT captured the same generalization across different phonological domains by encoding this using the same markedness constraint on output structure. Further, the invocation of a theory of Richness-of-the-base (ROTB) explicitly denied the existence of any constraints on lexical representations (that is, MSCs). Any such constraints are emergent properties from the grammar through constraint ranking, thereby placing the explanatory burden for both static phonotactic patterns and active processes into the grammar (i.e. CON). Thus in OT and its derivatives, phonotactic patterns and alternations are captured by the same mechanism. Thus the sibilant harmony pattern in Navajo above would be captured by a general constraint banning sequences of sibilants that do not agree for the feature [anterior]. This applies both in the lexicon as well as to morphologically complex forms engendering alternations. Yet, while the patterns in which phonotactic generalizations in the lexicon and alternations at the morpheme boundary, like Navajo harmony above, abound, there are many cases in the literature in which these two types of phonological generalizations do not accord with each other. In the remainder of this chapter,

I survey a set of cases in which these generalizations pull apart. I focus primarily on a set of patterns collectively known as derived-environment effects (a.k.a. non-derived environment blocking; Kiparsky 1993). In the latter part of this chapter, I present some toy simulations to illustrate why learning of these mismatch patterns is problematic.

2.2 Mismatches in static and active generalizations

Paster (2013) provides a survey of cases in which static phonotactic and dynamic alternation patterns pull apart from each other. She argues that if language change in OT involves constraint re-ranking, and a single constraint allows you to capture both static and active generalizations, then both static patterns and active processes should undergo changes *in tandem*. Paster uses these mismatch cases to argue for the fact that static phonotactic generalizations and alternations might need to be separately stated in the grammar, a position advanced earlier in by Anderson (1974). In the following three sections, I survey the taxonomy (Table 2.2) of observed interactions summarized in Paster (2013).

Static phonotactic generalization	Alternation	Example
✓	✓	Navajo sibilant harmony (Across-the-board)
✓	✗	Marash Armenian vowel harmony (Derived-environment blocking)
✗	✓	Korean /t/ palatalization (Derived-environment effect)
✓ but different characteristics	✓ but different characteristics	English Closed Syllable Shortening

Table 2.2: Summary of interactions between phonotactics in the lexicon and alternations

2.2.1 Derived-Environment Blocking: Static phonotactic generalization but no alternation

One common example of a mismatch are derived-environment blocking patterns (DEB). In these cases, a static phonotactic generalization that is true of the lexicon does not actually drive an alternation. (Paster 2013:86-87) cites an example from the Marash variety of Armenian (Vaux 1998). In this language, root vowels must be harmonic for both the features [back] and [round] (inventory: [i], [y], [e], [ø], [ɛ], [ɑ], [o] and [u])

(4) Harmonic roots

- | | | | |
|-----------------------|-----------------|----------|-----------|
| a. k ^h ɑnɑ | ‘how many/much’ | d. øsør | ‘today’ |
| b. ærin | ‘blood’ | e. ybyr | ‘when’ |
| c. hik ^h æ | ‘soul’ | f. urnog | ‘example’ |

Like Navajo, disharmonic roots are unattested (e.g. *k^hury), indicating a phonotactic constraint on harmony within roots. Unlike Navajo, however, this restriction does not hold across a morpheme boundary, as shown in (5).

(5) Disharmonic words

- | | |
|--------------------------------|---|
| a. ɑs-i-m | ‘say-theme.V-1SG’ |
| b. bɑrt ^h ɑgon-i-is | ‘owing-be-2SG’ |
| c. yr-iri | ‘interrog.pron-PL.NOM (from Adjarian 1954)’ |

Thus, despite the phonotactic restriction within roots, this restriction is not enforced by an alternation. This type of mismatch pattern also seems to commonly occur with laryngeal co-occurrence restrictions which often apply only to roots but not in a larger domain, such as across a root and suffix boundary (Gallagher 2010, MacEachern 1999). In Sanskrit, for example, roots (6) and roots with reduplicated material (7) can contain at most one aspirated stop (voiced or voiceless) (i.e. Grassmann’s law; data from Anderson 1970).

(6) Laryngeal co-occurrence restrictions in Sanskrit roots:

- a. /d^hauk/ ‘approach’ */d^hauk^h/
- b. /k^ha:d/ ‘chew’ */k^ha:d^h/

(7) Laryngeal co-occurrence restrictions in Sanskrit reduplicated forms:

- a. Root: /p^hal/ ‘burst’ → [pa=p^hala] ‘have burst’ *[p^hap^hala]
- b. Root: /d^hauk/ ‘approach’ → [d^hu=d^hauka] ‘have approached’ *[d^hud^hauka]

This phonotactic restriction, however, does not carry over across root-suffix boundaries, with more than one aspirated stop allowed to occur in these contexts (8):

(8) No laryngeal restriction across root-suffix boundaries

- a. Root: /b^hr/ ‘bear’ → 2nd pl. active present: [bib^hrt^ha] *[bibr^ha]
- b. Root: /d^ha:/ ‘put’ → 2nd dual middle present: [dad^ha:t^he] *[dada:t^he]

This contrasts with Ofo (MacEachern 1999) which has a similar laryngeal restriction on roots but in this case this restriction carries over to heteromorphemic sequences (9):

(9) /osk^ha + a^ha/ → [oska^ha] ‘the white or American egret’ (*osk^ha^ha)

Derived-environment blocking cases thus present one example in which static phonotactic generalizations and alternations pull apart, specifically where phonotactic generalizations are observed in the lexicon but these do not motivate alternations thus permitting certain structures only in underived environments.

2.2.2 Non-isomorphism between static generalization and active process

An intriguing example that Paster describes comes from English Closed Syllable Shortening (Myers 1987:CSS). Unlike in derived-environment blocking, here a generalization is true both

tautomorphemically and heteromorphemically, more or less. The exact characterization of each generalization, however, differs depending on whether or not it is within a root or across the morpheme boundary. Within a root, long vowels cannot occur before consonant clusters, unless *the cluster is coronal*. Across a morpheme boundary, on the other hand, long vowels cannot occur before *any* consonant clusters.

- (10) Static phonotactic constraint: No long vowel before CC, *unless coronal* (examples from (Paster 2013:87))
- a. *wild* [wâɪld] (**wilb* *[wâɪlb]) (coronal ✓, analogous labial ✗)
 - b. *saint* [sent] (**saimp* *[semp]) (coronal ✓, analogous labial ✗)
- (11) Alternation: No long vowels before *any* CC (examples from (Paster 2013:87))
- a. *convene* [kʰanvin]- *convention* [kʰanvɛnʃən] (coronal ✗)
 - b. *retain* [ritʰen] - *retentive* [ritʰɛntɪv] (coronal ✗)
 - c. *describe* [dɪskraɪb] - *descriptive* [dɪskrɪptɪv] (labial ✗)

In these cases, the phonotactic constraint motivating the alternation are more general in nature (any consonant cluster), whereas the phonotactic constraint governing roots only holds in a subset of contexts. Thus although the generalizations in both domains are similar, one would likely have to state this generalization using two different constraints.

2.2.3 Derived-environment effects: active alternation with no phonotactic support

In addition to the cases discussed above, prominent examples in which phonotactic generalizations and alternations pull apart from each other are cases which exhibit active alternations across morpheme boundaries but do not exhibit concomitant static generalizations within the lexicon. These patterns are known in the literature as derived-environment effects (DEE; also known as non-derived environment blocking, Kiparsky 1993, see also Wolf 2008 and Burzio 2011). These patterns will be the focus on this dissertation. One example that is often cited as

the paradigmatic case of a derived-environment effect is Finnish assibilation (Kiparsky 1973, 1993, Anttila 2006). In (12), stem-final /t/ assibilates to [s] before [i]:

- (12) Finnish assibilation: /t-i/ → [s-i]
- a. /halut-i/ → [halus-i] ‘want-3_{PSG.PRET}’ (*haluti)
 - b. cf. /halut-a/ → [halut-a] ‘want-INF’

While heteromorphemic /t-i/ sequences surface as [s-i], tautomorphemic /ti/ sequences are protected and surface faithfully as in (13). In fact, in (13b), we notice that there are two underlying sequences of /ti/, but only the heteromorphemic sequence assibilates whereas the stem-internal one does not, further suggesting that the generalization seems sensitive to morphological (or prosodic) structure.

- (13) /ti/ protected within stems in Finnish
- a. /aeiti/ → [aeiti] ‘mother’ (*aeisi)
 - b. /tilat-i/ → [tilasi] ‘order-3_{PSG.PRET}’ (*tilati, *silasi).
 - c. cf. /tilat-a/ → [tilata] ‘order-INF’ (*silata)

Interestingly, assibilation in Finnish is also fed by /e/-raising that applies to a set of /e/-final nouns in (14) (Kiparsky 1993). This case contrasts with those in (13) since the nominative form for ‘mother’ /aeiti/ does not show assibilation. Further, unlike the morphologically complex forms in (12), (14) is an example of a *phonologically* derived environment effect, since it seems that the prior application of /e/-raising allows for the assibilation rule to apply. For current purposes, we will not be discussing phonologically derived environment effects, and will concern ourselves primarily with morphologically derived environment effects. For a recent overview of both morphologically and phonologically derived-environment effects, I direct readers to Chapter 8 in Inkelas (2014).

- (14) /e/-raising feeds assibilation

- a. /vete/ → [veti] → [vesi] ‘water (NOM)’

Another well-known example is /t/-palatalization in Korean (Kiparsky 1993, Iverson & Wheeler 1988, Oh 1995, T. Cho 2001). Across a morpheme boundary, underlying /t, t^h/ palatalize to [c, c^h] before /i/ and /j/ as seen in (15). However, analogous monomorphemic forms in (16) do not palatalize. Here, there is a mismatch between tautomorphemic static generalizations ([ti] is allowed) and the heteromorphemic dynamic generalization ([t-i] is repaired by palatalization).

(15) Derived palatalization

- a. /mat-i/ → [maci] ‘eldest-NOM’¹
 b. /pat^h-i/ → [pac^hi] ‘field-NOM’

(16) Underived blocking of palatalization

- a. /mati/ → [mati] ‘knot, field’
 b. /titi-ta/ → [titita] ‘to tread’

Turkish velar deletion (Zimmer & Abbott 1978, Inkelas & Orgun 1995, Inkelas, Orgun, & Zoll 1997, Inkelas 2000) is yet another example. Turkish has an active process that deletes velar stops intervocalically. This occurs when vowel-initial suffixes are attached to the stem as in (17). Velar stops do occur intervocalically, however, as in (18). Interestingly, the forms in (18) contain two potential sites in which velar deletion can occur, but only the velar stop at the morpheme boundary undergoes deletion.

- (17) a. [bebek] ‘baby’ → [bebe-i] ‘baby-DAT’
 b. [katalog] ‘catalogue’ → [katalo-u] ‘catalogue-3sg.POSS’

- (18) a. [sokak] ‘street’ → [soka-a] ‘street-DAT’

¹In reality, stops are voiced intervocalically, so the surface form of /mat-i/ is actually [maji]. Since this does not have a bearing on the discussion of palatalization, I will use the voiceless counterparts in each of these cases for simplicity.

- b. [dakik] ‘precise’ → [daki-i] ‘precise-ACC’

Furthermore, mirroring the derived-environment blocking vowel harmony pattern in Western Armenian (§2.1), Turkish presents another case of derived-environment effects involving vowel harmony (Clements & Sezer 1982). In contrast the derived-environment blocking, here there is an active alternation enforcing harmony across a morpheme boundary (19) but stems themselves can be disharmonic (20). Heteromorphemic vowel harmony is enforced on disharmonic stems by agreeing with the final vowel of the stem.

(19) Turkish vowel harmony across morpheme boundaries

- a. /ip-lAr/² → [ip-ler] ‘rope-NOM.PL’
b. /sap-lAr/ → [sap-lar] ‘stalk-NOM.PL’

(20) Disharmonic stems in Turkish:

- a. takvim ‘calendar’
b. bobin ‘spool’
c. hesap ‘bank account’
d. silah ‘weapon’

In all of these cases so far, there is ostensibly an active alternation that is not supported by a phonotactic generalization in the lexicon. Malagasy shows an extreme case of this where the generalization enforced heteromorphemically is diametrically opposite of what is the statistical trend in the lexicon. In general, Malagasy exhibits backness dissimilation in vowels heteromorphemically. The examples in (21) show this alternation where the passive imperative suffix /-u/ alternates to [-i] if there is a back vowel in the stem (Zymet 2014).³

²Capital vowels indicate vowels that undergo harmony.

³The alternation is actually somewhat gradient based on distance and is blocked by an intervening front vowel (Zymet 2014)

- (21) Dissimilation locally and non-locally
- a. /babu-u/ → [babu-i] ‘plunder’
 - b. /tuv-u/ → [tuv-i] ‘fulfill’
 - c. /undan-u/ → [undan-i] ‘bolster’

Given this active alternation, we might expect that there should be a phonotactic generalization in the lexicon that favors disharmonic roots. In fact, as Zymet (2014) points out, there is actually the opposite, albeit slight, preference for harmonic sequences. Malagasy thus evinces an extreme case of derived-environment effect. Here the phonological alternation is not merely phonotactically unmotivated, there is an active preference for the opposite pattern. For further examples of derived-environment effects, I direct the reader to Chapter 8 in Wolf (2008).

Derived-environment patterns have proven to be a thorny problem for theoretical phonology precisely because of these mismatches in generalization, as first pointed out by Kiparsky (1973). Since then, a wealth of proposals have been put forth in the theoretical literature to account for these patterns. Some of these restrict where rules can apply, such as in the Revised Alternation Condition, which states that obligatory neutralization rules can only apply in a derived environment (Kiparsky 1973, 1982b). In Korean, for example, this means that coronal palatalization, a neutralizing rule, can only apply in a derived-environment (see also Ahn 1985, Iverson 1987), and thus stem-internal non-derived sequences are protected. This is also largely in the same spirit of Wolf’s (2008) proposal which interleaves morphological and phonological operations, with derived-environment patterns applying only if preceded by a morphological operation (i.e. affixation), or McCarthy’s (2003)’s Comparative Markedness proposal. Łubowicz’s (2002) conjoined constraint account shares similar intuitions as well. These proposals basically distinguish between “underlying” structures and new structures that are created due to morphological operations. In a similar vein, Inkelas & Orgun (1995) and Oh (1995) argue for an account in which stem-final segments which alternate (as in Turkish velar deletion or Korean coronal palatalization) are left unsyllabified and therefore are “visible” to the derived-environment rule, whereas stem-internal sequences are not. This prosodification

account argues essentially that stem-internal and stem-final sequences have different prosodic status and the fact that the latter alternates but the former does not falls out from this difference.

A second class of proposals deals with representational differences between tautomorphemic and heteromorphemic sequences. Kiparsky (1993) for example proposed that stem-final segments (such as in Korean coronal palatalization) are underspecified featurally and subject to different structure-building rules compared to stem-internal segments. T. Cho (2001) and Bradley (2007) provide a similar take but from the perspective of gestural timing, where the timing relation between two gestures created by morpheme concatenation are not specified and are thus more variable, hence the alternation due to gestural overlap. More recently, Inkelas (2015) has proposed that alternating segments are more weakly represented than non-alternating invariant segments within stems, resulting in two different faithfulness constraints: $\text{FAITH}_{strong} \gg \text{FAITH}_{weak}$. Inkelas proposes then that derived-environment effects result from the ranking of FAITH_{strong} above Markedness.

All in all, these numerous proposals share in common the notion that tautomorphemic sequences are in some way different from heteromorphemic sequences, either because they are subject to different kinds of process application or because they are represented differently. Implicit in these analyses is also the assumption that because they look similar structurally (i.e. active alternations without a concomitant phonotactic generalization in the lexicon), they form a class and are amenable to the same theoretical treatment. Recent work, however, has pointed out that some canonical cases of derived-environment effects are likely not as clear cut insofar as the alternation is less general than is predicted by most proposals (e.g. Finnish assibilation: Anttila 2006). I return to this question in Chapter 4 with an in-depth investigation of two well-known cases. What has often also been ignored in the theoretical literature is the question of how these phonological patterns are learnt and how productive these patterns actually are. Moreover, derived-environment effects provide a useful tool to investigate how phonotactics and alternations interact in phonological learning, as we shall in see in Chapter 3. In the next section, we first use toy simulations to illustrate why learning of alternations in these cases is problematic.

2.3 Learning mismatches: exploring possible outcomes

In enumerating different cases of mismatches in phonotactics and alternations, Paster (2013) argues that this is evidence that these generalizations can evolve independently of each other over time, therefore the close relationship between both types of knowledge is merely incidental. A more extreme view is that static generalizations over the lexicon play no role in the explanation of phonological alternations (e.g. Hale & Reiss 2008) or that they do not constitute phonological knowledge at all. Hale & Reiss (2008) take the particularly strong view point that accounting for static generalizations in the lexicon is to merely “state descriptive generalizations about the memorized content of the lexicon of a particular language. Even if we, as linguists, do find generalizations in our description of the lexicon, there is no reason to posit these generalizations as part of the speaker’s knowledge of their language, since they are computationally inert and thus irrelevant to the input-output mappings that the grammar is responsible for.” (Hale & Reiss 2008:17-18).

What these viewpoints ignore, however, is how exactly these generalizations are learnt when they mismatch. The implicit assumption is that a learner is able to arrive at the correct phonological generalization about phonotactics as well as alternations, even when the mismatch. Is this the case? In this section, we begin investigating this question by presenting toy simulations of learning in these mismatch cases. We will concentrate on three cases: (1) across-the-board patterns (phonotactics and alternations accord), (2) derived-environment effects (active alternation, no phonotactic generalization) and (3) derived-environment blocking (no alternation, but active phonotactic generalization in lexicon). These mismatch patterns present an interesting avenue into investigating the relationship between phonotactics and alternations in learning. In particular, if we are interested in alternation learning, the comparison between across-the-board patterns and derived-environment effects, where there are active alternations, should prove particularly informative, since the sole difference here is whether or not the static phonotactic generalizations accord with the alternation, thus providing a controlled way of examining the problem.

What are the predictions of the different models of the effect of phonotactics on alternation learning? In modular accounts, like rule-based phonology, a lack of a static generalization, specifically a mismatch, should have no impact on the learnability of a phonological alternation. This is not to say that static generalizations are not part of the grammar, although Hale & Reiss (2008) seem to argue for this, just that these two types of phonological knowledge are separately encoded and do not interact. In constraint-based models of phonological learning, it is hypothesized that phonotactic learning precedes, and so facilitates alternation learning (Hayes 2004, Prince & Tesar 2004, Tesar & Prince 2007, Jarosz 2006, Hayes & Wilson 2008). This predicts that derived-environment alternations should be more difficult to learn than an across-the-board pattern. Both types of models then have clearly distinct predictions as to the outcome of alternation learning, which we will explore in the rest of this dissertation. In the next section, I present the results of a series of toy simulations that illustrate what phonotactic generalizations are available at the outset of alternation learning given the different types of language patterns surveyed above.

2.3.1 What does the initial stages of alternation learning look like?

The general OT assumption that phonotactics precedes alternation learning does not actually completely answer the question of how exactly these two levels of generalization interact. Nor does it, on its own, answer the question of what the initial stages of alternation learning should look like given an acquired phonotactic grammar. Different types of constraint learning models predict differences in the way in which markedness constraints are ranked in the initial stages of phonotactic learning. Here, I examine the outcome of a number of toy simulations using two learning algorithms that utilize constraints: Constraint Demotion (Tesar 1997, Tesar & Smolensky 1998) and Maximum Entropy models (Hayes & Wilson 2008). I discuss these in turn in the following sections. Unfortunately, a learning algorithm for rule-based frameworks is not available to my knowledge.

Before describing each of these models, however, I first describe the simplifying assumptions made in the following simulations. It is also necessary at the outset to describe the different

choice points available for ranking (or weighting) constraints and what the content of the grammar is at the initial state of learning. In the simulations presented in this section, I adopt the following simplifying assumptions:

(22) *Simplifying assumptions:*

- a. Phonotactic learning is blind to morphological boundaries.
- b. Phonotactic learning treats surface forms as underlying forms.
- c. Phonotactic learning and alternation learning occur in separate stages.
- d. The outcome of pure phonotactic learning is the initial state for alternation learning.

These assumptions are necessarily simplifying and the real picture is most likely much more complex. For example, it is possibly the case that learners are able to do both phonotactic and alternation learning at the same time, as certain morphological boundaries become available. It suffices for the present purposes, however, to adopt these assumptions. Finally, I should emphasize that the goal of this section is not to show how one of the existing analyses (e.g. constraint conjunction) of derived environment effects could be learnt. Rather my goal is to show why phonotactic mismatches, assuming a simple set of constraints, are problematic for later phonological learning.

2.3.2 The learning data

The starting point for the learning data for the simulations below are three constructed toy cases that mimic natural language patterns. The first across-the-board case, shows a pattern in which a generalization is exceptionless both within and across a morpheme boundary. The second case is a derived-environment effect pattern, like in Korean, in which there is an active alternation. Finally, the third case, is a derived-environment blocking pattern, in which there is a static generalization without an active alternation. The summary of the type of observed forms given a Richness-of-the-Base-like input is given in Table 2.3.

	/..ti../	/..ci../	/..t+i/	/..c+i/
Across-the-board	[ci]		[c+i]	
Derived-environment effect	[ti]	[ci]	[c+i]	
Derived-environment blocking	[ci]		[t+i]	[c+i]

Table 2.3: Output forms given language type. Column headers are inputs assuming ROTB.

2.3.3 OT: Constraint demotion

The first model that I will be exploring is Recursive Constraint Demotion (CD; Tesar 1997, Tesar & Smolensky 1998). The input to CD is a set of *paired* underlying (UR) and surface representations (SR). In the initial stage of phonotactic learning, it is assumed that surface representations are mapped onto themselves via the identity map, such that the UR and SR are equivalent. For each input form, there is a set of rival outputs. In CD, it is assumed that the grammar has access to both markedness and faithfulness (Input-Output: IO-F) constraints from the outset of phonotactic learning (see Tesar & Smolensky 1998, Hayes 2004, Prince & Tesar 2004), together with the relevant violation profiles for each input-output pair.

CD evaluates input and output pairs. Specifically, CD generates a grammar that is consistent with a set of elementary ranking arguments: a competition between two candidates from a given input, where one is optimal and the other not. Each candidate is associated with their own constraint-violation profile. The optimal candidate is deemed the *winner* and the other the *loser*. The output of CD is a *stratified hierarchy* where within each strata, constraints do not conflict with each other. The goal of the learner then is to rank the constraints in such a way that all known ranking arguments derived from the learning data are satisfied simultaneously.

The algorithm can be summarized as follows (borrowing heavily from the summary in Hayes 2004: 169):

- (23)
- a. Start by collecting all *winner-only* preferring constraints (i.e. never violated by a winner) and placing these in the highest stratum (i.e. top of the hierarchy).
 - b. Candidates that violate constraints in the newly established stratum more times than the winner does are considered ‘explained’ and are taken out from the learning set.
 - c. Of the remaining unranked constraints, find those that do not prefer any losers in the remaining subset data and place them in the next highest stratum.
 - d. Repeat (b) and (c) until all constraints have been assigned to strata.

Both Hayes (2004) and Prince & Tesar (2004) argue for an additional bias for maintaining markedness constraints high in the hierarchy (see Gnanadesikan 2004). Put another way, the learner is biased against faithfulness constraints, wanting by default to rank faithfulness low. If a markedness constraint does not prefer any losers or winners (so it’s not active), it will nonetheless be placed in the highest stratum, precisely because it prefers no losers in the learning data. This is an important difference with a Maximum entropy model which we will come back to below. Further, at the outset of alternation learning, it is often argued by many that it is also assumed that Output-Output Faithfulness constraints (OO-F) are ranked above markedness constraints (Hayes 2004, Tessier 2012). One therefore has a “default ranking” of constraints as follows: OO-F \gg Markedness \gg IO-F.

The algorithm comes in at least three versions: Regular constraint demotion (CD), Biased constraint demotion (BCD: Prince & Tesar 2004) and Low-faithfulness constraint demotion (LCD: Hayes 2004). All three are implemented in OTSoft (Hayes, Tesar, & Zuraw 2013). The main difference between these models is what is done with the faithfulness constraints during ranking. Both BCD and LCD are biased to favor ranking markedness above faithfulness during the learning process. Regular CD does not have this bias. For the purposes of our current simulation goals, this difference does not matter and the results from each of these algorithms are the same. All the outcomes are reported from the BCD model.

I adopt the following structure-blind markedness constraints banning [ti] and banning [ci] in (24). I assume here that any constraints that are sensitive to morphological structure (e.g. *T+I that penalizes sequences that span a morphological boundary (see Martin 2011)) are unavailable during a stage of pure phonotactic learning.

- (24) Markedness constraints:
- a. *TI - No /-ti-/ sequences (structure-blind)
 - b. *CI - No /-ci-/ sequences (structure-blind)

Finally where necessary, the following general faithfulness constraint is assumed. We will set aside the issue of Output-Output faithfulness for current purposes.

- (25) Faithfulness constraints
- a. IDENT(T) - Do not change [t] (or [c])

To summarize, for each of the two learning algorithms examined, simulations were run with training from toy data that mimics either an Across-the-board, Derived-environment effect or a Derived-environment blocking language. In each case, the output of the initial stages of learning (constraint ranking or weights) were then used to predict output forms in the initial stages of alternation learning.

In the next subsections, I report the outcomes of the phonotactic learning stages for each of the language types. For each of these cases, I report what was fed into the model and what the learnt constraint ranking was after phonotactic learning. I then report what is predicted once alternation learning begins and morphological boundaries are available. All simulations were run using OTSoft (Hayes et al. 2013).

2.3.3.1 Across-the-board language

The BCD learner was asked to rank the following input in Table 2.4 in which [ti] sequences are never observed.

			*ti	*ci	Ident (t)
[ci]	[ti]	0	1		1
	[ci]	1		1	
[c+i]	[t+i]	0	1		1
	[c+i]	1		1	

Table 2.4: Across-the-board input

Given the input above, the learner learns the following constraint hierarchy in (26)

(26) Outcome of phonotactic learning: *TI \gg *CI \gg IDENT

Since the generalization is exceptionless, the difference between whether or not a structure-sensitive constraint is active in this stage of learning matters not. The relevant markedness constraints penalizing [ti] sequences remain high in the hierarchy. Note here that because the learner is biased to demote faithfulness constraints below markedness constraints, *CI is ranked above IDENT. When this hierarchy is applied to a novel /ti/ or /t+i/ sequence, the predicted output candidates are both palatalized [ci] and [c+i] respectively. Thus, the right output for alternations is trivially predicted which is unsurprising given the exceptionless pattern in the training data.

2.3.3.2 Derived-environment effect language

The learner was then trained on a derived-environment effect language in Table. 2.5. In this language, [ti] sequences are not found in across a morpheme boundary but are found within stems. Given the input above, the learner learns the following constraint hierarchy:

			*ti	*ci	Ident (t)
[ti]	[ti]	1	1		
	[ci]	0		1	1
[ci]	[ti]	0	1		1
	[ci]	1		1	
[c+i]	[t+i]	0	1		1
	[c+i]	1		1	

Table 2.5: Derived-environment effect input

(27) Outcome of phonotactic learning: IDENT \gg *TI,*CI

Since the learner does encounter [ti] in the input, it demotes the relevant constraint which penalizes this sequence. Thus when a learner encounters a form that requires alternations, it will predict that [t+i] will not alternate, and that stem-internal [ti] is well-formed, since IDENT outranks the relevant structure-blind markedness constraint. Thus it does not capture the fact that heteromorphic [t+i] should be repaired to [c+i].

2.3.3.3 Derived-environment blocking language

Finally, the learner is trained on the a derived-environment blocking language. In this language, [ti] sequences are not found in stems but are found across a morpheme boundary (Table. 2.6).

			*ti	*ci	Ident (t)
[ci]	[ti]	0	1		1
	[ci]	1		1	
[t+i]	[t+i]	1	1		
	[c+i]	0		1	1
[c+i]	[t+i]	0	1		1
	[c+i]	1		1	

Table 2.6: Derived-environment blocking input

Given the input above, the learner learns the following constraint hierarchy:

(28) Outcome of phonotactic learning: IDENT \gg *TI,*CI

The first thing to notice is that the constraint hierarchy with this language are exactly the same as those in the derived-environment effect case in the previous section. This is not surprising given our assumptions that the learner does not have access to morphological boundaries at this stage. Therefore, a morphologically complex form like /t+i/ and a monomorphemic form like /ti/ have the same violation profiles, despite having different morphological composition. While the learner successfully predicts that there should be no alternations, it fails to penalize stem-internal [ti] sequences, exhibiting the opposite problem from the Derived-environment effect language shown in the previous section. In both cases, the learner just learns that there must be a contrast between [ci] and [ti] somewhere in the word.

2.3.3.4 BCD: Summary

In this section, I have briefly shown how a constraint demotion algorithm would handle mismatch cases of phonotactics and alternations. Unsurprisingly, the learner learnt a constraint ranking that successfully captures the across-the-board generalization in a language in which [ti] sequences are never attested. Interestingly, since the phonotactic learner is blind to morphological structure during phonotactic learning, the constraint hierarchy learnt from both a derived-environment effect or derived-environment blocking language is the same. In both cases, the relevant markedness constraint, the structure-blind general *TI, is ranked below IDENT. The problem is that the learner cannot distinguish between where the [ti] sequences appear. Thus it would predict that, after the phonotactic learning stage, the static phonotactic pattern is not learnt in a derived-environment blocking language, and that the alternation is not learnt in the derived-environment effect language. Thus the generalizations are more difficult to learn in both of these cases compared to the across-the-board pattern.

2.3.4 Maximum entropy grammars

A Maximum entropy model (MaxEnt) refers to a type of stochastic model that has been used in the last decade or so, starting with Goldwater & Johnson (2003), to model phonological grammars. Since Goldwater & Johnson (2003), MaxEnt models have been used in a number of studies, for example in Hayes & Wilson's (2008) study of phonotactic learning (for other examples using MaxEnt see also Wilson 2006, Jäger 2007, Martin 2011, Pater, Staubs, Jesney, & Smith 2012, Pater & Moreton 2012, Hayes & White 2013, a.o.)

For any given input x , the MaxEnt model assigns a probability to each potential output candidate, y , using the formula in (29).

(29)

$$Pr(y|x) = \frac{\exp(-\sum_{i=1}^m w_i C_i(y, x))}{Z}$$
$$\text{where } Z = \sum_{y \in Y(x)} \exp(-\sum_{i=1}^m w_i C_i(y, x))$$

The weights can be seen as representing the scaling of importance of a given constraint *relative* to others in the grammar, with higher weights having a stronger role in decreasing the probability of a given candidate output. Put simply, a *penalty score* (Hayes & Wilson 2008) is calculated for each candidate by first multiplying, for each constraint, its weight w_i by the number of times a given input/output pair violates the constraint, $C_i(y, x)$, then summing over all the constraints in the grammar (i.e. $\sum_{i=1}^m w_i C_i(y, x)$). A *Maxent value* (Hayes & Wilson 2008) is then calculated by negating the score then raising the base of the natural logarithm e to the result. This gets you the numerator in the above equation (i.e. $\exp(-\sum_{i=1}^m w_i C_i(y, x))$). Finally, the probability of the candidate is calculated by dividing the Maxent value by the the sum of MaxEnt values of all possible output candidates, Z , which is the denominator.

It is worth noting that the implementation of MaxEnt, being a species of OT, has an assumed GEN component that generates possible output candidates of a given input. The nature of the constraints in CON can also be assumed. The difference, then, lies in the nature of EVAL.

While in classical OT, as in CD, evaluation was based on strict ranking of constraints, predicting categorical outcomes, MaxEnt generates a *probability distribution* over all possible output candidates, allowing for the total probability to be unevenly divided across different candidates. This particular property allows for the modeling of variable and gradient patterns in phonology (Jäger 2007, Hayes et al. 2009, Hayes & Moore-Cantwell 2011, Daland et al. 2011, Hayes, Wilson, & Shisko 2012, Hayes & White 2013, Kager & Pater 2012, Moore-Cantwell & Pater 2017).

One particularly attractive feature of MaxEnt models is that they are associated with a learning algorithm (Berger, Pietra, & Pietra 1996) that provably converges on the “best” grammar, given the data and constraints. The learner seeks to maximize the (log) probability of the *observed* data (30), which minimizes the probability of the *unobserved* data. Importantly, the calculation is sensitive to the number of observed tokens of input/output pairs during training. I direct the reader to §4.2.1 in Hayes & Wilson (2008) for a more thorough discussion of this. For our current purposes, it suffices to say that a MaxEnt model will be sensitive to the frequency distribution of forms in the input. Finally, the objective function in MaxEnt also often includes a “prior” which is subtracted from the probability of the data. This prior term, in (31), is a Gaussian distribution over each constraint weight, with two free parameters: the mean μ and standard deviation σ .

(30)

$$\log Pr(D) = \prod_{j=1}^n Pr(y_j|x_j)$$

(31)

$$\sum_{i=1}^m \frac{(w_i - \mu_i)^2}{2\sigma_i^2}$$

(32)

$$[\log \sum_{i=1}^m Pr(y_j|x_j)] - [\sum_{i=1}^m \frac{(w_i - \mu_i)^2}{2\sigma_i^2}]$$

The μ value for each constraint is the *a priori* preferred weight. This is subtracted from the learnt weight and the difference is then squared. So as the constraints vary more from their μ value, the penalty imposed by the prior increases. The σ^2 term determines the degree to which deviations from a constraint's μ value is penalized. Being in the denominator, a high value of σ^2 decreases the value of the prior, thereby allowing for more freedom to deviate from μ . A small value of σ^2 does the opposite and imposes a greater penalty for deviations. Thus with the priors, the goal of learning in a MaxEnt learner is to arrive the set of weights for a given set of constraints that maximizes the objective function given in (32), where the prior term (31) is subtracted from the log probability of the observed data (30). Together, the penalty imposed on the model increases the more the learnt constraint weights diverge from their *a priori* values. The crucial difference between CD and MaxEnt in this regard is that constraints which are not active are not assigned any weight, i.e. the weight stays at μ .

The simulations presented in the following sections are implemented in the MaxEnt Grammar Tool.⁴ As an initial assumption about the contents of the phonotactic grammar, I adopt the stance taken by Hayes & Wilson (2008) that the initial phonotactic stage of learning only has access to markedness constraints and not faithfulness constraints. The goal here is to just access the likelihood of surface forms. This is already a major point of difference between MaxEnt and CD, which requires the existence of IO-F constraints. The following parameters μ and σ are set at 0 and 1000 respectively for all the relevant markedness constraints, and will be kept constant for all simulations. We will also use the same markedness constraints from (24). In this case, the learner does not need to be furnished with faithfulness constraints, unlike in the BCD model.

For each language type, we ran a model using the same set of constraints as we used in the BCD learning model previously. For current purposes, we are primarily interested in the weight assigned to the markedness constraint after initial phonotactic learning, and we will set

⁴Software package developed by Colin Wilson with interface by Benjamin George. This is available for public use on Bruce Hayes's website at <http://www.linguistics.ucla.edu/people/hayes/MaxentGrammarTool/>.

aside the issue of Output-Output faithfulness. The goal is to examine to what extent the initial learning of the weight of the markedness constraint will allow for learning of an alternation.

For each language type, I present the input that the learner was trained on. I then report the weights learnt in each simulation and the predicted results in the alternation stage of learning. The probability of forms in the alternation stage were calculated in Excel using the learnt weights from the initial phonotactic stage of learning.

2.3.4.1 Across-the-board language

We will assume that the input contains an equal ratio of sequences which occur within a morpheme and across a morpheme boundary, and that this language only contains [ci] or [ti] sequences. The trained and predicted ratios for the Across-the-board language is shown in Table 2.7. In this language, there are no [ti] sequences, thus the ban against [ti] is a general one and is observed both tautomorphemically and heteromorphemically.

	Training proportion of total	Predicted proportion of total
[t+i]	0	0
[c+i]	0.50	0.50
[ti]	0	0
[ci]	0.50	0.50

Table 2.7: Across-the-board language

	*TI	*CI
Weights	4.67	0

Table 2.8: Weights learnt after phonotactic learning for Across-the-board language

The learner successfully matches the ratio of sequences in the training data (Table 2.7). The learnt weights of the relevant markedness constraints are given in Table 2.8. In this case, in the

alternation stage, the learner arrives at the correct generalizations about both tautomorphemic and heteromorphemic forms easily by weighing $*T(+)$ I highly (Table. 2.9). This is unsurprising given that the generalization is across-the-board.

