

**PHONOLOGICAL TRENDS IN THE LEXICON: THE ROLE OF
CONSTRAINTS**

A Dissertation Presented

by

MICHAEL BECKER

Submitted to the Graduate School of the
University of Massachusetts Amherst in partial fulfillment
of the requirements for the degree of

DOCTOR OF PHILOSOPHY

February 2009

Department of Linguistics

© Copyright by Michael Becker 2009

All Rights Reserved

PHONOLOGICAL TRENDS IN THE LEXICON: THE ROLE OF CONSTRAINTS

A Dissertation Presented

by

MICHAEL BECKER

Approved as to style and content by:

John J. McCarthy, Chair

Lyn Frazier, Member

Joe Pater, Member

Elisabeth Selkirk, Member

Shmuel Bolozky, Member

Adam Albright, Member

Elisabeth Selkirk, Department Chair
Department of Linguistics

In memory of my friend Dorit Markovich, 1961–2005

ACKNOWLEDGMENTS

Working on this dissertation was probably the longest and worst fun I've ever had. It was an enriching and humbling endeavor, which I only accomplished thanks to help from everybody I know.

First and foremost, to my committee chair, John J. McCarthy. I am yet to discover the bounds of his intellect, his patience, and his kindness. I feel greatly privileged to be counted among his many students. As for the other members of my committee, I thank Lyn Frazier for ruthlessly constructive criticism that helped make my work infinitely better. Joe Pater has been greatly influential in my interest in exceptions and variation, and I thank him for repeatedly amazing me with his ability to say the one sentence that untangles much confusion. I thank Lisa Selkirk for taking me on long journeys, showing me the world that lies beyond the Turkish possessive suffix, and helping me see the broader implications of my proposals. I thank Adam Albright for many long and deeply thought-provoking conversations. I tried to learn from Adam as much as I could. I thank Shmuel Bolozky for sharing his expertise and vast knowledge of Hebrew, and for giving me the opportunity to make an old dream of teaching Hebrew come true.

Beyond my committee, I thank John Kingston for saving my hide on multiple occasions. Thanks also to my other teachers at UMass, Kyle Johnson, Angelika Kratzer, Barbara Partee, Chris Potts, Peggy Speas, and Ellen Woolford. I thank my fellow monkeys, Kathryn Flack and Shigeto Kawahara, and the other UMass phonologists, Tim Beechey, Angela Carpenter, Della Chambless, Andries Coetzee, Emily Elfner, Ben Gelbart, Maria Gouskova, Karen Jesney, Mike Key, Wendell Kimper, Kathryn Pruitt, Takahito Shinya, Anne-Michelle Tessier, Adam Werle, and Matt Wolf for being the best thing about being a UMass student, for duck, for hugs, and for tea. Among the many others who made UMass

a great place, I would like to mention Jan Anderssen, Shai Cohen, Chris Davis, Amy Rose Deal, Annahita Farudi, Meg Grant, Minjoo Kim, Emerson Loustau, Andrew McKenzie, Aynat Rubinstein, Florian Schwartz, Cherlon Ussery, and Martin Walkow. Finally, special thanks are due to Kathy Adamczyk and to Sarah Vega-Liros.

At the Tel Aviv University, I thank Outi Bat-El for her constant support over the years, from my first day on campus and ever since. My interest in phonology owes much to her unique teaching style and her enthusiasm. I thank Chuck Kisseberth for showing me true commitment to linguistics. I also thank my other teachers, Galit Adam, Gabi Danon, Nirit Kadmon, Fred Landman, Aldo Sevi, and Tali Siloni. Thanks also to my colleagues at the Ben Gurion University, Nomi Erteschik-Shir, Idan Landau, and Yishai Tobin.

I thank Ram Frost of the Hebrew University's psychology department for his incomparable hospitality and generosity. By opening his lab to me and supporting me through many failed iterations of my experiments, he made my work on Hebrew possible. I also thank Hadas Velan, Boaz Tzur, Shlomit Brosh, and the other members of Ram's lab. They all deserve the finest chocolate.

I thank my collaborators in recent years, Lena Fainleib, Peter Jurgec, Nihan Ketrez, and Andrew Nevins, for their friendship and support, and for getting our prose out there. I thank Joey Sabbagh for showing me how co-teaching follows from general principles, and for helping me figure out the secret laws of Reed College. Thanks also to my other colleagues at Reed, Matt Pearson and Steven Hibbard.

For making linguistics as a field exciting and worth going into, I thank many people, including Diana Apoussidou, Sabrina Bendjaballah, Roy Becker (who is not my cousin), Adam Buchwald, Evan-Gary Cohen, Paul de Lacy, Noam Faust, Gillian Gallagher, Peter Graff, Bruce Hayes, Uri Horesh, Lary Hyman, Gaja Jarosz, Beste Kamali, Jonah Katz, Jaklin Kornfilt, Lior Laks, Susie Levi, Jean Lowenstamm, Bruce Morén, Elliott Moreton, Scott Myers, Sumru Özsoy, Omer Preminger, Curt Rice, Nathan Sanders, Donca Steriade, Bruce Tesar, Adam Ussishkin, Marc van Oostendorp, Andy Wedel, and Kie Zuraw.

For saving my life, I thank Dr. Deborah Smith. Her powers of healing far exceeded those of the medications she administered. I am forever indebted to Kathryn Flack and Matt Wolf, who were there for me when I needed them most. They are true friends. I am grateful to John McCarthy and Lisa Selkirk for their support in difficult times, and for helping minimize the impact of my illness on my graduate career.

I thank my parents Yigal and Ada Becker, for believing that with constant love and unconditional support I could parallel their level of education. Thanks also to my brothers, my aunts, my uncles, my cousins, and to Nir, Yoav, and Vered. I thank my grandmother Rivka Shershevsky for her love and above all for her patience.

Lastly, I thank Tom, for his love and companionship, in sickness and in health, and for giving me Tulip.

ABSTRACT

PHONOLOGICAL TRENDS IN THE LEXICON: THE ROLE OF CONSTRAINTS

FEBRUARY 2009

MICHAEL BECKER

M.A., TEL AVIV UNIVERSITY

Ph.D., UNIVERSITY OF MASSACHUSETTS AMHERST

Directed by: Professor John J. McCarthy

This dissertation shows that the generalizations that speakers project from the lexical exceptions of their language are biased to be natural and output-oriented, and it offers a model of the grammar that derives these biases by encoding lexical exceptions in terms of lexically-specific rankings of universal constraints in Optimality Theory (Prince & Smolensky 1993/2004). In this model, lexical trends, i.e. the trends created by the phonological patterning of lexical exceptions, are incorporated into a grammar that applies deterministically to known items, and the same grammar applies stochastically to novel items. The model is based on the Recursive Constraint Demotion algorithm (Tesar & Smolensky 1998, 2000; Tesar 1998; Prince 2002), augmented with a mechanism of constraint cloning (Pater 2006, 2008b).

Chapter 2 presents a study of Turkish voicing alternations, showing that speakers replicate the effects that place of articulation and phonological size have on the distribution

of voicing alternations in the lexicon, yet speakers ignore the effects of vowel height and backness. This behavior is tied to the absence of regular effects of vowel quality on obstruent voicing cross-linguistically, arguing for a model that derives regular phonology and irregular phonology from the same universal set of OT constraints.

Chapter 3 presents a study of Hebrew allomorph selection, where there is a trend for preferring the plural suffix [-ot] with stems that have [o] in them, which is analyzed as a markedness pressure. The analysis of the trend in terms of markedness, i.e. constraints on output forms, predicts that speakers look to the plural stem vowel in their choice of the plural suffix, and ignore the singular stem. Since real Hebrew stems that have [o] in the plural also have [o] in the singular, Hebrew speakers were taught artificial languages that paired the suffix [-ot] with stems that have [o] only in the singular or only in the plural. As predicted, speakers preferred the pairing of [-ot] with stems that have [o] in the plural, i.e. speakers prefer the surface-based, output-oriented generalization.

Chapter 4 develops the formal theory of cloning and its general application to lexical trends, and explores its fit with the typologically available data. One necessary aspect of the theory is the “inside out” analysis of paradigms (Hayes 1999), where the underlying representations of roots are always taken to be identical to their surface base form, and abstract underlying representations are limited to affixes. An algorithm for learning the proposed underlying representations is presented in a general form and is applied to a range of test cases.

TABLE OF CONTENTS

	Page
ACKNOWLEDGMENTS	v
ABSTRACT	viii
 CHAPTER	
1. INTRODUCTION	1
1.1 Lexical trends and constraint cloning	1
1.1.1 Identifying lexical trends	4
1.1.2 Lexical trends and conflicting grammars	4
1.1.3 Constraint cloning	7
1.1.4 Replicating lexical statistics	11
1.2 Structure of the dissertation	17
2. UNIVERSAL LEXICAL TRENDS IN TURKISH	19
2.1 Introduction	19
2.2 Turkish lexicon study	23
2.3 Speakers' knowledge of the lexicon	32
2.3.1 Materials and method	33
2.3.1.1 Speakers	33
2.3.1.2 Materials	33
2.3.1.3 Procedure	35
2.3.2 Results	36
2.3.3 Discussion	41
2.4 Analysis with cloned constraints	42
2.4.1 Constraint cloning	43
2.4.2 The place effect	47

2.4.3	The size effect	51
2.4.4	Combining place and size	56
2.4.5	The complex coda effect	59
2.4.6	Voicing alternations and $k \sim \emptyset$ alternations	61
2.4.7	Summary of the analysis	64
2.5	General-purpose learning with the MGL	65
2.5.1	Materials and method	66
2.5.2	Results	67
2.5.3	Discussion	68
2.6	UR-based approaches to final devoicing	69
2.7	Conclusions	71
3.	SURFACE-BASED LEXICAL TRENDS IN HEBREW	74
3.1	Introduction	74
3.2	Hebrew plurals: Lexicon study	76
3.3	Speakers' knowledge of lexical trends	82
3.3.1	Materials and methods	82
3.3.2	Participants	84
3.3.3	Results	84
3.3.4	Discussion	88
3.4	Using markedness constraints to learn lexical trends	90
3.4.1	Analysis	91
3.4.2	Ranking conflicts trigger the formation of generalizations	99
3.4.3	Learning specific patterns first	104
3.4.4	Vowel harmony and [o]-licensing	109
3.5	Product-orientedness in an artificial language	112
3.5.1	Materials	115
3.5.2	Methods	116
3.5.3	Participants	119
3.5.4	Transcription and encoding	120
3.5.5	Results	121
3.5.5.1	Generalization differences between the groups	122
3.5.5.2	No memorization differences between the groups	126
3.5.5.3	Bias towards <i>-im</i>	128
3.5.5.4	Errors and vowel perception	130
3.5.5.5	Summary of the experimental results	131

3.6	Discussion and analysis	131
3.6.1	The role of Universal Grammar in learning alternations	132
3.6.2	Stem changes and allomorph selection	134
3.6.3	The limited role of phonotactics	136
3.6.4	Learning alternations without Universal Grammar	138
3.6.5	The role of the grammar of real Hebrew	140
3.6.6	Source-oriented generalizations?	141
3.7	Conclusions	142
4.	LEXICAL TRENDS AS OPTIMALITY THEORETIC GRAMMARS	145
4.1	Introduction	145
4.2	Choosing the constraint to clone	145
4.2.1	Minimal conflict	146
4.2.2	Two independent conflicts	149
4.2.3	Overlapping conflicts	150
4.2.4	Interim summary: Choosing the least populated column	153
4.2.5	General-specific relations between trends; masking	155
4.2.6	Remaining questions	159
4.3	The cloning algorithm	161
4.3.1	Background: RCD	162
4.3.2	Cloning RCD	165
4.3.3	The domain of cloned constraints	169
4.3.4	Applying a grammar with cloned constraints	175
4.3.5	Searching for the UR of affixes	177
4.3.6	Supplying losers	185
4.3.7	Exceptionality and variation	188
4.4	Moving hidden structure into the grammar	191
4.4.1	Hidden structure in the grammar: Turkish	192
4.4.2	Hidden structure in affixes: Korean	196
4.4.3	Interim summary: Generalizing across affixes	204
4.4.4	Hidden structure in roots: English	206
4.4.5	The naturalness of lexical trends: Dutch	208
4.5	Cloning and the typology of lexical trends	211
4.5.1	Contrast neutralization and creation	211
4.5.2	Competing repairs	213
4.5.3	Exceptional emergence of the unmarked	214

4.6	Conclusions	218
5.	CONCLUSIONS AND FUTURE DIRECTIONS	220
5.1	Summary of the dissertation	220
5.2	Future directions	222
5.2.1	Unnatural phonology	222
5.2.2	The search for underlying representations	224
5.2.3	Issues in lexical organization	226
5.2.4	Lexically-specific grammars and phonotactics	227
	BIBLIOGRAPHY	228

CHAPTER 1

INTRODUCTION

1.1 Lexical trends and constraint cloning

In a wide variety of languages, there are cases of morphological categories that are expressed in more than one way. In English, for instance, the past tense is expressed on the majority of verbs by adding *-ed*, but on some verbs, the past tense is expressed by changing a vowel to [ɛ], e.g. *feed* ~ *fed*, *hold* ~ *held*.

A common theme in such limited-scope processes is their reported applicability to novel words. English speakers, for instance, are willing to offer *pred* as the past tense of *preed*, productively extending the limited pattern of changing a root vowel to [ɛ] (Albright & Hayes 2003).

Furthermore, speakers' willingness to apply a limited process to some novel form X depends on the number of existing base forms like X that do and don't undergo the minority process. Speakers are aware of the proportion of the words that undergo a minority process out of the total number of eligible words, i.e. speakers identify a trend in the application of the process in their lexicon (henceforth, a lexical trend), and apply this trend to novel items. Results of this type are reported by Zuraw (2000), Albright & Hayes (2003), Hayes & Londe (2006), Becker, Ketrez & Nevins (2007), and several others.

The wish to account for lexical trends in grammatical terms goes back at least as far as SPE (Chomsky & Halle 1968), where some lexical trends were derived by minor rules, i.e. rules that are formulated using the same mechanisms that are used for regular rules, but with a limited lexical scope. Other grammatical mechanisms, such as stochastic grammars, were offered in Zuraw (2000) and Hayes & Londe (2006), among others. There

are several reasons for thinking about lexical trends in grammatical terms: One reason is that lexical trends are stated with reference to the same objects that are characteristic of regular grammatical phenomena, such as phonological elements (features, syllables, etc.) and morphological elements (noun, root, etc.). Another, related reason is that lexical trends in one language are often found as regular grammatical processes in other languages: For example, intervocalic voicing is regular in Korean, but is a trend in Turkish, affecting stem-final stops in some words but not others.