		$*T(+)$ I wt = 4.67	$*C(+)$ I wt = 0	IDENT-F wt = 0	\exp^{-H}	Probability
/ti/	[ci]	0	1	1	1	0.99
	[ti]	1	0	0	0.009	<0.01
/t+i/	[c+i]	0	1	1	1	0.99
	[t+i]	1	0	0	0.009	<0.01

Table 2.9: Predicted alternations for structure-blind learner: Across-the-board language

The learner predicts successful learning of the alternation (>99%). It also correctly predicts that stem-internal sequences are categorically ill-formed. Thus as expected, in an across-the-board language, phonotactic learning successfully sets the learner up for alternation learning.

2.3.4.2 Derived-environment effect language

We will now consider what phonotactic learning might look like in a derived-environment effect language. The learner was trained on an input similar to the one used for the across-the-board simulation. In this case, stem-internal [ti] and [ci] sequences were equally represented, but across the morpheme boundary only [c+i] is attested (Table 2.10).

	Training proportion of total	Predicted proportion of total
[ti]	0.25	0.13
[ci]	0.25	0.37
[t+i]	0	0.13
[c+i]	0.50	0.37

Table 2.10: DEE language

	*TI	*C(+I)
Weights	1.09	0

Table 2.11: Weights learnt after phonotactic learning for DEE language

A learner that is blind to the morphological structure of the input misses the generalization that [ti] only occurs within stems and not across morphological boundaries. All the learner sees is that 75% of word forms contain [ci] and the remainder 25% contain [ti]. Because of this, the learner still assigns some weight to the markedness constraint *TI. When this weight is applied to alternations, alternations are predicted 75% of the time, and for stem-internal /ti/ there is still a preference for [ci], when there should be no such preference, since [ti] should be completely well-formed.

		*T(+I) wt = 1.09	*C(+I) wt = 0	IDENT-F wt = 0	\exp^{-H}	Probability
/ti/	[ci]	0	1	1	1	0.75
	[ti]	1	0	0	0.34	0.25
/t+i/	[c+i]	0	1	1	1	0.75
	[t+i]	1	0	0	0.34	0.25

Table 2.12: Predicted alternations for structure-blind learner: DEE language

2.3.4.3 Derived-environment blocking language

In the last simulation, we consider how learning in a derived-environment blocking language might proceed. The input to the derived-environment blocking language is shown in Table 2.13. This is the mirror image of a derived-environment effect language. Here [ti] sequences do not appear within a morpheme but are allowed to appear across a morpheme boundary. The predicted ratios of word forms is also shown in Table 2.13, with the learnt constraint weights shown in Table 2.14. Learning here looks exactly the same as in the derived-environment language, both in terms of weight and the predictions for alternations (Tables ?? and 2.15). This

is because the learner does not differentiate between [t+i] and [ti] sequences. All it takes into account is the overall frequency of both any [ti] sequence in the input, and in our simulations, this is always 25%. Thus the outcome of learning in both a derived-environment effect and derived-environment blocking language turns out to be the same.

	Training proportion of total	Predicted proportion of total
[ti]	0	0.13
[ci]	0.50	0.37
[t+i]	0.25	0.13
[c+i]	0.25	0.37

Table 2.13: DEB language

	*TI	*CI
Weights	1.09	0

Table 2.14: Weights learnt after phonotactic learning for DEB language

		*TI wt = 1.09	*CI wt = 0	IDENT-F wt = 0	\exp^{-H}	Probability
/ti/	[ci]	0	1	1	1	0.75
	[ti]	1	0	0	0.34	0.25
/t+i/	[c+i]	0	1	1	1	0.75
	[t+i]	1	0	0	0.34	0.25

Table 2.15: Predicted alternations for structure-blind learner: DEB language

2.3.4.4 MaxEnt: Summary

The toy simulations presented in the previous sections provide another basic confirmation of the intuition that phonological patterns are easier to learn in an across-the-board language. The

MaxEnt learner successfully learns the tautomorphemic and heteromorphemic generalizations in this case. We also found that the learning outcomes were the same regardless of whether the learner was trained on a derived-environment effect or derived-environment blocking language, due to the fact that the learner has no initial access to morphological boundaries. In either case, the learner predicts an overall preference for [ci] forms over [ti] forms since there are more [ci] instances in the input (75% vs. 25% [ti]) regardless of morphological structure. Thus in comparing between the derived-environment effect and across-the-board languages, the simulations suggest that alternations in the derived-environment effect language are learnable, although they are more difficult, as indicated by the lower predicted rate of alternations. At the same time, the learner fails to learn that tautomorphemic [ti] is completely well-formed.

The toy simulations also show that alternations are wrongly predicted in the derived-environment blocking language, where there should not be any alternations. But this might not matter in the end since if Output-Output faithfulness is heavily weighted initially (Hayes 2004, Do 2013, Tessier 2012), the learner will be biased to not posit any alternations anyway. And if the learner does not see any alternations in the ambient language, OO-Faith will remain heavily weighted blocking alternations from occurring. Thus it is possible that the a derived-environment blocking pattern will be more learnable than a derived-environment effect pattern.

2.3.5 Summary: Preliminary modeling

In this section, I investigated the behavior of two constraint learning algorithms given three types of phonological patterns in which the relation between tautomorphemic and heteromorphemic generalizations differ. The goal was to examine how phonotactic learning differs in each of these cases under the assumption that the learner does not have access to word-internal morphological boundaries. There are a number of similarities that these two models share. Both models learn an across-the-board pattern equally well. Both models basically distinguish between an across-the-board language on the one hand, and derived-environment effect and derived-environment blocking languages on the other. The main result here is that across-the-board languages are predicted to be more successfully learnt from the outset than the other two which contain

mismatches. Importantly, both models predict that the output of learning whether the input language is a derived-environment effect and derived-environment blocking type, will be the same, since the learner has no initial access to morphological boundaries.

In BCD, once a [ti] sequence is encountered, the constraint penalizing is immediately demoted, predicting that these sequences will be entirely well-formed. MaxEnt, however, is sensitive to the relative frequency of forms in the input. In our simulations, since although the model learnt that [ti] sequences are not completely ill-formed, it still encoded the generalization that overall [ci] sequences are more frequent relative to [ti] sequences. Thus even in a derived-environment language, a MaxEnt model might nonetheless be able to encode a gradient penalty against [ti] sequences which might then aid later alternation learning due to the over-representation of [ci] relative to [ti]. This is not possible with a BCD model since it is not stochastic.

Thus in this section, I have sketched out why patterns with a mismatch in phonotactics and alternations might be difficult to learn. In the next chapter, I present the results of a series of experiments that provide evidence for this conjecture.

CHAPTER 3

Learning derived-environment patterns: experimental investigations

3.1 Introduction

In the study of phonology, we are often concerned with accounting for phonological alternations: when a particular morpheme is pronounced in one way in one context but in a different way in a different context. In English for example, the past tense suffix *-ed* (/ -d/) is pronounced as a [d] when suffixed to a verb like *lag* (/laeg-d/ → [laegd]). But when suffixed to a verb like *lack* /laek/ the suffix is pronounced as a [t], agreeing in voicing with the preceding stem-final [k] (/laek-d/ → [laekt]). Thus, there is an alternation between [t] and [d] to ensure that heteromorphemic instances of [kd] never surface. Of course, native speakers of English understand that these two categorically distinct forms are instances of the same morpheme, and tacitly know what context supports the occurrence of each of these phonological variants (e.g. Gaskell & Marslen-Wilson 1996). Faced with variation of this sort, a child's task must necessarily involve learning the alternations in their native language. What might aid in the learning of this?

A growing body of work in the last decade has examined the question of how alternation learning proceeds. Many of these insights have been gained by using artificial grammar learning experiments with adults (e.g. Wilson 2006) and infants (K. S. White et al. 2008, J. White & Sundara 2014), as well as through computational modeling (Peperkamp, Calvez, Nadal, & Dupoux 2006, Calamaro & Jarosz 2015). A useful starting point is for learners to track the distribution of sounds to look for complementary distribution between two different segments

(Peperkamp et al. 2006, Calamaro & Jarosz 2015). In fact, K. S. White et al. (2008) showed that 12-month-old can rely on distributional information in an artificial language to learn an alternation.

Previous work has also investigated different kinds of learning biases that might drive learning of phonological alternations. One such bias is what is known as a “substantive” bias which proposes that learners are guided by perceptual, articulatory and general phonetic knowledge when learning phonological alternations (Wilson 2006, J. White 2014, Baer-Henney & van de Vijver 2012, Finley & Badecker 2012), such that learners are biased to learn patterns which accord with phonetic naturalness. Additionally, learners in artificial grammar learning experiments have been shown to preferentially learn formally simple phonological patterns, in particular, in terms of the number of relevant features that are involved with stating the phonological generalization (Saffran & Thiessen 2003, Moreton 2008, see Moreton & Pater 2012a for a review).

In this chapter, I test whether learners knowledge of phonotactics (speakers’ tacit knowledge of the licit and illicit sequences of sounds that can occur a language) can also facilitate the learning of alternations. Taking an example again from English, *act* [ækt] is a word, but a word like *akd* [ækd] is not a possible word since the [k] and [d] in hypothetical *ækd* are illegal sequences. So if a child hears monomorphemic forms like [ækt] but never hears [ækd], she might come to the conclusion that [kd] sequences are illegal in her language. Thus first learning static phonotactic generalizations about monomorphemic forms will then aid in the learning of alternations in multimorphemic forms. When faced with a novel morphologically complex word that has a [kd] sequence due to morpheme concatenation, the child will already know that this sequence has to be changed. The fact that phonotactics and phonological alternations are consistent with each other has been observed as early as in Chomsky & Halle (1968:382). In fact, in constraint-based theoretical phonology frameworks such as Optimality Theory (OT; Prince & Smolensky 1993/2004), both generalizations are captured using a single mechanism. Such models capture the generalization about the close relationship between phonotactics and alternations by ensuring that both types of phonological knowledge are governed by the

same constraint. Going back to the example above, a single constraint against [kd] sequences is responsible for both ensuring that words like *akd* are not attested and that the past tense morpheme /-d/ surfaces as [t] when following a voiceless [k] in *lacked*. This is in contrast to rule-based models of phonology (Chomsky & Halle 1968) in which both of these were *distinct* parts in the phonological grammar.

Capitalizing on this close relationship, recent computational models of alternation learning assume that phonotactic learning occurs prior to alternation learning and that learning the former aids in learning the latter (Hayes 2004, Hayes & Wilson 2008, Prince & Tesar 2004, Tesar & Prince 2007, Jarosz 2006). The timeline in infant development of these two types of phonological knowledge offers some support for this hypothesis: phonotactic knowledge emerges first around 9 months of age (Jusczyk et al. 1994, Saffran & Thiessen 2003, Friederici & Wessels 1993), while alternation learning has only been demonstrated a few months later at 12 months (K. S. White et al. 2008, J. White & Sundara 2014).

Although this is a widely held assumption, experimental evidence in support this assumption is still unclear. In the current study, we examine whether phonotactic learning facilitates the learning of alternations using an artificial grammar learning paradigm. To do so, we compare the learning of alternations in two groups of learners who differ in whether or not they are exposed to a supporting phonotactic generalization within stems. In the remainder of this section, I first summarize the results of previous experimental investigations that have examined the relation between phonotactics and alternations in learning.

3.1.1 Previous experimental investigations: phonotactics and alternations

Typically experimental studies on phonological learning have primarily focused on either the learning of phonotactic knowledge alone (Skoruppa & Peperkamp 2011, Linzen & Gallagher 2014, Saffran & Thiessen 2003, Chambers et al. 2003, Chambers, Onishi, & Fisher 2010, 2011, Onishi et al. 2002, Richtsmeier 2011) or the learning of phonological alternations alone (Wilson 2006, Kapatsinski 2010, Cristià, Mielke, Daland, & Peperkamp 2013, J. White 2014). In both

sets of studies, the relationship between these two levels of generalization is ignored. Only three studies have directly probed the link between phonotactic and alternation learning with largely inconclusive results.

Pater & Tessier (2005) trained American English speakers on an alternation ([t]-epenthesis) that was motivated by one of two phonotactic generalizations. One set of participants were trained on an epenthesis rule that inserted a [t] when a lax vowel occurred finally (e.g. /bli/ →[blɪt] but /bli/ →[bli]; Language 1). Here the alternation serves to enforce a phonotactic generalization that is internalized by English speakers (Moreton 1999): lax vowels do not occur in final open syllables. A second set of participants were trained on a language in which epenthesis occurred following front vowels but not back vowels (e.g. /li/ →[lit] but /fu/ →[fu]; Language 2), a generalization that is not supported in the English lexicon. Conforming to Pater & Tessier's (2005) prediction, learners learnt the alternation better in the language with phonotactic support ([t]-epenthesis following lax vowels) than the one without ([t]-epenthesis following front vowels). However, the authors themselves point out that, while both languages are of equivalent formal complexity, Language 2 is also typologically unnatural and unattested. Given that previous studies have shown that learners show a dispreference against unnatural patterns (Hayes et al. 2009, Becker et al. 2011, Hayes & White 2013), it is therefore possible that the poorer performance in Language 2 could be explained by this alone. Thus although Pater & Tessier's (2005) results are consistent with a link between phonotactic and alternation learning, they do not provide unequivocal evidence in support of it.

Two recent investigations by Pizzo (2015) and Chong (2016) have also failed to find conclusive evidence for the link. Using a more controlled design, Chong (2016) trained participants on an artificial language in which coronal stops (/t/ and /d/) palatalized to [tʃ] and [dʒ] across a morpheme boundary before /i/. This palatalization rule is based on Korean (Ahn 1985). In the Match language, [ti] and [di] sequences did not appear within stems, ensuring a match between stem phonotactics and alternations. In the Mismatch language, [ti] and [di] sequences did appear within stems mismatching with the alternation. Contrary to expectations, participants in both languages learnt the alternation equally well. However,

learners in the Match language did not infer a phonotactic constraint against [ti] or [di], despite these sequences never appearing in their training data. Thus we cannot be sure whether phonotactics aids in alternation learning since there was no evidence of phonotactic learning in Chong's learning experiment. What Chong's study suggests though is that alternation-based knowledge does not have a strong impact on phonotactic knowledge, a finding echoed by Pizzo (2015). Using a poverty-of-stimulus design, Pizzo (2015) trained participants on an alternation that involving either voicing assimilation (e.g. *nemab+fa* → *nemapfa*) or place assimilation (e.g. *lobon+fa* → *lobomfa*) in consonant clusters. In the training data, there were no consonant clusters within stems thus the evidence in the lexicon was ambiguous as to the nature of the phonotactic generalization. In the test phase, participants were given a pair of novel stems with stem-internal consonant clusters and had to decide which word belonged to the language they had just learnt (e.g. voicing assimilation: *madfas* vs. *matfas*). For example, if participants were trained on voicing assimilation, when faced with a pair like *madfas* vs. *matfas*, they should prefer the latter since this conforms to the alternation pattern. Pizzo was interested in seeing if participants extended their alternation knowledge to static phonotactics. Participants only showed significant generalization from alternations to phonotactics when there was an intermediate feedback stage where they were given explicit feedback on the task. In an implicit learning task, Pizzo found no clear effect, although there was a numerical trend in the predicted direction. Taken together, both Pizzo and Chong's findings suggest that the relationship between both types of phonological knowledge, if it exists at all, might be unidirectional.

The current study addresses the shortcomings of previous studies by building on the experimental design of Chong (2016). In order to probe the effect of phonotactic learning on alternation learning, we need to compare two equally natural phonological patterns, with both the alternation and phonotactic generalization being learnable in the laboratory setting. In this paper, we manipulate the learning of a pattern of vowel harmony. Vowel harmony is a phonological pattern in which vowel sequences have to agree in particular phonological features (e.g. backness). Vowel harmony has many properties that make it a good process to test in an artificial language experiment with English speakers. It is not a phonological pattern active

in English, so we can control for first language phonotactic knowledge. Moreover, previous studies using artificial grammar learning experiments have shown that learners, with a short amount of exposure, are able to learn these harmony patterns and generalize to unseen words (Finley & Badecker 2009, 2012, Pycha et al. 2003). Importantly for the current study, Skoruppa & Peperkamp (2011) showed that participants are able to successfully learn a static phonotactic generalization regarding well-formedness of words in an implicit learning task with relatively short exposure. Finally, vowel harmony has been implicated in mismatches between stem phonotactics and alternations that occur naturally across the world's languages. Turkish, for example, shows a vowel harmony alternation across a morpheme boundary, whereby adjacent vowel sequences alternate to agree in backness with the final vowel of the root (Clements & Sezer 1982, Lewis 1967) as shown in (1), although roots themselves can either be harmonic or disharmonic as in (2).

(1) Turkish vowel harmony across morpheme boundaries (data from Clements & Sezer 1982: 216)

- a. /ip-lAr/ ¹ → [ip-ler] 'rope-NOM.PL'
- b. /sap-lAr/ → [sap-lar] 'stalk-NOM.PL'
- c. /son-In/² → [son-un] 'village-GEN.SG'
- d. /jyz-In/ → [jyz-yn] 'face-GEN.SG'

(2) Turkish vowel harmony and disharmony in stems (data from Crothers & Shibatani 1980: 64)

- a. /sekiz/ 'eight'
- b. /oda/ 'room'
- c. /mezat/ 'auction'
- d. /kitap/ 'book'

¹The vowels in upper case indicate vowels in the suffix that harmonize to the vowel in the root.

²The surface form of the vowel 'I' in the suffix is derived by both rounding and backness harmony

Thus vowel harmony is not only learnable in a laboratory setting in terms of both active phonological alternations as well as static phonotactic generalizations, some languages with vowel harmony do show mismatches between both of these generalizations. In this paper, we test whether alternation learning of a vowel harmony process is facilitated when by a matching phonotactic generalization.

3.2 Experiment 1: Full language

Experiment 1 was designed to compare alternation learning when there is a match or mismatch in phonotactic generalizations within stems. Participants in this experiment were randomly assigned to one of three artificial languages involving vowel harmony: Harmonic, Semi-Harmonic and Non-Harmonic. In all three languages, there was an exceptionless alternation pattern in which the vowel in the plural suffix [-mu~mi] alternated to agree in backness with the final vowel of the singular stem. Thus participants in all three languages were trained on the same alternation with the same amount of evidence. Where the languages differed was in how much stem phonotactic support there was for the alternation. In the Harmonic language, vowels in all stems always agreed in backness (e.g. ['pime] but *['pimo]), supporting the alternation pattern across the morpheme boundary. In the Non-Harmonic language, vowels in half the stems agreed in backness (e.g. ['pime]), whereas the other half did not (e.g. ['pimo]), resulting in a mismatch between the alternation pattern and stem phonotactics. In the Semi-Harmonic language, vowels in 3/4 of the stems in agreed in backness whereas 1/4 did not.

Based on Skoruppa & Peperkamp's (2011) findings, we expected Harmonic language learners to learn the phonotactic pattern successfully. Because there is no such generalization available for Non-Harmonic language learners, they should infer no phonotactic generalizations regarding harmony. Finally, in the Semi-Harmonic language, we expected the learning of the phonotactic constraint to be better than in the Non-Harmonic language, but worse than in the Harmonic language group, given that learners seem to be sensitive to gradient statistical generalizations

in their input language (Frisch & Zawaydeh 2001, Hayes et al. 2009, Becker et al. 2011, Hayes 2000).

What are the predictions for alternation learning? Learners in each language receive the same amount of evidence for the alternation. If phonotactic learning facilitates alternation learning, we expected that Non-Harmonic language learners will, despite the mismatch in phonotactics, successfully learn the alternation pattern, and we should see no differences across groups. If, however, phonotactic learning facilitates alternation learning then we expected a similar result as with phonotactic learning: Harmonic language learners will learn the alternation the best, Non-Harmonic learners the worst, and Semi-Harmonic language learners in between.

3.2.1 Methods

3.2.1.1 Participants

45 American English participants were recruited from the UCLA Psychology Pool. Participants were randomly put into one of the three artificial language groups (15 in each). 25 more were tested but were excluded due to having the wrong first language background ($n = 4$), knowing a language with vowel harmony (i.e. Armenian; $n = 1$), not completing the experiment ($n = 7$), recognizing the vowel harmony pattern ($n = 2$), taking notes ($n = 3$) and issues with playing the sound files ($n = 8$).

3.2.1.2 Procedure

Participants were tested over the Internet using Experigen (Becker & Levine 2014), and were told that they were going to learn words from a foreign language. They were asked to pay attention to what they were hearing but were told that they did not have to memorize any of the words. On each trial in the training phase, either singular or plural word forms were presented auditorily accompanied by a corresponding image. That is, singular and plural word forms were not presented side-by-side on the same trial. This ensured that the task did not overtly provide a means of comparison between singular and plural forms. There were 3 training blocks

with 64 trials each resulting in a total of 192 training trials. In each phase of the experiment, participants were able to hear a particular stimulus item just once.

We then probed participants' knowledge of stem phonotactics (blick test; Scholes 1966) using a two-alternative forced-choice task in the first test phase as in Skoruppa & Peperkamp (2011). Participants heard two novel CVCV words (one harmonic and one non-harmonic, e.g. ['gike] vs. ['giko]) and had to decide which belonged to the language they had learned. The order of presentation of the two CVCV options was randomized such that harmonic words and non-harmonic words occurred equally as the first member of the pair.

Finally, we tested participants' knowledge of the phonological alternation in a wug test (Berko 1958). Participants heard a novel CVCV singular stem with an accompanying image, then heard two possible forms for the plural ([-mi] vs. [-mu]) and had to pick the correct word. The order of presentation of each possible plural form was counterbalanced such that each plural form occurred equally as the first member of the pair.

3.2.1.3 Artificial languages

Three artificial languages were constructed that consisted of bisyllabic CVCV singular stems, along the lines of the artificial languages in Finley & Badecker (2009). CVCV singular stems were constructed using consonants {p, b, t, d, k, g, m, n} and vowels {i, e, u, o}. The plural was marked with a suffix that had two allomorphs [-mu] or [-mi] that agreed with the backness/roundness specification of the final vowel of the stem. The allomorph [-mu] appeared when the final vowel of the root was back/rounded [u, o], and the allomorph [-mi] appeared when the final vowel of the root was front/unrounded [i, e]. Across all three languages, the plural suffix always harmonized with the final vowel, with stems occurring equally frequently with the [-mu] allomorph and the [-mi] allomorph (half each). Thus all three languages had the same amount of evidence for the alternation (100%).

Where the languages differed was in the proportion of harmonic stems. In the Harmonic language, all singular stems contained vowels that were harmonic for backness/roundness (e.g.

[ˈbuno] but *[ˈpume]). All harmonic V-V sequences occurred equally frequently. Contrastively, in the Non-Harmonic language, half the stems contained harmonic vowel sequences (e.g. [ˈbuno]) and half non-harmonic (e.g. [ˈpume]), yielding a mismatch between phonotactics and alternations. All V-V sequences, both harmonic and non-harmonic, occurred equally frequently. This means that the total number of harmonic V-V sequences was half that in the Harmonic language. Finally, in the Semi-Harmonic language, three-quarters of stems contained harmonic vowel sequences, with the remaining quarter of stems containing non-harmonic vowel sequences. Each possible non-harmonic V-V sequences occurred just once. Non-Harmonic stems were created by changing one of the vowels in a harmonic stem, thereby ensuring that as much as possible was kept constant across all the languages. In total, 32 CVCV stems were created for each language. A summary of the proportion of harmonic and non-harmonic stems across all three languages is shown in Table 3.1, and a full list of training items can be found in the Appendix.

	Mismatch		Match
	‘Non-Harmonic’	‘Semi-Harmonic’	‘Harmonic’
	Alternation		
No. of harmonizing stems	32 (100%)	32 (100%)	32 (100%)
	Stem phonotactic		
No. of harmonic stems (e.g. [buno])	16 (50%)	24 (75%)	32 (100%)
No. of non-harmonic stems (e.g. [pume])	16 (50%)	8 (25%)	0 (0%)

Table 3.1: Summary of Artificial Languages

For the blick test, 16 pairs of novel test words were created using the same set of consonants and vowels as the training stimuli. Each pair of blick test words contained a harmonic word (e.g. [ˈgike]) and a non-harmonic word (e.g. [ˈgiko]). Non-Harmonic words were created by changing one of the vowels front/back specification while maintain the same height feature. Half the blick items differed in the first vowel, and the other half in the second vowel. Finally, for the wug test, 16 novel bisyllabic words were created in the same fashion except that only

harmonic words were used. The same novel test stimuli in both blick and wug tests were used with learners in all three language groups.

To confirm that the the constructed languages differed in terms of their inferrable stem internal phonotactic generalizations, the languages were fed through the UCLA Phonotactic Learner (Hayes & Wilson 2008). The learner was trained on the stimuli in the training phase in each language and tested on both the test items in the blick test. Based on the outcome of phonotactic learning, the learner assigns a penalty score to each test item, with a higher score indicating that a particular item is less well-formed. For each pair of harmonic and non-harmonic blick test items, a difference score was calculated by subtracting the penalty score assigned to the harmonic word from the score assigned to the non-harmonic word (Non-harmonic-Harmonic). A bigger difference score indicates a bigger preference for harmonic over non-harmonic blick test items. As expected, there was a much stronger preference for the harmonic word in the Harmonic language (Difference score = 8.76) than in both the Semi- (Difference score = 3.34) and Non-Harmonic languages (Difference score = 0.54). This derives from the fact that the learner trained on the Harmonic language successfully inferred two bigram constraints penalizing disharmonic vowel sequences, namely $*[+back][-back]$ and $*[-back][+back]$.

For completeness, we also report on the penalty scores assigned wug words. As with blick test items, a difference score was calculated by subtracting the penalty score assigned to the harmonic plural from the score assigned to the non-harmonic plural. The modeling simulations confirmed the intuition that the alternation is indeed learnable in all three languages with the learner learns a constraint that prefers harmonic plurals (e.g. [kobo-mu]) over non-harmonic plurals (e.g. [kobo-mi]). This is indicated by a large difference in penalty score between non-harmonic and harmonic wug words in all three languages (Harmonic: 7.13; Semi-Harmonic: 8.91; Non-Harmonic: 9.36). These simulations thus show that the languages do differ in terms of stem-internal bigram phonotactics. Yet, as intended, the evidence in favor of the harmonic plural (i.e. the alternation) is more or less equally available in the learning data across all three languages.

3.2.1.4 Audio stimuli

Audio stimuli were recorded by a female phonetically-trained speaker of American English who was naïve to the goal of the current study. Target words were always realized with declarative intonation with stress placed on the initial syllable of the target word, ensuring that stress was always on the same syllable in both singular and plural forms. Voiceless stops were always produced with aspiration and voiced stops were produced with voicing through the entire closure. Recordings were made using PCQuirer (Scicon R & D 2015) at a sampling rate of 22,050 Hz and were scaled to 70 dB.

3.2.1.5 Visual stimuli

Visual stimuli consisted of digital images of animals and everyday objects obtained freely over the Internet (260 X 200 pixels). Singular images always contained just one animal/object and plural images always contained two. An example pair of singular and plural images are shown in Figure 3.1.



(a) Singular



(b) Plural

Figure 3.1: Example pair of visual stimuli

3.2.2 Blick test: phonotactic generalizations

3.2.2.1 Predictions

We were interested in the nature of the phonotactic generalization learners arrived at after training. Specifically, we were interested in whether learners show a preference for harmonic over non-harmonic words. Given that harmonic and non-harmonic words occur with equal frequency in their lexicon, we expected Non-Harmonic language learners to show no preference for harmonic words. That is, they should not infer any phonotactic constraint from the learning data. Harmonic language learners, however, never hear non-harmonic words, and thus should show a strong preference for harmonic words over non-harmonic words. That is, Harmonic language learners should infer a phonotactic constraint against non-harmonic words. Finally, given that we know that speakers are sensitive to the statistics in the lexicon such that they show gradient well-formedness judgments given gradient phonotactic evidence, we expected Semi-Harmonic language learners to show a small preference for harmonic words over non-harmonic words, but not to the same extent as Harmonic language learners.

3.2.2.2 Results

Figure 3.2(a) shows the rate of choosing harmonic words in the blink test across the three language groups. We first analyze the rate of choosing harmonic words using mixed-effect logistic regression (Jaeger 2008) with Language (Harmonic, Non-Harmonic and Semi-Harmonic) as a categorical factor using the *glmer* function from the *lme4* package (Bates, Maechler, Bolker, & Walker 2015) in R (R Core Team 2015). The model also contained by-subject and by-item random intercepts, as well as random slopes for Language by item. We were interested in (a) whether participants' performance in each language group was significantly different from chance (i.e. is the intercept significant?), and (b) how the rate of choosing harmonic words differed between each language group. Pairwise comparisons were conducted using the *multcomp* package in R (Hothorn, Bretz, & Westfall 2008), and p-values were adjusted using Shaffer's correction for multiple comparisons (Shaffer 1995). Overall, Harmonic language

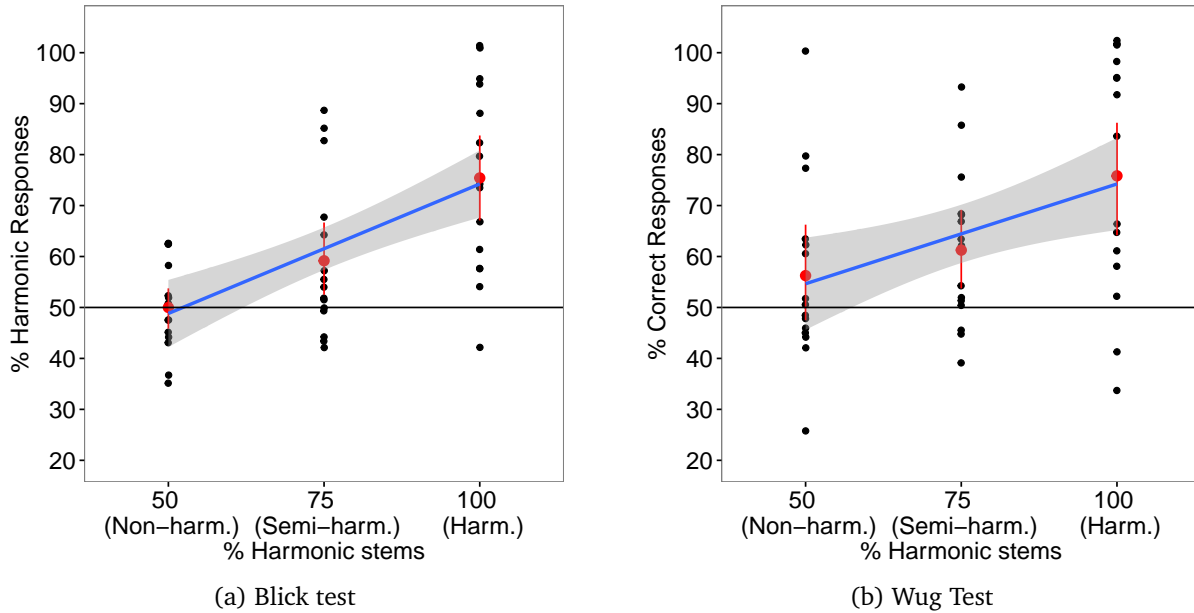


Figure 3.2: Experiment 1: Rate of choosing harmonic words in (a) blink test and (b) wug test. Each black dot represents a single participant with the large red-dot indicating mean rates for each language group. The red lines indicate 95% confidence intervals.

learners showed a significant preference for harmonic words ($\beta = 1.29$, $SE = 0.27$, $z = 4.76$, $p < 0.001$) whereas Non-Harmonic language learners selected harmonic words at chance ($\beta = -0.003$, $SE = 0.24$, $z = -0.01$, $p = 0.99$). Interestingly, Semi-Harmonic learners showed chance-level preferences for harmonic stems ($\beta = 0.43$, $SE = 0.27$, $z = 1.63$, $p = 0.31$), despite showing an overall numerical preference. Further, as predicted, Harmonic language learners chose harmonic words significantly more than both Non-Harmonic learners ($\beta = 1.30$, $SE = 0.28$, $z = 4.68$, $p < 0.001$) and Semi-Harmonic learners ($\beta = 0.86$, $SE = 0.36$, $z = 2.40$, $p = 0.03$). There was no significant difference between the rate of choosing harmonic words between Non-Harmonic and Semi-Harmonic learners ($\beta = 0.44$, $SE = 0.35$, $z = 1.23$, $p = 0.31$).

Although we failed to find a significant difference in preference for harmonic words between the Non- and Semi-Harmonic language learners, we were nonetheless interested in whether, on the whole, there was a significant relationship between the *proportion* of harmonic stems and the rate of choosing harmonic stems. That is, we were interested in whether there was a gradient

in the response preferences that corresponded with the amount of phonotactic evidence in the learning data. Another mixed effects logistic regression model was fit to the blick data with proportion of harmonic stems as a continuous variable. The model also contained by-subject and by-item random intercepts. This was the maximal model to converge. The rate of choosing harmonic words increased significantly as a function of the proportion of harmonic stems in the training language ($\beta = 0.024$, $SE = 0.005$, $z = 4.73$, $p < 0.001$). That is, the higher the proportion of harmonic stems in the language (in the Harmonic language) the more likely learners chose harmonic words over non-harmonic words in the blick test. Thus, overall, learners' rates of choosing harmonic plurals conformed to lexical statistics.

The results from the blick test confirm that the nature of phonotactic generalizations learnt from the training data differed given the differences in the lexical statistics in the input - the variable that was primarily manipulated. Given this, what kinds of generalizations did participants arrive at in terms of phonological alternations?

3.2.3 Wug test: phonological alternations

3.2.3.1 Predictions

Recall that in all three languages, there was consistent, exceptionless evidence for the phonological alternation across the morpheme boundary. So unlike in the blick test, the correct response in all three languages was the harmonic plural. For a singular like ['kobo], the correct plural should be ['kobomu] and not *['kobomi] in all three languages. If learning phonotactics does not influence alternation learning, we expect that learners in all three language groups should learn the alternation equally well. If learning phonotactics facilitates the learning of alternations, then we expected the strength of alternation learning to mirror that of phonotactic learning.

3.2.3.2 Results

The rate of choosing correct plurals in all three languages is shown in Figure 3.2(b). The rate of choosing the correct plurals (i.e. harmonic plurals) was modeled in the same manner as in the blick test. First, we analyzed the rate of choosing harmonic words using a mixed-effect logistic regression with Language (Harmonic, Non-Harmonic and Semi-Harmonic) as a categorical factor, with random intercepts by subject and item as well as an random slope for Language by item. Harmonic learners showed significantly above chance rates of choosing the correct plural ($\beta = 1.43$, $SE = 0.30$, $z = 4.78$, $p < 0.001$), indicating successful learning of the alternation, whereas Non-Harmonic learners did not learn the alternation and were at chance ($\beta = 0.52$, $SE = 0.27$, $z = 1.06$, $p = 0.29$). Semi-Harmonic learners showed at chance rates of choosing the correct plural ($\beta = 0.52$, $SE = 0.27$, $z = 1.94$, $p = 0.16$) as well. Pairwise comparisons revealed, unsurprisingly, that Harmonic learners learned the alternation significantly better than Non-Harmonic language learners ($\beta = 1.12$, $SE = 0.40$, $z = 2.76$, $p = 0.02$). Semi-Harmonic language learners' accuracy was significantly lower than Harmonic language learners ($\beta = 0.91$, $SE = 0.40$, $z = 2.28$, $p = 0.045$), and at the same time, not significantly different from Non-Harmonic language learners ($\beta = 0.21$, $SE = 0.40$, $z = 0.52$, $p = 0.60$). Thus Harmonic learners learnt the alternation, whereas those in the other two groups, on the whole, did not.

As in the blick test, we were also interested in examining whether the rate of choosing harmonic plurals was proportional to proportion of harmonic stems. The rate of choosing correct plurals was analyzed using a mixed-effects logistic regression model with proportion of harmonic stems in each language as a linear independent factor. The model also contained by-subject and by-item random intercepts. This was the maximal model to converge. Participants' performance in the wug test mirrored that of participants' in the blick test, with the rate of choosing harmonic plurals increasing significantly as a function of the proportion of harmonic stems in the training language ($\beta = 0.022$, $SE = 0.008$, $z = 2.82$, $p = 0.005$). Thus as in the wug test, there was gradient learning of the phonological generalization.

3.2.4 Correlation between phonotactics and alternations

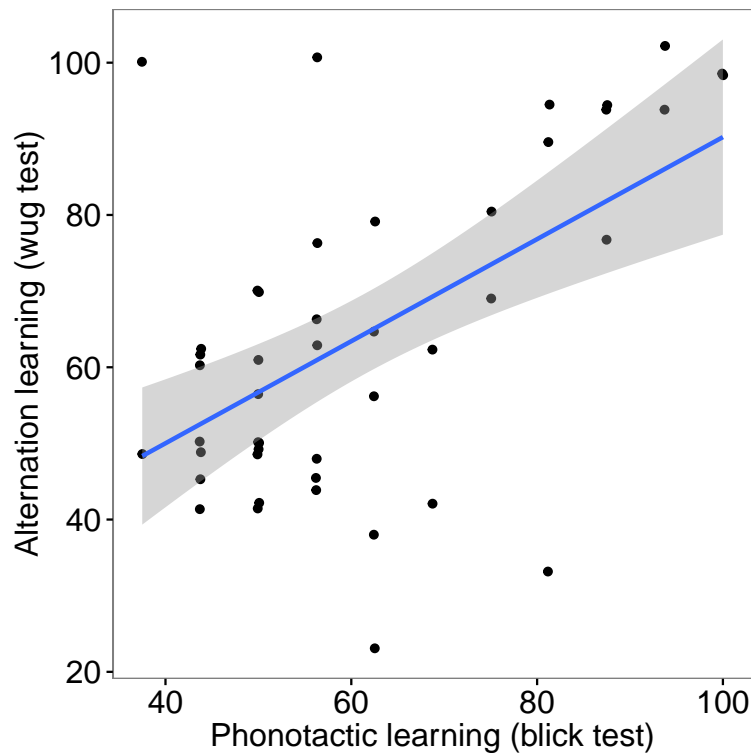


Figure 3.3: Experiment 1: Correlation of between performance in blick and wug tests

As seen in the previous sections, overall performance of participants in the wug test mirrored performance on the blick test. To further investigate the relationship between phonotactic and alternation learning, we inspected correlation between individual learner's performance on the blick test and their performance on the wug test. Figure 3.3 shows each learners' strength of alternation learning (wug test) as a function of phonotactic learning (blick test) collapsed across all three language groups. There was a significant positive correlation between the rate of choosing harmonic words in the blick test and the rate of choosing correct harmonic plurals ($R^2 = 0.32$, $r(43)=0.56$, $p < 0.001$), suggesting that even on an individual participant's level, alternation learning was correlated with phonotactic learning.

3.2.5 Discussion

The results of Experiment 1 show that alternations are more difficult to learn when stem phonotactic generalizations mismatch with the dynamic generalization. I also replicated the finding that participants show gradient phonotactic learning depending on the proportion of harmonic stems in the training lexicon (Frisch & Zawaydeh 2001, Hayes et al. 2009, Becker et al. 2011, Hayes 2000). Importantly, our results suggest that gradient learning of phonotactics leads to the gradient learning of alternations. Thus, our results provide clear evidence that matching phonotactics results in better learning of an alternation. In fact, participants in the Non-Harmonic language groups failed to learn the alternation despite there being exceptionless evidence for the alternation in the learning data. Performance by Semi-Harmonic learners was intermediate between both Harmonic and Non-Harmonic learners. Given that the main difference across the language groups was in terms of the proportion of harmonic stems, it seems that the alternation learning task is made more difficult if stem phonotactics do not support the dynamic generalization about alternations. More generally, on an individual level, the accuracy of alternation learning correlated with the degree to which each participant inferred a phonotactic constraint regarding harmonic stems, suggesting further that phonotactic learning aids alternation learning.

While we found a significant overall positive linear relation between proportion of harmonic stems and strength of both phonotactic and alternation learning, we failed to find a significant difference between the Non-Harmonic and Semi-Harmonic learners when these groups were directly compared. The difference in harmonic stem proportions between the Non- and Semi-Harmonic languages and the Semi-Harmonic and Harmonic languages were both 25%, with the Semi-Harmonic language being a mid-point between the two extremes. Thus the fact that there is a larger difference between the latter than the former pairs suggests that the degree to which exceptionality will affect the productivity of a particular phonotactic constraint or phonological alternation is not completely proportional to the amount of phonotactic evidence in the lexicon. While some recent work by Moore-Cantwell (2017) and Moore-Cantwell & Pater (2017) has investigated this question from a computational perspective, it is as yet still an open question

as to how exactly the degree of exceptionality affects productivity, an area that warrants further examination.

The basic finding in Experiment 1 was that participants were not successful at learning the alternation when the static phonotactic generalization in stems did not support the alternation. When exposed to a language in which both types of generalizations accorded with each other, participants successfully learnt the alternation. It is unclear what is *necessary* for this however. That is, what kind of information in the input facilitates the learning of an alternation? In the following series of experiments, we examine this question by varying what participants are exposed to in training. In Experiment 2, we were interested in ascertaining if morphological/semantic information was necessary for learners to learn the alternation.

3.3 Experiment 2: Distributional learning of full language

In Experiment 2, participants were only exposed to either the Harmonic or Non-Harmonic languages, since we were primarily concerned with the learning outcomes in the two languages which were most distinct from each other. Unlike in Experiment 1, in Experiment 2, participants were trained on singulars and plurals *without accompanying images*. We can conceive this type of learning as encouraging ‘pure phonotactic learning’ since participants are not given overt semantic associations that might facilitate a morphological parse of the stimuli. Of course it is possible that participants can still use transitional probabilities to segment the plural words into stem and suffix (Saffran, Aslin, & Newport 1996), but in the absence of accompanying images this would have to be implicit. Given that Finley & Badecker’s (2009) adopted a training paradigm that did not have accompanying images, and found successful learning of the harmony alternation, we expect participants exposed to the Harmonic language in Experiment 2 to successfully learn the alternation here as well. Note that the current task is more difficult than Finley & Badecker’s (2009) since in their experiment participants were trained on both the ‘unaffixed’ (e.g. [bidi]) and ‘affixed’ forms (e.g. [bidi-mi]) side-by-side on the same training trial. In our experiment these forms are randomly ordered such that a direct comparison of the

two forms is not available. Thus, if we obtain the same result here in a slightly more difficult learning task, we can be confident of the robustness of learning.

3.3.1 Methods

3.3.1.1 Participants

30 participants were recruited from the UCLA Psychology Pool. Participants were randomly put into one of the two artificial language groups (15 in each). 14 more were tested but were excluded due to wrong first language background ($n = 4$), knowing a language with vowel harmony (i.e. Armenian; $n = 1$), taking notes ($n = 4$), not completing ($n = 1$), one for having consistent issues playing the sound files ($n = 4$).

3.3.1.2 Stimuli

Only stimuli from the Harmonic and Non-Harmonic languages in Experiment 1 were used in Experiment 2.

3.3.1.3 Procedure

The general procedure was the same in Experiment 2 as in Experiment 1. There was a training phase followed by two test phases, blick and wug tests. In Experiment 2, participants were told that they were going to learn words from a new language, as in Experiment 1, but unlike in Experiment 1, participants were not shown images that accompanied each word during the training phase. Thus, participants did not have any semantic association and thus were not able to parse the training stimuli as singular or plural. Like in Experiment 1, there were 3 training blocks with a total of 192 training trials. After completing the training phase, participants moved on to the blick phase as in the previous experiment. The wug phase differed from Experiment 1, however. Since participants did not get any images in training, we were concerned that introducing images in the test phase would make the task more complex. Thus the 'wug' phase

in this Experiment 2 resembled the blink test in that participants just heard two possible words, corresponding to the two possible plurals, and had to decide which of those belonged to the language they had just learnt without any images.

3.3.2 Results - Blink test: phonotactic generalizations

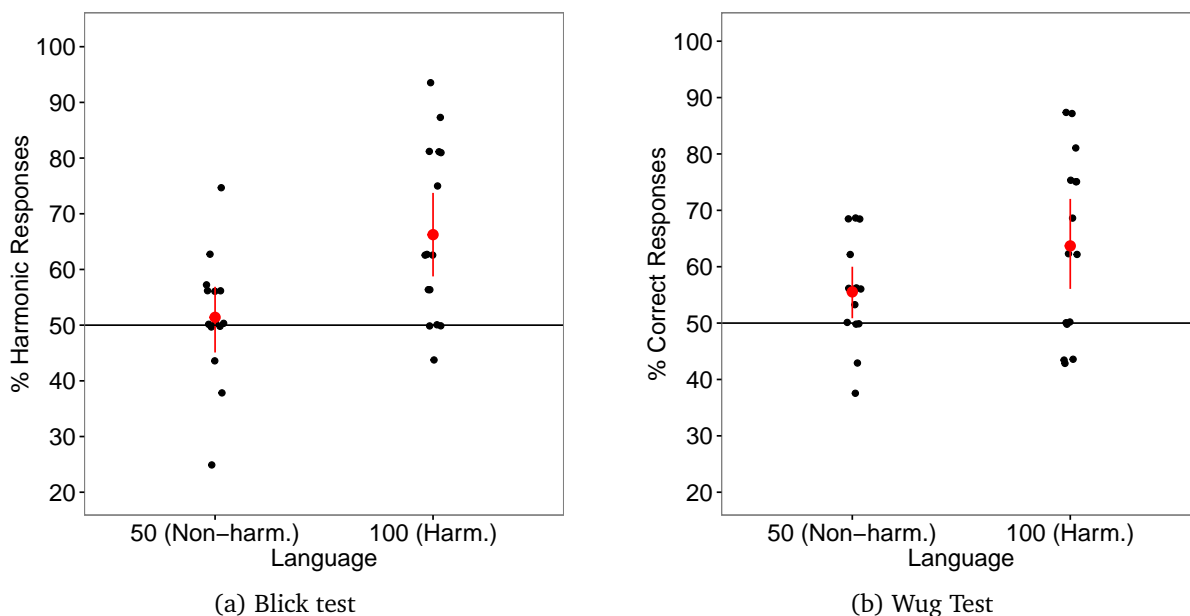


Figure 3.4: Experiment 2: Rate of choosing Semi-Harmonic harmonic words in (a) blink test and (b) wug test. Each black dot represents a single participant with the large red-dot indicating mean rates for each language group. The red lines indicate 95% confidence intervals.

The rate of choosing harmonic words in the blink test in Experiment 2 is shown in Figure 3.4(a). The rate of choosing harmonic words was analyzed using a mixed-effect logistic regression with Language (Harmonic and Non-Harmonic) as a categorical factor. The model also contained by-subject and by-item random intercepts, with Language as a random slope for item. As in Experiment 1, we were interested in ascertaining whether each language group's performance significantly differed from chance and each other, and we did this using the *multcomp* package which allowed for multiple simultaneous comparisons. Replicating the results in Experiment 1, Harmonic language chose harmonic words significantly above chance ($\beta = 0.73$, $SE = 0.21$, $z = 3.47$, $p = 0.002$), whereas Non-Harmonic language learners

were at chance ($\beta = -0.10$, $SE = 0.25$, $z = -0.38$, $p = 0.70$). Not surprisingly then, Harmonic language learners chose harmonic words at a significantly higher rate than Non-Harmonic language learners ($\beta = 0.83$, $SE = 0.25$, $z = 3.31$, $p = 0.002$). Harmonic language learners thus successfully learnt a phonotactic generalization that word forms had to contain harmonic vowel sequences. Contrastively, Non-Harmonic language learners, as expected, did not infer any vowel harmony constraint and accepted both harmonic and non-harmonic words at more or less equal rates.

3.3.3 Wug test: phonological alternations

3.3.3.1 Predictions

Given Finley & Badecker's (2009) result, as well as the results in Experiment 1, we expect that Harmonic language learners, despite not having semantic associations or a morphological parse in training, should nonetheless successfully learn the alternation, whereas those in the Non-Harmonic language group will fail to learn the alternation.

3.3.3.2 Results

The rate of choosing correct plurals in the wug test in Experiment 2 is shown in Figure 3.4(b). As in Experiment 1 and the analysis of blick test performance previously, the rate of choosing correct plurals was analyzed using a mixed-effect logistic regression with Language (Harmonic and Non-Harmonic) as a categorical factor. The model also contained by-subject and by-item random intercepts, with Language as a random slope for item. Mirroring again the blick test performance, Harmonic language learners chose correct plurals significantly more often than chance ($\beta = 0.62$, $SE = 0.17$, $z = 3.60$, $p < 0.001$), and at a marginally significantly higher rate than Non-Harmonic language learners ($\beta = 0.37$, $SE = 0.23$, $z = 1.64$, $p = 0.10$), who were at chance ($\beta = 0.25$, $SE = 0.15$, $z = 1.65$, $p = 0.10$). This indicates that, replicating Finley & Badecker's (2009) finding, that Harmonic learners successfully learnt the alternation in a

purely distributional learning training paradigm, without explicit access to a morphological parse or semantic information.

3.3.4 Discussion

The results in Experiment 2 replicated the results of Experiment 1. The phonotactic generalizations participants arrived at differed depending on the language they were trained on, with the learners trained on the Harmonic language inferring a vowel harmony constraint, whereas those trained on the Non-Harmonic language did not. The performance on alternation learning here also mirrored participants' performance in Experiment 1 with Harmonic language learners successfully learning the alternation and Non-Harmonic learners failing to do so, as indicated by their chance level performance on picking harmonic plurals. Recall that the sole difference between Experiments 1 and 2 is the fact that in Experiment 2, participants did not have explicit access to a morphological parse due to the absence of images in training. Nonetheless, participants picked the correct allomorph of the plural in the wug task. Thus, participants were still able to learn an alternation.

In Experiment 3, we extended this line of inquiry by examining whether learners are *biased* to have alternations reflect stem phonotactics without any exposure to the alternation at all. That is, do learners spontaneously expect alternations to reflect stem phonotactics in the absence of any evidence. If phonotactics and alternations are indeed encoded using a single mechanism, as in Optimality Theory, then we might expect participants to enforce a learnt phonotactic constraint on unseen potentially alternating forms. This would be the strongest evidence of a relationship between the two types of phonological knowledge.

3.4 Experiment 3: Stem-only training

In Experiment 3, using a poverty of stimulus design, we were interested in examining whether or not learners spontaneously expect alternations to reflect a learned phonotactic generalization. Participants were trained on the same languages as in Experiment 1. In Experiment 3,

however, learners were only presented with singular CVCV stems and did not get exposed to plurals in training. Thus, learners only had evidence for a static phonotactic generalization but did not get evidence regarding phonological alternations. If learners spontaneously expect alternations to reflect phonotactic generalizations, then we expect that learners should replicate the performance in the wug test of Experiments 1 and 2, learning the alternation successfully even without any evidence for this.

3.4.1 Methods

3.4.1.1 Participants

30 participants were recruited from the UCLA Psychology Pool. Participants were randomly put into one of the two artificial language groups (15 in each). 6 more were tested but were excluded due to having the wrong first language background ($n = 3$), not completing ($n = 1$), having consistent technical issues playing the sound files and for taking notes in training ($n = 1$).

3.4.1.2 Stimuli

Only stimuli from the Harmonic and Non-Harmonic languages in Experiment 1 were used in Experiment 3.

3.4.1.3 Procedure

The procedure in Experiment 3 was largely the same as in Experiment 1. In training, however, participants were only trained on bisyllabic singular stems (with accompanying images). Participants were never familiarized on trisyllabic plural words. In order to provide the same amount of learning data as in Experiment 1 and 2, there were 6 blocks of training instead of the previous 3, since initial piloting revealed that phonotactic learning did not occur on just 3 blocks of training. This ensured that participants were exposed to 192 training trials (as in

Experiments 1 and 2). The two test phases were the same as in Experiment 1. As in the wug test in previous experiments, participants were told that they were going to hear two words for a given image and that they had to pick the word they thought was correct for the language they had just learnt. Thus, in the wug test, learners had to generalize a learnt phonotactic generalization about bisyllabic singular stems to unseen alternations (trisyllabic plurals).

3.4.2 Results

3.4.2.1 Blick test: phonotactic generalizations

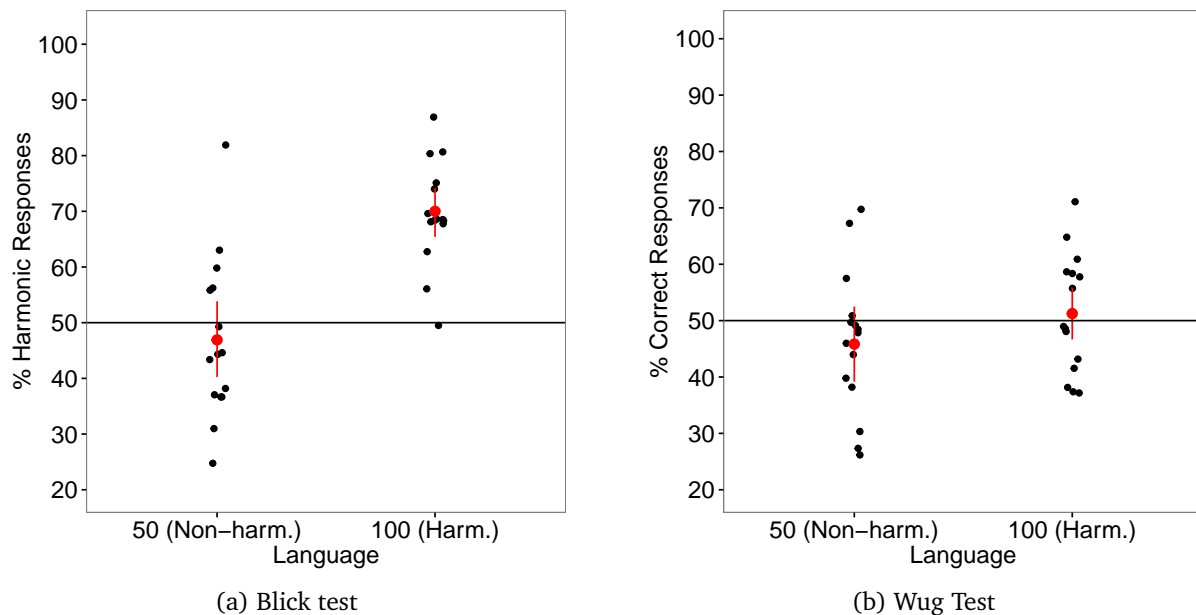


Figure 3.5: Experiment 3: Rate of choosing harmonic words in (a) blink test and (b) wug test. Each black dot represents a single participant with the large red-dot indicating mean rates for each language group. The red lines indicate 95% confidence intervals.

The rate of choosing harmonic words in the blink test in Experiment 3 was analyzed in the same way as in Experiment 2 (Figure 3.5(a)). Replicating the results of both Experiments 1 and 2, Harmonic language learners chose harmonic words significantly more often than chance ($\beta = 0.92$, $SE = 0.22$, $z = 4.18$, $p < 0.001$), and significantly more often than Non-Harmonic language learners ($\beta = 1.06$, $SE = 0.21$, $z = 5.13$, $p < 0.001$). As in both previous experiments, Harmonic language learners inferred a phonotactic preference for harmonic words whereas

Non-Harmonic language learners did not ($\beta = -0.14$, $SE = 0.19$, $z = -0.74$, $p = 0.46$), following the available evidence in the training data.

3.4.2.2 Wug test: phonological alternations

The rate of choosing correct plurals in the wug test in Experiment 3 was analyzed in the same ways in Experiment 2 (Figure 3.5(b)). Unlike in Experiments 1 and 2, learners in both languages chose correct plurals at chance (Harmonic: $\beta = 0.05$, $SE = 0.21$, $z = 0.25$, $p = 0.80$; Non-Harmonic: $\beta = -0.22$, $SE = 0.14$, $z = -1.53$, $p = 0.38$), and there was no significant difference in the rate of choosing correct plurals across both languages ($\beta = 0.27$, $SE = 0.21$, $z = 1.29$, $p = 0.38$). This indicates that despite inferring a phonotactic generalization preferring harmonic stems, Harmonic language learners nevertheless did not spontaneously extend this generalization to a novel morphological domain.

Given the lack of a difference in alternation learning between the two language groups, we were interested in ascertaining if participants showed a preference for one of the allomorphs over the other, as a strategy for responding in the wug test. The rate of choosing the [-mu] allomorph was calculated (Figure 3.6(a)) and analyzed in the same fashion as the preceding analyses. Participants in the Harmonic group showed a preference for the [-mu] allomorph as indicated by above chance rates of selecting it ($\beta = 0.73$, $SE = 0.31$, $z = 2.38$, $p = 0.03$), whereas those in the Non-Harmonic group showed no such preference ($\beta = 0.34$, $SE = 0.30$, $z = 1.14$, $p = 0.25$). Further, there was no correlation between their performance on the blick test and the wug test in the Harmonic language in which participants showed learning of the phonotactic constraint (Figure 3.6(b): $R^2 < 0.1$, $r(13) = -0.08$, $p = 0.78$).

3.4.3 Discussion

In Experiment 3, Harmonic language learners succeeded again in internalizing a phonotactic generalization favoring harmonic bisyllabic stems. However, they failed to extend a learnt

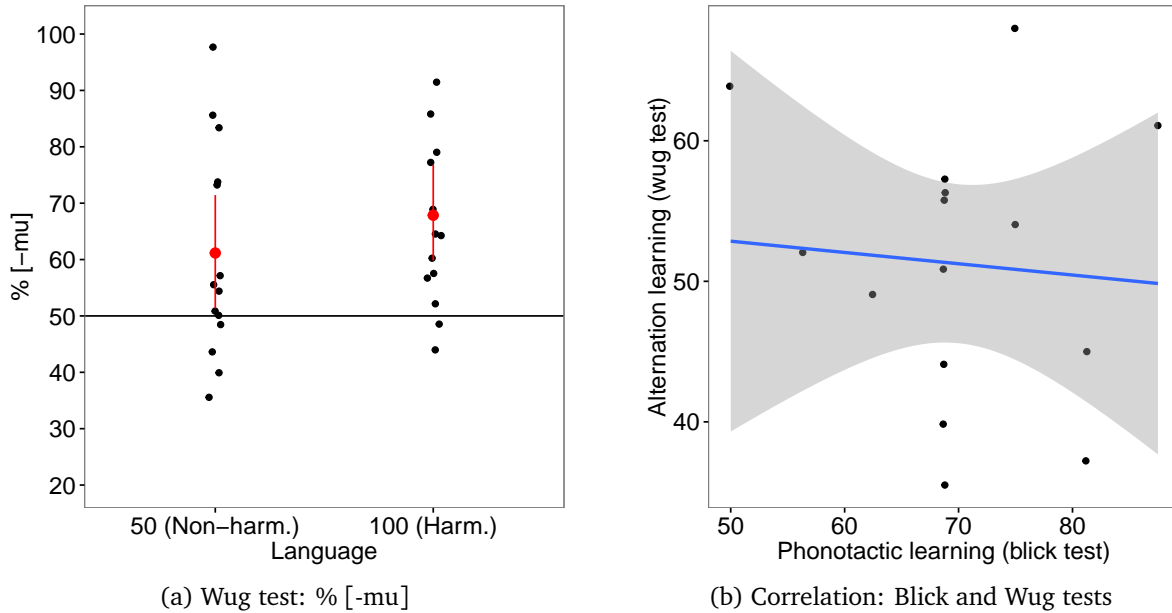


Figure 3.6: Experiment 3: (a) Rate of choosing [-mu] allomorph in Wug test. Each black dot represents a single participant with the large red-dot indicating mean rates for each language group. The red lines indicate 95% confidence intervals. (b) Correlation between phonotactic and alternation learning in Harmonic language.

phonotactic generalization to an unseen plural alternation, suggesting that learners are highly conservative in extending phonological generalizations to a novel morphological domain.

One reason that participants failed to extend the phonotactic to the alternation could be basically that these two types of phonological knowledge are separately acquired, and therefore exposure to alternations is required in order for an alternation to be learnt in the first place. But there remains one other alternative possibility, that might have to do with task effects. In Experiment 3, participants only heard bisyllabic stems in training. Then in the wug test, participants had to generalize to novel trisyllabic forms since all plurals are trisyllabic, as well as to a novel morphological context (i.e. plurals vs. singulars). Thus their failure to extend the generalization could be due to the novel trisyllabic form of the plurals. In Experiment 4, we tested this hypothesis.

3.5 Experiment 4: Trisyllabic stems

To address the possibility that learners need to have exposure to trisyllabic forms in order to generalize to unseen trisyllabic plurals, in Experiment 4, half of the bisyllabic stems in the training stimuli were modified to be trisyllabic by adding another CV syllable at the end. Thus learners were exposed to both bisyllabic and trisyllabic stems in training, and were then asked to extend the generalization to unseen trisyllabic plurals. So if learners succeed in generalizing from a stem phonotactic to a phonological alternation, when exposed to trisyllabic stems in training, it would suggest that learners do expect alternations to reflect phonotactics but require experience with the relevant word shapes (i.e. trisyllabic words). If, however, participants still fail to generalize to plurals, this would suggest that they require explicit evidence for an alternation in order to learn it, but that when evidence for alternations are available that learning this is aided if the phonotactics match the alternation.

3.5.1 Methods

3.5.1.1 Participants

30 participants were recruited from the UCLA Psychology Pool. Participants were randomly put into one of the two artificial language groups (15 in each). 9 more were tested but were excluded due to having the wrong first language background ($n = 5$), not completing ($n = 1$), and having consistent issues with playing the sound files ($n = 3$).

3.5.1.2 Stimuli

In order to create trisyllabic stems, extra CV sequences were generated from the original set of Cs {p, b, t, d, k, g, m, n} and Vs {i, e, u, o}, with each C and V occurring equally frequently in this position (each C occurred twice and each V occurred four times), yielding 16 novel CVs. These were concatenated with half of the stems (16 stems) in the training data used in Experiment 3 to create trisyllabic stems (e.g. old stem ['kete] + new CV [be] → new stem ['ketebe]). Half

of this set of stems were non-harmonic in the Non-Harmonic language, but harmonic in the Harmonic language. The other half were all harmonic in both languages. Thus the resulting training set, contained half bisyllabic and half trisyllabic stems. Only two languages were used for this experiment: Harmonic and Non-Harmonic. As in previous experiments, in the Harmonic language, all stems were harmonic. In the Non-Harmonic language, however, half stems were harmonic and half non-harmonic, split evenly across each stem type (bisyllabic and trisyllabic). A full set of stimuli items are shown in the Appendix. We used the same test stimuli as in Experiments 1-3, to maintain the ability to compare results across experiments. So learners only needed to decide between bisyllabic words in the blick test.

New trisyllabic stems were recorded by the same speaker used for the original stimuli in Experiment 1 using PCQuirer (Scicon R & D 2015) at a sampling rate of 22,050 Hz and were rescaled to 70 dB. As in previous experiments, stress was always placed on the initial syllable, voiceless stops were always aspirated and voiced stops always voiced throughout the closure.

As with the previous set of training data, the new training languages were fed into the UCLA Phonotactic Learner. Recall that larger differences in penalty scores between disharmonic and harmonic words indicates a larger preference for Harmonic over Disharmonic words. For the blick words, there was a clear preference for harmonic words in the Harmonic language (difference score = 9.78) but not in the Non-Harmonic language (difference score = 0.10). Furthermore, there was a preference for harmonic plurals in the Harmonic language (difference score = 7.71) but no such preference in the Non-Harmonic language (difference score = 0.82), unlike what we saw previously. This is different from the training data used in Experiments 1 and 2 since there are no alternating plurals in this training set. The learner thus confirmed the fact that the two languages differ in their phonotactic generalizations in support of harmonic words. Thus if learners choose to extend their learnt phonotactic generalizations to wug plurals, we expect them to show a preference for harmonic (i.e. correct) plurals. It is important to note here that the learner does not have access to morphological boundaries, thus when faced with a plural form it assesses the phonotactic legality of this form in the same way as a stem that just contains one morpheme.

3.5.1.3 Procedure

Training in Experiment 4 proceeded as in Experiment 3, with the sole difference that there were trisyllabic stems in training (e.g. ['ketebe]). Once training was completed, participants proceeded on to the two test phases, as in Experiments 1 and 3. Thus, as in Experiment 3, participants had to generalize a phonotactic generalization about singular stems to unseen plurals.

3.5.2 Results

3.5.2.1 Blick test: phonotactic generalizations

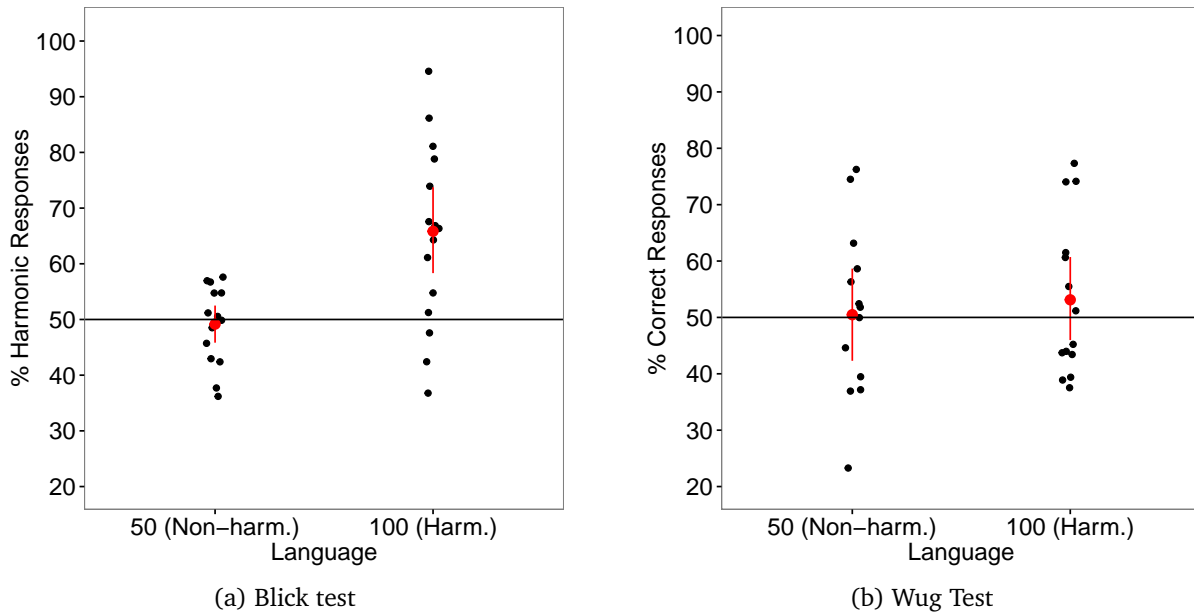


Figure 3.7: Experiment 4: Rate of choosing harmonic words in (a) blink test and (b) wug test. Each black dot represents a single participant with the large red-dot indicating mean rates for each language group. The red lines indicate 95% confidence intervals.

The rate of choosing harmonic words in the blink test in Experiment 4 was analyzed in the same way as in previous experiments (Figure 3.7(a)). Harmonic language learners chose harmonic words significantly above chance ($\beta = 0.68$, $SE = 0.17$, $z = 4.04$, $p < 0.001$), whereas Non-Harmonic language learners were at chance ($\beta = -0.03$, $SE = 0.17$, $z = -0.20$, $p =$

0.84). Importantly, Harmonic language learners showed a significantly stronger preference for harmonic words over non-harmonic words ($\beta = 0.71$, $SE = 0.19$, $z = 3.66$, $p < 0.001$) compared to Non-Harmonic language learners. This replicates results of Experiments 1-3 that show different phonotactic learning outcomes between both language groups depending on the presence of non-harmonic stems in the lexicon, but with training using trisyllabic as well as bisyllabic stems.

3.5.2.2 Wug test: phonological alternations

Participants' responses (Figure 3.7(b)) in the wug test were analyzed in the same manner as in previous experiments. Harmonic language learners chose correct (harmonic) plurals at a rate that was not significantly different from chance ($\beta = 0.04$, $SE = 0.26$, $z = 0.82$, $p > 0.50$), as did Non-Harmonic language learners ($\beta = -0.18$, $SE = 0.19$, $z = -0.95$, $p > 0.50$). Importantly, there was no significant difference in terms of accuracy between both language groups ($\beta = 0.21$, $SE = 0.26$, $z = 0.82$, $p > 0.50$). Thus, as in Experiment 3, participants did not extend the phonotactic generalization about stems, including trisyllabic stems, to novel unseen plurals that involved alternations.

As in Experiment 3, we further examined if participants in Experiment 4 showed a preference for one of the allomorphs over the other. The rate of choosing the [-mu] allomorph over the [-mi] allomorph was calculated and analyzed in the same manner as the rate of choosing correct plurals. Our results revealed that participants in both language groups, on the whole, did not choose one allomorph of the suffix exclusively over the other as seen in Figure 3.8(a), with both groups choosing the [-mu] allomorph at chance rates (Harmonic language: $\beta = 0.18$, $SE = 0.20$, $z = 0.92$, $p = 0.72$; Non-Harmonic language: $\beta = 0.04$, $SE = 0.22$, $z = 0.18$, $p = 0.86$). Furthermore, there was also no correlation between their performance on the blick test and the wug test in the Harmonic language in which participants showed learning of the phonotactic constraint (Figure 3.8(b): $R^2 = 0.16$, $r(13) = 0.40$, $p = 0.14$).

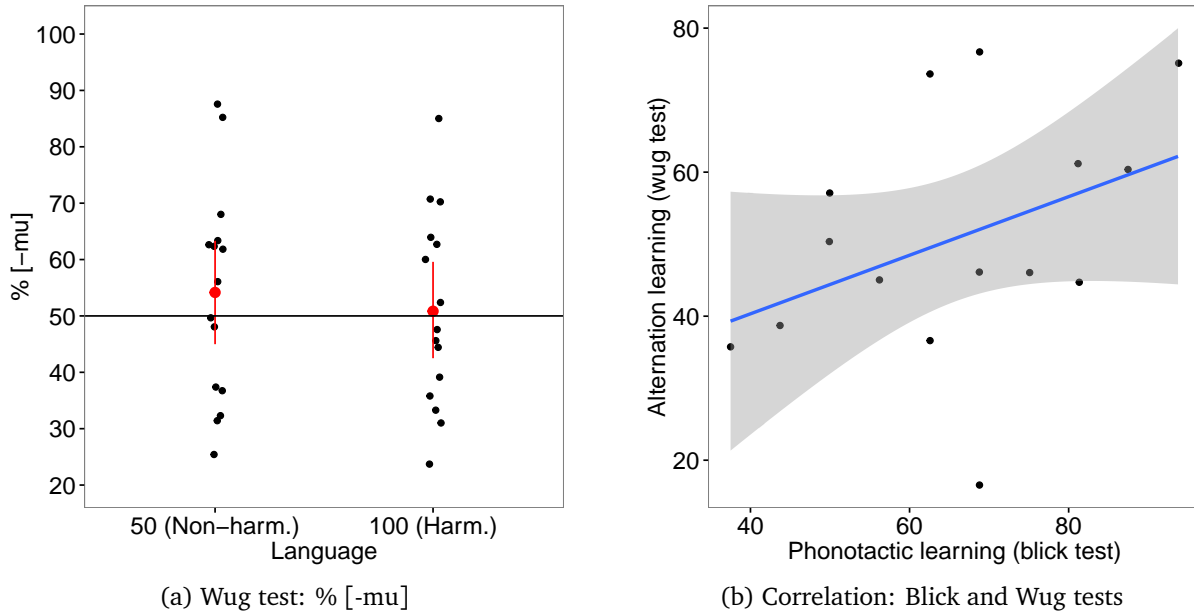


Figure 3.8: Experiment 4: (a) Rate of choosing [-mu] allomorph in Wug test. Each black dot represents a single participant with the large red-dot indicating mean rates for each language group. The red lines indicate 95% confidence intervals. (b) Correlation between phonotactic and alternation learning in Harmonic language.

3.5.3 Discussion

In Experiment 4, participants were trained on trisyllabic stems as well as bisyllabic stems, thus they had exposure to trisyllabic forms in training, and could have extended a learnt phonotactic generalization about trisyllabic forms to novel plurals that are all tri-syllabic. Harmonic language learners, as expected, showed successful internalization of phonotactics, at least as indicated by performance on the blick test with bisyllabic forms. Yet, contrary to expectations, Harmonic language learners failed to extend this generalization to the novel alternation, replicating the conservative behavior with regards to alternations in Experiment 3.