Much work on lexical trends assumes a grammar-external mechanism, such as Pinker & Prince's (1988) dual model. In this line of work, grammar (as constrained by Universal Grammar) is in charge of the "regular rules" of the language, while minority patterns are taken care of by associative networks. This view makes the prediction that Universal Grammar effects will not be visible in lexical trends – a prediction not borne out by observation.

A study of the distribution of voicing alternations in Turkish (chapter 2, see also Becker, Ketrez & Nevins 2007) shows that speakers are constrained by Universal Grammar when they learn this distribution. Turkish speakers replicate the effect of grammatical principles on the distribution, such as initial syllable faithfulness and place of articulation, and ignore non-grammatical principles, such as a relationship between vowel height and the voicing of a following consonant.

In work on plural selection in Hebrew (chapter 3), I show that speakers select plural suffixes based on the surface form of the plural stem rather than based on the stem's underlying representation, even though there is no evidence in the existing words of Hebrew for stating the generalization over surface forms. This preference is attributed to the markedness component of Universal Grammar, which is biased towards stating generalization over surface forms.

The product-oriented aspect of lexical trends was also noted in Albright & Hayes (2003). In the English past tense, several vowels in the present change to [o] in the past: [aɪ]

(e.g. *drive* ~ *drove*), [ei] (e.g. *break* ~ *broke*), [i] (e.g. *freeze* ~ *froze*), [u] (e.g. *choose* ~ *chose*). Speakers go beyond the observed mappings, and are willing to change any vowel in the present tense to [o] to make the past tense. Having several different phonological processes converge on the same output (a “conspiracy”, Kisseberth 1970) is a hallmark of grammatical behavior, and one of the central arguments in favor of using Optimality Theory (Prince & Smolensky 1993/2004).

Since speakers treat lexical trends as grammatical processes that have limited lexical scope, and since they are able to apply these processes to novel forms, one concludes that the grammar of a language needs to account for this behavior. Within the framework of Optimality Theory, a central approach in accounting for lexical trends is based on stochastic grammar (Boersma 1997), used in the analysis of lexical trends in Tagalog (Zuraw 2000) and in Hungarian (Hayes & Londe 2006). This approach and its relation to the proposal made here are discussed in §4.3.7.

To summarize, lexical trends show all the aspects of grammatical phenomena, and they should be described with the same mechanisms linguists use to describe regular grammatical phenomena. The desired theory will be able to take the existing words of the lexicon, extract statistical grammatical generalizations from them, and be able to project these generalizations onto novel words. Previous work in OT provided a way for projecting statistical grammatical generalizations onto novel words, but no mechanism was offered for extracting those generalizations from the existing words of the language. Work outside OT was able to extract generalizations from existing words, but those generalizations were not constrained by Universal Grammar, unlike the generalizations that humans extract from the words of their language.

I offer an OT-based model that uses constraint interaction to extract statistical generalizations from a lexicon and project them onto novel items. The model relies on the treatment of different processes within a single morphological category as a competition

between conflicting grammars, which give rise to competing constraint rankings in Optimality Theory.

1.1.1 Identifying lexical trends

When the expression of a single morphological category is unpredictable given the base form, lexical trends may arise. The past tense in English, for example, is not completely predictable given a verb root: The past tense may be expressed by suffixation of *-ed* (pronounced predictably as [d], [t] or [ɪd]), a change of a root vowel (e.g. *feed* ~ *fed*), or no change at all (e.g. *spread* ~ *spread*)¹. Results from Albright & Hayes (2003) clearly show that speakers identify partial generalizations, or trends, in the distribution of the different realizations of the past tense. For instance, among the real words of English, only verbs that end in [d] (e.g. *spread*, *rid*, *shed*) or [t] (e.g. *set*, *cut*, *split*, *burst*) can stay unchanged in the past. When given a novel verb, speakers replicate this lexical generalization, and only accept verbs as unchanged in the past when they end in [t] or [d] (e.g. *sned* can stay unchanged in the past, while *stib* cannot).

As discussed below and in chapter 4, speakers use ranking arguments to identify unpredictable patterns in the language they are exposed to, and they build information about lexical items into their constraint ranking. This lexically-enhanced grammar in turn allows speakers to replicate generalizations about their lexicon in dealing with novel items.

1.1.2 Lexical trends and conflicting grammars

The fact that English verbs can stay unchanged in the past only if they have a final [t] or [d] is not surprising given the presence of [d] in the regular *-ed* past, and an analysis that connects these two facts would seem like an insightful one. Optimality theory allows the generalization to be captured fairly easily: Given an underlying suffix [-d] and a constraint

¹Other expressions of the past tense include the unpredictable selection of [-t] after {n,l}-final roots (*learn* ~ *learn-t*, *spell* ~ *spel-t*), the change of a final [d] to [t] after {n,l} (*send* ~ *sent*, *build* ~ *built*), and the combination of a vowel change (most often to [ɛ]) and t-affixation (*sweep* ~ *sweep-t*).

that forbids clusters of alveolar stops, like [dt] and [dd], regular verbs resolve the cluster by epenthesis, and verbs that stay unchanged in the past resolve the cluster by deletion or fusion. Verbs that don't end in [d] or [t] don't violate the constraint on alveolar stop clusters, and thus have no reason to stay unchanged in the past.

The tableau in (1) shows the derivation of the verb [gɑrd] (*guide*). The first candidate in (1) is the winner, with an epenthetic vowel and hence a violation of DEP. The second candidate is zero-marked (i.e. it sounds identical to the root) by virtue of deleting the affixal [d], thus violating MAX². The final candidate is the faithful one, which violates a constraint on clusters of alveolar stops (*DD, see also Borowsky 1987), which is undominated in English.

(1)

/gɑrd + d/	*DD	MAX	DEP
a. ↵ gɑrdɪd			*
b. gɑrd		*!	
c. gɑrdd	*!		

The derivation of the zero-marked verb [sprɛd] (*spread*) is shown in (2). In order to make the zero-marked form the winner, DEP must dominate MAX, which is the opposite of the ranking required by *guide*.

²Alternatively, zero-marked verbs avoid a violation of *DD by fusing the root [t] or [d] and the suffixal [d], violating UNIFORMITY.

(2)

/sprɛd + d/	*DD	DEP	MAX
a. ☞ sprɛd			*
b. sprɛdid		*!	
c. sprɛdd	*!		

In terms of OT, then, zero-marked verbs are simply responding to a constraint ranking that's different from the constraint ranking that controls the regular verbs of the language. Regular *ed*-taking verbs that end in [t] or [d] require MAX to dominate DEP, whereas zero-marked verbs require the opposite ranking.

Verbs that do not end in [t] or [d], such as [star] (*star*), shown in (3), cannot be zero-marked using [d] as the underlying form of the past tense morpheme. The fully faithful form *starred* harmonically bounds the zero-marked form, since it doesn't violate any of the relevant constraints, including the one against clusters of alveolar stops. No ranking of these constraints can produce the zero-marked *star* as the past tense of *star*.

(3)

/star + d/	*DD	DEP	MAX
a. ☞ stard			
b. starɪd		*!	
c. star			*!

To summarize the result so far: Subjecting different verbs to different constraint rankings allows verbs to be zero-marked in the past only if they end in [t] or [d].

Furthermore, this result was derived from two other facts about English: (a) the language disallows final clusters of alveolar stops, and (b) the past tense is regularly marked by affixation of [d].

Zero-marking of the past tense was presented here as an alternative mechanism for satisfying a phonotactic constraint on English words, *DD. While regular verbs satisfy *DD by violating DEP, some verbs satisfy *DD by violating MAX. In other words, different verbs in English respond to different grammars: Verbs like *guide* respond to a grammar that requires $MAX \gg DEP$, while verbs like *spread* respond to a grammar that requires $DEP \gg MAX$. Verbs that don't end in [t] or [d], like *star*, are compatible with either ranking.

Learners can discover that different words of their language respond to different grammars, and then they can keep track of the grammar that each word requires. A mechanism for doing so depends on detecting inconsistency (Prince & Tesar 1999) and then solving the inconsistency by constraint cloning (Pater 2006, 2008b), as shown in the next section.

1.1.3 Constraint cloning

If English speakers are to recognize that the verbs *guide* and *spread* respond to different constraint rankings, they need to be able to extract ranking information from these words, and then discover that those rankings are mutually incompatible.

A simple way of doing this is by using winner-loser pairs (Tesar 1995 et seq.). For instance, the winner [gardid] from the tableau in (1), repeated as (4) below, can be paired with each of the two losers, [gard] and [gardd], to produce two winner-loser pairs (5). The result is a comparative tableau (Prince 2002), where a W means that a constraint prefers the winner (i.e. the constraint assigns less violation marks to the winner than it does to the loser), and an L means that a constraint prefers the loser (i.e. the constraint assigns less violation marks to the loser).

(4)

/gard + d/	*DD	MAX	DEP
a. \rightarrow gardɪd			*
b. gard		*!	
c. gardd	*!		

(5)

/gard + d/	*DD	MAX	DEP
a. gardɪd \succ gardd	W		L
b. gardɪd \succ gard		W	L

A row that has just one W and one L in it simply means that the constraint that assigned a W to the row must dominate the constraint that assigned an L to the row. Therefore, the first winner-loser pair reveals that $*DD \gg DEP$, and the second winner-loser pair reveals that $MAX \gg DEP$.

Making a comparative tableau out of the tableau in (2) yields (6). The first winner-loser pair reveals that $*DD \gg MAX$, and the second winner-loser pair reveals that $DEP \gg MAX$.

(6)

/sprɛd + d/	*DD	DEP	MAX
a. sprɛd \succ sprɛdd	W		L
b. sprɛd \succ sprɛɪd		W	L

One advantage of comparative tableaux over regular tableaux is that comparative tableaux can be combined, as in (7), which combines (5) and (6).

(7)

	*DD	MAX	DEP
a. gardɪd > gardd	W		L
b. gardɪd > gard		W	L
c. sprɛd > sprɛdd	W	L	
d. sprɛd > sprɛdɪd		L	W

The comparative tableau in (7) allows the ranking arguments from *guided* and *spread* to be compared and contrasted. Following the Recursive Constraint Demotion algorithm (RCD, Tesar & Smolensky 1998, 2000; Tesar 1998; Prince 2002), constraint rankings are discovered by identifying columns that only have W's and empty cells in them, "installing" them in a ranking, and then removing any winner-loser pairs that the installed constraints assigned W's to. In this case, *DD is such a constraint, since it only has W's in its column. It can be installed as the top-ranked constraint in the language, and winner-loser pairs (a) and (c) can be removed. The remaining comparative tableau is in (8).

(8)

	MAX	DEP
a. gardɪd > gard	W	L
b. sprɛd > sprɛdɪd	L	W

At this point, the ranking algorithm stalls, since there are no more constraints that have only W's in their columns. The information about the ranking of MAX and DEP is exactly contradictory: the first winner-loser pairs demands $\text{MAX} \gg \text{DEP}$, and the second winner-loser pair requires $\text{DEP} \gg \text{MAX}$. In the original RCD, inconsistency detection causes the ranking-finding process to stop, given RCD's focus on systems that can be described with a single consistent ranking. To extend this approach to systems that have exceptions, Pater (2006, 2008b) suggests that exceptional morphemes require a grammar that is inconsistent with the regular grammar of the language, and therefore inconsistency is a property of natural languages, and must be resolved. Pater suggests that a constraint be cloned, i.e. an extra copy of the constraint be made, and the new copy be made specific to the exceptional morpheme involved. In the English case at hand, either MAX or DEP will be cloned and the clone will be made specific to the root *spread*. Having exceptional morphemes be subject to lexically-specific clones and regular morphemes be subject to the general constraints, allows the different behavior of different morphemes to be captured in a single, consistent constraint ranking.

In the current proposal, constraint cloning does not result in one general constraint and one-lexically specific constraint, but rather two lexically-specific constraints. The reason for that will be made clear in §1.1.4.

In the English case, one of the constraints, either MAX or DEP, will be cloned. One clone will list verbs that end in [d] or [t] and take [-ɪd] in the past, like *guide*, and another clone will list zero-marked verbs like *spread*. The inconsistency in (8), then, triggers the cloning of one of the constraints. The result of cloning MAX is shown in (9), where each clone is specific to a lexical item.

(9)

	MAX_{gaid}	DEP	$MAX_{sprɛd}$
a. $gaidɪd \succ gaid$	W	L	
b. $sprɛd \succ sprɛɪd$		W	L

Since the comparative tableau in (9) contains a column that only has W's in it, the search for a ranking can continue³, and a consistent grammar for English can be obtained: First, the constraint MAX_{gaid} is installed, and the first winner-loser pair is removed.

With only the second winner-loser pair of (9) remaining, DEP can be installed. It will be added to the ranking below the last constraint to be installed, MAX_{gaid} , and the second winner-loser pair is removed. The remaining $MAX_{sprɛd}$ is left with no winner-loser pairs to deal with, so it is installed below DEP. The obtained grammar is $MAX_{gaid} \gg DEP \gg MAX_{sprɛd}$.

To motivate the lexical-specific nature of both clones, and discuss the exact nature of cloning, I turn to a discussion of lexical trends in Turkish.

1.1.4 Replicating lexical statistics

Identifying the existence of irregular patterns in a language is a necessary condition for learning a human language successfully, but it is not a sufficient condition. Language learners must also find the relative strength of competing patterns. When two behaviors compete for the same set of lexical items, such as the deletion and the epenthesis that compete for the *d*- and *t*-final verbs of English, as discussed above, speaker don't just recognize the existence of the two patterns, but also recognize how well-attested each

³Once a constraint is cloned, the search for a ranking can either starts from the beginning with the full set of winner-loser pairs, or equivalently, simply continue with the winner-loser pairs that were left over at the point of cloning. Starting the search for ranking from scratch only needs to happen when winner-loser pairs are added or removed, as discussed in §4.2.

pattern is. Speakers use their grammar to estimate the relative likelihood of the various behaviors that the grammar allows, and use this estimate to decide the fate of novel items they encounter. This section shows how constraint cloning can be used to extract the relative strength of an irregular pattern from the lexicon.

In Turkish, stem-final voiceless stops become voiced when an affix (such as the possessive) makes them intervocalic. This process applies to some words (10a), but not others (10b).

(10)	bare noun	possessive	
a.	tat	tad-i	‘taste’
	taɰ	taɰ-i	‘crown’
b.	at	at-i	‘horse’
	aɰ	aɰ-i	‘hunger’

The Turkish phenomenon is similar to the case of the English past tense: Different words of Turkish behave differently, and this difference can be captured in terms of constraint rankings. In Turkish, the relevant markedness constraints are those against intervocalic voiceless stops, such as *VtV and *VɰV⁴. In words like the ones in (10a), *VtV and *VɰV outrank faithfulness to voicing, causing a voiceless stop to become voiced. In words like the ones in (10b), faithfulness outranks *VtV and *VɰV, leaving the stem unchanged in the suffixed form. Note that faithfulness to voicing is violated in (10a) only if the stem-final stop is taken to be underlyingly voiceless, as it is in the bare noun. I will

⁴*VtV and *VɰV are not generally active in Turkish, and voiceless intervocalic stops occur freely in roots, e.g. *ata* ‘father’, *paɰa* ‘trotter’. The effect of *VtV and *VɰV must be limited in Turkish to derived environments, i.e. they must only affect stops that have become intervocalic under affixation. While this restriction could in principle be built into the definition of the constraints, e.g. *Vt]V, where the square bracket notes a morpheme boundary, a more attractive solution is offered in Wolf (2008b), who shows that principles of OT-CC (McCarthy 2007a) can be used to account for derived environment effects without hard-wiring these effects into the definition of the constraints.

assume that the learner takes the bare noun to be the underlying representation, a move that I discuss and motivate in §4.4.

In Turkish, the proportion of *t*-final nouns that exhibit the voicing alternation is low relative to the proportion of *tf*-final nouns that exhibit the voicing alternation. Speakers are aware of this difference, and when they are given novel *t*-final and *tf*-final nouns and are asked to add the possessive suffix, they choose voicing alternations more often with *tf*-final nouns than with *t*-final nouns. This replication of the relative strength of lexical trends in novel nouns is by no means restricted to Turkish, and it has been observed in a variety of languages, e.g. Tagalog (Zuraw 2000), Dutch (Ernestus & Baayen 2003), and many others.

The table in (11) shows counts of *t*-final and *tf*-final monosyllabic nouns in the Turkish Electronic Living Lexicon (TELL, Inkelas et al. 2000). The crucial point to notice here is that the 18 *t*-final nouns that alternate are more numerous than the 15 *tf*-final nouns that alternate, yet the alternating *t*-final nouns make only 15% of the total *t*-final nouns, relative to the larger 37% alternation rate among the *tf*-final nouns. So while *t*-final nouns show more alternation in absolute numbers, they show a smaller proportion of alternation. Since speakers prefer alternating [tʃ] to alternating [t], one can conclude that what speakers are attending to is not the number of alternating nouns for a given segment, but rather the number of alternating nouns relative to the number of non-alternating nouns for that segment.

(11)		alternating	non-alternating	% alternating
	t	18	102	15%
	tf	15	26	37%

It should also be pointed out that speakers must be able to keep track of alternation rates for [t] separately from [tʃ], rather than simply compute a single, global rate of alternations for all consonants. To achieve this result, speakers must come with a pre-existing propensity to keep track of the behavior of different segments separately, since

once two segments are merged into one category, there will be no overt evidence to suggest that they should be separated.

To achieve the intended result, i.e. to give the grammar a way to compare the relative numbers of alternating and non-alternating items, cloned constraints must keep track of both kinds of items. This is done by making all cloned constraints lexically-specific, rather than keep a general version of cloned constraints, as in Pater (2006, 2008b).

Turkish supplies conflicting evidence for the ranking of IDENT(voice), which penalizes voicing alternations, with respect to the ranking of *VtV and *VtʃV, which penalize intervocalic voiceless dental and pre-palatals stops, respectively. The comparative tableau in (12) shows the two kinds of *t*-final nouns.

(12)

	*VtV	IDENT(voice)
a. tad-i > tat-i	W	L
b. at-i > ad-i	L	W

Once the learner is exposed to the two kinds of *t*-final nouns, the ranking of *VtV relative to IDENT(voice) can no longer be found, since neither constraint has only W's in its column. The learner will then clone a constraint, in this case, *VtV (see §4.2 about choosing which constraint to clone). Both clones are made lexically-specific, and the result is the comparative tableau in (13), which gives rise to the grammar *VtV_{tat} >> IDENT(voice) >> *VtV_{at}.

(13)

	*VtV _{tat}	IDENT(voice)	*VtV _{at}
a. tad-i ∿ tat-i	W	L	
b. at-i ∿ ad-i		W	L

Since the general *VtV is no longer present in the grammar, the learner will have to list any new *t*-final nouns they encounter with one of the clones of *VtV. Items that get a W from *VtV will be listed with *VtV_{tat}, and items that get an L will be listed with *VtV_{at}. As the nouns tallied in (11) are gradually learned, the resulting grammar will approach the one in (14).

(14) *VtV_{18 items} ≫ IDENT(voice) ≫ *VtV_{102 items}

In this resulting grammar, most *t*-final nouns are listed with the clone of *VtV that ranks below IDENT(voice), meaning that their final [t] will surface unchanged in the suffixed form. Only 18 nouns are listed with the high-ranking clone of *VtV, making their [t] become a [d] intervocally. Since both kinds of nouns are listed in the grammar, the relative size of each group is available to the speaker, and the speaker can project the relative probability of alternation onto a novel word: When offered a novel *t*-final bare noun, and asked to derive its suffixed form, the speaker can randomly choose one of their listed *t*-final nouns and make the novel noun behave like it. Since only 15% of the listed nouns are listed above IDENT(voice), there is only a 15% chance for the novel noun to alternate. In effect, by choosing randomly from the nouns that are listed in the grammar, the speaker causes the likelihood of alternation of the novel noun to match the likelihood of alternation in the grammar, which in turn matches the likelihood of alternation in the lexicon.

Similarly for the *tf*-final nouns, once the speaker encounters *tf*-final nouns that do and do not alternate, they will clone *VtʃV, and eventually reach the grammar in (15).

$$(15) \quad *VtʃV_{15 \text{ items}} \gg \text{IDENT(voice)} \gg *VtʃV_{26 \text{ items}}$$

For the *tf*-final nouns, there are only 15 items listed with the clone of *VtʃV that ranks above IDENT(voice), compared to the 18 *t*-final nouns listed above IDENT(voice), but these 15 nouns make more than 40% of the total number of *tf*-final nouns, making the likelihood of an alternating [tʃ] higher than the likelihood of an alternating [t].

One of the responsibilities of the grammar is to estimate the relative likelihood of the various behaviors that it allows, letting speakers build on their knowledge of the lexicon when asked to use a novel item. The use of constraint cloning, as shown here, allows speakers to identify the existence of irregular patterns and also extract their relative strength from the lexicon.

The grammars in (14) and (15) are compatible with each other, as shown in (16), where they are combined. The two clones of *VtV list *t*-final nouns, while the two clones of *VtʃV list *tf*-final nouns.

$$(16) \quad *VtV_{18 \text{ items}}, *VtʃV_{15 \text{ items}} \gg \text{IDENT(voice)} \gg *VtV_{102 \text{ items}}, *VtʃV_{26 \text{ items}}$$

The grammar in (16) ensures that the listed items behave as expected, e.g. that the possessive form of *tatf* always comes out as *taçʃ-i* and never as **tatf-i*. Furthermore, the same grammar ensures that a novel *t*-final noun will probably keep its [t] voiceless in the possessive form, while a novel *tf*-final noun will be more likely to respect *VtʃV by alternating the [tʃ] with a [çʃ]. In other words, the same grammar derives the categorical behavior of listed items, and projects the trends that the listed items create onto novel items stochastically.

1.2 Structure of the dissertation

After the introduction to lexical trends and their treatment in OT using constraint cloning, two case studies are presented.

The first case study is Turkish voicing alternations, discussed in chapter 2. It presents a study of the Turkish lexicon, and compares it to results from a novel word task experiment, showing that speakers project lexical statistics onto novel items. Speakers use the size of words (mono- vs. poly-syllabic) and the identity of their final stop to define classes of similar lexical items, and project the behavior of each class onto novel items. Speakers do not use, however, the quality of the word-final vowel in calculating this similarity. I relate this language-specific observation to the cross-linguistic observation about speakers' reluctance to learn a relationship between vowel quality and the voicing of a neighboring consonant (Moreton 2008). The connection between language-specific lexical trends and cross-linguistic typological observations is formalized by deriving both kinds of phenomena from a single inventory of universal constraints, CON. The use of CON to express lexical trends means that only trends that can be expressed in terms of universal constraints can be learned. In other words, speakers use universal considerations when they assess the similarity of lexical items.

The second case study is Hebrew plural allomorphy, discussed in chapter 3. Again, a lexicon study is compared with results from a novel word task experiment, showing that speakers project a trend from their lexicon onto novel words. When choosing a plural suffix for masculine nouns, *-im* is chosen in the majority of cases, but the presence of an [o] in the stem significantly boosts the likelihood of choosing the plural allomorph *-ot*. In real Hebrew, every plural noun that has an [o] in its stem also has an [o] in the singular, so in real Hebrew, the connection between the presence of the [o] in the stem and the selection of the suffix *-ot* can be stated equally well over the singulars, the plurals, or the mapping between singulars and plurals. In an artificial mapping experiment, Hebrew speakers were asked to learn novel vowel mappings between singular and plural stems that put [o] only in

the singular or only in the plural. The speakers showed a preference for selecting the plural affix based on the vowel present in the plural stem. This preference doesn't come from real Hebrew, and I propose that it comes from universal grammar. I formalize this preference with the use of markedness constraints, which only assess output forms, in this case, plural forms.

With the support gathered in chapters 2 and 3 for the use of Optimality Theory to account for lexical trends, a formal theory is developed in chapter 4. I offer an extension of the Recursive Constraint Demotion algorithm (RCD, Tesar & Smolensky 1998, 2000; Tesar 1998; Prince 2002) with constraint cloning (Pater 2006, 2008b) that learns a grammar from language data that includes lexically-specific phonological processes. This chapter also offers a discussion of the revised assumptions about underlying representations in this model, specifically, the restriction of non-surface-true underlying representations to affixes, leaving roots necessarily surface-true underlyingly. Finally, the typology of lexical trends that the model predicts is examined.

CHAPTER 2

UNIVERSAL LEXICAL TRENDS IN TURKISH

2.1 Introduction

This chapter examines the phonology of voicing alternations in Turkish, and shows that Turkish speakers display a detailed, yet imperfect knowledge about trends in their lexicon. I propose that the source of the imperfection is Universal Grammar, which biases learners to notice some trends and ignore others.

Voicing alternations in Turkish are observed at the right edges of nouns, as in (17). Nouns that end in a voiceless stop in their bare form, such as the pre-palatal stop [tʃ], can either retain that [tʃ] in the possessive (17a-b), or the [tʃ] of the bare stem may alternate with the voiced [dʒ] in the possessive (17c-d).

(17)	bare stem	possessive	
a.	atʃ	atʃ-i	‘hunger’
b.	anaʃ	anaʃ-i	‘female cub’
c.	taʃ	taɟ-i	‘crown’
d.	amaʃ	amaɟ-i	‘target’

Whether the final stop of a given noun will or will not alternate is unpredictable. However, the noun’s size strongly correlates with its status: Most monosyllabic nouns do not alternate, while most poly-syllabic nouns do. Section §2.2 discusses several other factors that correlate with voicing alternations, and shows that Turkish speakers use only a subset of the available factors: They use the noun’s size and the place of articulation

of the final stop, but they do not use the quality of the vowel that precedes the word-final stop. A back vowel before a word-final [tʃ], for instance, correlates with more alternations, but Turkish speakers ignore this correlation in their treatment of novel nouns. This language-specific behavior can be understood from a cross-linguistic perspective: Typological observations commonly correlate the distribution of voice with a word's size and a consonant's place of articulation, but rarely or never with the quality of a neighboring vowel. Indeed, speakers are reluctant to learn patterns that correlate vowel height with the voicing of a neighboring consonant (Moreton 2008, see also Moreton & Thomas 2007).

From a cross-linguistic perspective, it is unsurprising that mono-syllabic nouns would behave differently from poly-syllabic nouns with respect to the voicing alternation. Initial syllables are often protected from markedness pressures, showing a wider range of contrasts and an immunity to alternations (Beckman 1998). Specifically in Turkish, the privileged status of the feature [voice] in initial syllables is not only seen in voicing alternations. Generally in the language, a coda stop followed by an onset stop will surface with the voicing feature of the onset stop (e.g. *is.tib.dat* 'despotism', **is.tip.dat*), but a coda stop in the initial syllable may disagree in voice with the following onset (e.g. *mak.bul* 'accepted', *eb.kem* 'mute').

The backness of a neighboring vowel, however, is never seen to interact with a consonant's voicing. While such a connection is mildly phonetically plausible (vowel backness correlates with tongue-root position, which in turn correlates with voicing), there is no known report of any language where consonant voicing changes depending on the backness of a neighboring vowel, or vice versa. Given this gap in the universal inventory of possible phonological interactions, it is no longer surprising that in Turkish, speakers show no sign of using vowel backness as a predictor of voicing alternations.

In Optimality Theory (Prince & Smolensky 1993/2004), typological observations are encoded in the structure of the universal inventory of constraints (CON). The constraints are crafted such that their interactions produce all and only the observed sound patterns

of the world's languages. The preferred status of initial syllables is encoded with a set of faithfulness constraints specific to initial syllables. The lack of interaction between vowel backness and voicing is encoded by the exclusion of constraints from CON that penalize some value of $[\pm\text{back}]$ next to some value of $[\pm\text{voice}]$, e.g. $*[+\text{back}][+\text{voice}]$. In the absence of such constraints, there is never a reason to change one of these features in the presence of the other, and the lack of interaction is predicted. The account of the Turkish facts offered here capitalizes on these aspects of CON, while remaining agnostic about the mechanism that excludes these constraints, be it by assuming an innate set of constraints (which is the regular assumption in OT since Prince & Smolensky 1993/2004, and in the context of learning in Tesar & Smolensky 1998, 2000; Tesar 1998; Hayes 2004; Jarosz 2006; Tesar & Prince 2006, among others), or by a mechanism of constraint induction (as in Hayes & Wilson 2008, Flack 2007a) that is purely phonotactic and therefore has no access to lexical trends.

A version of Optimality Theory is proposed where the behavior of individual lexical items is recorded in terms of lexically-specific constraint rankings (cf. Pater 2000, 2005, 2006, 2008b; Anttila 2002; Inkelas et al. 1997; Itô & Mester 1995). A noun with a non-alternating final stop, like *anɑtʃ* ~ *anɑtʃ-i*, is associated with the ranking IDENT(voice) \gg $*V\text{tʃ}V$, meaning that faithfulness to voicing outweighs the markedness pressure against intervocalic voiceless palatal stops. A noun with a final alternating stop, like *amaʃ* ~ *amaʃɰ-i*, is associated with the opposite ranking, i.e. $*V\text{tʃ}V \gg$ IDENT(voice). This assumes that the final stop in *amaʃ* is underlyingly voiceless, and that it surfaces unfaithfully in *amaʃɰ-i*, contrary to the traditional generative analysis of Turkish (Lees 1961; Inkelas & Orgun 1995; Inkelas et al. 1997), and in line with the suggestions in Hayes (1995b, 1999). This aspect of the analysis is discussed and motivated in §2.6.