At first glance, the failure to extend the static phonotactic generalization might be surprising if we consider the fact that learners in Finley & Badecker’s (2009) study were able to extend a learnt generalization about trisyllabic forms to a novel alternation. The difference, however, between our current experiment and Finley & Badecker’s (2009) is that in their experiment, participants were trained explicitly on one alternation and test on a different one which still

conformed to the same abstract generalization involving harmony. In our experiment, learners were not trained on any alternations but rather a static phonotactic generalization in singular stems. In the wug test, they then had to extend this generalization to a novel unseen alternation in plurals. Thus this presents a different task than that in Finley & Badecker (2009).

Together, the results of Experiment 3 and 4 show that learners are conservative in positing alternations when there is no evidence in the input for them.

3.6 General Discussion

Using an artificial grammar learning paradigm, we investigated how phonotactics and alternation learning interact. We specifically examined the question of whether the learning of phonological alternations is facilitated when the phonotactic generalization within stems matches the generalization that motivates alternations. In Experiment 1, we compared alternation learning across languages with three degrees of phonotactic match. Learners accurately learnt the phonotactic generalization in each case. Importantly, learning of the alternation was most successful in the Harmonic language where the phonotactics match the alternation. Learners in the Non-Harmonic language did not learn the alternation, while Semi-Harmonic language learners did. Further, in Experiment 2, learners showed the same learning behavior in the absence of semantic information in training. Alternation learning, therefore, is facilitated when phonotactics match. Learners, however, failed to extend a learnt phonotactic generalization to a novel unseen alternation regardless of whether they were trained on bisyllabic stems (Experiment 3) or trisyllabic stems (Experiment 4). Why did the results in Experiments 2 and 4 differ when learners in both experiments had trisyllabic forms in training? In Experiment 2, although learners did not have accompanying pictures, they nonetheless had implicit access to morphology since they were still exposed to plural word forms in training. Learners in Experiment 4 did not have this, and in fact, were trained explicitly on trisyllabic stems. Further, there were no pictures in the wug test in Experiment 2 so participants could have treated this as a blink test. Thus a combination of both of these factors might have facilitated learners'

performance on the wug task in Experiment 2. To sum up, although alternation learning is facilitated when phonotactics match, learners are, nonetheless, conservative in extending a phonotactic generalization to alternations.

What does this mean for the relationship between phonotactics and alternations in learning? Taken together, the results of both Experiments 1 and 2 show that phonotactic mismatches impede the learning of an alternation, suggesting that phonotactic learning facilitates alternation learning. The amount of evidence for alternations was kept consistent, and exceptionless, across the different language groups. Thus the failure of Non-Harmonic language learners to successfully learn the alternation was in spite of there being evidence to support this in the training data. A learning model in which phonotactic knowledge and alternation knowledge are acquired completely separately and independently of each other thus fails to account for this result. Our current study therefore provides experimental evidence in support of the hypothesis that learning phonotactics facilitates learning of alternations (Hayes 2004, Hayes & Wilson 2008, Prince & Tesar 2004, Tesar & Prince 2007, Jarosz 2006). Our results also further confirm that phonotactic learning is gradient and is consistent with the lexical statistics in the input (Frisch & Zawaydeh 2001, A. W. Coetzee & Pater 2008). Learners showed a gradient preference for harmonic words that was proportional to the lexical statistics of harmony in the training data. Interestingly, the success with which learners' learnt the alternation across all three languages mirrored their performance on the phonotactic learning task. Thus our results not only show that phonotactic learning is gradient, they also show that gradient phonotactic learning leads to gradient alternation learning.

While phonotactic learning can facilitate alternation learning, learners do not readily assume that alternations will reflect phonotactic generalizations. The lack of generalization in Experiments 3 and 4 suggest that learners are conservative in extending a learnt static phonotactic pattern to a novel alternation in across a different morphological domain. This likely reflects a general anti-alternation bias which is enforced early in morphophonological learning (Benua 2000, Tessier 2012, Do 2013, Hayes 2004, McCarthy 1998). Thus it is likely we would only see extension of the phonotactic generalization when learners experience alternations in the

training language. What is clear though is that when both types of generalizations are available, a matching phonotactic generalization aids the learning of the alternation.

Our study also contributes to the investigation of the *directionality* of influence between phonotactics and alternation learning. In our experiment it seems that the nature of phonotactic generalizations influences the ability to learn alternations, providing evidence for the flow of information from phonotactics to alternations (c.f. Pater & Tessier 2005). What our experiments suggest though is that the relation might be uni-directional and that alternation knowledge may not affect phonotactic knowledge. Although the alternation pattern in the Non-Harmonic language was always consistent, this did not affect learners' phonotactic knowledge. That is, since vowel sequences across a morpheme boundary always harmonized, even in the Non-Harmonic language there are actually more harmonic vowel sequences than disharmonic vowel sequences. Learners therefore had more overall evidence for harmonic vowel sequences, but the presence of the alternation failed to influence what learners' inferred about stem phonotactic generalizations. This echoes recent findings by both Pizzo (2015) and Chong (2016) who failed to robustly show an effect of alternation learning on the learning of stem-internal phonotactics. As Pizzo (2015) points out, that phonotactics should affect alternations is intuitive given that phonotactics seems to be acquired before alternations in infancy (Jusczyk et al. 1994, Saffran & Thiessen 2003, Friederici & Wessels 1993, K. S. White et al. 2008, J. White & Sundara 2014). The reverse relationship, with alternation learning affecting phonotactic knowledge, seems less motivated. Whether the two kinds of phonological knowledge are encoded by the same mechanism is still, however, an empirical question. At the very least, the fact that learners failed to spontaneously extend phonotactic generalizations to alternations shows that the simplest implementation is likely incorrect.

In addition to addressing the broader question regarding the relationship between phonotactics and alternations in learning, our study also has further implications for the learnability of phonological alternations with mismatching phonotactics. These types of patterns are known in the theoretical phonology literature as derived-environment effects (Kiparsky 1993). Traditional analyses of these patterns predict that the alternations should be productively learnt.

Yet, in Experiments 1 and 2, adults failed to learn the phonological alternation with cases of mismatching phonotactics, despite being exposed to evidence of an alternation in the learning data. This suggests the possibility that alternations in derived-environment patterns are more difficult to learn when compared to a language in which alternations and phonotactics match, making such language patterns typologically dispreferred. In fact, a ‘true’ derived-environment pattern akin to the mismatch Non-Harmonic language used in our experiments seems to be elusive.

Chong (*under revision*) shows that in one well-known example of derived-environment effects, Turkish velar deletion (Inkelas 2000, 2011, Inkelas & Orgun 1995, Inkelas et al. 1997, Lewis 1967, Sezer 1981, Zimmer & Abbott 1978), a careful inspection of the lexicon shows that there is no phonotactic constraint against intervocalic velars in the lexicon. Yet there is an alternation that deletes stem-final velars when these are followed by a vowel-initial suffix. Thus exactly the same sequences that are allowed to occur within a stem are nonetheless involved in an alternation across a morpheme boundary. Interestingly, however, the alternation is highly morphologically-conditioned (Sezer 1981, Inkelas 2011), applying productively only with polysyllabic nouns (Zimmer & Abbott 1978, Becker et al. 2011). Thus the alternation in this case seems highly circumscribed, when the phonotactics mismatch.

What about the learning of Turkish Vowel Harmony on which the Non-Harmonic language is based? Previous studies have shown that children as young as 2;0 show successful acquisition of vowel harmony in their productions (Aksu-Koç & Slobin 1985). Further both van Kampen, Parmaksiz, van de Vijver, & Höhle (2008) and Altan, Kaya, & Hohenberger (2016) show that Turkish infants show early preferences for harmonic words over disharmonic ones. An early study by Zimmer (1969) also showed that Turkish speakers show awareness of this harmony pattern. Given our finding that the success of alternation learning was proportional to the proportion of harmonic stems in the lexicon, what might the Turkish lexicon look like? To ascertain this, vowel-vowel co-occurrences in polysyllabic roots in the Turkish Electronic Living Lexicon (Inkelas, Küntay, Orgun, & Sprouse 2000) was calculated. There were a total of 12,491 polysyllabic roots. Table 3.2 shows the number of co-occurrences between front ([-back])

and back ([+back]) vowels. Observed/Expected (O/E) values were calculated for each cell. ‘Observed’ (O) values are the total number of times each VV combination is found in the corpus while ‘Expected’ (E) values are how often each VV combination is expected if each V1 and V2 sequence co-occurred based on chance. Expected values were then calculated by taking the product of the relevant marginal totals (row and column) and dividing it by the grand total (for other examples of the use of this measure see A. Coetzee 2008, Frisch & Zawaydeh 2001). O/E values are finally calculated by dividing the Observed by the Expected value. An O/E value of 1 indicates that a particular sequence occurs at the expected rate of occurrence, a value above 1 indicates over-representation of that particular VV and a value under 1 indicates underrepresentation. The percentages in bold indicate the proportion of V1s found in a particular V2 context and the percentages in italics indicate the proportion of V2s in a V1 context. The marginal percentages (bottom row and final column) indicate the expected proportions of each type of V1 and V2. We note that non-harmonic sequences (cells which are not shaded) are significantly under-represented ($\chi^2(1) = 3033.7, p < 0.001$) with O/E values under 1. In fact, harmonic VV sequences are overrepresented in the corpus, and thus the lexicon of Turkish shows a harmonic preference in VV sequences in the lexicon. It is possible then that the Turkish vowel harmony alternation is aided by the preference for harmonic stems in the lexicon. Thus the traditional description of Turkish as completely non-harmonic is not supported by the corpus counts. Instead, based on the larger frequency of harmonic roots in the Turkish lexicon, Turkish is better described as semi-harmonic.

Another well-known case of derived-environment effects is Korean palatalization (Kiparsky 1973, Iverson & Wheeler 1988, Kiparsky 1993, T. Cho 2001). There is some evidence that the alternation in this productive across a suffix boundary (Jun & Lee 2007). Analysis of Korean corpora also shows that this alternation is supported by a gradient phonotactic constraint in the lexicon (Chong, *under revision*; see Chapter 4). Taken together, these studies show that many of the well-known cases of derived-environment effects appear to be less than perfect examples (see also Anttila 2006, Inkelas 2011). A ‘true’ derived-environment language with a perfect mismatch between phonotactics and alternations, thus, has yet to be discovered, further

	V2: [-back]	V2: [+back]	Expected % of V2
V1: [-back]	5,649 (3,776) 67.80% / <i>65.85%</i> O/E = 1.50	2,930 (4,803) 32.20% / <i>25.92%</i> O/E = 0.61	44.01%
V1: [+back]	2,683 (4,556) 27.64% / <i>34.15%</i> O/E = 0.59	7,669 (5,796) 72.36% / <i>74.08%</i> O/E = 1.32	55.99%
Expected % of V1	45.32%	54.68%	

Table 3.2: Occurrence of VV combinations: by V1 type (front [-back] vs. back [+back]) and by V2 type (front [-back] vs. back [+back]) . Expected counts are in parentheses. Percentages in bold: row percentages; Percentages in italics: column percentages. The cells in gray indicate disharmonic sequences (i.e. [-back][+back] and [+back][-back])

supporting the idea that these patterns are dispreferred, likely due to the difficulty in learning them.

All in all, this study has provided evidence for the basic claim that phonotactic learning facilitates alternation learning. At the same time, we have also shown that learners are conservative in extending static phonotactic generalizations to novel alternations. Together, our results suggest that both types of phonological knowledge cannot be entirely independent of each other in a model of phonological learning, although the exact mechanism which links them is still an open question. We have further shown that patterns which show a mismatch between phonotactics and alternations are more difficult to learn, predicting that these patterns should be dispreferred cross-linguistically. While the results of the current study indicate that phonotactics and alternations interact in learning, the trajectory of learning across both types of phonological knowledge especially in infancy remains unknown. Given the timeline of phonological development in infancy, with phonotactic knowledge emerging before alternation knowledge, it would be illuminating to examine when different kinds of alternation knowledge emerge. Our prediction here is that phonotactically supported alternations should be learnt first.

Appendix

Table 3.3: Experiment 1 training lexicon: Harmonic language

No.	Singular	Plural	No.	Singular	Plural
1	beme	bememi	17	nedi	nedimi
2	pege	pegemi	18	gibe	gibemi
3	degi	degimi	19	nopu	nopumu
4	tipe	tipemi	20	kugo	kugomu
5	mine	minemi	21	gubu	gubumu
6	kipi	kipimi	22	neke	nekemi
7	dimi	dimimi	23	nibi	nibimi
8	podu	podomu	24	dopo	dopomu
9	dobu	dobomu	25	kete	ketemi
10	tonu	tonumu	26	peki	pekimi
11	muto	mutomu	27	tidi	tidimi
12	buno	bunomu	28	gomo	gomomu
13	gutu	gutumu	29	boku	bokumu
14	budu	budumu	30	pime	pimemi
15	tegi	tegimi	31	muko	mukomu
16	motu	motumu	32	kunu	kunumu

Table 3.4: Experiment 1 training lexicon: Semi-Harmonic language

No.	Singular	Plural	No.	Singular	Plural
1	beme	bememi	17	nedi	nedimi
2	pege	pegemi	18	gibe	gibemi
3	degi	degimi	19	nopu	nopumu
4	tipe	tipemi	20	kugo	kugomu
5	mine	minemi	21	gubu	gubumu
6	kipi	kipimi	22	neke	nekemi
7	dimi	dimimi	23	nibi	nibimi
8	podu	podomu	24	dopo	dopomu
9	dobo	dobomu	25	keto	ketomu
10	tonu	tonumu	26	peku	pekumu
11	muto	mutomu	27	tidu	tidumu
12	buno	bunomu	28	gome	gomemi
13	gutu	gutumu	29	boki	bokimi
14	budu	budumu	30	pume	pumemi
15	tegi	tegimi	31	miko	mikomu
16	motu	motumu	32	kuni	kunimi

Table 3.5: Experiment 1 training lexicon: Non-Harmonic language

No.	Singular	Plural	No.	Singular	Plural
1	beme	bememi	17	nodi	nodimi
2	pege	pegemi	18	gube	gubemi
3	degi	degimi	19	nepu	nepumu
4	tipe	tipemi	20	kigo	kigomu
5	mine	minemi	21	gibu	gibumu
6	kipi	kipimi	22	neko	nekomu
7	dimi	dimimi	23	nubi	nubimi
8	podu	podomu	24	dope	dopemi
9	dobu	dobomu	25	keto	ketomu
10	tonu	tonumu	26	peku	pekumu
11	muto	mutomu	27	tidu	tidumu
12	buno	bunomu	28	gome	gomemi
13	gutu	gutumu	29	boki	bokimi
14	budu	budumu	30	pume	pumemi
15	tegi	tegimi	31	miko	mikomu
16	motu	motumu	32	kuni	kunimi

Table 3.6: Experiment 4 training lexicon: Harmonic language

No.	Singular	No.	Singular
1	beme	17	ketebe
2	pegebi	18	pekiye
3	degin	19	nedi
4	tipege	20	gibe
5	mine	21	tidigi
6	kipi	22	gomonu
7	dimi	23	bokumo
8	podoku	24	nopu
9	dobo	25	kugo
10	tonuto	26	gubu
11	mutoko	27	nekepi
12	buno	28	pime
13	gutu	29	nibi
14	budutu	30	dopodo
15	tegime	31	muko
16	motu	32	kunudu

Table 3.7: Experiment 4 training lexicon: Non-Harmonic language

No.	Singular	No.	Singular
1	beme	17	ketobe
2	pegebi	18	pekupe
3	degin	19	nodi
4	tipege	20	gube
5	mine	21	tidugi
6	kipi	22	gomenu
7	dimi	23	bokimo
8	podoku	24	nepu
9	dobu	25	kigo
10	tonuto	26	gibu
11	mutoko	27	nekopi
12	buno	28	pume
13	gutu	29	nubi
14	budutu	30	dopedu
15	tegime	31	miko
16	motu	32	kunidu

Table 3.8: Blick test stimuli

No.	Harmonic	Non-Harmonic
1	deke	doke
2	nepe	nepo
3	pemi	pomi
4	tebi	tebu
5	kipe	kupe
6	gike	giko
7	dini	dinu
8	kibi	kubi
9	mogo	moge
10	gono	geno
11	podu	podi
12	bogu	begu
13	buto	bito
14	numo	nume
15	tudu	tudi
16	mutu	mitu

Table 3.9: Wug test stimuli

No.	Singular	[-mi] Plural	[-mu] Plural
1	mete	metemi	metemu
2	beke	bekemi	bekemu
3	neki	nekimi	nekimu
4	mipe	mipemi	mipemu
5	giti	gitimi	gitimu
6	pidi	pidimi	pidimu
7	kobo	kobomi	kobomu
8	konu	konumi	konumu
9	domu	domumi	domumu
10	tugo	tugomi	tugomu
11	tubu	tubumi	tubumu
12	gunu	gunumi	gunumu
13	bepi	bepimi	bepimu
14	dime	dimemi	dimemu
15	pugo	pugomi	pugomu
16	nodo	nodomi	nodomu

CHAPTER 4

Derived environments and the lexicon

In Chapter 3, using artificial grammar learning experiments, I showed how learning an alternation is more difficult when the phonotactic generalization in stems mismatch the alternation (i.e. a derived environment pattern).¹ What does this mean for languages that actually show these patterns? There is a hallowed literature on derived environment patterns dating back to Kiparsky (1973). These analyses have in common the assumption that the alternation pattern across the morpheme boundary is productive, with no phonotactic dispreference for morpheme-internal sequences which are not repaired. Given that we found alternation learning to be more difficult in the Non-Harmonic (derived environment) language in Chapter 3, are the alternations or phonotactic generalizations actually productive in natural language examples? In this chapter, I focus on examining whether there is phonotactic support in the lexicon for the alternation in a derived environment pattern. I examine the lexicon in two well-known cases of derived environment effects: Korean palatalization and Turkish velar deletion. I present corpus analyses as well as computational learning simulations of both patterns. I show that in both patterns the reported mismatches between phonotactics and alternations are superficial, undermining previous analytic assumptions related to analyses of derived environment effects.

4.1 Introduction

It has been observed that phonological alternations at morphological boundaries often reflect morpheme-internal static phonotactic patterns (Chomsky & Halle 1968, Kenstowicz & Kisseberth

¹A version of this chapter is under revision for publication.

1977, McCarthy 2002). Kenstowicz & Kisseberth (1977, 1979) give an example from Kirundi (Meeusen 1959, Rodegem 1970) where vowels preceding nasal+consonant (NC) clusters are [+long] as shown in (1):

(1) Vowels are [+long] before NC clusters within stems

- a. [umu-ru:ndi] 'a Rundi person' *[umu-rundi]
- b. [ku-ge:nd-a] 'to go' *[ku-gend-a]

Importantly, this static phonotactic generalization is also enforced across morpheme boundaries by a phonological alternation, vowel lengthening. Underlying short vowels in /ku-/, /ba-/ and /umu-/ lengthen when prefixed before a stem containing an initial NC cluster as in (2a), (2c) and (2e) but not a singleton in (2b), (2d) and (2f). Thus both the tautomorphemic static phonotactic generalization and the heteromorphemic generalization that motivates the phonological alternation can be captured using the same rule or constraint.

(2) Vowels lengthen before NC clusters across morpheme boundaries

- a. /ku-n-dor-a/ → [ku:ndora] 'to look at me'
- b. cf. /ku-ror-a/ → [kurora] 'to look at'
- c. /ba-n-taba:re/ → [ba:ntaba:re] 'that they help me'
- d. cf. /ba-taba:re/ → [bataba:re] 'that they help'
- e. /umu-ntu/ → [umu:ntu] 'person'
- f. cf. /umu-gabo/ → [umugabo] '(married) man'

However, these two types of generalizations (static phonotactic generalizations about the lexicon and dynamic generalizations about phonological alternations) do not always pattern alike. Morphologically derived environment effects (also known as non-derived environment blocking; e.g. Kiparsky 1973, 1993) are one such example of this mismatch (see Paster (2013) for a recent review of other examples). A textbook example of derived environment effects is

Korean palatalization. At morpheme boundaries, underlying stem-final coronal stops /t/ and /t^h/ palatalize to [c] and [c^h] respectively before a suffix-initial /i/ and /j/ as in (3) (Kiparsky 1973, Iverson & Wheeler 1988, Kiparsky 1993, T. Cho 2001).

- (3) Palatalization across morpheme boundaries: /t, t^h/ → /c, c^h/ before /i/ and /j/ ²
- a. /mat-i/ → [maci] 'eldest-NOM'
 - b. /pat^h-i/ → [pac^hi] 'field-NOM'
 - c. /pat-hje-jo/ → [pac^hejo]³ 'is butted'

However, palatalization fails to apply when the target consonant (/t/ or /t^h/) and /i/ or /j/ are within the stem (i.e. tautomorphemic). Thus /ti/ and /t^hi/ sequences which are repaired at the morpheme boundary are nonetheless attested within stems where they surface faithfully as in (4).

- (4) Blocking of palatalization tautomorphemically:
- a. /mati/ → [mati] 'knot, joint'
 - b. /t^him/ → [t^him] 'team'

Patterns such as these have continued to pose a challenge for phonological theory, starting with Kiparsky (1973) (for a recent review and proposal see Inkelas 2014, 2015). Previous analyses of such patterns in rule-based models (Chomsky & Halle 1968) or in Optimality Theory (Prince & Smolensky 1993/2004) have focused on protecting morpheme internal non-derived sequences (such as /ti/ in Korean) while ensuring that the very same sequences always alternate if they occur due to morpheme concatenation. This has been achieved through a number of theoretical

²The tense stop, /t*/, does not occur word-finally (e.g. Sohn 1999). In this paper, we transcribe the palatal consonants using the symbol for the palatal stop, although these are often transcribed using the symbol for the alveolo-palatal affricate /c/.

³An independent process ensures that the lax stop /t/ followed by /h/ becomes aspirated and that the onglide in /je/ deletes post-consonantly.

tools such as underspecification (Kiparsky 1993), interleaving morphological operations and phonological ones (Wolf 2008), conjoined constraints (Łubowicz 2002), and reference to new or old input (Comparative Markedness; McCarthy 2003), amongst others. What these analyses ignore, however, is the question of how productive these processes are and, relatedly, what generalizations about these patterns are actually encoded by speakers in the grammar. For example, Łubowicz (2002), examining Polish velar palatalization, concedes that the protected stem-internal sequences she cites all appear in loanwords. She argues, however, that these words have been wholly incorporated into the native grammar (Rubach 1984) since palatalization applies in these words across the morpheme boundary, arguing against these examples being purely exceptions to the phonological rule. Yet this argument ignores the question of how phonotactically well-formed such protected sequences actually are in the phonological grammar of Polish speakers.

In this chapter, I examine, using computational learning simulations, what static generalizations are available to learners in derived environment effect patterns, focusing on two well-known examples: Korean palatalization and Turkish velar deletion. I begin, in §2, by looking critically at the analytic assumptions adopted by existing accounts of morphologically derived environment effects. In §3, I provide the historical background regarding the palatalization patterns in Korean, including the origins of the current generalization mismatch. I then report the results of an in-depth corpus study of the Korean lexicon, as well as a phonotactic learning simulation of Korean. I end §3 by presenting a new analysis of the Korean pattern using lexically-specific constraints, arguing essentially that the stem-internal sequences protected from palatalization in Korean are examples of gradient exceptionality. These findings are compared to one other well-known case of morphologically derived environment effects, Turkish velar deletion, in §4. The implications of these results are taken up in §5, where I argue that various examples of morphologically derived environment effects, while structurally similar superficially, are actually quite different from each other.

To preview the results, non-palatalized sequences in Korean are underrepresented in the lexicon and this leads to the learning of a gradient phonotactic constraint that penalizes

such sequences across-the-board. In Turkish, however, intervocalic velars, while somewhat under-represented in the lexicon, are not sufficiently under-represented to be penalized by a phonotactic learner. Thus two putatively similar cases of derived environment effects are shown to have very different lexical statistics, providing evidence against the analysis of these patterns in a unified way. Instead, our investigation suggests that one case, Korean palatalization, is an example of gradient exceptionality, while the other case, Turkish velar deletion, is an example of morphologically-conditioned phonology. Further, the data support the conjecture that languages prefer to have static phonotactic generalizations and dynamic generalizations that are similar.

4.2 Morphological derived environment effects and analytic assumptions

4.2.1 The Derived environment condition

Analyses of derived environment patterns have largely made the assumption laid out in (5):

- (5) DERIVED-ENVIRONMENT CONDITION: Morphological derivedness is a necessary and sufficient condition for an applicable process to apply. (variously stated as the Strict Cycle Condition or the Revised Alternation condition; e.g. Kiparsky 1982a, 1993; see also Inkelas 2011, 2014)

On (5), a number of authors (e.g. Anttila 2006, Hammond 1992, Inkelas 2011, 2014) have argued that what are often seen as canonical cases of morphologically derived environment effects do not satisfy this condition, insofar as a derived environment (i.e. where the target and environment of a rule are from two different morphemes) does not actually guarantee that a particular process would apply. That is, a derived environment, while a necessary condition for a particular process to apply, is by no means a sufficient condition. One such case that has been examined in greater detail is Finnish assibilation (Anttila 2006, Hammond 1992, Kiparsky 1973, 1993, 2003). In Finnish, stem-final /t/ become [s] before /i/. This rule

is generally characterized as only occurring across a morpheme boundary as in (6); it fails to apply within stems (7).

(6) /t/ → [s] / __ across a morpheme boundary (*ti):

- a. /halut-i/ → [halusi] ‘want-PAST’
- b. cf. /halut-a/ → [haluta] ‘want-INF’
- c. /hakkat-i/ → [hakkasi] ‘beat-PAST’
- d. cf. /hakkat-a/ → [hakkata] ‘beat-INF’

(7) /ti/ sequences surface faithfully within stems:

- a. /tilat-i/ → [tilasi] ‘order-PAST’ *[silasi]
- b. /koti/ → [koti] ‘home’ *[kosi]

Yet the reality in the data is far more complex. Anttila (2006), citing Karlsson (1983), shows that not all /i/-initial suffixes actually trigger assibilation. Assibilation only occurs uniformly with the three suffixes shown in (8). Many other /i/-initial stem-level suffixes fail to trigger assibilation despite satisfying the phonological (and morphologically-derived) environment for process application as seen in (9). In at least one case, the suffix variably triggers assibilation as in (10). I refer the reader to Anttila (2006) for a full analysis of these patterns (all data in (8)-(10) are taken from Anttila 2006: 900-901). What is of note here is that upon closer inspection, the derived environment effect pattern in Finnish fails to conform to the assumption in (5) - having a derived environment does not guarantee that the rule will apply.

(8) Triggering suffixes.

- a. Plural /-i/: /vuote-i-nA/⁴ → vuosina ‘year-PL-ESS’
- b. Past tense /-i/: /huuta-i-vAt-kO/ → huusivatko ‘shout-PAST-3PPL-Q’
- c. Superlative /-impA/: /uute-impA-nA/ → uusimpana ‘new-SUP-ESS’

(9) Non-triggering suffixes.

- a. Instrumental /-ime/: /lentä-ime-n/ → lentimen ‘fly-INST-GEN’
(*lensimen)
- b. Conditional /-isi/: /tunte-isi/ → tuntisi ‘feel-COND’ (*tunsisi)

(10) Variable trigger.

- a. Adj. deriv. suffix /-inen/: /vete-inen/ → vesinen~vetinen ‘watery’

The data from Finnish suggest that an account of putative derived environment effect patterns cannot simply appeal to the derived condition, which although necessary is not a sufficient condition for a particular process to apply. This suggests that the assumption about derived environments in (5) is problematic, and instead argues for an account of these patterns as potentially morpheme-specific or morphologically-conditioned phonology.

4.2.2 Phonotactic ‘productivity’

Together with the derived-environment condition, analyses that aim to capture the fact that tautomorphic /ti/ sequences exist in Finnish, for example, also assume that such sequences are entirely phonotactically well-formed (11).

- (11) PHONOTACTIC PRODUCTIVITY: Static phonotactic patterns that violate the derived environment generalization are completely productive (i.e. morpheme-internal sequences are phonotactically well-formed).

⁴Vowels in upper case undergo vowel harmony.

Here I make a distinction between attestedness, which relates to whether particular phonotactic sequences exist in the lexicon, and well-formedness which relates to native speakers' grammatical judgments of legal and illegal structures in their native language. Attestedness is a categorical parameter since a particular sound sequence either exists or does not in the lexicon of a language, to the extent that having even just one lexical item is sufficient for a particular sound sequence to be attested. Well-formedness, however, is not a categorical notion: not all sound sequences (or words) are created equal (e.g. Coleman & Pierrehumbert 1997, Hayes 2000, Pierrehumbert 1994, Schütze 1996). In the domain of phonotactics, in particular, a large body of evidence has shown that speakers possess gradient intuitions (e.g. Bailey & Hahn 2001, A. Coetzee 2008, Coleman & Pierrehumbert 1997, Frisch, Pierrehumbert, & Broe 2004, Frisch & Zawaydeh 2001, Hay, Pierrehumbert, & Beckman 2003, Treiman, Kessler, Knewasser, Tincoff, & Bowman 2000) and that listeners even use gradient well-formedness constraints in speech processing (Frisch, Large, & Pisoni 2000, Kager & Shatzman 2007).

Thus, a second issue for morphologically derived environment effects is whether the sequences repaired by a phonological process really do count as fully well-formed in the static patterns in the lexicon. In the same way that these patterns are not as general as was previously thought in derived environments, it is possible that the static (non-derived) patterns in the lexicon are not as widespread or 'productive' as we assume. Taking Korean as an example, this would entail that not only is a sequence [ti] repaired at a morpheme boundary (due to a constraint like *ti), but speakers actually show some dispreference for such sequences even when they occur within stems. This question is taken up in the next section, using computational learning simulations.

4.3 Korean palatalization

4.3.1 Historical origins and further background

How did the current mismatch develop historically in Korean? The origin of the current ostensibly derived condition on palatalization dates back to Early Modern Korean (circa early

nineteenth century; Y.-M. Y. Cho 2009, Lee & Ramsey 2011, and references therein). By the start of the 19th century, palatalization was an obligatory process that was an across-the-board sound change that neutralized the coronal stops /t, t^h, t^{*}/ to their corresponding palatal affricate counterparts /c, c^h, c^{*}/ before the high front vowel /i/ and the palatal glide /j/. This was a process that applied both within morphemes and across morpheme boundaries. Thus for a time both the static phonotactic generalization about the lexicon as well as the dynamic generalization motivating alternations were in agreement.

The current mismatch in generalizations has three sources. The first was the monophthongization of /i/ sequences to [i] which occurred following the sound change that involved palatalization. Thus words that had underlying /t_i/ or /t^h_i/ became previously unattested [ti] and [t^hi] (12).

(12) Source 1: historical monophthongization of /i/.

- a. * Δ t_i > Δ ti ‘where’
- b. *mat_i > mati ‘joint’

The fact that surface [ti] or [t^hi] were not palatalized represents an example of diachronic opacity, specifically counterfeeding. In principle, the monophthongization of /i/ could have fed the palatalization process, but this did not occur, resulting in a generalization that was not entirely surface true. This same process is also reflected in synchronic monophthongization of /i/ sequences that were the result of morpheme concatenation (13).

(13) Source 2: synchronic monophthongization of /i/:

- a. /t^{*}_i-ita/ → [t^{*}_iita] ~ [t^{*}ita] ‘to become aware’
eye-PASS-PRED
- b. /t^h_i-ita/ → [t^h_iita] ~ [t^hita] ‘to be open’
open-PASS-PRED

The final source of /ti/ and /t^{hi}/ sequences are loanwords borrowed from English and other European languages which were systematically borrowed in faithfully. Some examples are given in (14).

(14) Source 3: Loanwords from English

- a. /sitilom/ from ‘CD-ROM’
- b. /anthikhi/ from ‘antique’

Cho’s (2009) broader observation in relation to morphologically derived environment effects is the fact that many putative cases have a similar historical origin to that outlined for Korean: a particular phonological process had historically applied across-the-board both tautomorphemically and heteromorphemically, but sequences which were previously unattested were reintroduced through borrowing into the language as well as other independent phonological processes. She points out that this is the case with other well-known examples of morphologically derived environment effects such as Chamorro Vowel Lowering (Chung 1983), Finnish Vowel Coalescence (Anttila 2009) and Polish First Velar Palatalization (Łubowicz 2002). Y.-M. Y. Cho (2009) further suggests that words that are exceptions to the more general palatalization rule are marginal in the lexicon because of their historical origins. In what follows, I examine Cho’s claim about the marginality of /ti/ and /t^{hi}/ sequences in more detail and its consequences for phonotactic learning.

4.3.2 Corpus study

In order to investigate the lexical trends pertaining to [ti] and [t^{hi}] sequences in Korean, I examined two corpora. The first is a corpus compiled by the National Academy of Korean Language (National Academy of Korean Language 2003).⁵ The NAKL corpus contains over 50,000 frequently used Korean words (native and Sino-Korean), including loanwords, as well as corre-

⁵Now the National Institute of Korean Language.

sponding frequencies of each word from various, usually print, sources. The second is a corpus of Child-Directed Speech (CDS) compiled from two sub-corpora from CHILDES (MacWhinney 2000). To facilitate analysis, both corpora were first pre-processed and each syllabary was split up into its component Korean letters or digraphs using the grapheme-to-phonetic conversion system of Kim, Lee, & Lee (2002). The system also applies regular neutralizing phonological rules at the appropriate morphological boundaries detected in the process of conversion. The corpus contains lexemes and thus any morphological boundaries would be due to derivational suffixes. Note that the conversion system only applies transformations if there is already an orthographic character available for the resulting sound, thus this does not reflect any purely allophonic changes (such as intervocalic voicing).

4.3.3 NAKL

4.3.3.1 Entire lexicon

The final corpus, after any duplicate items were excluded, included 53,196 lexical items. We were first interested in how many words in the corpus contained the consonant [t], [t^h] or [t*] followed by [i] or [j] (I will refer to these as [Ti] or [Tj] respectively, and collectively as [TI]). Although we do not have overt evidence of palatalization of /t*/ because these do not occur word-finally, we are treating the coronal stop series here as a natural class, since these all participated in the historical sound change. A count of [TI] entries in the corpus is given in Table 1. One notices that out of a total of 53,196 words in the corpus, only 436 contained [TI] sequences, less than 1% of the lexicon. One further notes that out of these words, 284 (65%) are loanwords (e.g. /t^him/ = team, etc.). Thus these sequences are rare in terms of absolute type frequency in the corpus.

Table 4.1: No. of words that contain [ti], [t^{hi}], [t*ⁱ], [tj], [t^{hj}] and [t*^j] in NAKL corpus (and by lexical strata).

CV Type	Entire Lexicon	Native	Sino-Korean	Loanword
[ti]	208	68	5	135
[t ^{hi}]	167	30	4	133
[t* ⁱ]	32	28	4	0
[tj]	14	5	0	9
[t ^{hj}]	15	4	4	7
[t* ^j]	0	0	0	0
Total	436	135	17	284

It is possible, though, that the rarity of words with such sequences is merely attributable to either the overall rarity of the coronal stop series or the high front vocoids. To ascertain if the rarity of [Ti] and [Tj] words is statistically significant given the independent frequency of its components segments, a two-by-two contingency table was constructed (Table 2) that compared the frequency of occurrence of these sequences compared to other CV combinations. Note here that we are counting the type frequency of each CV sequence and not the number of words that contain at these sequences as in Table 1 — for example, *pata* ‘sea’ contributes both to the upper right ([ta]) and the lower right cells ([pa]). Observed/Expected (O/E) values were calculated for each cell. ‘Observed’ (O) values are the total number of sequences of each CV combination found in the corpus. ‘Expected’ (E) values are how frequently each CV combination is expected if each C and V co-occurred based on chance. That is, given the independent occurrence of the coronal stop series and the independent occurrence of high front vocoids, how often do we expect to see them co-occur? Expected values were calculated by taking the product of the relevant marginal totals (row and column) and dividing it by the grand total (for other examples of the use of this heuristic, see A. Coetzee 2008, Frisch & Zawaydeh 2001). O/E values are then calculated by dividing the Observed over the Expected value. An O/E value of 1 indicates that a particular sequence occurs at the expected rate of occurrence. O/E values

above 1 indicate over-representation of that particular CV and O/E values under 1 indicate under-representation, and we can thus compare the degree of attestedness in the corpus of particular CV sequences. The percentages in bold indicate the proportion of vowels found in a particular consonantal context and the percentages in italics indicate the proportion of consonants in a particular vocalic context. The marginal percentages (bottom row and final column) indicate the expected proportions.