Given this approach, the behavior of mono-syllabic nouns, like *atʃ* ~ *aʃɰ-i*, can be recorded separately from the behavior of poly-syllabic nouns, by using a faithfulness constraint that protects the voicing feature of stops in the base's initial syllable,

IDENT(voice)_{σ1}. The existence of constraints in CON that are specific to initial syllables allows Turkish speakers to learn separate lexical trends for monosyllabic and polysyllabic nouns. On the other hand, in the absence of universal constraints that relate voicing and vowel backness, the backness of the stem-final vowel cannot be used in recording the behavior of any lexical items, and this aspect of the lexicon goes ignored by speakers.

To encode lexically-specific constraint rankings, the version of Optimality Theory used here is one augmented by a mechanism of constraint cloning (proposed in Pater 2006, 2008b, see also Mahanta 2007; Coetzee 2008). In this theory, language learners detect that their language requires opposite rankings of a pair of constraints, and then clone one of those constraints. In the Turkish case, speakers realize that some lexical items require IDENT(voice) \gg *VʃV and some lexical items require the opposite ranking. They clone one of the constraints, say IDENT(voice), and then non-alternating nouns are associated with the clone of IDENT(voice) that ranks over *VʃV, and alternating nouns are associated with the clone that ranks under *VʃV.

The resulting grammar contains two lists of nouns, as every *ʃ*-final noun of Turkish is listed under one of the clones of IDENT(voice). Since most *ʃ*-final nouns do alternate, most nouns will be listed with the clone that ranks below *VʃV. Now suppose a speaker encounters a novel noun in its bare form, and they are required to produce the possessive form. The grammar allows the final stop to either alternate or not alternate, but the alternating behavior is more likely, since more nouns are listed with the clone of IDENT(voice) that ranks below *VʃV. Cloned constraints allow speakers to reach a grammar that records the behavior of known items, and then project that behavior probabilistically onto novel items.

The full analysis of Turkish will involve the faithfulness constraints IDENT(voice) and IDENT(voice)_{σ1}, to protect final stops from becoming voiced, and additionally MAX and MAX_{σ1}, to protect final dorsals from deleting (see §2.4.6). These faithfulness constraints conflict with a family of markedness constraints against voiceless stops, either between two

vowels (*VpV, *VtV, *VtʃV, *VkV) or between a sonorant consonant and a vowel (*RpV, *RtV, *RtʃV, *RkV). Each stop-final noun of Turkish is listed under a pair of conflicting constraints, or equivalently, each pair of conflicting constraints accumulates a list of lexical items, and this listing allows the speaker to project the lexical statistics onto novel nouns.

This ability of speakers to project trends from their lexicon onto novel items is a well-established observation (see Zuraw 2000, Albright et al. 2001, Ernestus & Baayen 2003, Hayes & Londe 2006, among others). The theoretical contribution of this work is two-fold: (a) It relates the projection of language-specific lexical trends to cross-linguistic patterns of phonological interactions, by deriving both from the inventory of universal constraints in CON, and (b) it offers an OT-based grammar that applies deterministically to known items, and projects lexical trends directly from those items onto novel nouns.

2.2 Turkish lexicon study

The distribution of voicing alternations in the lexicon of Turkish depends heavily on the phonological shape of nouns. For instance, while the final stop in most mono-syllabic nouns does not alternate (18a), the final stop in most poly-syllabic words does alternate with its voiced counterpart (18b). This section offers a detailed quantitative survey of the Turkish lexicon, based on information from the Turkish Electronic Living Lexicon (TELL, Inkelas et al. 2000).

(18)	Bare stem	Possessive	
a.	atʃ	atʃ-i	‘hunger’
b.	amaʃ	amaɟ-i	‘target’

Several phonological properties of Turkish nouns will be discussed, showing that four of them correlate with stem-final alternations: (a) the noun’s size (mono-syllabic vs. poly-syllabic), (b) the place of articulation of the stem-final stop, (c) the height of the vowel that precedes the stem-final stop, and (d) the backness of that vowel.

Of the 3002 nouns in TELL whose bare stem ends in a voiceless stop, almost 90% are poly-syllabic, and in most of those, the final stop alternates¹ (19). The rate of alternation is much lower for monosyllables, especially in those with a simplex coda.

(19)	Size	<i>n</i>	% alternating
	Monosyllabic, simplex coda (CVC)	137	11.7%
	Monosyllabic, complex coda (CVCC)	164	25.9%
	Polysyllabic (CVCVC and bigger)	2701	58.9%

The distribution of alternating stops also varies by the place of articulation of the word-final stop (20). Most word-final labials, palatals and dorsals² do alternate, but only a small proportion of the final coronals do.

(20)	Place	<i>n</i>	% alternating
	Labial (p)	294	84.0%
	Coronal (t)	1255	17.1%
	Palatal (tʃ)	191	60.5%
	Dorsal (k)	1262	84.9%

While longer words correlate with a higher proportion of alternating nouns, size does not affect all places equally (21). In all places, CVC words alternate less than CVCVC words, but the behavior of CVCC words is not uniform. For labials and palatals, a majority of CVCC words alternate, patterning with the CVCVC words. For the dorsals, the CVCC words pattern together with the shorter CVC words, showing a modest proportion of

¹Some nouns in TELL are listed as both alternators and non-alternators. In calculating the percentage of alternating nouns, such nouns were counted as half alternators (although in reality it's entirely possible that the actual rate of alternation is different from 50%). Therefore, the proportion of alternating nouns is calculated by adding the number of alternating nouns and half the number of vacillating nouns, and dividing the sum by the total number of nouns.

²Dorsals delete post-vocally, see §2.4.6 for discussion.

alternators. Finally, the coronals show a very minor place effect, with CVCC words actually having a slightly higher proportion of alternators than either longer or shorter words.

(21)

	CVC		CVCC		CVCVC	
Place	<i>n</i>	% alt	<i>n</i>	% alt	<i>n</i>	% alt
p	30	26.7%	16	75.0%	248	91.5%
t	41	6.1%	79	19.0%	1135	17.3%
tʃ	23	17.4%	18	58.3%	150	67.3%
k	43	3.5%	51	9.8%	1168	91.2%

In other words, it is not the case that size and place each have a constant effect. Their effect on the distribution of voicing alternations cannot be accurately described separately. Anticipating the discussion in §2.3.2, it will be seen that indeed speakers treat each place/size combination separately.

Further study of TELL reveals a correlation between the quality of the vowel that precedes the word-final stop and the proportion of alternating nouns: high vowels correlate with a higher proportion of alternating stops relative to non-high vowels, and so do back vowels relative to front vowels. This correlation is rather surprising, since cross-linguistically, vowel quality is not known to influence the voicing of a neighboring obstruent³.

A noun-final stop is about 30% more likely to alternate when following a high vowel than when following a non-high vowel (22).

³Vowel length does correlate with voicing, with long vowels correlating universally with voiced consonants and short vowels with voiceless consonants (Lisker & Abramson 1964; Ohala 1983; Volatis & Miller 1992). In some cases, such as that of Canadian Raising, the change in vowel length causes a concomitant change in vowel quality. See §2.4.2 below for discussion.

(22)	Height of stem-final vowel	<i>n</i>	% alternating
	–high	1690	41.7%
	+high	1312	71.9%

The correlation with height, however, is not equally distributed among the different size and place combinations. The table in (23) shows that in most size/place combinations, there are only modest differences (less than 10%) between the proportions of alternating nouns given the height of the preceding vowel. A larger correlation in the opposite direction (53%) is seen for the CVCC *tf*-final words, but this is limited to a mere 18 nouns, which explains its negligible impact on the overall size correlation. The correlation with height is concentrated at the longer *t*-final nouns, where several hundred nouns show 24% more alternating stops following a high vowel.

(23)	CVC		CVCC		CVCVC	
	–high	+high	–high	+high	–high	+high
p	19	11	13	3	132	116
	26%	27%	77%	67%	85%	99%
t	24	17	55	24	796	339
	10%	0%	15%	29%	10%	34%
tʃ	14	9	8	10	91	59
	18%	17%	88%	35%	66%	69%
k	31	12	33	18	474	694
	2%	8%	12%	6%	87%	94%

A fourth and final phonological property that significantly correlates with the distribution of voicing alternations is the backness of the stem-final vowel (24). When preceded by a back vowel, a stem-final stop is about 10% more likely to alternate compared to a stop preceded by a front vowel.

(24)	Backness of stem-final vowel	<i>n</i>	% alternating
	–back	1495	49.5%
	+back	1507	60.3%

Just like vowel height, the correlation with vowel backness is not uniformly distributed in the lexicon. As seen in (25), the correlation with backness is small (at most 13%) for labial-, coronal- and dorsal-final nouns. A robust correlation with backness is seen in *tf*-final words of all sizes. Averaged over the 191 *tf*-final nouns, the proportion of alternating nouns is 30% higher following a back vowel relative to a front vowel.

(25)	CVC		CVCC		CVCVC	
	–back	+back	–back	+back	–back	+back
p	12	18	4	12	113	135
	33%	22%	75%	75%	96%	87%
t	18	23	34	45	673	462
	8%	4%	26%	13%	16%	19%
tf	11	12	10	8	66	84
	14%	21%	40%	81%	50%	81%
k	19	24	25	26	510	658
	8%	0%	16%	4%	90%	92%

In contrast to the four properties that were examined until now (size, place, height and backness), a phonological property that has but a negligible correlation with the distribution of voicing alternations is the rounding of the stem's final vowel (26).

(26)	Rounding of stem-final vowel	<i>n</i>	% alternating
	–round	2524	54.6%
	+round	478	56.4%

A closer examination of vowel rounding is no more revealing, and the details are omitted here for lack of interest. Other phonological properties that were checked and found to be equally unrevealing are the voicing features of consonants earlier in the word, such as the closest consonant to the root-final stop, the closest onset consonant, and the closest obstruent.

To sum up the discussion so far, four phonological properties of Turkish nouns were seen to correlate with stem-final voicing alternations in Turkish:

- Size: mono-syllables alternate less than poly-syllables, and among the mono-syllables, roots with simplex codas alternate more than roots with complex codas.
- Place (of articulation): Stem-final coronals alternate the least, while labials and dorsals alternate the most.
- Vowel height: stem-final stops are more likely to alternate following a high vowel compared to a non-high vowel.
- Vowel backness: stem-final stops are more likely to alternate following a back vowel compared to a front vowel.

All of these properties allow deeper insight when considered in pairs: Size and place have a non-uniform interaction, with CVCC words behaving like CVC words when dorsal-final and like CVCVC words when labial- or palatal-final. Height and backness interact with place non-uniformly: the correlation with height is concentrated in the coronal-final nouns, while the correlation with backness is concentrated in the palatal-final nouns.

In statistical parlance, the aforementioned properties can be understood as predictors in a regression analysis. Since TELL makes a three-way distinction in stop-final nouns (nouns that don't alternate, nouns that do, and "vacillators", i.e. nouns that allow either alternation or non-alternation), an ordinal logistic regression model was fitted to the lexicon using the *lrm()* function in R (R Development Core Team 2007). The dependent variable was a three-

level ordered factor, with non-alternation as the lowest level, alternation as the highest level, and optional alternation as the intermediate level.

Five independent variables were considered:

- **Size:** a three-level unordered factor, with levels corresponding to mono-syllables with a simplex coda (CVC), mono-syllables with a complex codas (CVCC), and poly-syllables (CVCVC). CVC was chosen as the base level.
- **Place:** a four-level unordered factor, with levels corresponding to coronal, palatal, labial and dorsal. Dorsal was chosen as the base level.
- **High, back and round:** each of the three features of the stem-final vowel was encoded as two-level unordered factor. The base levels chosen were non-high, front and unrounded.

First, each of these five predictors was tried in its own model, to assess each predictor's overall power in the lexicon (27). This power is measured by R^2 and by the model's likelihood ratio (Model L.R.), which comes with a number of degrees of freedom and a p-value. It turns out that *place*, *high*, *size*, and *back* are highly predictive of alternations, in that order, and *round* isn't⁴.

(27)	R^2	Model L.R.	df	p
place	.482	1469	3	<.001
high	.113	284	1	<.001
size	.078	193	2	<.001
back	.015	37	1	<.001
round	0	0	1	.489

⁴Another method for assessing the predictive power of each feature separately is a TiMBL simulation (Daelemans et al. 2002). Given the data in TELL, this system creates a number called "information gain" for every predictor that it is given. The system confirmed the verdict in (27), assigning the five predictors the following information gain values, respectively: .367, .071, .047, .009 and .0004.

While *high* has a larger R^2 than *size*, the interaction of *high* and *place* is less powerful than the interaction of *size* and *place*. The interaction of *place* with each of *size*, *high*, and *back* were tested in separate models, summarized in (28).

(28)	R^2	Model L.R.	df	p
place* <i>size</i>	.588	1920	11	<.001
place* <i>high</i>	.519	1621	7	<.001
place* <i>back</i>	.488	1496	7	<.001

When a base model that has *place*size* as a predictor is augmented with *place*high*, R^2 goes up to .616. Augmenting the base model with *place*back* only brings R^2 up to .594. Finally, model with all three of the interactions in (28) as predictors reaches an R^2 of .622, with a model L.R. of 2078 for 19 degrees of freedom. This final model is given in (29)⁵.

The model in (29) hardly contains any surprises, as it confirms the validity of the observations made earlier in this section. It simply restates the numerical observations as differences in the propensity to alternate relative to the arbitrarily chosen baseline levels of the predictors, namely CVC size, dorsal place, non-high vowels and front vowels. The size effect is mostly limited to the difference between CVC and CVCVC, with none of the CVCC levels reaching significance relative to CVC. In the CVCVC size, the coronal and palatal places alternate significantly less than the baseline dorsal, and labial place only approaches significance at this size. The vowel features reach significance for the interaction of high and coronal, and for the interaction of back and palatal.

⁵The model in (29) was validated with the fast backwards step-down method of the *validate()* function, and the predictor *back* was the only one deleted. Since the interaction of *back* with *place* was retained, I did not remove *back* from the model, so as not to leave an interaction in the model without its components. In 200 bootstrap runs, seven factors were considered: the three interaction factors, and the four basic factors they were made of. At least 5 of the 7 factors were retained in 197 of the runs, and in the vast majority of the runs, the three interaction factors were among the ones retained. The R^2 of the model was adjusted slightly from .6213 to .6117. An additional step of model criticism was taken with the *pentrace()* function, which penalizes large coefficients. With a penalty of .3, The penalized model was left essentially unchanged from the original model in (29), with slight improvements of the p-values of the vowel-place interactions at the fourth decimal place.

(29)	Coefficient	SE	Wald z	p
(y>=vacillator)	-3.502	0.745	-4.70	>0.001
(y>=alternating)	-3.822	0.746	-5.13	>0.001
COR	-0.102	0.976	-0.10	0.917
LAB	2.201	0.954	2.31	0.021
PAL	1.249	0.950	1.31	0.189
CVCC	0.783	0.869	0.90	0.367
CVCVC	5.488	0.735	7.47	0.000
high	0.874	0.205	4.27	0.000
back	0.288	0.204	1.41	0.158
CVCC * COR	0.703	1.102	0.64	0.523
CVCC * LAB	2.022	1.157	1.75	0.081
CVCC * PAL	1.269	1.129	1.12	0.261
CVCVC * COR	-4.011	0.959	-4.18	>0.001
CVCVC * LAB	-1.737	0.901	-1.93	0.054
CVCVC * PAL	-3.110	0.919	-3.38	0.001
COR * high	0.620	0.254	2.45	0.014
LAB * high	0.533	0.539	0.99	0.323
PAL * high	-0.754	0.387	-1.95	0.051
COR * back	0.077	0.254	0.30	0.762
LAB * back	-0.755	0.490	-1.54	0.123
PAL * back	1.136	0.386	2.95	0.003

To summarize the study of the Turkish lexicon, it was found that both size and place are excellent predictors of the alternation status of nouns. Larger nouns are more likely to alternate, and coronal-final nouns are less likely to alternate. In addition, the height and

backness of final stem vowels are also good predictors in combination with place: High vowels promote the alternation of coronals, and back vowels promote the alternation of palatals. All of these generalizations were confirmed to be highly statistically significant in a logistic regression model. In other words, the size of nouns, the place of their final stop, and the height and backness of their final vowels all strongly correlate with voicing alternations in a way that is statistically unlikely to be accidental.

2.3 Speakers' knowledge of the lexicon

In the previous section, the distribution of voicing alternations in the Turkish lexicon was examined and shown to be rather skewed. The distribution of alternating and non-alternating noun-final stops is not uniform relative to other phonological properties that nouns have: Size, place, height and backness were identified as statistically powerful predictors of alternation.

What the humans who are native speakers of Turkish know about the distribution of voicing alternations, however, is a separate question, which is taken on in this section. It will turn out that native speakers identify generalizations about the distribution of voicing alternations relative to the size of nouns and the place of articulation of their final stops. However, speakers ignore, or fail to reproduce, correlations between the voicing of final stops and the quality of the vowels that precede them.

A novel word task (Berko 1958) was used to find out which statistical generalizations native speakers extract from their lexicon. This kind of task has been shown to elicit responses that, when averaged over several speakers, replicate distributional facts about the lexicon (e.g. Zuraw 2000 and many others).

2.3.1 Materials and method

2.3.1.1 Speakers

Participants were adult native speakers of Turkish ($n = 24$; 13 males, 11 females, age range: 18-45) living in the United States. Some of the speakers were paid \$5 for their time, and others volunteered their time. The experiment was delivered as a web questionnaire, with some speakers doing the experiment remotely. For those speakers, reaction times were indicative of the speakers taking the questionnaire in one sitting, with no discernible distractions or pauses.

2.3.1.2 Materials

A male speaker of Turkish, a graduate student from the economics department, recorded the bare form and two possible possessive forms for each noun, repeated three times. Each stimulus was normalized for peak intensity and pitch and inspected by a native speaker to be natural and acceptable. One of the possessive forms was completely faithful to the base, with the addition of a final high vowel that harmonized with the stem, following the regular vowel harmony principles of the language. In the other possessive form, the stem final stop was substituted with its voiced counterpart, except for post-vocalic *k*'s, which were deleted.

Creating stimuli that exemplify all size, place and vowel quality combinations would have come up to 96 (four places * three sizes * eight vowel qualities). Since the lexical distribution of voicing alternations among palatals and labials is fairly similar, and in the interest of reducing the number of stimuli, the palatal and labial categories were collapsed into one category, using 12 words of each place, compared to 24 for the coronal- and dorsal-final words. The total number of stimuli, then, was 72 (three place categories * three sizes * eight vowel qualities).

Additionally, native Turkish nouns disallow the round nonhigh vowels *o*, *ö* in non-initial position. To make the stimuli more Turkish sounding, non-high round vowels in the second

syllable of the CVCVC words were replaced with the corresponding high vowels *u*, *ü*. The nouns that were used are presented in (30).

The non-final consonants were chosen such that the resulting nouns all sound plausibly native, with neighborhood densities equalized among the stimuli as much as possible.

(30)

		CVC		CVCC		CVCVC	
		–high	+high	–high	+high	–high	+high
p/tʃ	–back	gep	yiʃ	telp	gintʃ	hevetʃ	çisip
	–round						
	+back	dap	niʃ	pantʃ	dirp	yyap	ma.ıʃ
	+round						
t	–back	köʃ	züp	yöntʃ	kürp		bölütʃ türütʃ
	–round						
	+back	poʃ	tup	solp	munʃ		konup guyup
	+round						
k	–back	pet	hit	zelt	ʃint	niket	gevit
	–round						
	+back	fat	mit	hant	ʃirt	ya.at	pisit
	+round						
k	–back	söt	çüt	gönt	nürt		sölüt bünüt
	–round						
	+back	yot	nut	çolt	bunt		ʃforut muyut
	+round						
k	–back	vek	zik	helk	tink	mesek	perik
	–round						
	+back	çak	pik	vank	nirk	tatak	banik
	+round						
k	–back	hök	sük	sönk	pürk		nönük düyük
	–round						
	+back	mok	nuk	bolc	dunk		zoruk yuluk
	+round						

Finally, 36 fillers were included. All the fillers ended in either fricatives or sonorant consonants. To give speakers a meaningful task to perform with the fillers, two lexically-specific processes of Turkish were chosen: vowel-length alternations (e.g. *ruh* ~ *ru:h-u* ‘spirit’) and vowel-∅ alternations (e.g. *burun* ~ *burn-u* ‘nose’). Eighteen fillers displayed vowel-length alternations with a CVC base, and the other eighteen displayed vowel-∅ alternations with a CVCVC base. All of the fillers were chosen from a dictionary of

Turkish, some of them being very familiar words, and some being obsolete words that were not familiar to the speakers we consulted.

The materials were recorded in a sound attenuated booth into a Macintosh computer at a 44.1 KHz sampling rate. Using Praat (Boersma & Weenink 2008), the token judged best of each suffixed form was spliced and normalized for peak intensity and pitch. Peak intensity was normalized using Praat's "scale peak" function set to 0.6. For pitch normalization, three points were manually labeled in each affixed form: the onset of the word, the onset of the root's final segment (the onset of the burst in the case of stops), and the offset of the word. Then, a reversed V-shaped pitch contour was superimposed on the materials, with a pitch of 110 Hz at the onset of the word, 170 Hz at the onset of the root-final segment, and 70 Hz at the offset of the word. These values were chosen in order to best fit most of the speaker's actual productions, such that changes would be minimal.

Finally, for each stimulus, two .wav files were created by concatenating the two suffixed forms with a 0.8-second silence between the two, once with the voiceless form followed by the voiced form, and once with the voiced followed by the voiceless. A linguist who is a native speaker of Turkish verified that the final materials were of satisfactory quality. While she had some concerns about stress being perceived non-finally in a few of the filler items, no problems were found with the stimuli.

2.3.1.3 Procedure

Before the beginning of the experiment, speakers were reminded that voicing alternations are lexically-specific by presenting a familiar non-alternating paradigm (*top* ~ *top-u* 'ball') next to a familiar alternating paradigm (*çep* ~ *çeb-i* 'pocket'). Then, speakers were asked to choose the possessive form of two familiar alternating nouns (*dolap* 'cupboard' and *aatf* 'tree'), and feedback was given on their choices.

The stimuli were presented in a self-paced forced-choice task. The base form, e.g. *fet* was presented in Turkish orthography, which reflects the relevant aspects of the phonology

faithfully. The participants saw an overt possessor with genitive case followed by a blank, to provide the syntactic context for a possessive suffix, e.g. *Ali'nin* _____ “Ali’s _____”, and they heard two possible possessed forms, e.g. *fet-i* and *fed-i*. Speakers pressed “F” or “J” to choose the first or the second possessive form they heard. Most speakers took 15-20 minutes to complete the experiment.

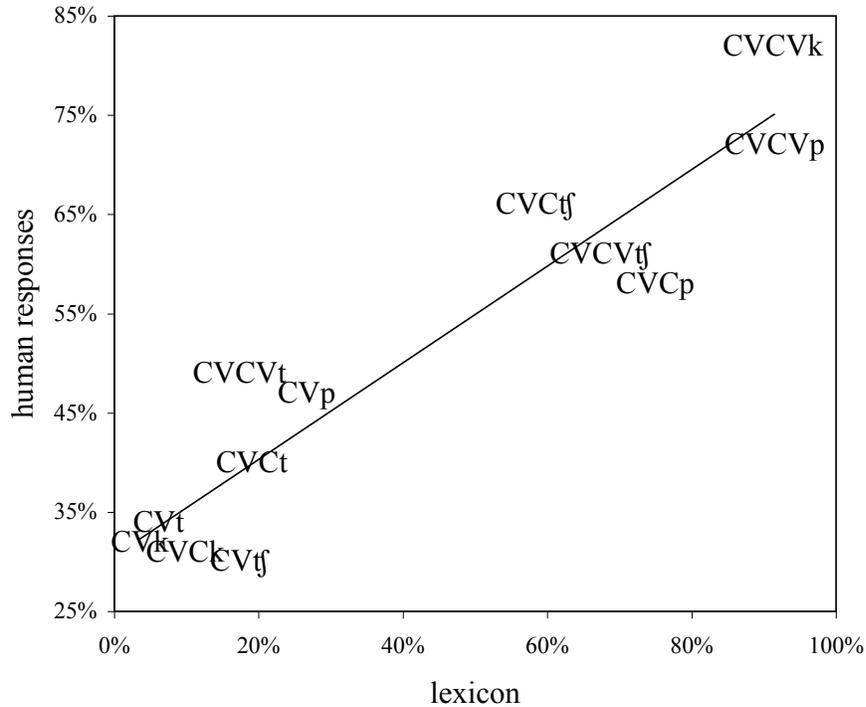
The order of the stimuli and the order of the choices were randomized. Additionally, the fillers were randomly distributed among the first three quarters of the stimuli.

2.3.2 Results

The experimental results are plotted in (31), grouped by size and place, plotted against the percent of alternating words in the lexicon with the matching size and place. The correlation is excellent (Spearman’s rank correlation test, $S = 46$, $\rho = .839$, $p < .005$), showing that speakers have accurately matched the percentages of alternating words in the lexicon. On average, the proportion of alternating responses ranges from 30% to 82%, as opposed to a wider range of 6% to 92% in the lexicon. Nevertheless, this compressed range of responses⁶ correlates with the lexicon very well.

⁶The source of the compression of the human results comes both from between-speaker and within-speaker sources. Some participants showed a strong preference for alternating responses, and some showed the opposite preference, resulting in at least 3 and at most 22 alternating responses per item, thus covering only 79% of the range of 0 to 24 alternating responses possible with 24 participants. Additionally, individual participants varied as to how strong the size and place effects were in their responses, with weak-effect participants causing further compression. The strength of these effects did not correlate with participants’ overall preference for alternation or non-alternation.

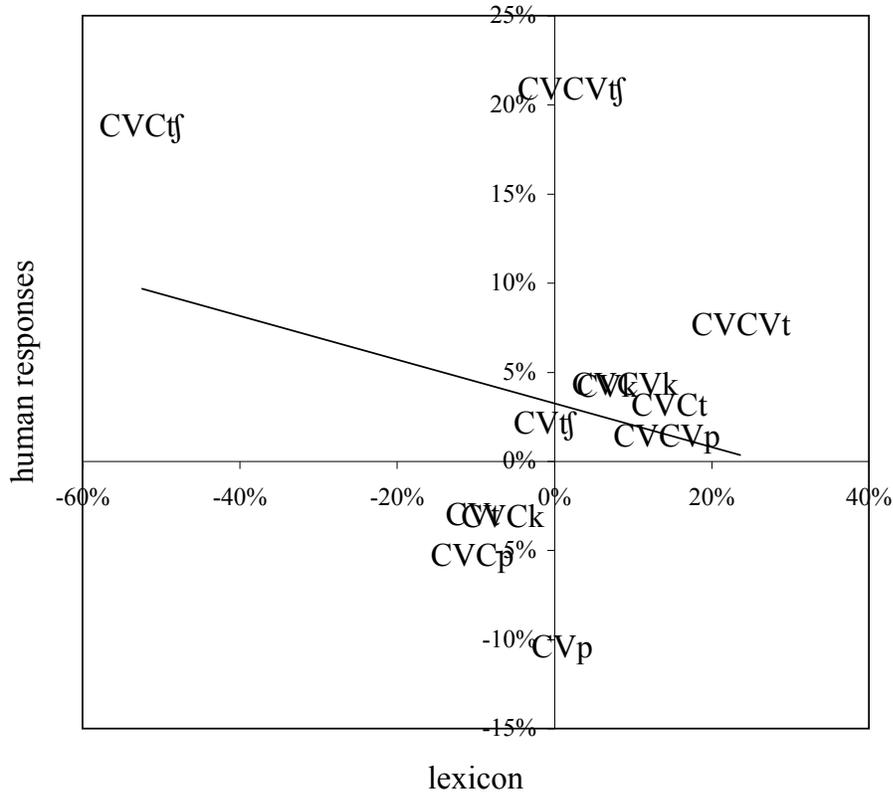
(31) Proportions of nouns with voicing alternations in the lexicon vs. the percent of alternating choices in the experiment, by size and place.



In stark contrast to the tight correlation between the experimental results and the lexicon for place and size effects, as seen in (31), there is no pattern when the height or backness effects are considered. The chart in (32) shows the results of the height factor. Each point in this chart shows the difference in rates of alternation between high and non-high vowels, by size and place. Positive values indicate more alternations with [+high] vowels, and negative values indicate more alternations with [−high] vowels.

There is no correlation between the lexicon and speakers' performance when vowel height is considered (Spearman's rank correlation test, $S = 196.8$, $\rho = .312$, $p > .1$). The chart in (32) shows that speakers' behavior was essentially random with respect to vowel height.