	[i , jV] ([I])	Other Vs	Expected % of Cs
[t, t ^h , t*] ([T])	454 (5,798) 1.63% / <i>1.43%</i> O/E = 0.08	27,424 (22,073) 98.37% / <i>22.75%</i> O/E = 1.24	<i>18.31%</i>
Other Cs	31,247 (25,903) 25.12% / <i>98.57%</i> O/E = 1.21	93,112 (98,672) 74.87% / <i>77.25%</i> O/E = 0.95	<i>81.69%</i>
Expected % of Cs	20.82%	79.18%	

Table 4.2: Occurrence of CV combinations: by consonant type (T vs. other Cs) and vowel type (i, j vs. other Vs). Expected counts are in parentheses. Percentages in bold: row percentages; Percentages in italics: column percentages.

What is clear from Table 4.2 is that the actual observed number of [TI] sequences in the NAKL corpus is about a tenth of what would be expected due to chance (randomly combining each C and V), and this underrepresentation is statistically significant ($\chi^2(1) = 7625.1$, $p < 0.001$). While we expect about 18% of consonants to be /T/, only 1.43% of Cs in the [I] context are [T], indicated by the percentages in italics. Similarly, while we expect about 21% of vowels to be /I/, only 1.63% of [I]s occur with [T] (the percentages in bold). That is, [TI] sequences occur at about a tenth the rate that we would expect them to occur given the independent occurrence of [T] and [I].

	[i , jV]	Other Vs	Expected % of Cs
[c, c ^h , c*]	5,944 (5,051) 22.50% / 18.75% O/E = 1.08	20,473 (20,916) 77.50% / 16.98% O/E = 0.98	17.35%
Other Cs	25,757 (26,200) 20.47% / 81.25% O/E = 0.98	100,063 (99,620) 79.53% / 83.02% O/E = 1.00	82.65%
Expected % of Cs	20.82%	79.18%	

Table 4.3: Occurrence of CV combinations: by consonant type (CH vs. other Cs) and vowel type (i, j vs. other Vs) in the entire NAKL. Expected counts are in parentheses. Percentages in bold: row percentages; Percentages in italics: column percentages.

As a comparison, Table 4.3 shows the same calculations for [c, c^h, c*] and [i] or [jV] sequences ([CHi] and [CHj] respectively). In this case, we see instead a small, statistically significant ($\chi^2(1) = 54.41$, $p < 0.001$) over-representation of [CHi] and [CHj] sequences in the corpus. This is perhaps expected given that historically /Ti/ and /Tj/ sequences palatalized to [CHi] and [CHj] across-the-board, as was discussed above.

Thus, our corpus investigation so far supports the hypothesis that [Ti] and [Tj] sequences, while attested in the Korean lexicon, are actually exceedingly rare and marginal, with a majority of such words being loanwords. Even in loanwords, however, we might expect that such sequences might be dispreferred. This would provide further evidence of a general constraint against [TI] sequences. In the next section, I investigate this further by exploring the distribution of [Ti] and [Tj] sequences in each stratum of the Korean lexicon.

4.3.3.2 Generalizations across different strata

The Korean lexicon can be divided into three strata. Beyond the native stratum, a large proportion (as much as 60%) of the current Korean vocabulary is Sino-Korean (Sohn 1999).

Although the exact time Chinese vocabulary entered the Korean lexicon is difficult to determine, it is generally assumed that this occurred before 900 AD (Sohn 1999). Thus, when historical palatalization occurred in the nineteenth century, it would have applied to both native and Sino-Korean strata of the lexicon. Most loanwords from English and other European languages, however, were likely only borrowed in relatively recently in the twentieth century and thus did not undergo the historical palatalization process (Y.-M. Y. Cho 2009). In this section, we are interested in examining whether the under-representation of [Ti] and [Tj] in the NAKL taken as a whole is due purely to the native and Sino-Korean lexicon or whether or not this extends to loanwords as well.

The entire NAKL corpus was divided into three sub-corpora corresponding to the different lexical strata. This was done by identifying whether Roman orthography or Chinese orthography were included in the comments field of the NAKL corpus. Tagging was done in the following order. First, any entries with English orthography in the comments field were tagged as loanwords. This was done even when there might have been a native or Sino-Korean derivational suffix. Next, any of the remaining words tagged with Chinese orthography were tagged as Sino-Korean words, even if these had native derivational suffixes. Whatever remained after the two rounds of exclusion were coded as native words. Of the 53,196 words in the corpus, 13,459 were native, 36,504 were Sino-Korean and 3,233 were loanwords. Two-by-two contingency tables were constructed and O/E values were calculated in the same fashion as in section 4.1.1.

As can be seen in Table 4.4, in both the native and Sino-Korean strata of the lexicon, [TI] are significantly under-represented (native: $\chi^2(1) = 2829.3$, $p < 0.001$; Sino-Korean: $\chi^2(1) = 5026$, $p < 0.001$). In fact, the general distribution is comparable in both lexicons. In the native stratum, we expect between 20-24% of CV sequences to be [TI] but we see less than 2% of such cases, and in the Sino-Korean stratum, while we expect between 15-20% of CVs to be [TI], less 1% of CVs actually are. When we turn to the loanword stratum (Table 4.6), we see that the O/E for [TI] is much higher than in the native or Sino-Korean lexicons (0.70 vs. 0.07 or 0.005), although it is still statistically under-represented ($\chi^2(1) = 57.63$, $p < 0.001$). Thus, while it seems that there are some qualitative differences in the native and Sino-Korean lexicons

compared to the loanword lexicon, there is still a small dispreference for [TI] in loanwords, despite English sequences /t, d, θ, ð/+/ɪ, i/ (TI) being borrowed faithfully into Korean as [ti] or [t^hi] (Y.-M. Y. Cho 2009). This suggests that Korean speakers extend an, albeit weaker, native statistical dispreference to new loans.

	[i , jV]	Other Vs	Expected % of Cs
[t, t ^h , t*]	147 (1995) 1.48% / <i>1.79%</i> O/E = 0.07	9,761 (7,913) 98.52% / <i>29.94%</i> O/E = 1.23	<i>24.27%</i>
Other Cs	8,073 (6225) 26.12% / <i>98.21%</i> O/E = 1.30	22,839 (24,687) 73.88% / <i>70.06%</i> O/E = 0.93	<i>75.73%</i>
Expected % of Vs	20.14%	79.86%	

Table 4.4: Occurrence of CV combinations: by consonant type (T vs. other Cs) and vowel type (i, j vs. other Vs) - Native lexicon. Percentages in bold: row percentages; Percentages in italics: column percentages.

	[i , jV]	Other Vs	Expected % of Cs
[t, t ^h , t*]	17 (3,380) 0.11% / <i>0.08%</i> O/E = 0.005	16,118 (12,755) 99.89% / <i>19.98%</i> O/E = 1.26	<i>15.81%</i>
Other Cs	21,358 (17,995) 24.86% / <i>99.92%</i> O/E = 1.19	64,727 (67,912) 75.14% / <i>80.02%</i> O/E = 0.95	<i>84.19%</i>
Expected % of Vs	20.95%	79.05%	

Table 4.5: Occurrence of CV combinations: by consonant type (T vs. other Cs) and vowel type (i, j vs. other Vs) - Sino-Korean lexicon. Percentages in bold: row percentages; Percentages in italics: column percentages.

	[i , jV]	Other Vs	Expected % of Cs
[t, t ^h , t*]	290 (412) 15.80% / 13.77% O/E = 0.70	1,554 (1,423) 84.20% / 21.25% O/E = 1.09	19.57%
Other Cs	1,825 (1,694) 24.08% / 86.23% O/E = 1.07	5,728 (5,846) 75.92% / 78.75% O/E = 0.98	80.43%
Expected % of Vs	22.46%	77.54%	

Table 4.6: Occurrence of CV combinations: by consonant type (T vs. other Cs) and vowel type (i, j vs. other Vs) - Loanwords. Percentages in bold: row percentages; Percentages in italics: column percentages.

An alternative explanation for the under-representation of [TI] in the loanword sub-corpus might be that Koreans are just matching the statistics of such sequences in English. That is, the reason we see the lexical statistics in Korean is because Korean speakers are just borrowing in [TI] words matching the proportion of [TI] words in English. Thus, the under-representation here might reflect a generalization about the English lexicon and not the Korean one. To rule out this possibility, we examined the rate of borrowing into Korean from English. For this analysis, we consulted a separate corpus of Korean loanwords collated as well by the National Academy of Korean Language (2001) which already contained the source of the Korean word in English orthography and accompanying phonetic transcriptions. This corpus proved more appropriate since we are interested here in the statistics of the English lexicon. An edited version of the Carnegie Mellon University Pronunciation (CMU) dictionary was used here as a stand-in for the entire English lexicon. This version of the CMU was edited by Hayes & White (2013) and contains words with a CELEX (Baayen, Piepenbrock, & Gulikers 1995) frequency of ≥ 1 , thus

it excludes very low frequency items.⁶ The CMU corpus was tagged for whether an English word was loaned into Korean or not based on identifying CMU words that also appeared in the NAKL loanword corpus. Out of the 2486 unique entries in the NAKL loanword corpus (2001), only 1709 of these also appeared in the CMU. These were the only ones included in the counts below in Table 4.7. Loanwords that did not appear in the edited CMU were not included since we assumed that these were generally lower frequency items, such as acetal, adenosine etc.

	/t, d, θ, ð + /l, i/ ([TI])	Other CVs	Expected % of Loaned/Not Loaned
LOANED	153 (242) 4.01% / 5.53% O/E = 0.63	3,659 (3,570) 95.99% / 8.98% O/E = 1.02	8.76%
NOT LOANED	2,612 (2,523) 6.58% / 94.47% O/E = 1.04	37,097 (37,186) 93.42% / 91.02% O/E = 1.00	91.24%
Expected % of CVs	6.35%	93.65%	

Table 4.7: Observed/Expected counts of English CV sequences (loaned/not loaned against TI/CV). Percentages in bold: row percentages; Percentages in italics: column percentages.

As can be see in Table 4.7, we expect, all else being equal, 8.76% of the English CV sequences to be borrowed in to Korean, but only 5.53% of possible [TI] words are borrowed in. Similarly, while we expect 6.35% of CVs to contain [TI], only 4.01% of loaned CVs do. Thus, given the general expected rate of loaning English CV sequences into Korean, the number of English [TI] sequences loaned in is significantly less than what we expect ($\chi^2(1) = 38.01$, $p < 0.001$). This suggests then that the under-representation of [TI] sequences in the loanword stratum in Table

⁶Original version downloadable at: <http://www.speech.cs.cmu.edu/cgi-bin/cmudict>. Modified CMU version available online at: http://www.linguistics.ucla.edu/people/hayes/PhonologicalNaturalness/TrainingData_HayesAndWhite.txt

6 is not a result of frequency matching with the lexicon of English. So even though Korean speakers can borrow [TI] sequences faithfully, they seem to avoid doing so, suggestive of a dispreference for word with such sequences.

4.3.4 Child-directed speech

In the previous section, we showed how [Ti] and [Tj] sequences are under-attested across all the strata of the Korean lexicon. The NAKL corpus, however, represents a sample drawn from adult sources (such as newspapers etc.). It is possible that these lexical statistics do not actually reflect what a typical Korean-learning child might hear in acquisition, and is idiosyncratic to the specific corpus we used. Thus if we are ultimately interested in the question of what is available in the learning data for a child then our case is bolstered by an investigation of child-directed speech (CDS). To ascertain how robust the statistical distributions we found in NAKL are, we examined a corpus of Child-Directed Speech (CDS). A large corpus of CDS was created by combining the two available Korean CDS corpora on CHILDES (MacWhinney 2000): the Jiwon corpus (Ghim 2005) and the Ryu corpus (Ryu 2012, Ryu & Yasuhiro 2014). The Jiwon corpus consists of recordings of a mother-child interaction for a single child from ages 2;0 to 2;3. The Ryu corpus consists of longitudinal data from three children aged 1;3 to 3;9 with recorded interactions with caregivers: mother, father, grandmother, and grandfather. Utterances from all caregivers were included in the analysis. Since the CDS corpora involve play sessions with the child, there are a number of transcribed utterances that contain “reduplicated” forms which often contain repeated sequences of syllables or part-words. It is unclear if a child would actually encode these “words” in their lexicon, so such instances were filtered out as much as possible. An arbitrary cut-off was adopted such that words with more than 20 segments were excluded from the corpus. Almost all of the transcribed “words” with 20 segments or more fit into the category of repetitions described above. That said, many of these cases which may involve fewer repetitions fall under this criterion, but to filter out each of these forms would require checking each word in the lexicon. The total resulting corpus for analysis had 40,317 words. O/E values were calculated in the same way as in previous analyses above.

	[i , jV]	Other Vs	Expected % of Cs
[t, t ^h , t*]	524 (3,607) 3.18% / <i>1.95%</i> O/E = 0.15	15,937 (12,853) 96.82% / <i>16.60%</i> O/E = 1.24	<i>13.39%</i>
Other Cs	26,415 (23,332) 24.81% / <i>98.05%</i> O/E = 1.13	80,058 (83,140) 75.19% / <i>83.40%</i> O/E = 0.96	<i>86.61%</i>
Expected % of Vs	21.91%	78.09%	

Table 4.8: Occurrence of CV combinations: by consonant type (T vs. other Cs) and vowel type (i, j vs. other Vs) - CDS corpus. Percentages in bold: row percentages; Percentages in italics: column percentages.

What we see in Table 4.8 then is a similarly, statistically significant ($\chi^2(1) = 3895$, $p < 0.001$), under-representation of [Ti] in the CDS corpus as we saw in the NAKL corpus. To the extent that CDS contains more frequent forms in the language, this result bolsters the argument that even amongst more frequent forms, the lexical statistics we found for the lexicon as a whole still hold.

4.3.5 Corpus Summary

To summarize, so far, we have examined two different Korean corpora, one of adult printed sources and another of CDS. In both studies, [Ti] and [Tj] sequences, while attested, occur significantly less than what we would expect given the independent occurrence of coronal stops and high front vocoids. Interestingly, this statistical dispreference extends, albeit more weakly, into loanwords as well where new borrowings contain fewer than expected such sequences, even though these sequences are usually borrowed faithfully (Y.-M. Y. Cho 2009). Our initial hypothesis then that [Ti] and [Tj] are marginal in the Korean lexicon is supported. In the

next section, I describe a series of learning simulations that are aimed at ascertaining whether learners could easily arrive at a markedness constraint along the line of *[Ti] or *[Tj].

4.3.6 Learning a phonotactic grammar of Korean

While the previous sections uncovered a statistical under-attestation of the relevant sequences in Korean, this does not necessarily entail that such forms would be penalized by the grammar. In this section, we are interested in ascertaining whether or not the underrepresentation of [Ti] and [Tj] would actually translate into the learning of a constraint penalizing such sequences in a phonotactic grammar. Both the NAKL and CDS corpora were used as learning data for the UCLA Phonotactic Learner (Hayes & Wilson 2008). The entire NAKL corpus was included, including loanwords. The motivation for including loanwords here was due to the fact that we are assuming that in an early stage of phonotactic learning, a child does not have explicit knowledge of lexical strata, thus as far as she is concerned there is no difference between a native, Sino-Korean word or loanword. We assumed the following features in Table A1 in Appendix A for the Korean phoneme inventory. Note that glide-vowel sequences were assumed to be two separate segments for the purposes of this simulation. Also since /e/ and /ɛ/ are now merged in most speakers' productions (Eychenne & Jang 2015, Shin, Kiaer, & Cha 2013), these categories were both collapsed to /e/.

The phonotactic learner was asked to only find bigram constraints, with a maximum number of constraints set at 180. The O/E accuracy threshold for constraints was set at 0.30, following the simulations done by Hayes & Wilson (2008). Since type frequency is typically implicated in the learning of phonotactic constraints over the lexicon (e.g. Pierrehumbert 2003, Edwards et al. 2004, Richtsmeier 2011), each input had a frequency of 1. All other parameters were set at default. The key point to take note of is that we are making fairly uncontroversial assumptions in the parameters of this simulation and the model is not being biased towards finding the relevant constraints.

No.	Constraint	Weight	Description
1.	*[+spread_glott][+word_boundary]	7.14	No aspirated stops word-finally
2.	*[+word_boundary][+nasal,+dorsal]	6.67	No word-initial ɲ
3.	*[+const_glott][-approx.]	6.36	No tense stops preceding non-approx. consonants
4.	*[-lateral,-syllabic][+lateral]	6.16	No non-lateral consonant before a lateral
5.	*[-syllabic][+nasal,+dorsal]	5.99	ɲ cannot occur following consonants
6.	*[-cons.,+front,-syll.][-labial,+high]	5.81	*[jɪ], *[ji]

Table 4.9: Top weighted constraints learned from NAKL corpus. Grey cells indicate constraints that also fall into the top six in the simulation with Child-Directed Speech (Table 4.10)

4.3.6.1 Results: NAKL

As a first check of the results of the simulations, the highest weighted constraints were inspected to see if these corresponded to well-known phonotactic constraints in Korean. The top six constraints are shown in Table 4.9.

The learner discovers sensible constraints: the top six of these conform to what we know about the phonotactic restrictions in Korean, such as *[+spread_glott][+word_boundary] which ensures that aspirated stops do not occur word-finally (these neutralize to the lenis stops) as well as constraints that require [ɲ] to follow a vowel, *[-syllabic][+nasal,+dorsal] and *[+word_boundary][+nasal,+dorsal].

Crucially, the learner discovers a constraint penalizing [Ti] and [Tj] sequences: *[-sonorant,-strident][-spread_glott,-const_glott,+high,+front]⁷ (from here on *TI) and gives this a weight of 1.916. This constraint is ranked 55th in weight out of 134 constraints that the learner discovered. As a comparison the learner was run on a modified NAKL corpus in which [Ti] and [Tj] sequences were excluded (so the generalization is artificially made exceptionless). In this case, the learner assigns *TI a weight of 6.169 (ranked 4th out of 127). Thus, *TI in the Korean

⁷Note that the feature matrix here groups the vowels [i, j] together with the lenis [c] as a natural class. A simulation run with an initial constraint that excludes [c] (i.e. just the vowels [i, j]) arrives at the same result.

lexicon is not a categorical constraint, unsurprisingly, but it nonetheless is assigned a sizable weight.

4.3.6.2 Results: Child-Directed Speech

No.	Constraint	Weight	Description
1.	*[+cont.,-lat.,-syll.][+word_bound.]	6.783	No word-final glides or fricatives
2.	*[+spread_glott][-approximant]	6.599	No asp. stops preceding non-approx. cons.
3.	*[+const_glott][-approximant]	6.521	No tense stops preceding non-approx. cons.
4.	*[+const_glott][+word_boundary]	6.303	No tense stops word-finally
5.	*[+spread_glott][+word_boundary]	6.204	No aspirated stops word-finally
6.	*[+word_boundary][+nasal,+dorsal]	6.141	No word-initial [ŋ]

Table 4.10: Top weighted constraints learned from CDS corpus. Grey cells indicate constraints that also fall into the top six in the simulation with the NAKL corpus (Table 4.9)

As with the NAKL simulation we first inspected which constraints received the largest weights. Reassuringly, the learner arrives at a similar list of high-ranked constraints (they share 3 out of the top 6 constraints), and all six reflect well-known phonotactic restrictions in Korean (Table 4.10). Thus, the learner is arriving at sensible constraints that are similar to what was learned with the adult corpus above.

As with the NAKL corpus previously, the phonotactic learner also assigns a considerable weight, 1.599, to a constraint (*[-sonorant,-strident][-spread_glott,-const_glott,+high,+front]) that penalizes [Ti] and [Tj] sequences. This constraint is ranked 77 out of the 103 bigram constraints posited by the model. As a comparison, a learner trained on a corpus without any [Ti] and [Tj] sequences at all learns the same constraint but assigns it a weight of 6.015 (ranked 6th out of 107 constraints). The examination of the CDS corpus thus replicates our findings with the adult NAKL corpus.

4.3.7 Modeling Summary

Using default parameters, we have shown how a probabilistic phonotactic learning model assigns a sizable weight to a constraint penalizing [Ti] and [Tj] sequences despite these forms actually existing in the lexicon of Korean. This constraint was robustly attested in both a corpus of words from printed sources and one of Child Directed Speech. To the extent that we assume CDS contains highly frequent word forms, the replication of the basic finding in CDS provides strong support for this constraint occurring in the learning data. We can conclude, therefore, that the statistical under-representation we found does indeed translate into a well-formedness penalty for words with [Ti] and [Tj]. This suggests that the implicit assumption in analyses of morphologically derived environment effects that such sequences occur freely in the lexicon is empirically not the case.

4.3.8 A new analysis

Both the corpus studies and phonotactic modeling simulations above suggest strongly that there is a general structure-blind (Martin, 2011) markedness constraint, *TI, which is active in the Korean grammar. That is, regardless of whether a morpheme boundary intervenes between T and I, the constraint still penalizes the sequence. Yet we also know that words with such forms exist in the Korean lexicon, e.g. /mati/ ‘joint’. Thus an analysis of the Korean palatalization pattern needs to predict that existing [TI] words should surface faithfully but it should also penalize novel words that do contain such sequences. The analysis presented in this paper builds on the new empirical data presented in the previous sections. An analysis of Korean palatalization should:

1. make use of a general markedness constraint following the results of phonotactic learning.
2. capture the fact that [TI] sequences are under-represented (= less phonotactically well-formed).

3. Stem-internal sequences are essentially exceptions since there are few enough of such cases when we take the Korean lexicon as a whole
4. capture the fact that existing words have fixed outputs containing stem-internal [TI].
5. predict categorical alternations at morpheme boundaries.
6. predict that a nonce word with stem-internal [TI] should be dispreferred.

To the best of my knowledge, palatalization is both general in that it applies to all suffixes where the phonological conditions are met (15) and extends to loanwords (Jun & Lee 2007).

- (15) Palatalization occurs in both inflectional and derivational suffixes (data from Cho, 2009)
- a. /hɛ tot-i/ → [hɛ toci] ‘sun-rise’
sun rise-NML
 - b. /kut-i/ → [kuci] ‘firmly’
be.firm-ADV
 - c. /put^h-i/ → [puc^hi] ‘to affix’
adhere-CAUS
 - d. /pat^h-ita/ → [pac^hita] ‘to be the field’
field-COP

(i) and (iv) can be captured using a single markedness constraint in (16) that bans [TI] sequences in the output.

- (16) *[-sonorant,-strident][+high,+front,+tense] (*TI): Assign one violation mark to every sequence of [t, t^h, t*][i, j] in the output.⁸

⁸I am assuming here that [j] is [+tense].

This markedness constraint dominates a faithfulness constraint that bans spreading of palatal features. I assume here that a number of features will be involved and [PAL] is a stand-in for these. Palatalization therefore is modeled as the spreading of [PAL] from trigger to target, assuming association lines as in Autosegmental Phonology (Goldsmith 1976, 1990) (Goldsmith, 1976, 1990). So (17) essentially blocks the association of [PAL] features to adjacent root nodes (equivalent of a constraint like *SPREAD). In tableau (18), candidate (18)a violates the higher-ranked markedness constraint *TI. Candidate (18)b, however, violates DEPLINK since [PAL] is newly associated to the root node of the stem-final consonant, but does not violate higher-ranked *TI.

(17) DEPLINK[PAL]: (following Jurgec 2011; cf. Itô, Mester, & Padgett 1995, a.o.) Let x_i be an input root node and x_o its output correspondent. Assign one violation mark iff x_o is associated with the feature [PAL] and x_i is not. (Abbreviated: DEPLINK)

(18) /mat-i/ 'eldest-NOM'

/mat-i/ 'eldest-NOM'	*TI	DEPLINK
a. [mat]i [pal]	*!	
b. [mac]i \ [pal]		*

How might one capture the fact that existing words should have fixed outputs containing [TI], but nonce words might be phonotactically dispreferred? In work on gradient well-formedness, studies often test the predictions of a model on novel data (from well-formedness ratings or rates of alternation from a wug test; e.g. Hayes et al. 2009). But as Moore-Cantwell and Pater (2017; for a similar discussion of this issue, see Pater 2010, Zuraw 2000) point out, such studies fail to account for the fact that existing words, which would be penalized by some markedness constraint, surface faithfully.

To account for these facts, Moore-Cantwell & Pater (2017) propose that the grammar contains both general as well as lexically-specific versions of constraints, with every lexical item associated with its own instantiation of a given constraint. While Moore-Cantwell and Pater implement their model to capture gradient exceptionality using Maximum Entropy Grammar (Goldwater & Johnson 2003, Hayes & Wilson 2008) it suffices presently to approximate this using regular strict domination Optimality Theory (Prince & Smolensky 1993/2004). We want the analysis to capture the fact that palatalization in Korean does apply reliably across morpheme boundaries, while words like /mati/ ‘joint’ are really exceptions but surface faithfully and reliably with [TI] sequences. I posit a lexically-specific version of the faithfulness constraint DEPLINK, in (19). As with the general schema of indexed constraint, the indexed faithfulness constraint is ranked higher than the markedness constraint and the general faithfulness constraint. The constraint in (19) penalizes the spreading of [PAL] within the specific lexical item mati ‘joint’. In tableau (20), candidate (20)b violates the higher-ranked indexed constraint, allowing the faithful candidate (20)a to win despite violating the general markedness constraint *TI.

(19) $DEPLINK_{joint}$: Let x_i be an input root node and x_o its output correspondent. Assign one violation mark iff x_o is associated with the feature [PAL] and x_i is not, and iff x_o and [PAL] are both within the root/stem ‘joint’.

(20) /mati/ ‘joint’

/mati/ ‘joint’	$DEPLINK_{joint}$	*TI	DEPLINK
☞ a. [mati] [pal]		*	
b. [maci] \ [pal]	*!		*

The tableau in (20) contrasts with that in (21). Here the target root node and the trigger [PAL] are not within the same stem mat ‘eldest’. Thus the lexically-specific constraint is not violated

by candidate (21)b which has the spreading of the [PAL] feature. The brackets are shown to indicate the domain of the stem. Notice that the difference in (20) and (21) is whether or not the trigger and target are in the same stem.

(21) /mat-i/ 'eldest-NOM'

/mat-i/ 'eldest-NOM'	DEPLINK _{eldest}	*TI	DEPLINK
a. [mat]i [pal]		*!	
a b. [mac]i [pal]			*

This analysis also extends to cases where there are two potential targets in a single stem, e.g. /tikit/ 'letter t'. The relevant lexically-specific constraint is given in (22), and the workings of the constraints are shown in tableau (23). There are four possible candidates presented. The indexed faithfulness constraint allows us to differentiate between the stem-internal [ti] sequence and the [ti] sequence which are not in the same stem. So palatalization of the stem-internal [t] in candidates (23)c and (23)d violate DEPLINK_{tikit}, but palatalization of stem-final [t] does not as in (23)b. (23)a violates *TI more than (23)b, leaving (23)b as the correct output.

(22) DEPLINK_{tikit}: Let x_i be an input root node and x_o its output correspondent. Assign one violation mark iff x_o is associated with the feature [PAL] and x_i is not, and iff x_o and [PAL] are both within the root/stem $\hat{O}tikit'$.

(23) /tikit-i/ 'letter 't'-NOM'

/tikit-i/	DEPLINK _{'tikit'}	*TI	DEPLINK
a. [tikit̥i] [p] [p]		*!*	
b. b. [tikic̥i] \ [p] [p]		*	*
c. [cikit̥i] \ [p] [p]	*!	*	*
d. [cikic̥i] \ \ [p] [p]	*!		**

With these constraints a factorial typology predicts that candidate (23)c is never attested, since this candidate is harmonically bounded (although this is not true in MaxEnt Grammar, since harmonically bounded candidates are predicted to be possible). The predicted winners based on three different rankings are shown in (24).

(24) Factorial Typology

Ranking	Predicted Candidate	Pattern
DEPLINK, DEPLINK _{tikit} \gg *TI	[tikiti]	Across-the-board non-alternation
*TI \gg DEPLINK, DEPLINK _{tikit}	[cikici]	Across-the-board alternation
DEPLINK _{tikit} \gg *TI \gg DEPLINK	[tikici]	Derived-env. pattern but specific to particular lexical items

The analysis presented here differs from previous proposals (e.g. constraint conjunction, Łubowicz (2002); sequential faithfulness, Burzio 1997, Itô & Mester 2003, Wolf 2008; under-specification, Kiparsky 1973; strength scales, Inkelas 2015; gestural timing, Bradley 2007, T. Cho 2001). In particular, the claim of the current analysis is that the stem-internal [TI] sequences are exceptions to the more general markedness constraint that disprefers such sequences. That is, the derived environment effect in Korean is really an example of exceptionality. The main theoretical prediction that departs from previous analyses relates to how acceptable nonce words with [TI] are to Korean speakers. Previous accounts, in ensuring stem-internal sequences were completely well-formed, predict that such nonce words should be perfectly acceptable. The current account, however, predicts that Korean speakers will rate words nonce words with TI as less acceptable than other comparable sequences such as [ta] or [ci], since the general constraint against such sequences (*TI) is ranked high, and there are no constraints that penalize the latter. Further confirmation of this prediction awaits future work testing native speaker knowledge.

The analysis of Korean palatalization as case of exceptionality leads us to the question of whether or not all other cases of derived environment effects fall into this category. We take up this question in the next section, focusing on another putative derived environment effect pattern, this time in Turkish.

4.4 Turkish velar deletion

Turkish velar deletion (Inkelas 2000, Inkelas & Orgun 1995, Inkelas et al. 1997, Lewis 1967, Sezer 1981, Zimmer & Abbott 1978) is another oft-cited case of a derived environment effect. Velars delete intervocalically if at a morpheme boundary (25), but are protected from deletion within morphemes (26) (data from Inkelas, 2011, 2015):

- (25) Suffix-boundary deletion
- a. /bebek-In/⁹ → [bebein] ‘baby-GEN’
 - b. cf. /bebek/ → [bebek] ‘baby-NOM’
 - c. /ipek-A/ → [ipee] ‘cotton-DAT’
 - d. cf. /ipek/ → [ipek] ‘cotton-NOM’
- (26) Deletion blocked morpheme-internally (from Inkelas 2011, 2015)
- a. /hareket/ → [hareket] ‘motion’
 - b. /sigorta/ → [sigorta] ‘insurance’
 - c. /sokak-A/ → [sokaa] ‘street-DAT’
 - d. /mekik-A/ → [mekie] ‘(weaver’s) shuttle’

Inkelas (2011, 2014; see also Sezer 1981) points out that the specifics of the process are far more nuanced. While velar deletion occurs in both native words and loanwords as well as in morphologically simplex and complex stems, it fails to apply to verb roots, although the phonological conditions are met as in (27). Here we have a minimal pair /gerek/ which can either be a noun or verb. While the noun undergoes deletion (27a), the verb does not (27b) despite satisfying the phonological conditions for velar deletion to occur.

- (27) Verbal roots
- a. /gerek-Ijor/ → [gerekijor] ‘is necessary-PROG’ (verb)
 - b. cf. /gerek-i/ → [gerei] ‘need-ACC’ (noun)

Furthermore, deletion does not seem to apply when the velar consonant is suffix-initial as in (28) compared to when it is stem-final as seen in (25) despite both context being morphologically derived.

⁹Capital vowels indicate vowels that undergo vowel harmony.

- (28) /-ki/ suffix
- a. /sene-ki/ → [seneki] 'year-REL' (*senei)
 - b. /ada-dA-ki/ → [adadaki] 'island-LOC-REL' (*adadai)

The application of velar deletion seems to be confined to polysyllabic nouns, and is usually blocked from occurring with monosyllabic nouns. Polysyllabic nouns in the Turkish Electronic Living Lexicon (TELL: Inkelas et al. 2000) corpus have a deletion rate of overall 93% whereas monosyllables have a deletion rate of only 3% (Becker et al. 2011) and Turkish speakers seem to extend this trend to nonce words in wug testing (Becker et al. 2011, Zimmer & Abbott 1978). Whether a particular lexeme alternates, however, is unpredictable. Becker et al. (2011) posit an analysis which relies on lexically-specific constraint cloning for each lexical item where the cloned faithfulness constraint (e.g. MAX) that blocks deletion is ranked above the markedness constraint *VKV. It should be noted then that even in the context in which velar deletion most readily applies, an analysis still necessitates lexically-specific constraints. Thus Turkish velar deletion appears to be a less-than-canonical derived environment effect pattern.

It should already be clear at this point that although often described as morphologically derived environment effects, Korean palatalization and Turkish velar deletion do not evince the same phonological patterns. In Korean, palatalization is, to the best of my knowledge, productive and general insofar as it applies to all suffixes which provide the appropriate phonological environment. In Turkish, however, velar deletion is morphologically restricted to certain word categories amongst other factors. This raises the question of how exactly these languages might differ in terms of their static phonotactic generalizations. In Korean, while there were indeed stem-internal exceptions to the constraint *TI, there was nonetheless a significant under-representation of such sequences such that a reliable phonotactic generalization could be learned. But how strongly is the constraint motivating velar deletion represented in the lexicon of Turkish? In the following sections, I report on the results of an investigation of two corpora parallel to what was examined in Korean.

4.4.1 TELL Corpus

First, we examine the lexical statistics in TELL. TELL contains approximately 30,000 lexemes that were compiled from a variety of existing dictionaries as well as transcribed pronunciations from two speakers of a large proportion of these in various verbal and nominal inflected forms. For current purposes, we will be querying the transcribed roots in the database. In total, the resulting corpus contained 16,757 transcribed roots.¹⁰ O/E values were calculated in the same way as in section 4. This time we were interested in comparing the velar stops /k, g/ vs. other stops and affricates in the intervocalic context vs. all other contexts. Table 4.11 shows this calculation, indicating that intervocalic velar stops occur significantly less than is expected in intervocalic position, ($\chi^2(1) = 23.16, p < 0.001$). Yet while there seems to be a statistical under-representation of velars intervocalically, one notes that that degree of under-representation (roughly indicated by the size O/E values) is different from what we saw in Korean. Recall that an O/E value of 1 indicates that a particular combination of segments co-occurs at essentially an expected rate of co-occurrence given their independent frequency of occurrence; an O/E of 0 indicates that a particular combination does not occur at all. In the Korean case, the O/E value of [TI] sequences was much smaller, and in fact closer to 0, suggesting that the rarity of such sequences. In Turkish, however, the O/E value for intervocalic velars is much closer to 1, the expected rate of co-occurrence. Thus, we might expect then that these cases might play out differently phonotactic learning (see section 7.3).

Given the grammatical category differences in terms of where velar deletion applies (broadly speaking: nouns vs. verbs; Inkelas 2011, 2014, Sezer 1981), the corpus was next split into two different sub-corpora based on whether lexical items could be classified as nouns or verbs. Although the entries in TELL are not actually tagged for lexical category, we can infer a particular entries lexical category by consulting which inflected forms are included. If the aorist, infinite or causative cells were filled in, the entry was coded as a verb; if the predicative, accusative;

¹⁰We have not filtered out roots with duplicates that correspond to different lexemes in the analysis presented here. Calculations using unique roots yield the same qualitative results. The current analysis is presented to allow for the use of lexical category information.

	V_V	Other contexts	Expected % of Cs
/k, g/	1,389 (1,524) 17.77% / 33.25% O/E = 0.91	6,426 (6,291) 82.23% / 37.25% O/E = 1.02	36.47%
Other stops/affricates	2,789 (2,654) 20.50% / 66.75% O/E = 1.05	10,823 (10,958) 79.51% / 62.75% O/E = 0.99	63.53%
Expected % of occurrence in context	19.50%	80.50%	

Table 4.11: Occurrence of /k,g/ compared to other stops/affricates in V_V vs. other contexts. Expected counts are in parentheses. Percentages in bold: row percentages; Percentages in italics: column percentages.