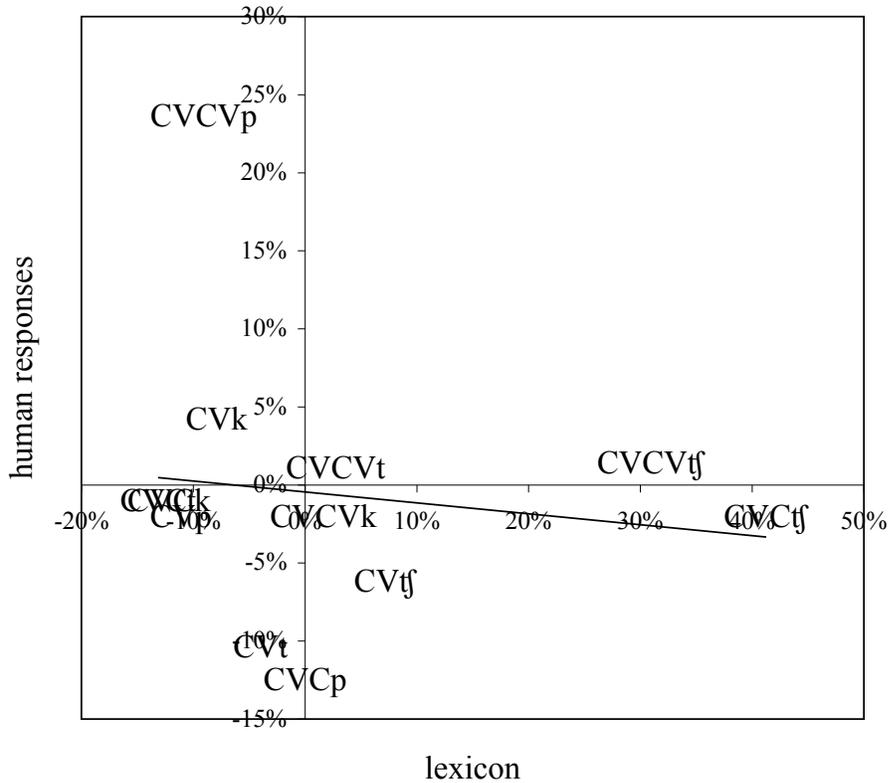
(32) Differences between high and non-high stem-final vowels in the lexicon vs. the differences between high and non-high vowels in the experiment, by size and place.



The lack of correlation in (32) is probably only due to a subset of the points, most noticeably CVCtʃ, CVCVtʃ, and CVp. There is no sense, however, in which these are “outliers”, as they represent a sizable proportion of the data. The data for the CVCtʃ point comes from 18 lexical items and from 96 experimental responses (4 items * 24 participants). The regression analysis below confirms the lack of correlation.

When vowel backness is considered (33), the result is essentially the same: There is no correlation between the lexicon and speakers’ responses when the results are categorized by size, place and backness (Spearman’s rank correlation test, $S = 326.1$, $\rho = -.140$, $p > .1$). Each point in (33) shows the difference in rates of alternation between back and front vowels, by size and place. Positive values indicate more alternations with back vowels, and negative values indicate more alternations with front vowels.

- (33) Differences between back and front stem-final vowels in the lexicon vs. the differences between back and front vowels in the experiment, by size and place.



The contrast between the strong correlation in (31) and the lack of correlation in (32-33) shows that speakers' behavior is best understood as replicating the lexicon's size and place effects, but not replicating its height or backness effects. This contrast is seen in the statistical analysis below.

The results were analyzed with a mixed-effects logistic regression in R (R Development Core Team 2007) using the *lmer()* function of the LME4 package, with *participant* and *item* as random effect variables. The fixed effect variables were the same ones used in the analysis of the lexicon: *size*, *place*, *high*, *back* and *round*.

An initial model was fitted to the data using only *size* and *place* as predictors. Adding their interaction to the model made a significant improvement (sequential ANOVA model comparison, $\chi^2(6) = 50.58, p < .001$). The improved model with the interaction term is given in (34). This model shows that labial place and CVCVC size are more conducive

to alternating responses than the baseline dorsal place and CVC size, respectively. As for interactions, for the CVCC size, palatal place is more conducive to voicing than the baseline dorsal place with the same CVCC size. Additionally, in the CVCVC size, all places are less conducive to alternating responses than the baseline dorsal place with the same CVCVC size. All of these effects mirror the lexical effects as presented in §2.2. The model stays essentially unchanged when validated by the *pvals.fnc()* function (Baayen 2008).

(34)	Estimate	SE	<i>z</i>	<i>p</i>
(Intercept)	-0.864	0.283	-3.056	0.002
COR	0.111	0.256	0.434	0.665
LAB	0.744	0.304	2.451	0.014
PAL	-0.119	0.320	-0.372	0.710
CVCC	-0.089	0.260	-0.341	0.733
CVCVC	2.694	0.285	9.469	< 0.001
CVCC:COR	0.385	0.361	1.065	0.287
CVCC:LAB	0.641	0.431	1.487	0.137
CVCC:PAL	1.867	0.447	4.173	< 0.001
CVCVC:COR	-1.936	0.377	-5.142	< 0.001
CVCVC:LAB	-1.436	0.455	-3.154	0.002
CVCVC:PAL	-1.126	0.457	-2.463	0.014

The addition of any vowel feature to the baseline model (*high*, *back* or *round*) made no significant improvement ($p > .1$). No vowel feature approached significance, either on its own or by its interaction with *place*. For example, adding the interaction *place*high* to the model in (34) gives a new model where the interaction of coronal place and *high* is almost exactly at chance level ($p = .981$). Adding *place*back* to the baseline model gives an interaction of palatal place and *back* that is non-significant ($p = .661$) and its coefficient is negative, i.e. going in the opposite direction from the lexicon.

In other words, *size* and *place* had statistically significant power in predicting the choice of alternation vs. non-alternation of stem-final stops. Crucially, however, none of the vowel features had a significant effect on the participants' choices.

To summarize the findings, Turkish speakers reproduced the distribution of voicing alternations in the lexicon by paying attention to the size of the nouns and the place of the final stops, while ignoring the quality of the vowel that precedes the stem-final stop.

2.3.3 Discussion

The experimental results show that Turkish speakers generalize their knowledge of the voicing alternations in their lexicon. Not contenting themselves with memorizing the alternating or non-alternating status of single nouns, speakers have access to the relative proportion of alternating nouns categorized by size and place. Using size and place as factors, speakers must somehow project their lexical statistics onto novel items. Although the height and backness of stem-final vowels are strongly correlated with alternations in the lexicon, speakers' treatment of stem-final vowels in novel words is random, showing no significant interaction with their choice of alternating or non-alternating forms.

Speakers failed to reproduce the correlation between vowels and voicing alternations in spite of an abundance of overt evidence, while learning the size and place effects even with very little evidence. For instance, the difference in alternation rates between *tf*-final CVC and CVCC nouns was successfully reproduced in the experiment results, even though the evidence comes from 23 and 18 actual nouns, respectively. The evidence for the vowel effects, however, comes from hundreds of nouns.

The proposal advanced here is that the results are best understood in light of a theory of universally possible phonological interactions, as encoded in a set of universal constraints. Only factors that can be expressed in terms of constraint interaction can be identified by language learners, with other lexical generalizations going unnoticed. This model is

contrasted with general-purpose statistical learners that can learn any robust distributional generalization, as discussed in §2.5.

2.4 Analysis with cloned constraints

Turkish speakers evidence a detailed knowledge of trends in their lexicon that regulate the choice of alternation or non-alternation of stem-final stops. Furthermore, speakers are biased by Universal Grammar to learn only lexical trends that can be captured in terms of cross-linguistically observed interactions between phonological elements. This section shows how an OT-based model can be used to learn the trends that humans learn. The model reads in the lexicon of Turkish and projects a probabilistic grammar from it, a grammar that can in turn be used to derive novel words in a way that correlates with the experimental results shown in §2.3.

Given a stop-final novel noun and asked to choose a possessive form for it, Turkish speakers consult a subset of their lexicon: For instance, given the noun *dap*, speakers identify it as a mono-syllabic *p*-final simplex-coda noun, and they compare it to the other mono-syllabic *p*-final simplex-coda nouns in their lexicon. If they have 30 such nouns, of which 8 alternate and 22 don't alternate, as in TELL, then the likelihood that *dap* will exhibit a voicing alternation is 8 out of 30, or 27%.

In other words, Turkish speakers partition their lexicon based on phonological principles. The mass of stop-final nouns is partitioned by the size of each noun (mono- vs. polysyllabic), by the place of articulation of the final stop (p, t, tʃ, k), and by the complexity of the final coda, and within each such group, alternating nouns are separated from non-alternating nouns. This creates a total of $2 * 4 * 2 * 2 = 32$ partitions. Nouns that don't end in a stop are all lumped together in the “elsewhere” partition.

Constraint cloning is a mechanism for partitioning the lexicon and listing the words that belong in each partition. The partitions are defined by the set of universal constraints

in CON, which ensures that nouns are only categorized based on universal grammatical principles.

2.4.1 Constraint cloning

The OT-based model proposed here makes crucial use of the concept of Inconsistency Resolution, offered by Pater (2006, 2008b), which relies on the Recursive Constraint Demotion Algorithm (RCD, Prince & Tesar 1999).

In RCD, the speaker learns from “errors”, or mismatches between the words of the language they are exposed to and the words that are produced by their current grammar. Suppose the learner hears the adult form [kanat] ‘wing’, but their grammar produces [kana], because the markedness constraint *CODA out-ranks faithfulness in their grammar (35).

(35)

[kanat]	*CODA	MAX
a. ☹ kanat	*!	
b. ☺ kana		*

Since the current winner, [kana], is different from the adult form, the speaker constructs a winner-loser pair, as in (36). The tableau in (36) is a comparative tableau (Prince 2002), where W means “winner-preferring” (i.e. the constraint assigns less violations to the winner) and L means “loser-preferring (i.e. the constraint assigns less violations to the loser).

(36)

	*CODA	MAX
a. kanat \succ kana	L	W

RCD takes winner-loser pairs such as the one in (36) and extracts a grammar from them by identifying columns that don't have L's in them and "installing" them. In this simple case, MAX can be installed, meaning that it is added to the grammar below any other previously installed constraints (which would be at the top of the grammar in this case, since no constraints were previously installed), and winner-loser pairs that MAX assigns a W to are removed from the tableau. Once MAX is thus installed, the tableau is emptied out, and the remaining constraints, in this case just *CODA, are added at the bottom of the grammar. The resulting grammar is now MAX \gg *CODA, which allows codas to be produced, as in adult Turkish.

There is no guarantee, however, that RCD will always be able to install any constraints and remove all of the winner-loser pairs from the tableau. If all of the available columns have L's in them, RCD will stall. This situation arises when the language provides the learner with conflicting data, as in (37). In some words, a stem-final stop is voiceless throughout the paradigm (37a-b), and in others, a final stop shows up voiceless in the bare stem and voiced in the possessive (37c-d).

(37)	bare stem	possessive	
a.	atʃ	atʃ-i	'hunger'
b.	anaʃ	anaʃ-i	'female cub'
c.	taʃ	taɖ-i	'crown'
d.	amaʃ	amaɖ-i	'target'

Assuming the bare stem with its voiceless stop as the underlying form,⁷ the non-alternating forms rank faithfulness to the underlying representations above the markedness

⁷Assuming the bare stem as the underlying representation goes against the tradition in generative linguistic theory, which assumes that alternating stops and non-alternating stops have different specifications for voice underlyingly (Inkelas & Orgun 1995; Inkelas et al. 1997, yet cf. Hayes 1995b). The empirical shortcomings of the traditional approach are addressed in §2.6.

pressure against intervocalic voiceless stops (38), while alternating forms require ranking faithfulness below markedness (39).

(38)

/ anatʃ + i /	IDENT(voice)	*VʃV
a. ↗ anatʃ-i		*
b. anaɕʒ-i	*!	

(39)

/ amatʃ + i /	*VʃV	IDENT(voice)
a. ↗ amaɕʒ-i		*
b. amatʃ-i	*!	

With this understanding of the situation, the ranking between the faithfulness constraint IDENT(voice) and the markedness constraint *VʃV cannot be determined for the language as a whole. Pairing the winners in (38) and (39) with their respective losers allows the ranking arguments to be compared, as in (40).

(40)

	IDENT(voice)	*VʃV
a. anatʃ-i > anaɕʒ-i	W	L
b. amaɕʒ-i > amatʃ-i	L	W

Since the ranking arguments in (40) are inconsistent, there are no rows with no L's in them, and therefore no constraints can be installed, and a grammar cannot be found using

RCD. Pater (2006, 2008b) proposes a mechanism for resolving such inconsistencies by cloning. In cloning, the speaker replaces a universal constraint of general applicability with two copies, or clones, of the universal constraint that are lexically-specific, with each clone listing the lexical items it applies to⁸.

Given the situation in (40), the speaker can clone IDENT(voice), making one clone specific to the root *anaɫf* (and any other lexical items that IDENT(voice) assigns a W to), and the other clone specific to the root *amaɫf* (and any other lexical items that IDENT(voice) assigns an L to). The resulting grammar is no longer inconsistent:

(41)

	IDENT (voice) _{anaɫf}	IDENT (voice) _{amaɫf}	*VɫfV
a. anaɫf-i > anaɫɟ-i	W		L
b. amaɫɟ-i > amaɫf-i		L	W

Now RCD can be successfully applied to (41): First, IDENT(voice)_{anaɫf} is installed, and the first winner-loser pair is removed. This leaves the column of *VɫfV with no L's in it, so *VɫfV is installed below IDENT(voice)_{anaɫf}, and the second winner-loser pair is removed. The remaining constraint, IDENT(voice)_{amaɫf} is added to the ranking below *VɫfV. The resulting grammar is IDENT(voice)_{anaɫf} ≫ *VɫfV ≫ IDENT(voice)_{amaɫf}, which correctly blocks the voicing alternation in *anaɫf-i* but allows it in *amaɫɟ-i*. In the case of (40), choosing to clone IDENT(voice) solved the inconsistency, but cloning *VɫfV would have been equally useful. The question of which constraint to clone is addressed systematically in §4.2.

⁸Pater (2006, 2008b) suggests a slightly different mechanism, where one clone is lexically specific and the other clone stays general. I argue in §2.4.2 below that both clones must be lexically specific to account for the behavior of Turkish speakers.

The cloning of IDENT(voice), and the listing of lexical items with its clones, divided the lexicon into three partitions: One partition contains the items listed with the high-ranking clone of IDENT(voice), another partition contains the items listed with the low-ranking clone of IDENT(voice), and a third partition contains all the lexical items that are not listed with either clone. These partitions are not arbitrary, but rather determined by the the mark that IDENT(voice) assigns to each winner-loser pair: W, L, or none.

Once a constraint is cloned, its clones accumulate lists of the stems they apply to. This approach allows for two sub-grammars to coexist in a language, while keeping track of the number of lexical items that belong to each sub-grammar. Since the number of lexical items of each kind becomes available in the grammar, the speaker can estimate the likelihood of each behavior.

The rest of this section shows how constraint cloning creates a grammar of Turkish that reflects speakers' knowledge of the lexicon, as determined by the experimental findings in §2.3.

2.4.2 The place effect

As discussed in §2.2, all stops are not equally likely to alternate: While the stops in most *tʃ*-final and *p*-final nouns alternate, the stops in most *t*-final nouns do not. The table in (42), repeated from (20) above, lists the numbers of alternating and non-alternating (faithful) paradigms by the place of articulation of the final stop, as found in TELL (Inkelas et al. 2000).