	V_V	Other contexts	Expected % of Cs
/k, g/	932 (1,049) 18.10% / 31.30% O/E = 0.89	4,218 (4,101) 81.90% / 36.23% O/E = 1.03	35.23%
Other stops/affricates	2,045 (1,928) 21.60% / 68.69% O/E = 1.06	7,423 (7,540) 78.40% / 63.77% O/E = 0.99	64.77%
Expected % of occurrence in context	25.57%	79.63%	

Table 4.12: Occurrence of /k, g/ compared to other stops/affricates in V_V vs. other contexts in Turkish nouns. Expected counts are in parentheses. Percentages in bold: row percentages; Percentages in italics: column percentages.

professional or possessive cells were filled in, these were coded as a noun. Roots in the corpus for which lexical category could not be determined by the method above were excluded. These usually involved only having a transcribed root with no other cells transcribed. In total, there were 11,272 identified nouns and 1786 identified verbs. Since these forms were phonetically transcribed, these transcriptions contained palatal stops [c] and [ɟ] which are allophones of /k/ and /g/ respectively (e.g. Göksel & Kerslake 2005, Lewis 1967). These were recoded as /k/ and /g/ respectively. The same calculations were then conducted on the noun corpus (Table 4.12) and verb corpus (Table 4.13) separately. There was a significant statistical under-representation

of intervocalic velars in nouns ($\chi^2(1) = 25.01$, $p < 0.001$), but not with verbs ($\chi^2(1) = 0.24$, $p = 0.63$). The difference in distribution of intervocalic velars by lexical category might be related to the fact that velar deletion at the morpheme boundary does not occur with verb but largely occurs with nouns.

	V_V	Other contexts	Expected % of Cs
/k, g/	89 (85) 12.34% / <i>41.98%</i> O/E = 1.04	632 (636) 87.66% / <i>39.97%</i> O/E = 0.99	<i>40.21%</i>
Other stops/affricates	123 (127) 11.47% / <i>58.02%</i> O/E = 0.97	949 (945) 88.53% / <i>60.02%</i> O/E = 1.00	<i>59.79%</i>
Expected % of occurrence in context	13.41%	88.18%	

Table 4.13: Occurrence of /k, g/ compared to other stops/affricates in V_V vs. other contexts in Turkish verbs. Expected counts are in parentheses. Percentages in bold: row percentages; Percentages in italics: column percentages.

4.4.2 Turkish CDS

Parallel with our investigation of Korean, in this section we examine two small corpora of CDS in Turkish available also from CHILDES (MacWhinney 2000): the Aksu-Koçcorpus (Slobin 1982) and the Turkey corpus (Turkey 2012). The Aksu-Koç corpus contains caregiver-child interactions with 34 children between the ages of 2;0 and 4;4. The Turkey corpus contains caregiver-child interactions with just one child between the ages of 1;4 and 2;4. The corpora were combined and cleaned of any English words. The resulting small corpus of 6,107 unique “lexeme” (split on spaces in the transcript) presumably contains morphologically complex forms. Note that we are not using a phonetically-transcribed corpus in this case and our analysis here is based on orthography. However, we assume here that a child might plausibly store the entire morphologically complex form at some stage of acquisition, and the lexical statistics are then calculated over this proto-lexicon. Only one of the corpora is tagged which makes it difficult to analyze lexical statistics by lexical category with these corpora, thus nouns and verbs were not

separated out. As shown in Table 4.14, there is a statistically significant under-representation ($\chi^2(1) = 311.74, p < 0.001$) of intervocalic velars, and the O/E value is much lower than what we observed with the much larger TELL corpus.

	V_V	Other contexts	Expected % of Cs
/k, g/	369 (699) 13.84% / 17.08% O/E = 0.53	2,297 (1,967) 86.16% / 37.81% O/E = 1.17	32.39%
Other stops/affricates	1,791 (1,461) 32.16% / 82.92% O/E = 1.23	3,778 (4,108) 67.84% / 62.19% O/E = 0.92	67.61%
Expected % of occurrence in context	26.23%	73.70%	

Table 4.14: Occurrence of /k, g/ compared to other stops/affricates in V_V vs. other contexts in Turkish CDS. Expected counts are in parentheses. Percentages in bold: row percentages; Percentages in italics: column percentages.

4.4.3 Modeling a Turkish grammar

Although we found evidence of statistical under-representation of intervocalic velars in both TELL and the CDS corpora, it is not a given that these would be under-represented enough to be penalized by a phonotactic learner. We note that unlike in Korean where O/E values were all under 0.2, the O/E values for Turkish intervocalic velars are all above 0.5, suggesting a smaller degree of under-representation. In this section, I present the results of two phonotactic modeling simulations using the UCLA Phonotactic Learner, one trained on the entire TELL root lexicon, another trained on the entire CDS corpus. To allow comparison with the results of the Korean simulation, the O/E accuracy criterion for the learner was set at 0.3 as well in these simulations. However, instead of being asking to find bigram constraints, the algorithm was in this case asked to find trigram constraints since the relevant phonotactic constraint governs a sequence of three segments. We assume a featural inventory in Table A2 in Appendix A.

The simulation with Child-Directed Speech utilizes corpora that are orthographically transcribed which includes what is traditionally described as soft ‘g’ <ğ>. In speech, words with

these characters are usually pronounced with lengthening of the preceding vowel instead. Although Zimmer & Orgun (1999) transcribed this as [ɣ], (Inkelas et al., 2000) report that for their speakers in TELL this is usually realized as length on the previous vowel or produced as a glide intervocalically (Göksel & Kerslake 2005, Lewis 1967). For the current purposes, all orthographic instances of <ğ> were deleted, thus leaving vowel-vowel sequences where <ğ> had occurred intervocalically to reflect as much as possible the actual phonetic realization. Otherwise, because the Turkish orthographic system is fairly accurate vis-à-vis phonetic pronunciation, no other transformations were conducted on the corpora.

Across both simulations, the phonotactic learner failed to learn a constraint that penalizes intervocalic velars (*VKV; Inkelas & Orgun 1995) specifically. Recall that the parameters of these simulations were the same as those in the Korean simulations, except for the fact that the learner was asked to also find trigram constraints (not just bigram). The fact that the phonotactic learner failed to discover a *VKV constraint in spite of the statistical under-representation of such sequences suggests that finding statistical under-representation in the input does not entail that a learner will actually notice this under-representation.¹¹ That is, it seems the degree of under-representation matters. The relevant O/E values were much higher in Turkish compared to Korean, in most cases closer to 1 than 0. Thus even if intervocalic velars are under-represented in Turkish, they are still relatively common compared to /TI/ sequences in Korean. Whatever the case, making the same modelling assumptions, the learner does not learn a phonotactic constraint against intervocalic velars in Turkish, but it does in Korean for [TI] sequences.

¹¹In fact, even when the constraint *VKV was specified ahead of time, the model failed to assign it any weight, although it assigns a very small weight when trained on the CDS corpus (0.37). This was the case even when we allowed the learner to find more constraints (250). We might worry that the difference between Korean, where the learner did discover *TI, and Turkish, where it did not discover *VKV, is due to the difference in complexity between the two constraints. That is, for Korean, it was only necessary to search the space of bigram constraints, while Turkish requires searching the larger space of trigram constraints. As a comparison, a trigram model was run with the Korean NAKL corpus. While the model failed to discover the simple *TI constraint, instead breaking the pattern into more specific constraints like *tti and *tj that combine to penalize TI, when the more general *TI constraint was specified at the start of the simulation, the model always assigned this constraint a sizable weight (~1.5).

4.4.4 Island of reliability: polysyllabic nouns?

Given Zimmer & Abbott (1978) finding that Turkish speakers extend velar deletion to nonsense polysyllabic words, there is the possibility that there is an island of reliability for a phonotactic generalization (Albright 2002). That is, there might possibly be a cophonology (Inkelas & Zoll 2007) or sublexicon (Becker & Gouskova 2016) of polysyllabic nouns in which there is a strong and reliable phonotactic generalization (*VKV). To investigate this, we identified polysyllabic nouns from our noun corpus described in §7.1 and calculated the same O/E values as above shown in Table 4.15. We essentially find the same results as we found in Table 4.12 when we considered all nouns. Given the failure to find a phonotactic constraint in the previous section, it is unlikely that we would find one here either. Thus, it seems that although Turkish speakers do seem to generalize velar deletion to novel polysyllabic nouns (Zimmer & Abbott 1978)), they do not seem to be relying on a phonotactic generalization within the sublexicon of polysyllabic nouns. In fact, a phonotactic learner similarly fails to penalize intervocalic velars in this case as well, although if the relevant constraint is fed at the initialization of the model, it is given a miniscule weight after learning (~ 0.09).

	V_V	Other contexts	Expected % of Cs
/k, g/	932 (1,044) 20.05% / <i>31.31%</i> O/E = 0.89	3,716 (3,604) 79.95% / <i>36.14%</i> O/E = 1.03	<i>35.05%</i>
Other stops/affricates	2,045 (1,933) 23.75% / <i>68.69%</i> O/E = 1.06	6,567 (6,679) 76.25% / <i>63.86%</i> O/E = 0.98	<i>64.95%</i>
Expected % of occurrence in context	22.45%	77.55%	

Table 4.15: Occurrence of /k, g/ compared to other stops/affricates in V_V vs. other contexts in Turkish polysyllabic nouns. Expected counts are in parentheses. Percentages in bold: row percentages; Percentages in italics: column percentages.

4.4.5 Summary

In contrast to what we found with Korean, velar deletion in Turkish does not have strong phonotactic support from the lexicon. We should note that although there is a statistically significant under-representation of intervocalic velars in the lexicon as a whole, the set of nouns as well as in CDS, this under-representation does not actually translate into the robust learning of a phonotactic constraint that penalizes such sequences. It should be pointed out though that the CDS corpus was much smaller than any of the other corpora examined in both Korean and Turkish with just over 6,000 items.

Of course, with a richer phonological representation, it might be possible for a learner to arrive at such a constraint. Indeed this is what Inkelas & Orgun (1995) propose, arguing that velars which are syllabified at an earlier level of the morphological derivation (i.e. stem-internal velars), are immune from deletion. Only unsyllabified, i.e. stem-final velars, are subject to the *VKV constraint which triggers deletion. Their analysis further protects deletion from occurring with verbal suffixes since these are attached at an earlier level of the derivation where *VKV is not applicable. The important difference here is that there is not a simple phonotactic solution to Turkish velar deletion, unlike in Korean where a constraint is readily arrived at without reference to the derivational cycle. As Inkelas (2011) points out, velar deletion in Turkish is essentially a case of morphologically-conditioned phonology.

4.5 Another brief example: Finnish Assibilation

In this section, I briefly present the results of a small corpus study of the lexicon of Finnish. Recall that in Finnish, /t/ assibilates to [s] before [i] but only across a morpheme boundary (see §2.1). The basic pattern is repeated below with the alternation occur across a morpheme boundary in (29) but not within stems (30). As Anttila (2006) and Karlsson (1983) have pointed out, however, not all suffixes engender the alternation despite meeting the phonological requirements (31). In fact, the alternation seems to be limited to three suffixes in (32).

(29) /t/ → [s] /__ across a morpheme boundary (*ti):

- a. /halut-i/ → [halusi] ‘want-PAST’
- b. cf. /halut-a/ → [haluta] ‘want-INF’
- c. /hakkat-i/ → [hakkasi] ‘beat-PAST’
- d. cf. /hakkat-a/ → [hakkata] ‘beat-INF’

(30) /ti/ sequences surface faithfully within stems:

- a. /tilat-i/ → [tilasi] ‘order-PAST’ *[silasi]
- b. /koti/ → [koti] ‘home’ *[kosi]

(31) Non-triggering suffixes.

- a. Instrumental /-ime/: /lentä-ime-n/ → lentimen ‘fly-INST-GEN’
(*lensimen)
- b. Conditional /-isi/: /tunte-isi/ → tuntisi ‘feel-COND’ (*tunsisi)

(32) Triggering suffixes for assibilation

- a. Plural /-i/: /vuote-i-nA/¹² → vuosina ‘year-PL-ESS’
- b. Past tense /-i/: /huuta-i-vAt-kO/ → huusivatko ‘shout-PAST-3RPL-Q’
- c. Superlative /-impA/: /uute-impA-nA/ → uusimpana ‘new-SUP-ESS’

This pattern shares in common the morphologically restricted nature as the Turkish velar deletion process described in the previous sections. The question then is whether the /ti/ sequences in the lexicon, which are repaired across a morpheme boundary, are underrepresented or not. I consulted the “The Frequency Lexicon of the Finnish Newspaper Language” (CSC - IT Center for Science 2004) which contains 9,996 of the most common lemmas taken from Finnish newspapers. CV sequences were tabulated as in the previous sections above 4.16. We notice that the O/E value for /ti/ sequences is 1 indicating that in this corpus /ti/ sequences occur

¹²Vowels in upper case undergo vowel harmony.

at the expected rate. A chi-square test confirms that there is no skew in the distribution ($\chi^2(1) = 0.065, p = 0.79$), suggesting that there is no underrepresentation of /ti/ sequences in Finnish. Therefore, there does not seem to be phonotactic support for the assibilation process that occurs across a morpheme boundary in Finnish, making this parallel to the Turkish velar deletion case described above. In both cases, there is no strong phonotactic support for the alternation, and the alternation is highly restricted to specific morphological contexts.

	/i/	Other Vs	Expected % of Cs
/t/	795 (790) 16.12% / 46.19% O/E = 1.00	4,138 (4143) 83.88% / 45.83% O/E = 1.00	45.88%
Other Stops	926 (931) 15.92% / 53.08% O/E = 1.00	4,892 (4,887) 84.008% / 54.17% O/E = 1.00	54.12%
Expected % of Vs	16.01%	83.99%	

Table 4.16: Occurrence of /t/ compared to other stops before /i/ vs. other contexts in Finnish. Expected counts are in parentheses. Percentages in bold: row percentages; Percentages in italics: column percentages.

4.6 General Discussion

The present chapter set out with the aim of examining in closer detail the quality of the phonotactic generalizations available in languages with derived-environment effect patterns. We investigated in detail two well-known examples of derived-environment effects: Korean palatalization and Turkish velar deletion. I also briefly examined the lexical statistics as it pertains to the well-known derived environment pattern in Finnish assibilation. We found differing results with respect to the strength of the phonotactic constraint motivating the phonological alternation in each case. These results were consistent across different types of corpora within each language. Specifically, while there is a robust, albeit gradient, phonotactic constraint dispreferring [TI] that is able to motivate alternations in Korean, no such constraint dispreferring intervocalic velars is readily available in Turkish. Here I discuss the implications

of these results for morphologically derived environment effects in general, as well as for the relation between static phonotactic generalizations and phonological alterations, and the learning of such generalizations.

4.6.1 Derived environment effects as a unified phenomenon?

A major finding of the current study is that despite structural similarities, patterns previously described together as examples of morphologically derived environment effects are by no means a unified phenomenon. On the surface, both Korean palatalization and Turkish velar deletion share structural similarities in that the phonological process is purported to only apply when the environment is achieved by virtue of the concatenation of two morphemes, but not within morphemes. This surface similarity belies stark differences once we start looking more closely at the quantitative patterns of alternations and the lexicon. On the one hand, in Korean, there is a productive general phonological constraint that drives palatalization such that even stem-internal forms violating this constraint are rare. On the other hand, in Turkish, velar deletion is morphologically conditioned and not general, and the stem-internal forms that contain intervocalic velars while occurring slightly less than expected, are nonetheless frequent enough that there is no reliable penalty against them. Thus, these two cases are in some respects opposites of each other.

More broadly speaking, our results call into question the traditional notion of morphologically derived environment effects. In Korean and Turkish, both considered uncontroversial canonical examples of morphologically derived environment effects, the actual patterns do not hold up to scrutiny, especially when one takes into consideration the assumptions laid out in (5) and (11) repeated here as (29) and (30).

- (33) DERIVED-ENVIRONMENT CONDITION: Morphological derivedness is a necessary and sufficient condition for a process to occur (variously stated as the Strict Cycle Condition or the Revised Alternation condition; Kiparsky 1973, 1982b).

- (34) PHONOTACTIC PRODUCTIVITY: Static phonotactic patterns are completely productive (i.e. morpheme-internal sequences are phonotactically well-formed).

Each of the cases examined here fails on one of these assumptions, Turkish on (33) (Inkelas 2011, 2014, Sezer 1981), and Korean on (34).¹³ A cursory survey of the literature on morphologically derived environment effects indicates that many other well-known cases of derived environment effects are similarly murky. Finnish assibilation (Anttila 2006) discussed in §2.1 and §4.6 is another example, which like Turkish velar deletion, fails on (33): a derived-environment provides a necessary but not sufficient condition for the phonological process to apply.

Thus it seems that the most well-known cases of derived-environment effects are less canonical than had previously thought, an observation previously articulated by Inkelas (2011, 2014). Whereas Inkelas (2011, 2014) conjectures that many morphologically derived environment effects might be instances of morphologically-conditioned phonology, our examination of Korean here suggests that at least some putative derived environment effect cases are really instances of gradient exceptionality. In Korean, stem-internal [TI] sequences are exceedingly rare and come mostly from loanwords, the latter of which was first observed by Y.-M. Y. Cho (2009), and confirmed quantitatively in more detail here. As Y.-M. Y. Cho (2009) further points out, a number of other cases are similar in having exceptions that are mostly loanwords (e.g. Finnish Vowel Coalescence (Anttila 2009), Polish First Velar Palatalization (Łubowicz 2002), see Y.-M. Y. Cho (2009) for other examples). Thus, a closer inspection of the lexicon of other cases of derived environment effects might well reveal a similar picture as in Korean where we observe a dispreference tautomorphemically for sequences repaired heteromorphemically.

Taken together, these results argue for the assertion that patterns classified traditionally as morphologically derived environment effects might in reality be cases of either gradient

¹³Korean arguably also violates (29) if we take into account compounding, where underlying /ti/ sequences can also occur across compound (or prefix) boundaries (Oh 1995). But there is a conspiracy here where n-insertion can variably occur to fix such sequences. It is an open question then as to the extent to which n-insertion applies to prevent underlying /ti/ from surfacing (see also Martin 2011).

exceptionality or morphologically conditioned phonology. Given this, one wonders what utility there is to the very notion of a ‘derived environment effect’.

4.6.2 Derived environment effects and the relationship between static and dynamic generalizations

In the previous section, I argued for the fact that the two instances of morphologically derived environment effects examined in this paper are really two different patterns. But while they are indeed different from each other, they are similar insofar as the static generalizations and dynamic ones motivating phonological alternations are much more closely linked than the surface facts would suggest.

Morphologically derived environment effects are instances in which there is a loss of a static phonotactic generalization, although the dynamic one motivating the alternation is still active (Paster, 2013). Paster (2013) further suggests that these two levels of generalization are somewhat independent to the extent they can develop historically in different ways. Under this view, however, we would expect to find that alternations are entirely productive heteromorphemically and analogous segmental sequences are not in any way dispreferred tautomorphemically (i.e. the patterns should conform to (29) and (30)). Yet as argued for above, this does not seem to be the case. In fact, while there are indeed exceptions to the general constraint motivating palatalization in Korean, constraint-violating stem-internal sequences are rare enough for a reliable static phonotactic constraint against these same sequences to be learned. In Turkish, contrastively, there is no strong phonotactic support for velar deletion - stem-internal VKV sequences while less frequent than expected are nonetheless well-formed - and the alternation is highly morphologically conditioned. Although the alternation is productive (Becker et al. 2011, Zimmer & Abbott 1978), it is unclear that there’s a straightforward way to account for this without referring to morphological information or word category. Thus these cases of mismatch in the static and dynamic generalizations turn out to only be apparent (Table 4.17). In Korean, where the alternation is productive, there is also a gradient, static phonotactic generalization. In Turkish, where there is no static phonotactic generalization, the alternation is much more

constrained. This general pattern is one that is shared by both Turkish and Finnish in that there is no static phonotactic pattern in the lexicon that supports the alternation, and in both cases the alternation is highly morphologically conditioned. These facts perhaps suggest a bias for the maintenance of broader, more general, phonological generalizations (see also Martin 2011), or at the very least a bias to maintain similar generalizations across different levels. So while it is true that static phonotactic patterns and alternations can pull apart historically (as happened in Korean through independent phonological processes and borrowings), the two do not seem to be entirely independent. Thus any theory that posits complete independence of stem phonotactic generalizations and alternations misses this relationship.

Table 4.17: Static phonotactics vs. alternations

	Static generalization	Alternation
Korean	Yes (weaker)	Yes
Turkish	No	No-ish (Yes but very much constrained)

Our results also have interesting implications from the phonological learning standpoint. Constraint-based learning models (Hayes 2004, Hayes & Wilson 2008, Prince & Tesar 2004, Tesar & Prince 2007) predict that patterns like morphologically derived environment effects should be more difficult to learn. If phonotactic learning occurs prior to alternation learning, then a learner might initially learn the wrong phonotactic generalization. For example, in Turkish, a learner might initially learn to accept [VkV] although she would need to learn to delete [k] in polysyllabic nouns later. In fact, Turkish learners do not seem to have access to a robust phonotactic constraint that motivates velar deletion, even when we only consider the part of the lexicon (polysyllabic nouns) where velar deletion reliably occurs. It is maybe not surprising that the alternation is highly constrained in this case. Contrastively, if offending words are rare enough, a probabilistic learner (e.g. Hayes & Wilson 2008, Jarosz 2006, 2011) will nonetheless learn an, albeit weaker, more gradient phonotactic constraint. This is exactly the situation in Korean with regards to palatalization. Such a finding thus argues against

non-probabilistic learners such as Biased Constraint Demotion (Prince & Tesar 2004) which would likely be too brittle to handle gradient well-formedness.

4.6.3 Conclusion

In this paper, we examined the relationship between static phonotactic patterns and phonological alternations in two paradigmatic examples of derived environment effects: Korean palatalization and Turkish velar deletion. We found that neither case conformed to the usual assumptions regarding derived environment effects. They are in fact mirror-images of each other. In Korean, there is an active alternation with a gradient static phonotactic constraint, whereas in Turkish, there is a constrained alternation with no static phonotactic constraint. The significance of our result lies in the fact that these supposedly structurally similar cases are not canonical cases of derived environment effects, and moreover, they call into question the notion of morphologically derived environment effects. Our results also suggest some bias to maintain similar generalizations across different domains (tautomorphemically vs. heteromorphemically).

A number of avenues of investigation remain open. For one, it is a puzzle as to why Korean allowed [TI] sequences to be borrowed in faithfully in the first place. Borrowing occurred after the counterfeeding diachronic sound change was complete (Y.-M. Y. Cho 2009), producing novel [TI] sequences. Yet as we saw from the corpus results, TI was strongly under-represented in the native and Sino-Korean lexicons. So what allowed for faithful borrowings of [TI]? It is possible here that other considerations regarding loanword adaptation are at play here which prefer such forms (e.g. orthographic effects: Daland, Oh, & Kim 2015).

Further, given our finding of two different “types” of putative derived environment cases, one wonders whether we would be able to construct a typology of derived environment patterns that actually captures what these patterns look like. That is, do all derived environment patterns reduce to either gradient exceptionality or morphologically conditioned phonology? Finally, the findings of our study lead us to the conjecture that a derived environment effect is likely to

only be productive if it is supported by a phonotactic generalization in the lexicon. If this is unavailable, then learning breaks down, resulting in morphologically-condition phonology.

Appendix

	syllabic	consonantal	sonorant	continuant	del. rel.	approximant	nasal	voice	spread glot.	const. glot	labial	round	coronal	anterior	distributed	strident	lateral	dorsal	high	low	front	back	tense
p	-	+	-	-	-	-	-	-	-	-	+	-	-	0	0	0	-	0	0	0	0	0	0
t	-	+	-	-	-	-	-	-	-	-	-	-	+	+	-	-	-	0	0	0	0	0	0
k	-	+	-	-	-	-	-	-	-	-	-	-	-	0	0	0	-	+	+	-	0	0	0
c	-	+	-	-	+	-	-	-	-	-	-	-	+	+	+	+	-	+	+	-	+	-	0
p ^h	-	+	-	-	-	-	-	-	+	-	+	-	-	0	0	0	-	0	0	0	0	0	0
t ^h	-	+	-	-	-	-	-	-	+	-	-	-	+	+	-	-	-	0	0	0	0	0	0
k ^h	-	+	-	-	-	-	-	-	+	-	-	-	-	0	0	0	-	+	+	-	0	0	0
c ^h	-	+	-	-	+	-	-	-	+	-	-	-	+	+	+	+	-	+	+	-	+	-	0
p*	-	+	-	-	-	-	-	-	-	+	+	-	-	0	0	0	-	0	0	0	0	0	0
t*	-	+	-	-	-	-	-	-	-	+	-	-	+	+	-	-	-	0	0	0	0	0	0
k*	-	+	-	-	-	-	-	-	-	+	-	-	-	0	0	0	-	+	+	-	0	0	0
c*	-	+	-	-	+	-	-	-	-	+	-	-	+	+	+	+	-	+	+	-	+	-	0
s	-	+	-	+	+	-	-	-	+	-	-	-	+	+	-	+	-	0	0	0	0	0	0
s*	-	+	-	+	+	-	-	-	-	+	-	-	+	+	-	+	-	0	0	0	0	0	0
h	-	-	-	+	+	-	-	-	+	-	-	-	-	0	0	0	-	0	0	0	0	0	0
m	-	+	+	-	0	-	+	+	-	-	+	-	-	0	0	0	-	0	0	0	0	0	0
n	-	+	+	-	0	-	+	+	-	-	-	-	+	+	-	-	-	0	0	0	0	0	0
ŋ	-	+	+	-	0	-	+	+	-	-	-	-	-	0	0	0	-	+	+	-	0	0	0
l	-	+	+	+	0	+	-	+	-	-	-	-	+	+	-	-	+	0	0	0	0	0	0
j	-	-	+	+	0	+	-	+	-	-	-	-	-	0	0	0	-	+	+	-	+	-	+
w	-	-	+	+	0	+	-	+	-	-	+	+	-	0	0	0	-	+	+	-	-	+	+
i	+	-	+	+	0	+	-	+	-	-	-	-	-	0	0	0	-	+	+	-	+	-	+
i	+	-	+	+	0	+	-	+	-	-	-	-	-	0	0	0	-	+	+	-	-	+	+
u	+	-	+	+	0	+	-	+	-	-	+	+	-	0	0	0	-	+	+	-	-	+	+
e	+	-	+	+	0	+	-	+	-	-	-	-	-	0	0	0	-	+	-	-	+	-	+
ø	+	-	+	+	0	+	-	+	-	-	+	+	-	0	0	0	-	+	-	-	+	-	+
o	+	-	+	+	0	+	-	+	-	-	+	+	-	0	0	0	-	+	-	-	-	+	+
ʌ	+	-	+	+	0	+	-	+	-	-	-	-	-	0	0	0	-	+	-	-	-	+	-
a	+	-	+	+	0	+	-	+	-	-	-	-	-	0	0	0	-	+	-	+	-	-	+

Table 4.18: Korean Features

	syllabic	long	consonantal	sonorant	continuant	del. rel.	approximant	tap	trill	nasal	voice	spread glot.	const. glot	labial	round	labiodental	coronal	anterior	distributed	strident	lateral	dorsal	high	low	front	back	tense
p	-	-	+	-	-	-	-	-	-	-	-	-	-	+	-	-	-	0	0	0	-	-	0	0	0	0	0
t	-	-	+	-	-	-	-	-	-	-	-	-	-	-	-	-	+	+	-	-	-	-	0	0	0	0	0
k	-	-	+	-	-	-	-	-	-	-	-	-	-	-	-	-	-	0	0	0	-	+	+	-	0	0	0
b	-	-	+	-	-	-	-	-	-	-	+	-	-	+	-	-	-	0	0	0	-	-	0	0	0	0	0
d	-	-	+	-	-	-	-	-	-	-	+	-	-	-	-	-	+	+	-	-	-	-	0	0	0	0	0
g	-	-	+	-	-	-	-	-	-	-	+	-	-	-	-	-	-	0	0	0	-	+	+	-	0	0	0
ʃ	-	-	+	-	-	+	-	-	-	-	-	-	-	-	-	-	+	-	+	+	-	-	0	0	0	0	0
ʒ	-	-	+	-	-	+	-	-	-	-	+	-	-	-	-	-	+	-	+	+	-	-	0	0	0	0	0
f	-	-	+	-	+	+	-	-	-	-	-	-	-	+	-	+	-	0	0	0	-	-	0	0	0	0	0
s	-	-	+	-	+	+	-	-	-	-	-	-	-	-	-	-	+	+	-	+	-	-	0	0	0	0	0
ʃ	-	-	+	-	+	+	-	-	-	-	-	-	-	-	-	-	+	-	+	+	-	-	0	0	0	0	0
h	-	-	+	-	+	+	-	-	-	-	-	+	-	-	-	-	-	0	0	0	-	-	0	0	0	0	0
v	-	-	+	-	+	+	-	-	-	-	+	-	-	+	-	+	-	0	0	0	-	-	0	0	0	0	0
z	-	-	+	-	+	+	-	-	-	-	+	-	-	-	-	-	+	+	-	+	-	-	0	0	0	0	0
ʒ	-	-	+	-	+	+	-	-	-	-	+	-	-	-	-	-	+	-	+	+	-	-	0	0	0	0	0
m	-	-	+	+	-	0	-	-	-	+	+	-	-	+	-	-	-	0	0	0	-	-	0	0	0	0	0
n	-	-	+	+	-	0	-	-	-	+	+	-	-	-	-	-	+	+	-	-	-	-	0	0	0	0	0
r	-	-	+	+	+	0	+	+	-	-	+	-	-	-	-	-	+	+	-	-	-	-	0	0	0	0	0
j	-	-	-	+	+	0	+	-	-	-	+	-	-	-	-	-	-	0	0	0	-	+	+	-	+	-	+
l	-	-	+	+	+	0	+	-	-	-	+	-	-	-	-	-	+	+	-	-	+	-	0	0	0	0	0
i	+	-	-	+	+	0	+	-	-	-	+	-	-	-	-	-	-	0	0	0	-	+	+	-	+	-	+
y	+	-	-	+	+	0	+	-	-	-	+	-	-	+	+	-	-	0	0	0	-	+	+	-	+	-	+
u	+	-	-	+	+	0	+	-	-	-	+	-	-	+	+	-	-	0	0	0	-	+	+	-	-	+	+
e	+	-	-	+	+	0	+	-	-	-	+	-	-	-	-	-	-	0	0	0	-	+	-	-	+	-	+
ø	+	-	-	+	+	0	+	-	-	-	+	-	-	+	+	-	-	0	0	0	-	+	-	-	+	-	+
o	+	-	-	+	+	0	+	-	-	-	+	-	-	+	+	-	-	0	0	0	-	+	-	-	-	+	+
a	+	-	-	+	+	0	+	-	-	-	+	-	-	-	-	-	-	0	0	0	-	+	-	+	-	+	0

Table 4.19: Turkish Features

CHAPTER 5

Simulating learning in derived-environment effects

In the previous chapter, we examined the phonotactic generalizations that were available in two cases of derived-environment effects: Korean palatalization and Turkish velar deletion. While previous analyses (e.g. Wolf 2008, Łubowicz 2002) had largely assumed that stem-internal sequences that went against the cross-morpheme generalization were well-formed, I showed that at least in one case, Korean palatalization, this assumption does not seem to hold. I then proposed a sketch of a phonological analysis that argued for the purported derived-environment pattern in Korean as a case of lexical exceptionality: words with surface [ti] in them are encoded as lexical exceptions due to their rarity, thereby allowing for the general markedness to drive the categorical alternation. It was also argued that another well-known derived-environment pattern, Turkish velar deletion, is a case of morphologically-conditioned phonology insofar as the alternation is only licensed in certain morphological contexts, although in these specific contexts, it is largely productive. I further suggested that in both cases the purported mismatches in terms of phonotactics and alternations belie generalizations that are actually more similar, potentially indicating a bias for more general constraints (Martin 2011).

In this chapter, I sketch out how a bias in favor of more general phonological constraints can be achieved in a MaxEnt grammar (see Chapter 2). I implement this by capitalizing on the priors imposed on constraints and weights in MaxEnt learners which I discuss further in §1. In §2-5, I present toy simulations of five different languages that differ in the degree to which phonotactics and alternations mismatch.

5.1 Favoring generality: Implementing with a prior

As described in Chapter 2, MaxEnt learners contain two free parameters for each constraint (priors). The first is μ which is the preferred weight for a given constraint. In our simulations below, these will be determined by the outcome of phonotactic learning of the various simulations presented in the previous section. The second parameter is σ^2 which determines how much the weights are allowed depart from their preferred value μ . When σ^2 is small, it is more difficult for the weights to depart from their preferred value. In MaxEnt, when two constraints are able to explain the same data, and they share the same μ and σ^2 values, weight is uniformly distributed across both constraints. Our simulations will make use of two types of markedness constraints: structure-blind and structure-sensitive constraints. Structure-blind constraints do not refer to any morphological structure, whereas structure-sensitive constraints do. So for example, both structure-blind $*_{TI}$, which penalizes *any* [ti] sequence, and structure-sensitive $*_{T+I}$, which penalizes [ti] sequences spanning a morpheme boundary, are able to explain the categorical alternation of /t+i/ to [c+i]. All else being equal, a learner will assign the burden of explanation (in terms of weight) equally across both constraints. Here, however, I will present a model that encodes a bias for weights to be preferentially placed on the more general structure-blind constraint $*_{TI}$ over the more specific structure-sensitive $*_{T+I}$. So constraints in these simulations will have different σ^2 values. This perhaps captures the intuition that, all else being equal, learners prefer more general, less complex constraints (Moreton & Pater 2012a, Hayes & Wilson 2008), where $*_{T+I}$ is more complex than $*_{TI}$ since the former references morphological structure whereas the latter does not. The working assumption here is that it is more difficult and more complex to learn structure-sensitive constraints, thus learners will prefer the explanatory burden to be placed on the general constraint and not the structure-sensitive one.

We will assume the following constraints for this simulation:

- (1) Constraints:
 - a. *TI: No TI sequences (anywhere)
 - b. *T+I: No TI sequences across a morpheme boundary
 - c. FAITH: Faithfulness constraint preventing palatalization
 - d. FAITH_{LexC1}: Lexically-specific faithfulness constraint preventing palatalization within a stem (for Lexical Class 1) (see Chapter 4 for a definition).

Here, we will assume that exceptional words with [ti] sequences that do not palatalize within stems are part of a lexical class: Lexical Class 1 (LexC1). Thus the constraint FAITH_{LexC1} applies only to words in the input, and not to novel /ti/ words. For the constraints *TI, *FAITH and *FAITH_{LexC1} will have σ^2 values set at a constant 500. What we will vary is the σ^2 values for the structure-sensitive constraint: 0.1, 10, 50, 100, 200, 300, 400 and 500. When σ^2 is small, we expect little weight to be assigned to the structure-sensitive *T+I constraint, and that the weight assigned to this constraint will increase as the σ^2 value increases to be the same as the other constraints. The μ value for this constraint is set at 0. For each language, we will examine the model's prediction on untrained data. Specifically, we're interested in examining what the learner's behavior is as the σ^2 value changes.

5.2 Across-the-board language

We start the examination with a baseline language modeled on Korean palatalization except that the ban on [ti] is across-the-board, so there is no mismatch between phonotactics and alternations. The model was initialized with a μ value of 6.169 for *TI which was the weight assigned to this constraint in the learning simulation in Chapter 4 in which the data set did not contain any [ti] sequences. All other constraints are assigned a μ of 0. The learner is trained on the input in Table 5.1. In this case, the lexically-specific FAITH_{LexC1} is not relevant, since

there are no lexical exceptions to palatalization. This is left in the tableaux in Table 5.1 for completeness and to aid comparison with the other simulations in this chapter that do make use of this constraint. The learner is then tested on novel forms in Table 5.2. These will be the same test items used in all the simulations reported below.

Input	Ouput	Freq.	FAITH _{LexC1}	*TI	*T+I	FAITH
/t+i/	[t+i]	0		1	1	
	[c+i]	1				1
/ti/ _{LexC1}	[ti]	0		1		
	[ci]	0	1			1
/c+i/	[t+i]	0		1	1	1
	[c+i]	1				
/ci/	[ti]	0		1		1
	[ci]	1				

Table 5.1: Input for across-the-board language

Input	Ouput	Freq.	FAITH _{LexC1}	*TI	*T+I	FAITH
/t+i/ _{novel}	[t+i]	0		1	1	
	[c+i]	0				1
/ti/ _{novel}	[ti]	0		1		
	[ci]	0				1
/c+i/ _{novel}	[t+i]	0		1	1	1
	[c+i]	0				
/ci/ _{novel}	[ti]	0		1		1
	[ci]	0				

Table 5.2: Novel test items

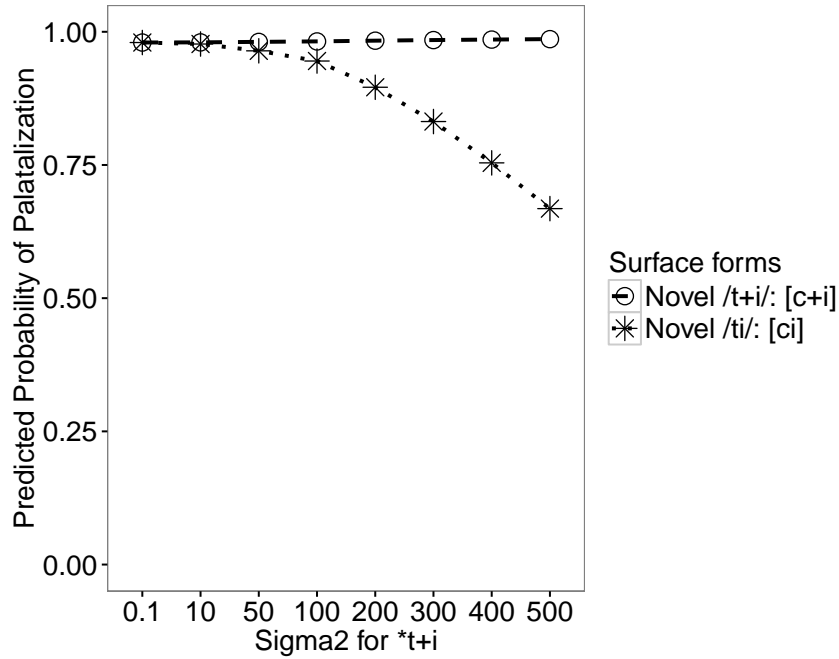
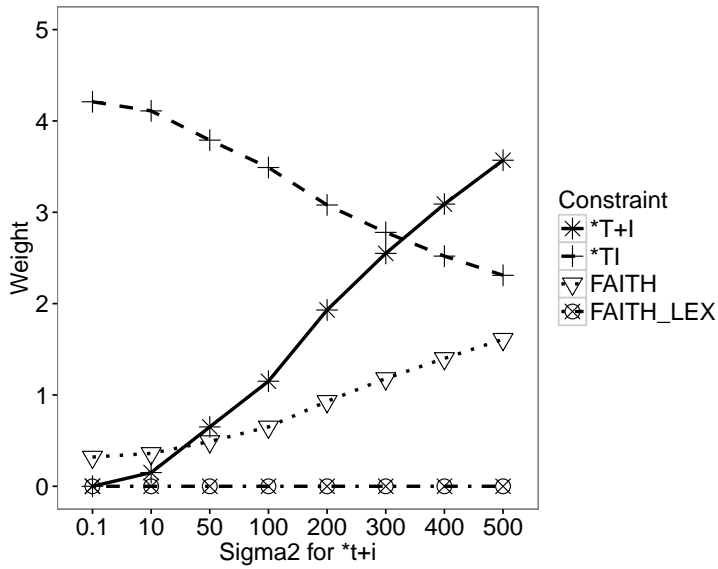



Figure 5.1: Predicted probability of palatalization



(a) Learnt constraint weights



(b) Hasse Diagram of desired grammar

Figure 5.2: Model outputs: Constraint weights

We expect in this language that the σ^2 value of the $*_{T+I}$ constraint does not need to be very high in order to learn the language since the generalization does not have to be accounted for using the structure-sensitive $*_{T+I}$ constraint. Figures 5.1 and 5.2 show the predicted probability of palatalization of the inputs $/ti/_{novel}$ and $/t+i/_{novel}$ with a range of σ^2 values, and the learnt weights for each of the constraints respectively. Here we concentrate on the forms with underlying $/ti/$ sequences, since underlying $/ci/$ sequences are always predicted to surface faithfully. We notice that as the value of σ^2 increases, the weight assigned to $*_{T+I}$ increases, and so does the rate of palatalization for $/t+i/$ (dashed line in Figure 5.1). As the weight of $*_{T+I}$ increases, this means that the weight of structure-blind $*_{TI}$ can be lower without sacrificing any accuracy on the training data. But this has the unintended consequence of lowering the rate of palatalization for $/ti/_{novel}$ which was not in the training data. As far as the learner is concerned, this is of no consequence since it is still modeling the training data accurately. But typically, we generally expect human learners to reject data that is not presented in training (e.g. experiments in Chapter 3, Skoruppa & Peperkamp 2011, Linzen & Gallagher 2014). The learner is, therefore, not accurately modeling human behavior, since $[ti]$ sequences should be rejected. Counterintuitively, the model with a lower σ^2 value is preferred in this case.

The simulations, perhaps unsurprisingly, shows how using just a structure-blind constraint, with a very low σ^2 value for the structure-sensitive constraint, accurately models the input data in an across-the-board language. Somewhat counterintuitively, however, the learner actually requires that σ^2 be low in order for it to model the training data in a way that conforms to what we expect human learners to do, in particular it predicts a weakening of the stem-internal phonotactic constraint as the σ^2 value increases.

5.3 Toy Korean

In this section we present a simulation of Toy Korean. In this toy language (Table 5.3), there is a frequency difference between stem-internal $[ti]$ and $[ci]$ sequences in the training data. This aims to capture the fact that Korean has an under-representation of $[ti]$ sequences in the

lexicon as we found in Chapter 4. We will assume the same constraints as in §2 although in this language *TI is initialized with a μ of 1.916, which was the weight learnt in Chapter 4 for Korean. All other constraints were initialized with a μ value of 0.

Input	Ouput	Freq.	FAITH _{LexC1}	*TI	*T+I	FAITH
/t+i/	[t+i]	0		1	1	
	[c+i]	1				1
/ti/ _{LexC1}	[ti]	0.1		1		
	[ci]	0	1			1
/c+i/	[t+i]	0		1	1	1
	[c+i]	1				
/ci/	[ti]	0		1		1
	[ci]	0.9				

Table 5.3: Input for Toy Korean

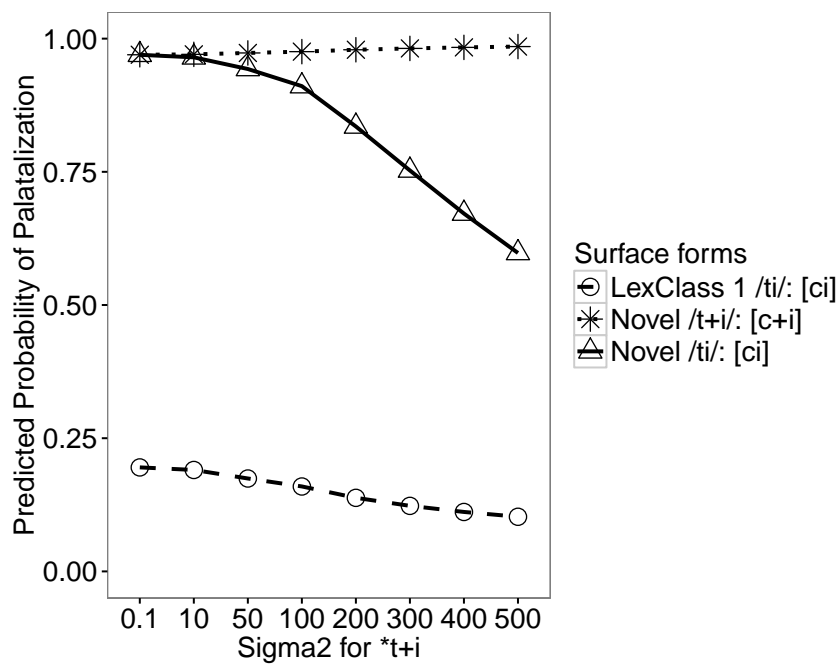


Figure 5.3: Model outputs: Predicted probability of palatalization

Figures 5.3 and 5.4 shows the outcome of learning with different σ^2 values, with the predicted probabilities of palatalization for inputs $/ti/_{LexC1}$, $/ti/_{novel}$ and $/t+i/_{novel}$. Overall, the learner is accurate on the alternation, regardless of the σ^2 value. Interestingly, the learner overpredicts palatalization of $/ti/_{LexC1}$, suggesting a constant pressure to palatalize, even for existing exceptional words. The model does best when σ^2 is around 100. Here the learner is able to learn a more or less categorical phonological alternation (98%), while maintaining a strong preference for [ci] (91%) for an input $/ti/_{novel}$. Thus, when σ^2 is at 100, the language is able to maintain, on the one hand, an essentially categorial alternation, but at the same a gradient phonotactic constraint in stems.




Figure 5.4: Model outputs: Constraint weights

5.4 A "true" derived-environment effect

Having modeled Toy Korean, we now present a simulation of what is required to learn a "true" derived-environment language. By "true" derived-environment language, I mean here a language in which stem-internal [ti] sequences are completely well-formed, but there is

nonetheless a categorical phonological alternation. The learner was trained on the input in Table 5.4, with an equal frequency of /ci/ and /ti/. In this language, the goal is to learn to accept [ti] sequences when they are stem-internal, so the learner should predict that for a novel input /ti/_{novel}, the preferred output is faithful [ti] and not [ci]. The μ value for all constraints was set at 0 for this simulation

Input	Ouput	Freq.	FAITH _{LexC1}	*TI	*T+I	FAITH
/t+i/	[t+i]	0		1	1	
	[c+i]	1				1
/ti/ _{LexC1}	[ti]	0.5		1		
	[ci]	0	1			1
/c+i/	[t+i]	0		1	1	1
	[c+i]	1				
/ci/	[ti]	0		1		1
	[ci]	0.5				

Table 5.4: Input for a “true” derived-environment language

The results of the simulations are shown in Figures 5.5 and 5.6. Even when the σ^2 value is equal across all the constraints (at 500), the learner nonetheless never quite predicts a preference for [ti] over [ci] as an output for /ti/_{novel}. That is, even without a bias against placing weight on the structure-sensitive constraint, the learner still shows a persistent preference for [ci] over [ti], in accords with the alternation. The issue is that in order to get a categorical alternation while preserving stem-internal [ti] sequences, the learner must weight $*T+I \gg FAITH$ and $FAITH \gg *TI$. Even when the σ^2 values are uniform across all the constraints, the learner still cannot achieve the correct dominance relationship between the relevant constraints. Only when the σ^2 value for the structure-sensitive constraint is higher ($\sigma^2 = 1000$; Figures 5.5 and 5.6) - that is, when the complex structure-sensitive constraint is treated as preferable to the general structure-blind constraint - does the learner then show a preference for the faithful [ti]


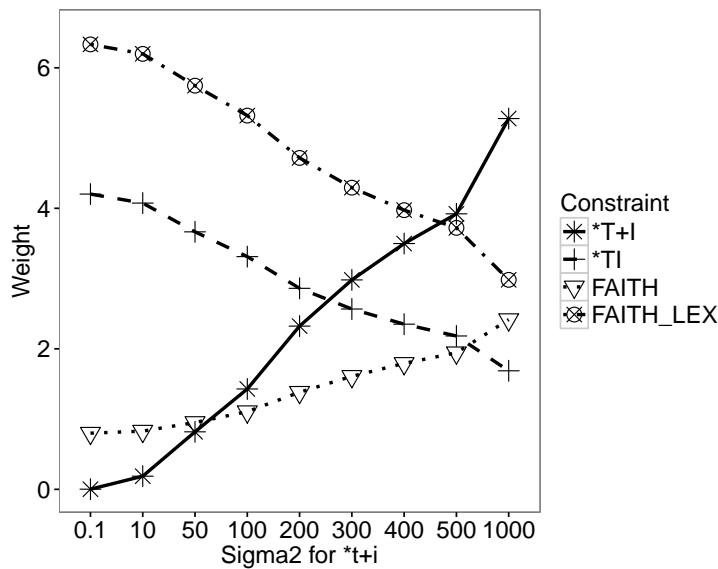



Figure 5.5: Model outputs: Predicted probability of palatalization



(a) Learnt constraint weights



(b) Hasse Diagram of desired grammar

Figure 5.6: Model outputs: Constraint weights

output. This suggests that a perfect derived environment language, in which tautomorphic sequences are completely phonotactically well-formed but at the same time heteromorphic

sequences are ill-formed and have to be repaired, is likely to be much more difficult to learn. This further suggests that learning in “true” derived-environment patterns might never be entirely accurate due to “grammatical leakage” (Martin 2011), since responsibility for the alternation will always be assigned to both constraints, and it is unlikely that the general *TI constraint will ever reach 0, with constant pressure for novel /ti/ sequences to palatalize.

5.5 No phonotactic generalization and free variation

So far, in each of the toy simulations, we have kept the rate of alternations constant at 1, with alternations always occurring. What I have manipulated is the evidence for the phonotactic generalization in three ways: (1) exceptionless constraint against [ti] (§2), (2) small preference for constraint against [ti] (§3) and (3) no constraint against [ti] (§4). In this section, I examine what occurs when there is no constraint against [ti] stem-internally and the alternation is in free variation (i.e. each candidate is equally probable). This is like the language modeled in §3 but with a much lower rate of alternation. This is a language that is inspired by the Turkish velar deletion pattern that was discussed in Chapter 4 in which the alternation is morphologically conditioned. Thus this might represent an earlier stage of learning Turkish in which a learner might not have differentiated the contexts in which deletion applies or not, represented here with the alternation rate at 50%. The learner was trained on the input shown in Table 5.5 using the same series of σ^2 values as was done previously. The μ value for each of the constraints was set at 0 for this simulation.

Input	Ouput	Freq.	FAITH _{LexC1}	*TI	*T+I	FAITH
/t+i/	[t+i]	0.5		1	1	
	[c+i]	0.5				1
/ti/ _{LexC1}	[ti]	0.5		1		
	[ci]	0	1			1
/c+i/	[t+i]	0		1	1	1
	[c+i]	1				
/ci/	[ti]	0		1		1
	[ci]	0.5				

Table 5.5: Input for a free variation language




Figure 5.7: Model outputs: Predicted probability of palatalization

The results of the simulations are shown in Figures 5.7 and 5.8. The learner matches, more or less, the rate of alternations when σ^2 is very low at 0.1, and remains constant even as the σ^2 value for *T+I increases. When σ^2 is at 0.1, the learner accurately models the alternation,

but at the cost of only showing a small tendency to be faithful to $/ti/_{LexC1}$, since $*TI$ is still weighted very high. As the σ^2 value increases, more weight is placed on the structure-sensitive constraint, and eventually, the structure-sensitive constraint takes up more of the burden in explaining the alternation (when $\sigma^2 = 300$). What we note here is that the learner shows a consistent preference for surface $[ti]$ forms, although this is not quite categorical. Overall, the learner manages to learn a preference for faithful $[ti]$ stems (lower rate of palatalization for $/ti/_{novel}$) while maintaining the free variation with the alternation.

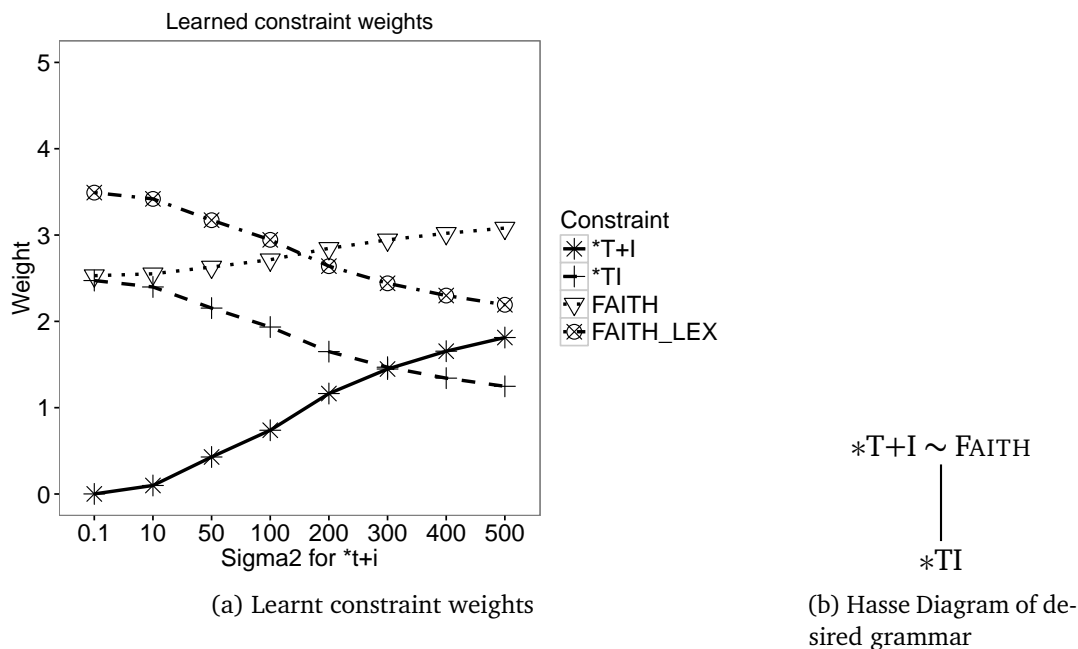


Figure 5.8: Model outputs: Constraint weights

5.6 Summary

In this chapter, I presented some simulations of how learning of mismatched generalizations across phonotactics and alternations might occur, depending on the mismatch profiles. In our simulations, I implemented a bias for alternations to reflect general structure-blind constraints by constraining how much weight could be assigned onto the structure-sensitive constraints. Table 5.6 shows a summary of the subjectively “best” σ^2 value for modeling each language. The

across-the-board language needed just the structure-blind constraint to accurately model the learning data, although the learner actually got worse in approximating human behavior as σ^2 increased. Toy Korean with a gradient phonotactic and categorical alternation along was also not too difficult to learn although it only achieved an approximation of the data at a higher σ^2 value than the across-the-board language. Yet, the learner suffered from the same issue as in the across-the-board language, where as the σ^2 value increased, the learner predicted lower rates of palatalization for /ti/_{novel}. A “true” derived-environment pattern was not learnt very well. In particular, the learner never learnt to loosen the stem-internal phonotactic constraint against [ti] sequences, and only did so when the σ^2 value was higher than every other constraint (Table 5.6). That is, only when it was preferred over the structure-blind constraint. Finally, the learner quite easily approximates the learning data in which there is no phonotactic constraint and the alternation is in free variation.

	σ^2 for *T+I	σ^2 for other constraints
Across-the-board	0.1	500
Toy Korean	100	500
Derived-environment effect	1000	500
Free variation	500	500

Table 5.6: Summary: “Best” σ^2 values for *T+I for each language.

I have attempted to show how constraining the weight of the structure-sensitive constraint might allows us to implement a generality bias. The proposal sketched out here is far from complete and, of course, needs to be tested against a larger, more realistic data set in which we have clear evidence from native speaker judgments as to how the generalizations in both domains differ. In fact, the models here actually produce rather counter-intuitive results. Across all four simulations, since the learner was trained on alternations, it was always accurate at predicting the rate of palatalization for the alternation, but the generalization about stem-internal [ti] sequences was not always predicted, even in the across-the-board language. Here as σ^2 increases, the learner overpredicts unattested [ti] sequences to surface. Thus it seems

that the learner is overfitting the alternations, at the expense of accurately modeling the stem-internal phonotactic generalization. The issue is possibly due to the way in which the learning data is specified, in particular how the frequencies of the stem-internal /ti/ _{LexC1} are specified. Alternatively this might be due to the constraint set that we have used. How to remedy this will be examined in future extensions of the current sketch proposal. One possibility might be that the learner entertains rich-base type mappings like /ti/ → [ci] as well as identity mapping /ti/ → [ti]. If this is the case, the learner must match the rates of these mappings in the learning data in order to be accurate, so an increase in σ^2 will mean a closer match to that rate. However, even in such a scheme, the derived-environment effect language will still likely require a counterintuitive σ^2 value in order to accurately model the pattern, since the lack of a phonotactic skew in stems means that, as far as the learner is concerned, there is no reason to assume any unfaithful mappings for /ti/.

Regardless, the basic proposal is in the spirit of other proposals that use of the free parameters in the MaxEnt model to encode biases in learning has previously been proposed in Wilson (2006) and J. White (2017). Both were primarily concerned with implementing a substantive bias for alternations to occur between sounds which are more perceptually similar or the P-map (Steriade 2001/2008). Using *MAP constraints (Zuraw 2007, n.d.) which penalize the correspondence between a given pair of sounds, J. White (2017) implemented this bias by setting the μ (preferred weight) values for each *MAP constraint based on perceptual confusion data. Thus, in his implementation of the substantive bias, there was an *a priori* ranking of constraints, with constraints penalizing alternations between more distant sounds having a higher weight, thus requiring more evidence to overturn. Wilson (2006) encodes a similar bias in his model, but makes use of the σ^2 values instead. Under his implementation, markedness constraints (such as *[ki], *[ka]) have different σ^2 values with constraints against less phonetically motivated constraints having a smaller σ^2 value. Because of this they are under more pressure to remain close to their preferred weights.

The proposal also has some similarities to the model presented in Martin (2011). Martin (2011) was primarily interested in modeling the fact that geminates (e.g. [pp]) do not exist

within a morpheme in English, but nonetheless appear heteromorphemically though at a rate significantly lower than chance. He argued, much like I did in Chapter 4, that this indicates a gradient dispreference for such forms. In addition to a structure-blind constraint $*[p(+)]p$ which penalizes all geminate sequences anywhere in a morphologically complex word, Martin's model also contained two kinds of structure-sensitive constraints: $*[pp]$ and $*[p+p]$. The former only penalizes tautomorphemic geminates, whereas the latter penalizes heteromorphemic ones. Having both kinds of constraints in his model allowed for weight to be assigned to both the structure-sensitive $*[pp]$ as well as the structure-blind $*[p(+)]p$ to capture the fact that geminate sequences within words were categorically banned. And because weight was assigned to the more general constraint, this then predicted an increased penalty for heteromorphemic $[p+p]$. The grammatical leakage here is allowed to occur precisely because the algorithm chooses to uniformly distribute weights across multiple constraints if those constraints are equally able to explain particular patterns in the data, here $*[pp]$ and $*[p(+)]p$. Martin initialized his model by enforcing a uniform prior across his constraints, thereby allowing leakage from tautomorphemic domain to a heteromorphemic one, and it's crucially the presence of the structure-blind constraint that allows for this.

The learner that I have sketched out above differs in a number of ways from Martin's simulations. For one, we have adopted the use of faithfulness constraints as well as markedness constraints, where Martin only used markedness constraints. This is motivated by the fact that we are modeling an alternation, and not just a purely static phonotactic generalization. Martin was primarily concerned with phonotactic learning and specifically in whether phonotactic learning in one morphological domain extends to another one. In the current set of simulations, however, we are concerned with modeling a categorical alternation with a gradient phonotactic generalization. In Martin's case there is a categorical morpheme internal generalization that spreads in a gradient manner to the heteromorphemic case. Thus, at least in the Korean palatalization case, the 'grammatical leakage' here is also in a different direction. In Martin's simulations, the prior on constraint weights, σ^2 , was set at a uniform value across all constraints, in fact, leakage across the different phonological domains relied precisely on the fact that weight

was shared uniformly across constraints whenever two constraints were able to account for the same pattern. In my own simulations, however, I have attempted to show how adjusting the prior on the more complex structure-specific constraint gives us different quantitative predictions for the well-formedness of novel stem-internal sequences that are not protected by the lexically-specific faithfulness constraints. Thus, unlike Martin, where the weight is shared uniformly across multiple constraints, I have attempted to show how favoring the assignment of weight to the more general structure-blind constraint allows for simultaneous gradient well-formedness, with lexically-specific faithfulness constraints, as well as categorical alternations.

All in all, what is presented here is a first step to trying to model mismatches across phonotactics and alternations. A future step would be to utilize models with an implementation that incorporates both phonotactic and alternation learning, such as the model elaborated in Jarosz (2006, 2011). Examining the learning behavior in such a model would be particularly enlightening on how these mismatches are learnt.

CHAPTER 6

General Discussion & Conclusion

6.1 Summary of dissertation

The goal of this dissertation was to examine a central assumption in the literature on phonological learning: that phonotactics facilitates alternation learning. I examined this question by comparing alternation learning in a language with matching phonotactics to one in which the stem-internal phonotactic generalization mismatches the alternation. These latter types of patterns are known as derived-environment effects. In particular, I was interested in shedding light on the following questions:

1. How does phonotactic learning interact with alternation learning - does phonotactic learning facilitate alternation learning?
2. How might a learning perspective on phonological mismatches shed light on how to theoretically account for these kinds of phonological patterns?

I began, in Chapter 2, by describing cases in the literature that show mismatches between alternations and stem phonotactic generalizations. In particular, I concentrated on derived-environment effects, a well-known class of phonological patterns that have proven thorny to account for in our theoretical models. Although these mismatch patterns exist, when compared to patterns which show a match in phonotactic generalizations and alternations, they seem to be more difficult to learn. I illustrated how phonotactic learning in two of these mismatch patterns, derived-environment blocking and derived-environment effects, might proceed in two different computational models of learning: a Biased Constraint Demotion

(BCD) Model and a MaxEnt model. Without morphological information in the training data, neither model distinguished phonotactic learning in a derived-environment effect language and a derived-environment blocking language, resulting in a difficulty in either learning the right root phonotactic generalization or alternation. The BCD learner failed to learn the appropriate constraint that would motivate the alternation (e.g. *TI) in a derived-environment effect language, and failed to learn the appropriate root phonotactic constraint that would penalize stem-internal sequences in a derived-environment blocking language. The MaxEnt learner shows an overall preference for alternation in both types of languages, although this was at a lower rate than when trained on a language that showed an across-the-board generalization. This predicted a lower rate of alternations for a derived-environment effect pattern, but in so doing, wrongly predicted alternations for a derived-environment blocking pattern. It also predicted that both languages should show a gradient root phonotactic constraint which does not accurately fit either language. Thus preliminary learning simulations in this chapter showed how phonotactic learning in these mismatch patterns makes it difficult to then arrive at the correct generalization about alternations.

Using an artificial grammar learning paradigm, in Chapter 3, I compared the learning of alternations in languages with a phonotactic match and mismatch. Learners trained on the mismatch language (Non-harmonic) language were unable to learn the alternation despite being trained on the alternation. Contrastively, those trained on the match (Harmonic) language successfully learnt the alternation. We also saw the intermediate learning of the alternation when the phonotactic generalization of stems only partially accorded with the alternation (Semi-harmonic language). When trained only on phonotactic generalizations without any exposure to alternations, however, learners do not spontaneously extend the learnt phonotactic generalization to novel alternations. I also conjectured that our results suggest that languages with a phonotactic mismatch, i.e. derived-environment effects, should be typologically dispreferred.

In Chapter 4, I examined in greater detail what the empirical patterns, in terms of the lexicon, are in two well-known cases of derived-environment effects: Korean palatalization and Turkish velar deletion. Pairing corpus analysis with computational learning of phonotactic generaliza-

tions using a MaxEnt learner, I showed that in Korean, contrary to previous assumptions about the data, there is a gradient phonotactic penalty against sequences which are repaired across a morpheme boundary. Thus there is phonotactic support for the alternation. Contrastively, in Turkish, we found no evidence for a phonotactic constraint penalizing intervocalic velar sequences in the lexicon, indicating that the alternation is not phonotactically supported. This accords with the fact that the alternation in Turkish is morphologically circumscribed (Inkelas 2011, 2014, Sezer 1981). I also showed briefly how Finnish assibilation showed a similar statistical pattern as Turkish velar deletion. I then presented an analysis of the Korean alternation pattern using lexical faithfulness constraints to capture the fact that the sequences in the lexicon which show the phonotactic mismatch are rare enough to be encoded as lexical exceptions to a more general phonological generalization.

Finally, in Chapter 5, I sketched a proposal of how a bias for a general constraint might be implemented using differing σ^2 values for a general structure-blind constraint and a more specific structure-sensitive constraint. I proposed that this preference could be encoded by having a smaller σ^2 value for the more complex structure-sensitive constraint, thus preferring weight to be assigned to the more general structure-blind constraint. I showed that in order to learn a toy derived environment language successfully (novel stems with the sequence /ti/ treated faithfully more often than not), it was necessary to give a larger σ^2 to the complex, structure-sensitive constraint, treating it as easier to learn than the simple, structure-blind constraint. This would be an unrealistic assumption, which supports the decreased learnability of derived-environment patterns.

Overall our results suggest that (a) learners preferentially learn alternations that accord with static phonotactic patterns, and that (b) derived-environment effects as a class of patterns are potentially suspect, likely due to the fact that they are much more difficult to learn relative to patterns which are across-the-board. Here I conclude by considering the potential implications of this work for theoretical descriptions of derived-environment effects and for a model of phonological learning.

6.2 Implications for derived-environment effects

The converging results from our studies indicate that derived-environment effect alternations are more difficult to learn relative to those in an across-the-board pattern. Moreover, upon closer inspection, canonical cases of derived-environment effects in the literature are revealed to be less than perfect examples of these patterns. Neither Turkish velar deletion nor Korean palatalization, the patterns investigated in detail here, hold up to the assumptions that the alternation is productive and that stem-internal sequences are phonotactically well-formed. In fact, in Korean, there is a productive alternation, but there is also a gradient phonotactic constraint in the lexicon. In Turkish, conversely, there is no phonotactic constraint in the lexicon, but the alternation is not productive, and is confined to certain morphological contexts.

Why might derived-environment mismatches arise in some languages? In at least some cases, the answer lies in foreign loanwords or the interaction of independent sound changes. In fact, Łubowicz (2002) has suggested this for Polish palatalization, and Y.-M. Y. Cho (2009) has also suggested this was the case for Korean palatalization. But in both cases, they present analyses which predict stem-internal sequences to be phonotactically well-formed. Our current study offers the first quantitative investigation of the lexicon to provide evidence that this is not always the case. Similarly, others have also pointed out previously that some cases of derived-environment effects are confined to certain morphological contexts. In fact, the most famous case of derived-environment effects, Finnish assibilation (Kiparsky 1973), is really only true of three suffixes (Anttila 2009), and as I showed briefly in Chapter 4, the sequence that is repaired heteromorphemically is not underattested in the lexicon. In the same vein, Turkish velar deletion only applies to polysyllabic nouns (Sezer 1981, Inkelas 2011), and even here is not a completely categorical process (Becker et al. 2011). Further, Malagasy vowel dissimilation (Zymet 2014), which I suggested was an interesting case in which the the alternation enforces a generalization that is the opposite of the mild statistical tendency of harmony in the lexicon, only occurs with one suffix (J. Zymet, p.c.). Thus no case so far exhibits the qualities of a ‘true’

derived-environment effect pattern in which there is a productive alternation that eliminates a structure (such as /ti/) that is nonetheless phonotactically well-formed within stems.

What does this mean for the very notion of a “derived-environment effect”? The fact that no case examined so far conforms to the general conception of what a derived-environment effect pattern should look like suggests that this very notion is suspect. I argue that these patterns have only been described in this way due to the lack of a proper consideration of the complexity of the actual patterns in the data. Once we consider the quantitative patterns in both the lexicon and alternations, it clearly emerges that patterns which have been previously described as instances of derived-environment effects are not the same beast, despite the surface structural similarity.

Can the notion of a derived-environment effect be salvaged? The strongest answer that one might make based on our current results is: no. A less rigid answer might be that derived-environment effects might only be useful as a broad class of patterns which show an alternation of some sort but with some kind of mismatch in phonotactics. But in this case, this notion is purely a descriptive label for a diverse set of patterns. What I think might be more insightful, perhaps, is to construct a typology of derived-environment effects that pays close attention to the quantitative patterns and relative productivity of both the phonotactic generalizations and alternations. So far, we have at least two types of cases: Korean palatalization is a case of gradient phonotactics and categorical alternations, while Turkish velar deletion is a case where there is no phonotactic constraint with a morphologically conditioned alternation. Further examination of other cases will shed light on the extent to which other cases conform to these two patterns or, as it likely, that these patterns exist on a continuum of increasing alternation productivity and decreasing phonotactic acceptability.

A piece of the puzzle that is missing from the current study is what native speakers’ internalized knowledge of these patterns are. Thus our modeling results await confirmation from wug testing with native speakers of Korean and Turkish. The predictions are that Korean speakers should show categorical behavior with alternations but gradient phonotactic dispreference for wug words that contain [ti]. Turkish speakers on the other hand should not show any

dispreference for intervocalic velar sequences in stems and should show limited generalization of the alternation beyond the contexts in which it occurs in the language (e.g. novel suffixes that could in principle trigger alternations).

6.3 Implications for phonological learning: phonotactics and alternations

The results of the artificial language experiments in Chapter 3 provide evidence that phonotactic learning and alternation learning are not completely separate processes. The design of our experiments ensured that the evidence for alternations was consistent across the different language groups, while manipulating the amount of static phonotactic evidence in the training data. Thus participants' failure to learn the alternation in the derived-environment language is particularly striking since there was ample evidence for it. Thus by showing that phonotactic mismatches impede the learning of a phonological alternation, I have shown that phonotactic knowledge does have an impact on alternation knowledge. Moreover, I have shown that gradient learning of a phonotactic generalization (as in the Semi-Harmonic language in Experiment 1 of Chapter 3) leads to gradient learning of the alternation.

That said, the phonotactic learning simulation in Chapter 3 of three artificial languages revealed that the alternation was in principle learnable in all three languages despite the differences in stem phonotactics. Interestingly, the learner achieves this in different ways. In the Harmonic language, the learner learns two bigram constraint that penalize [+back][-back] and [-back][+back] vowel sequences. This is the equivalent of learning a single constraint that would be formalized as *[-αback][αback]. In both the Semi- and Non-harmonic languages, however, the learner had to resort to utilizing two trigram constraints that ban the last two vowels of a three vowel sequence from disagreeing in backness features: *[+syllabic][-back][+back] and *[+syllabic][+back][-back]. Two things should be of note here. The first is the fact that in all three languages the learner can actually arrive at the correct generalization about plurals in the learning data and therefore we might expect that they learn all the patterns equally well.

This was not the case for adult learners as we have shown in Experiments 1 and 2 in Chapter 3. The second is that even though the learners are both able to learn the alternation pattern from a purely phonotactic generalization, the Harmonic learners are able to do this with a much simpler bigram constraint than both the Semi- and Non-harmonic language learners. The latter groups have to resort to a more complex trigram constraint to get the correct generalization. Thus, although there was in fact a trigram phonotactic solution to learning the alternation in the Non-harmonic language, learners failed to discover this generalization, potentially providing another piece of support for the idea that learners are biased towards simpler constraints.

So far, we have only investigated what happens with alternation learning if it is not supported by phonotactics. But what about cases when there are no alternations yet there seems to be an phonotactic generalization in the lexicon - derived-environment blocking? In Chapter 2, our initial modeling of phonotactic learning of this type of language showed a similar result to learning in a derived-environment effect language. In a stochastic MaxEnt model, the learner predicts the encoding of a gradient phonotactic constraint due to the overall preponderance for one type of sequence (e.g. [ci]) over another (e.g. [ti]). In our simulations the overall statistical distribution of these sequences was the same in both the derived-environment effect and derived-environment blocking languages, thus a learner that is not sensitive to morphological structure infers the same generalizations in either case. The mismatch here between alternations and phonotactics, however, does not seem too impede the ability for speakers to infer the root phonotactic constraints which do not engender alternations. Using a repetition task, Gallagher (2013) for example showed that Cochabamba Quecha speakers encode the laryngeal cooccurrence restrictions regarding ejectives and plain stops in a root; this phonotactic constraint, however, does not engender alternations. Thus Cochabamba Quechua speakers internalize root phonotactic knowledge but do not extend this to alternations.

Unlike a derived-environment effect pattern, a derived-environment blocking pattern might perhaps be less difficult to learn. This would suggest that while an alternation is easier to learn if it is supported by phonotactics, just having a phonotactic generalization does not necessarily lead to the learning of an alternation, nor the failure to internalize that phonotactic

generalization about the lexicon. So there might be an asymmetric effect on learning depending on the locus of the mismatch. Note that this is from the perspective of a human learner, since a computational model treats these equivalently as we saw in Chapter 2. In fact, we saw that learners in the experiments presented in Chapter 3 failed to extend a learnt phonotactic generalization to novel alternations that they were not trained on. If learners are already biased against positing alternations, as we might infer from the results of the Experiments 3 and 4 in Chapter 3, it might be the case that derived-environment blocking cases are actually the default.