(42)	Place	Alternating	Faithful	Total	% alternating
	p	247	47	294	84%
	t	214	1041	1255	17%
	tʃ	117	74	191	61%
	k	1071	191	1262	85%

To replicate the effect that place has over the distribution of voicing alternations, the language learner must separately keep track of words that end in different stops. The fact that voicing affects stops of different places of articulation differently is well documented (e.g. Lisker & Abramson 1964; Ohala 1983; Volatis & Miller 1992). Additionally, the lenition of voiceless stops to voiced stops between vowels is also very well documented. Kirchner (1998) surveys numerous languages that lenite all of their voiceless stops between vowels, and several that lenite some of their voiceless stops, but his survey also has languages that lenite only labials (e.g. Gitksan, Hoard 1978), only coronals (e.g. Liverpool English, Wells 1982) or only dorsals (e.g. Apalai, Koehn & Koehn 1986). This typology can not only motivate a general constraint against intervocalic stops, but also a family of constraints that penalize voiceless stops between vowels: *VpV, *VtV, *VʃV, *VkV. The interaction of each of these constraints with IDENT(voice) will allow the speaker to discover the proportion of the stop-final nouns of Turkish that alternate in each place of articulation.

Note that for each place of articulation, the speaker has to keep track of both the number of words that alternate and the number of words that do not. Simply keeping a count of words that alternate leads to a wrong prediction: Compare, for instance, *t*-final words and *ʃ*-final words. There are 214 *t*-final words that alternate, but only 117 *ʃ*-final words that do. If the speaker were to only keep a count of alternating words, they would reach the conclusion that *t*-final words are more likely to alternate. But in fact, speakers choose alternating responses with *ʃ*-final words more often than they do with *t*-final words, reflecting the relative proportions of alternating and non-alternating nouns, not the absolute number of alternating nouns.

Similarly, keeping track of just the non-alternating nouns will also make the wrong prediction. Comparing *ʃ*-final words and *k*-final words, we see that there are more than twice as many *k*-final non-alternators than there *ʃ*-final non-alternators. Speakers, however, choose non-alternating responses with *k*-final words less often than they do with *ʃ*-final

words. In order to match the proportion of alternating stops in each place, both alternating and non-alternating words will need to be tracked.

Imagine a learner that has learned just two paradigms, *amatf* ~ *amaɕ-i* and *sepet* ~ *seped-i*. While one alternates and the other doesn't, no inconsistency is detected yet, since IDENT(voice) interacts with two different markedness constraints (43).

(43)

	IDENT(voice)	*VtV	*VɸV
a. <i>amaɕ-i</i> > <i>amatf-i</i>	L		W
b. <i>sepet-i</i> > <i>seped-i</i>	W	L	

Running RCD on (43) yields the clone-free grammar $*V\text{ɸ}V \gg \text{IDENT(voice)} \gg *VtV$. If the speaker learns the word *anatf* ~ *anatf-i*, however, the grammar becomes inconsistent (44).

(44)

	IDENT(voice)	*VtV	*VɸV
a. <i>amaɕ-i</i> > <i>amatf-i</i>	L		W
b. <i>anatf-i</i> > <i>anaɕ-i</i>	W		L
c. <i>sepet-i</i> > <i>seped-i</i>	W	L	

Since there are no columns in (44) that don't have L's in them, RCD stalls. Cloning either $*V\text{ɸ}V$ or IDENT(voice) can resolve the inconsistency. In this case, $*V\text{ɸ}V$ is chosen since its column has the least number of non-empty cells (choosing a constraint to clone based on the number of non-empty cells is discussed in §4.2). The result of cloning $*V\text{ɸ}V$ is shown below:

(45)

	ID(voice)	*VtV	*VtʃV _{amatʃ}	*VtʃV _{anatʃ}
a. amaɕʃ-i > amatʃ-i	L		W	
b. anatʃ-i > anaɕʃ-i	W			L
c. sepet-i > seped-i	W	L		

Installing *VtʃV_{amatʃ} removes the first winner-loser pair. This leaves IDENT(voice) with no L's in its column, so it is installed, and the last two winner-loser pairs are removed. Then, *VtV and *VtʃV_{anatʃ} are installed, yielding the ranking in (46).

(46) *VtʃV_{amatʃ} >> IDENT(voice) >> *VtV, *VtʃV_{anatʃ}

The resulting grammar has successfully partitioned the data available to the learner: Lexical items that end in *tʃ* are listed with the two clones of *VtʃV, and the *t*-final noun was not listed, since *t*-final nouns behave consistently in this limited set of data.

Cloning of *VtV will only become necessary once the speaker encounters a word with an alternating *t*, e.g. *kanat* ~ *kanad-i* 'wing', as in (47). Note that whenever the speaker learns a new paradigm, information about constraint conflicts may change; therefore, constraint cloning always starts from square one with the addition of a new winner-loser pair.

(47)

	ID(voice)	*VtV	*VʃV
a. amaɟ-i > amatʃ-i	L		W
b. anatʃ-i > anaɟ-i	W		L
c. kanad-i > kanat-i	L	W	
d. sepet-i > seped-i	W	L	

Given (47), cloning *VʃV will not suffice to make the grammar consistent. If *VʃV is cloned first, the learner will install *VʃV_{amatʃ} and remove the first winner-loser pair, but then they will still have a tableau with no columns that have no L's in them. Cloning *VtV as well will solve the inconsistency, and the resulting grammar would be as in (48).

(48) *VʃV_{amatʃ}, *VtV_{kanat} >> IDENT(voice) >> *VtV_{sepet}, *VʃV_{anatʃ}

The resulting grammar in (48) successfully partitions the lexicon: *t*-final nouns are listed with clones of *VtV, and *ʃ*-final nouns are listed with clones of *VʃV. These partitions are defined by the constraints that distinguish winners from losers. The language learner's ability to treat each place separately is a consequence of the availability of universal constraints that relate voicing and place of articulation. These constraints let the speaker detect inconsistency in each place separately, and create lists of lexical items in each place.

2.4.3 The size effect

Both the lexicon (§2.2) and the experimental results (§2.3) show a higher preference for alternations in poly-syllabic nouns relative to mono-syllabic, in every place of articulation. The size effect is not equal across the different places, however. Mono-syllabic nouns

generally don't alternate, regardless of the place of articulation of their final stop. Poly-syllabic nouns usually do alternate if they are *p*-final or *tf*-final, but not if they are *t*-final. Speakers have replicated this pattern of differential treatment of poly-syllabic nouns. In statistical terms, the size and place effects have a significant interaction, and the implication for the learner is that the proportion of alternating nouns is learned separately in each place-size combination.

The proposed account of this size effect relies on the position of the alternating final stop relative to the initial syllable of the root. In a mono-syllabic noun, the unfaithful mapping from a voiceless stop to a voiced one affects the initial syllable of the base, while a voicing alternation in a poly-syllabic noun doesn't affect the initial syllable. Initial syllables are known to enjoy greater faithfulness cross-linguistically, as formalized by Beckman (1997). The availability of a faithfulness constraint that protects only mono-syllabic roots allows the speaker to partition the lexicon along this dimension, putting mono-syllables in one partition, and leaving the other nouns, which are therefore poly-syllabic, in another partition.

The role of the word-initial syllable in the distribution of voice in Turkish is not limited to voicing alternations. Generally in the language, a coda stop followed by an onset stop will surface with the voicing feature of the onset stop (also known as regressive voicing assimilation, e.g. *is.tib.dat* 'despotism', **is.tip.dat*), but a coda stop in the initial syllable may surface with its independent voicing specification (e.g. *mak.bul* 'accepted', *eb.kem* 'mute').

For concreteness, this section focuses on learning the *tf*-final nouns of Turkish with simple codas. The relevant lexical counts are in (49).

(49)

	CVtʃ	CVCVtʃ	Total
Faithful	18	44	62
Alternating	3	96	99
Total	21	140	161

Given both mono-syllabic and poly-syllabic nouns that do and do not alternate, as in (50), the learner can successfully separate mono-syllabic roots from poly-syllabic ones by cloning the specific IDENT(voice)_{σ1} first.

(50)

	IDENT	IDENT _{σ1}	*VtʃV
a. satʃ-1 > saɕ-1	W	W	L
b. taɕ-1 > tatʃ-1	L	L	W
c. anatʃ-1 > anaɕ-1	W		L
d. amaɕ-1 > amatʃ-1	L		W

IDENT(voice)_{σ1} can be identified as more specific than IDENT(voice) by examining the number of W's and L's in each column, since the more specific constraint will necessarily assign a subset of the W's and L's that the general constraint assigns. The result of cloning IDENT(voice)_{σ1} is in (51). Since only mono-syllabic stems are assigned W's or L's by IDENT(voice)_{σ1}, only mono-syllables get listed by clones at this point.

(51)

	IDENT	IDENT _{σ1 saf}	IDENT _{σ1 taf}	*VtjV
a. satf-1 > sacɕ-1	W	W		L
b. taɕ-1 > tatf-1	L		L	W
c. anaf-1 > anaɕ-1	W			L
d. amaɕ-1 > amatf-1	L			W

The column of IDENT(voice)_{σ1 saf} has no L's in it, so it can be installed, and the first winner-loser pair can be removed from the tableau. While the mono-syllabic *tf*-final nouns were successfully listed by clones of IDENT(voice)_{σ1}, the learner is not quite ready to discover the rest of the *tf*-final nouns. Given the tableau in (51), there are no constraints to install after the installation of IDENT(voice)_{σ1 saf}, so either IDENT(voice) or *VtjV will need to be cloned. Once either of them is cloned, *tatf* and *amach* will be listed with one clone, and *anaf* will be listed with the other. Assuming it is IDENT(voice) that is cloned, the resulting grammar will be the one in (52).

(52) IDENT(voice)_{σ1 saf} ≫ IDENT(voice)_{anaf} ≫ *VtjV ≫
 IDENT(voice)_{σ1 taf}, IDENT(voice)_{tatf, amatf}

The problem with the grammar in (52) is that the lexicon is not neatly partitioned in the way the learner needs it to be: The specific IDENT(voice)_{σ1} correctly lists all and only the mono-syllables, but the general IDENT(voice), in addition to correctly listing all the poly-syllabic *tf*-final nouns, also incorrectly lists the mono-syllabic *tf*-final alternators.

The problem is that the general IDENT(voice) assigns W's and L's to all nouns, regardless of size, potentially allowing some nouns to “double dip”, as seen in (52). To ensure that nouns are not listed multiple times, the learner needs to make sure that

when they clone a specific constraint and list words with the clones, they also ignore any W's or L's that a more general constraint assigns to these listed words. In the case of (51), the learner needs to notice that IDENT(voice) is more general than IDENT(voice)_{σ1} (as determined by the fact that IDENT(voice) assigns a superset of the W's and L's that IDENT(voice)_{σ1} assigns), and ignore (or “mask”) the W's and L's that IDENT(voice) assigns to the nouns that are listed by IDENT(voice)_{σ1}⁹. The correct tableau, with the masking of the W that IDENT(voice) assigns to *satf-t* and the L that it assigns to *taç-t*, is in (53).

(53)

	IDENT	IDENT _{σ1 satf}	IDENT _{σ1 taf}	*VtjV
a. satf-1 > sacç-1	⊗ W	W		⊗ L
b. taçç-1 > taf-1	⊗ L		L	W
c. anaf-1 > anaçç-1	W			L
d. amaçç-1 > amatf-1	L			W

Given the tableau in (53), the column of IDENT(voice) has the fewest W's and L's, so IDENT(voice) will be chosen for cloning. The learner will clone IDENT(voice) and successfully list just the poly-syllables with it. The resulting grammar will be the one in (54). This grammar achieves the intended partitioning of the lexicon: The *tj*-final nouns are divided into mono-syllables and poly-syllables, and within each category, the nouns are further divided into alternators and non-alternators.

(54) IDENT(voice)_{σ1 satf} ≫ IDENT(voice)_{anaf} ≫ *VtjV ≫
 IDENT(voice)_{σ1 taf}, IDENT(voice)_{amatf}

⁹The masking operation can also be defined to operate only on L's, since the W's will be removed by the installation of a clone of the specific constraint, and masking of W's will turn out to be vacuous.

To summarize, the analysis of the size effect in Turkish relies on the availability of a specific version of IDENT(voice) that only assesses voicing alternations in mono-syllables. The speaker uses the specific IDENT(voice)_{σ1} to list the mono-syllables, leaving the poly-syllables to the care of the general IDENT(voice). The intended result relies on two principles: (a) the selection of the constraint to clone by identifying the column with the fewest non-empty cells, and (b) the masking of W's and L's from general constraints upon the listing of items with a specific constraint.

2.4.4 Combining place and size

The distribution of the voicing alternations in Turkish is analyzed here as affected by two factors: The place of articulation of the final stop, which was attributed to the markedness of different stops between vowels, and the size, which was attributed to specific faithfulness to voicing in mono-syllables. The two effects have a significant interaction, where the size effect is strong in labials and palatals and much smaller for coronals. This section will show how the learner can model this interaction by using pairs of constraints to list lexical items.

The tableau in (55) shows the full range of possible winner-loser pairs given two places (*t* and *tʃ*), two sizes (mono-syllabic and poly-syllabic) and two alternation patterns (faithful and alternating). The intended result is for the speaker to partition their lexicon by size and place, making four partitions, and within each of the four, further partition and list alternating and non-alternating items separately. Using the cloning technique that was offered in §2.4.2 and §2.4.3 above, no constraint will lead to the correct partitioning: For instance, cloning IDENT(voice)_{σ1} will separate the alternating mono-syllabic nouns from the non-alternating mono-syllabic nouns, so *satʃ* and *at* will be listed with one clone and *tatʃ* and *tat* will be listed with the other clone. But this listing collapses the place distinction, putting *tʃ*-final nouns and *t*-final nouns in the same partition.

(55)

	IDENT	IDENT _{σ1}	*VtjV	*VtV
a. satj-1 > saçj-1	W	W	L	
b. taçj-1 > tatj-1	L	L	W	
c. anatj-1 > anaçj-1	W		L	
d. amaçj-1 > amatj-1	L		W	
e. at-1 > ad-1	W	W		L
f. tad-1 > tat-1	L	L		W
g. sepet-i > seped-i	W			L
h. kanad-1 > kanat-1	L			W

The mechanism of cloning must be made sensitive to the various sources of conflict in the data: The column of IDENT(voice)_{σ1} indeed contains W's and L's, but these conflict with different constraints. Some W's that IDENT(voice)_{σ1} assigns are offset by L's from *VtV, and some are offset by L's from *VtjV. Similarly, the L's that IDENT(voice)_{σ1} assigns are offset by W's from *VtV and from *VtjV.