While our results are suggestive of an effect of phonotactic knowledge on alternation knowledge in learning, whether they are encoded by a shared mechanism as is suggested by the constraint architecture of OT remains unclear. Learners are not inherently biased to extend a phonotactic generalization to unseen alternations as we have shown in Experiments 3 and 4 in Chapter 3. This is also consistent with Pizzo's findings. In Pizzo's (2015) case it was from alternations to phonotactics, and in Chapter 3, from phonotactics to alternations. An eventual model must therefore be able to account for the fact that phonotactic matches aid alternation learning, but that learners do not spontaneously expect alternation knowledge to reflect phonotactic knowledge or vice versa.

Finally, our understanding of the relative trajectory of phonotactic and alternation learning in infancy is still unclear. Infants show a precocious ability to learn phonotactics solely from distributional data by 8.5-9 months (Saffran & Thiessen 2003, Friederici & Wessels 1993, Jusczyk et al. 1993, K. S. White et al. 2008), but only show knowledge of alternations a little later at 12 months (K. S. White et al. 2008, J. White & Sundara 2014). Yet we do not actually have any evidence that alternations which are phonotactically motivated are learnt earlier than those which are not, a prediction of phonological models and the artificial grammar learning studies from this dissertation. Thus examining specific cases that would bear on this question would provide a fruitful avenue for future research.

REFERENCES

- Adjarian, H. (1954). *Complete grammar of the Armenian language*. Yerevan: Haykakan SSH GA Kratarakchutyun.
- Adriaans, F., & Kager, R. (2010). Adding generalization to statistical learning: The induction of phonotactics from continuous speech. *Journal of Memory and Language*, 62(3), 311–331.
- Ahn, S.-C. (1985). *The interplay of phonology and morphology in Korean* (Ph.D. Dissertation). University of Illinois, Urbana-Champaign, Urbana-Champaign.
- Aksu-Koç, A. A., & Slobin, D. I. (1985). The acquisition of Turkish. In D. I. Slobin (Ed.), *Cross-linguistics study of language acquisition: Vol. 1* (pp. 839–878). New Jersey: Lawrence Erlbaum Associates.
- Albright, A. (2002). Islands of reliability for regular morphology: Evidence from Italian. *Language*, 78(4), 684–709.
- Altan, A., Kaya, U., & Hohenberger, A. (2016). Sensitivity of Turkish infants to vowel harmony in stem-suffix sequences: Preference shift from familiarity to novelty. In J. Scott & D. Waughtal (Eds.), *Proceedings of the 40th Boston University Conference on Language Development*.
- Anderson, S. R. (1970). On Grassmann's Law in Sanskrit. *Linguistic Inquiry*, 1(4), 387–396.
- Anderson, S. R. (1974). *The organization of phonology*. New York: Academic Press.
- Anttila, A. (2006). Variation and opacity. *Natural Language & Linguistic Theory*, 24, 893–944.
- Anttila, A. (2009). Derived environment effects in Colloquial Helsinki Finnish. In K. Hanson & S. Inkelas (Eds.), *The Nature of the Word: Studies in Honor of Paul Kiparsky* (pp. 433–460). Cambridge: MIT Press.
- Baayen, R. H., Piepenbrock, R., & Gulikers, L. (1995). *The CELEX lexical database (Release 2) [CD-ROM]*. Philadelphia, PA: University of Pennsylvania, Linguistic Data Consortium [Distributor].

- Baer-Henney, D., & van de Vijver, R. (2012). On the role of substance, locality, and amount of exposure in the acquisition of morphophonemic alternations. *Laboratory Phonology*, 3(2), 221–250.
- Bailey, T. M., & Hahn, U. (2001). Determinants of wordlikeness: Phonotactics or lexical neighborhoods? *Journal of Memory and Language*, 44(4), 568–591. doi: 10.1006/jmla.2000.2756
- Bates, D., Maechler, M., Bolker, B., & Walker, S. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, 67(1), 1–48. doi: <http://dx.doi.org/10.18637/jss.v067.i01>
- Becker, M., & Gouskova, M. (2016). Source-oriented generalizations as grammar inference in Russian vowel deletion. *Linguistic Inquiry*, 47(3).
- Becker, M., Ketrez, N., & Nevins, A. (2011). The surfeit of the stimulus: Analytic biases filter lexical statistics in Turkish laryngeal alternations. *Language*, 87(1), 84–125. doi: <http://dx.doi.org/10.1353/lan.2011.0016>
- Becker, M., & Levine, J. (2014). *Experigen - an online experiment platform*. (Available at: <http://becker.phonologist.org/experigen>)
- Benua, L. (2000). *Phonological relations between words*. New York: Garland Press.
- Berger, A. L., Pietra, V. J. D., & Pietra, S. A. D. (1996). A Maximum Entropy approach to natural language processing. *Computational Linguistics*, 22(1), 39–71.
- Berko, J. (1958). The child's learning of english morphology. *WORD*, 14(2-3), 150–177. doi: <http://dx.doi.org/10.1080/00437956.1958.11659661>
- Bradley, T. J. (2007). Morphological derived-environment effects in gestural coordination: A case study of Norwegian clusters. *Lingua*, 117, 950–985.

- Burzio, L. (1997). Cycles, non-derived environment blocking, and correspondence. In J. Dekkers, F. van der Leeuw, & J. van de Weijer (Eds.), *The Pointing Finger: Conceptual Studies in Optimality Theory* (pp. 47–87). Oxford: Oxford University Press.
- Burzio, L. (2011). Derived environment effects. In M. van Oostendorp, C. J. Ewen, E. Hume, & K. Rice (Eds.), *The Blackwell Companion to Phonology*. Blackwell.
- Calamaro, S., & Jarosz, G. (2015). Learning general phonological rules from distributional information: A computational model. *Cogn Sci*, 39(3), 647–666. doi: 10.1111/cogs.12167
- Chambers, K. E., Onishi, K. H., & Fisher, C. (2003). Infants learn phonotactic regularities from brief auditory experience. *Cognition*, 87, B69–B77.
- Chambers, K. E., Onishi, K. H., & Fisher, C. (2010). A vowel is a vowel: Generalizing newly-learned phonotactic constraints to new contexts. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 36(3), 821–828.
- Chambers, K. E., Onishi, K. H., & Fisher, C. (2011). Representations for phonotactic learning in infancy. *Language Learning and Development*, 7(4), 287–308.
- Cho, T. (2001). Effects of morpheme boundaries on intergestural timing: Evidence from Korean. *Phonetica*, 58(3), 129–162. doi: <http://dx.doi.org/10.1159/000056196>
- Cho, Y.-M. Y. (2009). A historical perspective on nonderived environment blocking: The case of Korean palatalization. In K. Hanson & S. Inkelas (Eds.), *The Nature of the Word: Studies in Honor of Paul Kiparsky* (pp. 461–486). MIT Press.
- Chomsky, N., & Halle, M. (1968). *The sound pattern of English*. New York, NY: Harper and Row.
- Chong, A. J. (2016). Learning consequences of derived-environment effects. In *Proceedings of the Linguistic Society of America* (Vol. 1).
- Chung, S. (1983). Transderivational relationships in Chamorro phonology. *Language*, 59(1), 35. doi: 10.2307/414060

- Clements, G. N., & Sezer, E. (1982). Vowel and consonant disharmony in Turkish. In H. van der Hulst & N. Smith (Eds.), *The structure of phonological representations II*. Dordrecht: Foris.
- Coetzee, A. (2008). Phonological variation and lexical frequency. Ms.: University of Michigan. Available as ROA-952-0208, Rutgers Optimality Archive, <http://roa.rutgers.edu>.
- Coetzee, A. W., & Pater, J. (2008). Weighted constraints and gradient restrictions on place co-occurrence in Muna and Arabic. *Natural Language & Linguistic Theory*, 26(2), 289–337. doi: 10.1007/s11049-008-9039-z
- Coleman, J., & Pierrehumbert, J. B. (1997). Stochastic phonological grammars and acceptability. In *Computational Phonology: Third meeting of the ACL special interest group in computational phonology* (pp. 49–56). Association for Computational Linguistics.
- Cristià, A., Mielke, J., Daland, R., & Peperkamp, S. (2013). Similarity in the generalization of implicitly learned sound patterns. *Laboratory Phonology*, 4(2), 259-285. doi: <http://dx.doi.org/10.1515/lp-2013-0010>
- Crothers, J., & Shibatani, M. (1980). Issues in the description of Turkish vowel harmony. In R. Vago (Ed.), *Issues in vowel harmony* (pp. 63–88). Amsterdam: John Benjamins Publishing Company.
- CSC - IT Center for Science. (2004). *Frequency lexicon of the Finnish newspaper language [text corpus]*. Kielipankki. Retrieved online from <http://urn.fi/urn:nbn:fi:lb-201405272>.
- Daland, R., Hayes, B., White, J., Garellek, M., Davis, A., & Normann, I. (2011). Explaining sonority projection effects. *Phonology*, 28, 197–234.
- Daland, R., Oh, M., & Kim, S. (2015, May). When in doubt, read the instructions: Orthographic effects in loanword adaptation. *Lingua*, 159, 70–92. doi: 10.1016/j.lingua.2015.03.002
- Do, Y. (2013). *Biased learning of phonological alternations* (Unpublished doctoral dissertation). Massachusetts Institute of Technology, Cambridge, MA.

- Edwards, J., Beckman, M. E., & Munson, B. (2004). The interaction between vocabulary size and phonotactic probability effects on children's production accuracy and fluency in nonword repetition. *Journal of Speech, Language, and Hearing Research*, 47, 421–436.
- Eychenne, J., & Jang, T.-Y. (2015). On the merger of Korean mid front vowels: Phonetic and phonological evidence. *Journal of the Korean Society of Speech Sciences*, 7(2), 119–129.
- Finley, S., & Badecker, W. (2009). Artificial language learning and feature-based generalization. *Journal of Memory and Language*, 61, 423–437.
- Finley, S., & Badecker, W. (2012). Learning biases for vowel height harmony. *Cognitive Science*, 13, 287–327.
- Fountain, A. V. (1998). *An Optimality Theoretic account of Navajo prefixal syllables* (Ph.D. Dissertation). University of Arizona, Tucson.
- Friederici, A. D., & Wessels, J. M. I. (1993). Phonotactic knowledge of word boundaries and its use in infant speech perception. *Perception & Psychophysics*, 54(3), 287–295.
- Frisch, S. A., Large, N. R., & Pisoni, D. B. (2000). Perception of wordlikeness: Effects of segment probability and length on the processing of nonwords. *Journal of Memory and Language*, 42(4), 481–496. doi: 10.1006/jmla.1999.2692
- Frisch, S. A., Pierrehumbert, J. B., & Broe, M. B. (2004). Similarity avoidance and the OCP. *Natural Language & Linguistic Theory*, 22(1), 179–228.
- Frisch, S. A., & Zawaydeh, B. A. (2001). The psychological reality of OCP-Place in Arabic. *Language*, 77(1), 91–106.
- Gallagher, G. (2010). Perceptual distinctness and long-distance laryngeal restrictions. *Phonology*, 27, 435–480.
- Gallagher, G. (2013). Speaker awareness of non-local ejective phonotactics in Cochabamba Quechua. *Natural Language & Linguistic Theory*, 31(1067-1099).

- Gaskell, M. G., & Marslen-Wilson, W. D. (1996). Phonological variation and inference in lexical access. *Journal of Experimental Psychology: Human Perception and Performance*, 22, 144–158.
- Ghim, H.-R. (2005). *Korean Jiwon Corpus*. Retrieved online from CHILDES: <http://childes.talkbank.org/access/EastAsian/Korean/Jiwon.html>.
- Gnanadesikan, A. (2004). Markedness and faithfulness constraints in child phonology. In R. Kager, J. Pater, & W. Zonneveld (Eds.), *Constraints in phonological acquisition* (pp. 73–108). Cambridge: Cambridge University Press.
- Göksel, A., & Kerslake, C. (2005). *Turkish: A comprehensive grammar*. London: Routledge. (OCLC: 249107163)
- Goldsmith, J. (1976). *Autosegmental Phonology* (Ph.D. Dissertation). Massachusetts Institute of Technology, Cambridge, MA.
- Goldsmith, J. (1990). *Autosegmental and Metrical Phonology*. Oxford: Blackwell.
- Goldwater, S., & Johnson, M. (2003). Learning OT constraint rankings using a Maximum Entropy model. In *Proceedings of the Workshop on Variation within Optimality Theory*. Stockholm University.
- Hale, M., & Reiss, C. (2008). *The phonological enterprise*. Oxford: Oxford University Press.
- Halle, M. (1959). *The sound pattern of Russian*. The Hague: Mouton.
- Hamann, S., Apoussidou, D., & Boersma, P. (2012). Modelling the formation of phonotactic restrictions across the mental lexicon. In *Proceedings from the Annual Meeting of the Chicago Linguistics Society 45* (pp. 193–206).
- Hammond, M. (1992). Deriving the Strict Cycle Condition. In *Papers from the 28th Regional Meeting of the Chicago Linguistic Society* (pp. 126–140). Chicago: Chicago Linguistic Society.

- Hay, J., Pierrehumbert, J. B., & Beckman, M. (2003). Speech perception, well-formedness, and the statistics of the lexicon. In J. Local, R. Ogden, & R. Temple (Eds.), *Papers in laboratory phonology VI* (pp. 58–74). Cambridge: Cambridge University Press.
- Hayes, B. (2000). Gradient well-formedness in Optimality Theory. In J. Dekkers, F. van der Leeuw, & J. van de Weijer (Eds.), *Optimality Theory: Phonology, syntax, and acquisition* (pp. 88–120). Oxford: Oxford University Press.
- Hayes, B. (2004). Phonological acquisition in Optimality Theory: The early stages. In R. Kager, J. Pater, & W. Zonneveld (Eds.), *Constraints in phonological acquisition* (pp. 158–203). Cambridge: Cambridge University Press.
- Hayes, B., Kirchner, R., & Steriade, D. (2004). *Phonetically-based phonology*. Cambridge: Cambridge University Press.
- Hayes, B., & Londe, Z. C. (2006). Stochastic phonological knowledge: The case of Hungarian vowel harmony. *Phonology*, 23, 59–104.
- Hayes, B., & Moore-Cantwell, C. (2011). Gerard Manley Hopkins's sprung rhythm: Corpus study and stochastic grammar. *Phonology*, 28, 235–282.
- Hayes, B., Siptár, P., Zuraw, K., & Londe, Z. (2009). Natural and unnatural constraints in Hungarian vowel harmony. *Language*, 85(4), 822–863. doi: 10.1353/lan.0.0169
- Hayes, B., Tesar, B., & Zuraw, K. (2013). *OTSoft*. (Published: software package, <http://www.linguistics.ucla.edu/people/hayes/otsoft/>)
- Hayes, B., & White, J. (2013). Phonological naturalness and phonotactic learning. *Linguistic Inquiry*, 44(1), 45–75. doi: http://dx.doi.org/10.1162/LING_a_00119
- Hayes, B., & Wilson, C. (2008). A Maximum Entropy model of phonotactics and phonotactic learning. *Linguistic Inquiry*, 39(3), 379–440. doi: <http://dx.doi.org/10.1162/ling.2008.39.3.379>

- Hayes, B., Wilson, C., & Shisko, A. (2012). Maxent grammars for the metrics of Shakespeare and Milton. *Language*, 88(4), 691–731.
- Hothorn, T., Bretz, F., & Westfall, P. (2008). Simultaneous Inference in General Parametric Models. *Biom. J.*, 50(3), 346–363. doi: <http://dx.doi.org/10.1002/bimj.200810425>
- Inkelas, S. (2000). Phonotactic blocking through structural immunity. In B. Stiebels & D. Wunderlich (Eds.), *Lexicon in focus* (pp. 7–40). Berlin: Akademie Verlag.
- Inkelas, S. (2011). Another look at velar deletion in Turkish, with special attention to the derived environment condition. In E. E. Taylan & B. Rona (Eds.), *Puzzles of Language: Essays in honour of Karl Zimmer*. Wiesbaden: Harrassowitz.
- Inkelas, S. (2014). *The interplay of morphology and phonology*. Oxford: Oxford University Press.
- Inkelas, S. (2015). Confidence scales: A new approach to derived environment effects. In Y. E. Hsiao & L.-H. Wee (Eds.), *Capturing phonological Shades*. Cambridge: Cambridge Scholar Press.
- Inkelas, S., Küntay, A., Orgun, C. O., & Sprouse, R. (2000). Turkish electronic living lexicon (TELL): A lexical database. In M. G. et al (Ed.), *Proceedings of the Second International Conference on Language Resources and Evaluation*. Paris: European Languages Resources Association.
- Inkelas, S., & Orgun, C. O. (1995). Level ordering and economy in the lexical phonology of Turkish. *Language*, 71, 763–793.
- Inkelas, S., Orgun, C. O., & Zoll, C. (1997). Implications of lexical exceptions for the nature of the grammar. In I. Roca (Ed.), *Derivations and constraints in phonology*. Oxford: Clarendon Press.

- Inkelas, S., & Zoll, C. (2007). Is grammar dependence real? A comparison between cophological and indexed constraint approaches to morphologically conditioned phonology. *Linguistics*, 45, 133–171.
- Itô, J., & Mester, A. (2003). *Japanese morphophonemics: Markedness and word structure*. Cambridge, MA: MIT Press.
- Itô, J., Mester, A., & Padgett, J. (1995). Licensing and underspecification in Optimality Theory. *Linguistic Inquiry*, 26, 571–631.
- Iverson, G. K. (1987). The revised alternation condition in Lexical Phonology. *Nordic Journal of Linguistics*, 10, 151–164.
- Iverson, G. K., & Wheeler, D. W. (1988). Blocking and the Elsewhere condition. In M. Hammond & M. Noonan (Eds.), *Theoretical morphology: Approaches in modern linguistics* (pp. 325–338). San Diego: Academic Press.
- Jaeger, T. F. (2008). Categorical data analysis: Away from ANOVAs (transformation or not) and towards Logit Mixed Models. *J Mem Lang*, 59(4), 434–446. doi: <http://dx.doi.org/10.1016/j.jml.2007.11.007>
- Jäger, G. (2007). Maximum Entropy models and Stochastic Optimality Theory. In A. Zaenen, J. Simpson, T. H. King, J. Grimshaw, J. Maling, & C. Manning (Eds.), *Architecture, rules and preferences: Variations on themes by Joan W. Bresnan* (pp. 467–479). CSLI Publications.
- Jarosz, G. (2006). *Rich lexicons and restrictive grammars - Maximum likelihood learning in Optimality Theory*. (Ph.D. Dissertation). Johns Hopkins University, Baltimore, MD.
- Jarosz, G. (2011). The roles of phonotactics and frequency in the learning of alternations. In *BUCLD 35: Proceedings of the 35th annual Boston University Conference on Language Development* (pp. 321–333). Somerville, MA: Cascadilla Press.
- Jun, J., & Lee, J. (2007). Multiple stem-final variants in Korean native nouns and loanwords. *Journal of the Linguistic Society of Korea [Eoneohag]*, 47, 159–187.

- Jurgec, P. (2011). *Feature spreading 2.0: A unified theory of assimilation* (Ph.D. Dissertation). University of Tromsø, Tromsø.
- Jusczyk, P. W., Friederici, A. D., Wessels, J. M. I., Svenkerud, V. Y., & Jusczyk, A. M. (1993). Infants' sensitivity to the sound patterns of native language words. *Journal of Memory and Language*, 32, 402–420.
- Jusczyk, P. W., Hohne, E. A., & Bauman, A. (1999). Infants' sensitivity to allophonic cues for word segmentation. *Perception & Psychophysics*, 61(8), 1456–1476.
- Jusczyk, P. W., Luce, P. A., & Charles-Luce, J. (1994). Infants' sensitivity to phonotactic patterns in the native language. *Journal of Memory and Language*, 33, 630–645.
- Kager, R., & Pater, J. (2012). Phonotactics as phonology: Knowledge of a complex restriction in Dutch. *Phonology*, 29(1), 81–111.
- Kager, R., & Shatzman, K. (2007). Phonological constraints in speech processing. *Linguistics in the Netherlands*, 24, 99–111. doi: 10.1075/avt.24.11kag
- Kang, Y. (2010). The emergence of phonological adaptation from phonetic adaptation: English loanwords in Korean. *Phonology*, 27, 225–253.
- Kapatsinski, V. (2010). Velar palatalization in Russian and artificial grammar: Constraints on models of morphophonology. *Laboratory Phonology*, 1(2), 361–393.
- Kari, J. M. (1976). *Navajo verb prefix phonology*. New York: Garland Press.
- Karlsson, F. (1983). *Suomen kielen äänne- ja muotorakenne [The Phonological and Morphological Structure of Finnish]*. Helsinki: Werner Söderström Osakeyhtiö.
- Kenstowicz, M., & Kisseberth, C. (1977). *Topics in phonological theory*. New York: Academic Press.
- Kenstowicz, M., & Kisseberth, C. (1979). *Generative Phonology: Description and theory*. New York, NY: Academic Press.

- Kim, B., Lee, G., & Lee, J.-H. (2002). Morpheme-based grapheme to phoneme conversion using phonetic patterns and morphophonemic connectivity information. In *ACM Trans. Asian Lang. Inf. Process* (Vol. 1, pp. 65–82).
- Kiparsky, P. (1973). Phonological representations. In O. Fujimura (Ed.), *Three dimensions of linguistic theory* (pp. 1–135). Tokyo: TEC.
- Kiparsky, P. (1982a). *Explanation in phonology*. Dordrecht: Foris.
- Kiparsky, P. (1982b). From cyclic phonology to lexical phonology. In H. van der Hulst & N. Smith (Eds.), *The structure of phonological representations I* (pp. 131–175). Dordrecht: Foris.
- Kiparsky, P. (1993). Blocking in nonderived environments. In E. Kaisse & S. Hargus (Eds.), *Phonetics and Phonology 4: Studies in lexical phonology* (pp. 277–313). San Diego: Academic Press.
- Kiparsky, P. (2003). Finnish noun inflection. In D. Nelson & S. Manninen (Eds.), *Generative Approaches to Finnic Linguistics* (pp. 109–161). Stanford: CSLI Publications.
- Lee, K.-M., & Ramsey, S. R. (2011). *A history of the Korean Language*. Cambridge: Cambridge University Press.
- Lewis, G. (1967). *Turkish grammar*. Oxford: Oxford University Press.
- Linzen, T., & Gallagher, G. (2014). The timecourse of generalization in phonotactic learning. In J. Kingston, C. Moore-Cantwell, J. Pater, & R. Staub (Eds.), *Proceedings of Phonology 2013*. Washington, DC: Linguistic Society of America.
- Łubowicz, A. (2002). Derived environment effects in Optimality Theory. *Lingua*, 112(4), 243–280. doi: 10.1016/S0024-3841(01)00043-2
- MacEachern, M. (1999). *Laryngeal co-occurrence restrictions*. New York: Garland Press.
- MacWhinney, B. (2000). *The CHILDES Project: Tools for analyzing talk. 3rd Edition. Vol. 2: The Database*. Mahwah, NJ: Lawrence Erlbaum Associates.

- Martin, A. (2011). Grammars leak: Modeling how phonotactic generalizations interact within the grammar. *Language*, 87(4), 751–770. doi: <http://dx.doi.org/10.1353/lan.2011.0096>
- Mattys, S. L., & Jusczyk, P. W. (2001). Phonotactic cues for segmentation of fluent speech by infants. *Cognition*, 78, 91–121.
- Mattys, S. L., Jusczyk, P. W., Luce, P. A., & Morgan, J. L. (1999). Phonotactic and prosodic effects on word segmentation in infants. *Cognitive Psychology*, 38, 465–494.
- McCarthy, J. J. (1998). Morpheme structure constraints and paradigm occultation. In M. C. Gruber, D. Higgins, K. Olson, & T. Wysocki (Eds.), *Proceedings of the Chicago Linguistic Society 35: Vol. II The panels* (p. 123-150). Chicago: Chicago Linguistic Society.
- McCarthy, J. J. (2002). *A thematic guide to Optimality Theory*. Cambridge: Cambridge University Press.
- McCarthy, J. J. (2003). Comparative markedness. *Theoretical Linguistics*, 29, 1–51.
- McDonough, J. (1991). On the representation of consonant harmony in Navajo. In *Proceedings of the West Coast Conference on Formal Linguistics (WCCFL)* (Vol. 10, pp. 319–335). Somerville, MA: Cascadilla Press.
- McDonough, J. (2003). *The Navajo sound system*. Dordrecht: Kluwer.
- Meeusen, A. E. (1959). *Essai de grammaire rundi*. Tervuren: Musée Royal du Congo Belge.
- Moore-Cantwell, C. (2017). *Concurrent learning of the lexicon and phonology*. Talk given at Annual Meeting of Linguistic Society of America 2017.
- Moore-Cantwell, C., & Pater, J. (2017). Gradient exceptionality in maximum entropy grammar with lexically specific constraints. *Catalan Journal of Linguistics*, 15.
- Moore-Cantwell, C., & Sanders, L. (2014). *Two types of implicit knowledge of probabilistic phonotactics*. (Talk given at Annual Meeting of Phonology 2014)

- Moreton, E. (1999). Evidence for phonotactic grammar in speech perception. In *Proceedings of the 14th International Congress of Phonetic Sciences*. San Francisco.
- Moreton, E. (2008). Analytic bias and phonological typology. *Phonology*, 25(2), 83–127.
- Moreton, E., & Pater, J. (2012a). Structure and substance in artificial-phonology learning. Part II: Substance. *Language and Linguistics Compass*, 6(11), 702–718.
- Moreton, E., & Pater, J. (2012b). Structure and substance in artificial-phonology learning. Part I: Structure. *Language and Linguistics Compass*, 6(11), 686–701.
- Myers, S. (1987). Vowel shortening in English. *Natural Language & Linguistic Theory*, 5, 485–518.
- National Academy of Korean Language. (2001). *Oylaye sayong siltay cosa: 1990 nyendo [Survey of the state of loanword usage: 1990]*. Seoul: National Academy of Korean Language.
- National Academy of Korean Language. (2003). *Hyeondae gug-eo sayong bindo josa gyeolgwa pail [Frequently used words in modern language]*. Seoul: National Academy of Korean Language.
- Oh, M. (1995). A prosodic analysis of nonderived-environment blocking. *Journal of East Asian Linguistics*, 4(4), 261–279.
- Onishi, K. H., Chambers, K. E., & Fisher, C. (2002). Learning phonotactic constraints from brief auditory experience. *Cognition*, 83, B13–B23.
- Paradis, C., & Lacharité, D. (1997). Preservation and minimality in loanword adaptation. *Journal of Linguistics*, 33(2), 379–430.
- Paster, M. (2013). Rethinking the ‘duplication problem’. *Lingua*, 126, 78–91. doi: <http://dx.doi.org/10.1016/j.lingua.2012.11.015>
- Pater, J. (2010). Morpheme-specific phonology: Constraint indexation and inconsistency resolution. In S. Paker (Ed.), *Phonological Argumentation: Essays on Evidence and Motivation* (pp. 123–154). London: Equinox.

- Pater, J., & Moreton, E. (2012). Structurally biased phonology: Complexity in learning and topology. *The EFL Journal*, 3, 1-44.
- Pater, J., Staubs, R., Jesney, K., & Smith, B. (2012). Learning probabilities over underlying representations. In *Proceedings of the Twelfth Meeting of ACL-SIGMORPHON: Computational Research in Phonetics, Phonology, and Morphology*.
- Pater, J., & Tessier, A.-M. (2005). Phonotactics and alternations: Testing the connection with artificial language learning. In K. Flack & S. Kawahara (Eds.), *University of Massachusetts Occasional Papers in Linguistics [UMOP]* (Vol. 31, pp. 1–16).
- Peperkamp, S., Calvez, R. L., Nadal, J.-P., & Dupoux, E. (2006). The acquisition of allophonic rules: Statistical learning with linguistic constraints. *Cognition*, 101, B31–B41.
- Pierrehumbert, J. B. (1994). Syllable structure and word structure: A study of triconsonantal clusters in English. In P. A. Keating (Ed.), *Phonological Structure and Phonetic Form: Papers in Laboratory Phonology II* (pp. 168–188). Cambridge: Cambridge University Press.
- Pierrehumbert, J. B. (2003). Phonetic diversity, statistical learning, and acquisition of phonology. *Language and Speech*(46), 115–154.
- Pizzo, P. (2015). *Investigating properties of phonotactic knowledge through web-based experimentation* (Ph.D. Dissertation). University of Massachusetts Amherst, Amherst.
- Prince, A., & Smolensky, P. (1993/2004). *Optimality Theory: Constraint interaction in Generative Grammar*. Cambridge, MA: Blackwell Publishers.
- Prince, A., & Tesar, B. (2004). Learning phonotactic distributions. In R. Kager, J. Pater, & W. Zonneveld (Eds.), *Constraints in phonological acquisition* (pp. 245–291). Cambridge University Press.
- Pycha, A., Nowak, P., Shin, E., & Shosted, R. (2003). Phonological rule-learning and its implications for a theory of vowel harmony. In M. Tsujimura & G. Garding (Eds.), *Proceedings*

- of the West Coast Conference on Formal Linguistics (WCCFL) 22 (pp. 101–114). Somerville, MA: Cascadilla Press.
- R Core Team. (2015). *R: A language and environment for statistical computing*. Vienna, Austria: R Foundation for Statistical Computing.
- Richtsmeier, P. T. (2011). Word-types, not word-tokens, facilitate extraction of phonotactic sequences by adults. *Laboratory Phonology*, 2(1), 157–183. doi: 10.1515/labphon.2011.005
- Rodegem, F. M. (1970). *Dictionnaire Rundi-Français*. Tervuren: Musée Royal de l'Afrique Central.
- Rubach, J. (1984). *Cyclic and lexical phonology*. Dordrecht: Foris.
- Ryu, J. (2012). The L1 acquisition of the imperfective aspect markers in Korean: A comparison with Japanese. *International Journal of Asian Language Processing*, 22, 1–13.
- Ryu, J., & Yasuhiro, S. (2014). The first language acquisition of the Korean imperfective aspect markers -ko iss-/-a iss-. In S. Nam, H. Ko, & J. Jun (Eds.), *Japanese/Korean Linguistics 21* (pp. 265–279). Stanford: CSLI Publications.
- Saffran, J. R., Aslin, R. N., & Newport, E. L. (1996). Statistical learning by 8-month-old infants. *Science*, 274(5294), 1926–1928. doi: <http://dx.doi.org/10.1126/science.274.5294.1926>
- Saffran, J. R., & Thiessen, E. D. (2003). Pattern induction by infant language learners. *Developmental Psychology*, 39, 484-494.
- Sapir, E., & Hoijer, H. (1967). *The phonology and morphology of the Navaho language*. Berkeley: University of California Press.
- Scholes, R. (1966). *Phonotactic grammaticality*. The Hague: Mouton.
- Schütze, C. (1996). *The empirical base of linguistics: Grammaticality and linguistic methodology*. Chicago: University of Chicago Press.

- Scicon R & D. (2015). *PCQuirer*. Los Angeles, CA.
- Sezer, E. (1981). The k/ø alternation in Turkish. In G. N. Clements (Ed.), *Harvard Studies in Phonology* (pp. 354–382). Bloomington: IULC.
- Shaffer, J. P. (1995). Multiple hypothesis testing. *Annual Review of Psychology*, 46(1), 561–584. doi: 10.1146/annurev.ps.46.020195.003021
- Shin, J., Kiaer, J., & Cha, J. (2013). *The sounds of Korean*. Cambridge: Cambridge University Press.
- Skoruppa, K., & Peperkamp, S. (2011). Adaptation to novel accents: Feature-based learning of context-sensitive phonological regularities. *Cognitive Science*, 35(2), 348–366. doi: <http://dx.doi.org/10.1111/j.1551-6709.2010.01152.x>
- Slobin, D. (1982). Universal and particular in the acquisition of language. In E. Wanner & L. Gleitman (Eds.), *Language acquisition: The state of the art* (pp. 128–172). New York: Cambridge University Press.
- Sohn, H.-M. (1999). *The Korean language*. Cambridge: Cambridge University Press.
- Stanley, R. (1967). Redundancy rules in phonology. *Language*, 43, 393–436.
- Steriade, D. (2001/2008). The phonology of perceptibility effect: The P-Map and its consequences for constraint organization. In K. Hanson & S. Inkelas (Eds.), *The nature of the word* (pp. 151–179). Cambridge, MA: MIT Press.
- Sundara, M., Kim, Y., White, J., & Chong, A. J. (2014). *There is no pat in patting: Acquisition of phonological alternations by English-learning 12-month-olds*. Talk given at Boston University Conference on Language Development 38.
- Tesar, B. (1997). *Multi-recursive constraint demotion*. (Ms., Rutgers University. ROA-197)

- Tesar, B., & Prince, A. (2007). Using phonotactics to learn phonological alternations. In *Proceedings of the Thirty-Ninth Conference of the Chicago Linguistics Society, vol. II: The panels* (pp. 209–237).
- Tesar, B., & Smolensky, P. (1998). The learnability of Optimality Theory. *Linguistic Inquiry*, 29(2), 229–268.
- Tessier, A.-M. (2012). Testing for OO-Faithfulness in the acquisition of consonant clusters. *Language Acquisition*, 19(2), 144–173. doi: 10.1080/10489223.2012.660552
- Treiman, R., Kessler, B., Knewasser, S., Tincoff, R., & Bowman, M. (2000). English speakers' sensitivity to phonotactic patterns. In M. B. Broe & J. Pierrehumbert (Eds.), *Papers in Laboratory Phonology V: Acquisition and the Lexicon* (pp. 269–282). Cambridge: Cambridge University Press.
- Turkay, F. (2012). *Turkay corpus*. CHILDES.
- van Kampen, A., Parmaksiz, G., van de Vijver, R., & Höhle, B. (2008). Metrical and statistical cues for word segmentation: The use of vowel harmony and word stress as a cue to word boundaries by 6- and 9-month-old turkish learners. In A. Gavarró & M. J. Freitas (Eds.), *Language acquisition and development : proceedings of GALA 2007* (pp. 313–324). Newcastle: Cambridge Scholars Publishing.
- Vaux, B. (1998). *The phonology of Armenian*. Oxford University Press.
- White, J. (2014). Evidence for a learning bias against saltatory phonological alternations. *Cognition*, 130(1), 96–115. doi: <http://dx.doi.org/10.1016/j.cognition.2013.09.008>
- White, J. (2017). Accounting for the learnability of saltation in phonological theory: a Maximum Entropy model with a P-Map bias. *Language*, 93(1), 1–36.
- White, J., & Sundara, M. (2014). Biased generalization of newly learned phonological alternations by 12-month-old infants. *Cognition*, 133(1), 85–90.

- White, K. S., Peperkamp, S., Kirk, C., & Morgan, J. L. (2008). Rapid acquisition of phonological alternations by infants. *Cognition*, *107*, 238–265.
- Wilson, C. (2006). Learning phonology with substantive bias: An experimental and computational study of velar palatalization. *Cognitive science*, *30*(5), 945–982.
- Wolf, M. (2008). *Optimal interleaving: Serial phonology-morphology interaction in a constraint-based model* (Ph.D. Dissertation). University of Massachusetts, Amherst.
- Zimmer, K. E. (1969). Psychological correlates of some Turkish morpheme structure conditions. *Language*, *45*(2), 309–321.
- Zimmer, K. E., & Abbott, B. (1978). The k/ø alternation in Turkish: Some experimental evidence for its productivity. *Journal of Psycholinguistic Research*, *7*(1), 35–46.
- Zimmer, K. E., & Orgun, C. O. (1999). Turkish. In *Handbook of the International Phonetic Association: A guide to the use of the International Phonetic Alphabet* (pp. 154–158). Cambridge: Cambridge University Press.
- Zuraw, K. (n.d.). **MAP constraints*. (Unpublished Ms., UCLA)
- Zuraw, K. (2000). *Patterned exceptions in phonology* (Ph.D. Dissertation). University of California, Los Angeles, Los Angeles, CA.
- Zuraw, K. (2007). The role of phonetic knowledge in phonological patterning: Corpus and survey evidence from Tagalog. *Language*, *83*, 277–316.
- Zymet, J. (2014). *Distance-based decay in long-distance phonological processes: Probabilistic models for Malagasy, Latin, English, and Hungarian* (Master's Thesis). University of California, Los Angeles, Los Angeles, CA.