To capture the different sources of conflict in the data, lexical items that are listed with clones of IDENT(voice)_{σ1} must also mention which constraint they conflict with: If a lexical item gets a W from IDENT(voice)_{σ1}, this W must be offset by an L from some other constraint, and vice versa. The clones of IDENT(voice)_{σ1} don't simply list lexical items, but rather list lexical items by the constraint they conflict with, or more formally, clones list ⟨constraint, {lexical items}⟩ pairs. This is shown in (56). As before, the listing of items

with clones of the specific $\text{IDENT}(\text{voice})_{\sigma 1}$ causes the masking of W's and L's from the column of the more general IDENT.

(56)

	IDENT	$\text{IDENT}_{\sigma 1}$ $\langle *VfV, \text{saf} \rangle,$ $\langle *VtV, \text{at} \rangle$	$\text{IDENT}_{\sigma 1}$ $\langle *VfV, \text{taf} \rangle,$ $\langle *VtV, \text{tat} \rangle$	*VfV	*VtV
a. $\text{saf-1} \succ \text{saf-1}$	Ⓜ	W		L	
b. $\text{taf-1} \succ \text{taf-1}$	Ⓛ		L	W	
c. $\text{anaf-1} \succ \text{anaf-1}$	W			L	
d. $\text{amaf-1} \succ \text{amaf-1}$	L			W	
e. $\text{at-1} \succ \text{at-1}$	Ⓜ	W			L
f. $\text{tat-1} \succ \text{tat-1}$	Ⓛ		L		W
g. $\text{sepet-i} \succ \text{sepet-i}$	W				L
h. $\text{kanat-1} \succ \text{kanat-1}$	L				W

Next, the learner is ready to clone $\text{IDENT}(\text{voice})$, which will again list items by the constraints they conflict with. The resulting grammar is in (57).

$$(57) \text{IDENT}(\text{voice})_{\sigma 1} \langle *VfV, \text{saf} \rangle \langle *VtV, \text{at} \rangle \gg \text{IDENT}(\text{voice}) \langle *VfV, \text{anaf} \rangle \langle *VtV, \text{sepet} \rangle \gg *VfV, *VtV \gg$$

$$\text{IDENT}(\text{voice})_{\sigma 1} \langle *VfV, \text{taf} \rangle \langle *VtV, \text{tat} \rangle, \text{IDENT}(\text{voice}) \langle *VfV, \text{amaf} \rangle \langle *VtV, \text{kanat} \rangle$$

This grammar correctly partitions the lexicon: Clones of $\text{IDENT}(\text{voice})_{\sigma 1}$ list all the mono-syllabic stop-final nouns that the speaker has, and those are further divided by

markedness constraints into *t*-final and *ʈ*-final nouns. Of course, the full grammar also lists *p*-final nouns under *VpV, and those *k*-final nouns that show a voicing alternation are listed under *VkV (for more on *k*-final nouns, see §2.4.6). The nouns that were assessed neither W's nor L's by IDENT(voice)_{σ1}, which are therefore poly-syllabic, are listed by clones of the general IDENT(voice). These again are listed by the markedness constraint that IDENT(voice) conflicts with, correctly separating the poly-syllabic nouns according to the place of articulation of their final stop.

This grammar allows the speaker to learn the proportion of alternating nouns in each size and place combination, with these combinations made available by listing lexical items with pairs of constraints.

2.4.5 The complex coda effect

As discussed in §2.2 and §2.3, stop-final CVC nouns have a lower proportion of alternators relative to CVCC nouns. The complexity of the coda does not have the same effect in all places of articulation, e.g. CVCC nouns have a proportion of alternators that's similar to the proportion of alternators among the poly-syllables when *p*-final and *ʈ*-final nouns are considered, but *k*-final CVCC nouns pattern with the mono-syllabic *k*-final nouns, which have a low proportion of alternators.

Of the 354 stop-final nouns in TELL that have a complex coda, 244 have a sonorant before the final stop, and 39% of those 244 nouns alternate. Of the 110 nouns that have an obstruent before their final stop, only 3% alternate. Since only sonorants lead to a non-negligible proportion of alternators, only sonorants were used in the experiment in §2.3, and hence only nouns with a sonorant before their final stop will be considered below.

The alternation of nouns with simple codas was attributed in §2.4.2 to a family of markedness constraints that penalize intervocalic voiceless stops: *VpV, *VtV, *VʈV, and *VkV. Similarly, the alternations of nouns with complex codas is attributed here to markedness constraints that penalize voiceless stops between a sonorant consonant and a

vowel, namely *RpV, *RtV, *RtʃV, and *RkV. This formulation of the constraints collapses the distinction between the nasal sonorants {m, n} and the oral sonorants {l, λ, r, y}, which might be an over-simplification. In the lexicon, stops are more likely to alternate following nasals than following oral sonorants (47.6% vs. 29.3%), a tendency that was also found in the experimental results (49.0% vs. 39.6%).

The behavior of alternating and non-alternating *tʃ*-final nouns with final complex codas is shown in (58). The markedness constraint *RtʃV prefers alternation, while the familiar IDENT(voice) and IDENT(voice)_{σ1} prefer a faithfully voiceless root-final stop.

(58)

	IDENT	IDENT _{σ1}	*RtʃV
a. göntʃ-ü > göndʒ-ü	W	W	L
b. genʒ-i > gentʃ-i	L	L	W
c. gülüntʃ-ü > gülünʒ-ü	W		L
d. güvenʒ-i > güventʃ-i	L		W

With different markedness constraints regulating voicing alternations in nouns with simplex codas and complex codas, the learner can easily partition the lexicon by the complexity of the final coda. Adding the nouns with complex codas in (58) to the grammar in (57) gives rise to the more complete grammar in (59).

(59) IDENT(voice)_{σ1} ⟨*VtʃV, satʃ⟩ >> IDENT(voice) ⟨*VtʃV, anatʃ⟩
 ⟨*VtV, at⟩ ⟨*VtV, sepet⟩
 ⟨*RtʃV, göntʃ⟩ ⟨*RtʃV, gülüntʃ⟩
 >> *RtʃV, *VtʃV, *VtV >>
 IDENT(voice)_{σ1} ⟨*VtʃV, tatʃ⟩ , IDENT(voice) ⟨*VtʃV, amatʃ⟩
 ⟨*VtV, tat⟩ ⟨*VtV, kanat⟩
 ⟨*RtʃV, gentʃ⟩ ⟨*RtʃV, güventʃ⟩

The grammar in (59) allows the speaker to partition their *tf*-final nouns by their mono- or poly-syllabicity, and within each length, by the complexity of their coda. Within each of the four kinds of *tf*-final nouns, alternators are separated from non-alternators, giving the speaker access to the relative proportion of alternating nouns in each partition. The stimuli with complex codas that were used in the experiment in §2.3 were all mono-syllabic, and for those nouns, speakers successfully replicated the proportion of alternators from the lexicon.

Poly-syllabic nouns with complex codas were not treated separately in the statistical analyses in §2.2 due to their small number relative to the poly-syllabic nouns with simple codas. Of the 301 mono-syllabic nouns in TELL, the 164 nouns that have a complex coda make a respectable 54.5%. However, the 190 poly-syllabic nouns with a complex coda make a mere 7% of the 2701 poly-syllabic nouns in TELL. Consequently, poly-syllabic nouns with complex codas are not very representative of the Turkish lexicon as a whole, nor are they representative of the poly-syllabic nouns of Turkish, and therefore they were not tested in the experiment in §2.3. They are included in the analysis here for the sake of completeness only.

2.4.6 Voicing alternations and $k \sim \emptyset$ alternations

The discussion of voicing alternations in §2.2 and §2.3 abstracted away from the fact that post-vocalic dorsals delete, rather than become voiced. The crucial observation in this context is that the voicing of stem-final stops and the deletion of stem-final dorsals are in *complementary distribution*. This is seen in (60) below, where post-vocalic dorsals either surface faithfully in the possessive (a-b) or delete (c-d), whereas post-consonantal dorsals either surface faithfully (e-f) or voice (g-h).

(60)	bare stem	possessive	
a.	ok	ok-u	‘arrow’
b.	tfekik	tfekik-i	‘slanting’
c.	gök	gö-ü	‘sky’
d.	tfilek	tfile-i	‘strawberry’
e.	mülk	mülk-ü	‘real estate’
f.	mehenk	mehenk-i	‘measure’
g.	renk	reng-i	‘color’
h.	kepenk	kepeng-i	‘rolling shutter’

Given a *k*-final noun in Turkish, it is not predictable whether it will surface faithfully or unfaithfully, but if it is known to surface unfaithfully, it is predictable whether the final [k] will voice (following a consonant) or delete (following a vowel). If dorsal deletion were in some sense an independent process of Turkish, its complementary distribution with respect to voicing would be left unexplained.

Both the voicing and the deletion of final dorsals show a size effect in TELL (61). While the size effect is dramatic for the post-vocalic dorsals (3% vs. 93%), there is also a noticeable size effect for the post-consonantal dorsals (10% vs. 41%).

(61)	Size	Faithful	Alternating	% alternating
Deletion	mono-syllabic	42	1	3%
	poly-syllabic	79	1048	93%
Voicing	mono-syllabic	45	5	10%
	poly-syllabic	19	13	41%

The deletion of a final dorsal does not violate IDENT(voice), but rather violates MAX, a faithfulness constraint that penalizes deletion. To learn the size effect, the learner will need

to use the general MAX and the specific $MAX_{\sigma 1}$, which penalizes the deletion of material from the initial syllable of the stem.

The complementary distribution of voicing alternation and dorsal deletion is apparent from the summary of the ranking arguments, exemplified with mono-syllabic nouns in (62). There is a conflict between $IDENT(voice)_{\sigma 1}$ and *RkV, and there is a separate conflict between $MAX_{\sigma 1}$ and *VkV. The learner is free to discover each conflict separately.

(62)

	$IDENT_{\sigma 1}$	*RkV	$MAX_{\sigma 1}$	*VkV
a. mülk-ü > mülğ-ü	W	L		
b. reng-i > renk-i	L	W		
c. ok-u > o-u			W	L
d. gö-ü > gök-ü			L	W

If $IDENT_{\sigma 1}$ is cloned first, $IDENT(voice)_{\sigma 1} \langle *RkV, mülk \rangle$ will be installed, followed by the installation of *RkV. Then, either $MAX_{\sigma 1}$ or *VkV will need to be cloned. If $MAX_{\sigma 1}$ is cloned, the resulting grammar will be as in (63).

(63) $IDENT(voice)_{\sigma 1} \langle *RkV, mülk \rangle \gg *RkV \gg MAX_{\sigma 1} \langle *VkV, ok \rangle \gg *VkV$
 $\gg IDENT(voice)_{\sigma 1} \langle *RkV, renk \rangle, MAX_{\sigma 1} \langle *VkV, gök \rangle$

Equivalently, If $MAX_{\sigma 1}$ is cloned first, followed by the cloning of $IDENT(voice)_{\sigma 1}$, the resulting grammar, in (64), is just as good as the grammar in (63) in accounting for the available data.

(64) $MAX_{\sigma 1} \langle *VkV, ok \rangle \gg *VkV \gg IDENT(voice)_{\sigma 1} \langle *RkV, mülk \rangle \gg *RkV$
 $\gg IDENT(voice)_{\sigma 1} \langle *RkV, renk \rangle, MAX_{\sigma 1} \langle *VkV, gök \rangle$

Since the deleting dorsals and the voicing dorsals are in complementary distribution, and controlled by separate constraints, it doesn't matter which trend leads to cloning first.

2.4.7 Summary of the analysis

This section offered an OT-based model that allows speakers to detect inconsistent behavior in their lexicon, and encode the inconsistency in terms of lexically-specific constraint clones. Each cloned constraint lists the items that it applies to, with each item listed with the constraint that triggered the inconsistency. This lexically-enriched grammar can be applied to novel items, with clones that list more items more likely to exert their influence, thus projecting the lexical trend unto the novel items.

The listing of lexical items with clones can also be seen as partitioning the lexicon: Each item is classified according to its behavior, getting listed with an appropriate clone if it participates in a lexical trend, or going unlisted if it isn't.

In Turkish, voicing alternations are irregular. Stem-final voiceless stops become voiced before vowel-initial suffixes in some words due to markedness constraints that favor lenition, and stay voiceless in other words due to faithfulness to their base form, which is also assumed to be their underlying form. The availability of a family of markedness constraints that affect each place of articulation separately, (viz. *VpV, *VtV, *VʃV, *VkV) allows speakers to partition the stop-final nouns of Turkish according to the place of articulation of the final stop. With access to the relative number of items in each partition of the lexicon, speakers can project this aspect of the lexical statistics onto novel forms. The availability of faithfulness constraints that are specific to initial syllables (viz. IDENT(voice)_{σ1} and MAX_{σ1}) and general faithfulness constraints allows speaker to partition the stop-final nouns of Turkish according to their size: Alternations in mono-syllabic nouns can be identified as causing unfaithfulness to the only, and hence to the initial syllable of the base, whereas alternations in longer nouns do not affect the initial syllable. This lets

speakers partition the lexicon by the size of its nouns, and then project the lexical statistics onto novel items.

In the proposed model, the language learner identified the most specific lexical trend that can be expressed with constraint interaction. Whenever the behavior of lexical items causes ranking conflicts, lexical items are recorded with reference to two most specific conflicting constraints: One of the constraints is cloned, and items are listed under a clone, paired with the other constraint that was involved in the conflict. In Turkish, this allows speakers to combine the place effect and the size effect, listing nouns according to their size and the place of their final stop.

Since the model only uses the Universal constraints in CON to record lexical trends, it ignores facts about the lexicon that cannot be expressed with universal constraints. Since languages are not observed to have interactions of obstruent voicing with the height or backness of neighboring vowels, there are no constraints that penalize combinations of voicing with neighboring vowel qualities. In the absence of such constraints, Turkish speakers cannot record the effect that vowel height and backness have on the distribution of voicing alternations.

2.5 General-purpose learning with the MGL

The Minimal Generalization Learner (MGL) of Albright & Hayes (2002, 2003, 2006) is an information-theoretic algorithm that generalizes patterns over classes of words that undergo similar alternations. MGL provides a reflection of trends in the lexicon and has the potential to generalize them to novel outputs. The MGL has been shown to successfully model humans' experimental results in novel word-formation tasks with the past tense in English and with similar tasks in other languages, and is thus a good representative of a class of models that access lexical patterns without any bias against generalizing from phonologically unnatural trends.

The MGL works by reading in pairs of surface forms that are morphologically related, such as a bare noun and its possessive form in Turkish, creating a rule for each pair, and then generalizing over those rules to make more general rules. These more general rules can be applied to novel bare nouns, giving a set of possible derived forms with a confidence score assigned to each.

2.5.1 Materials and method

To simulate the behavior of the human participants as described in the experiment in §2.3, the MGL was provided with all the stop-final words in TELL as training data, and with the stimuli of the experiment as test items. In addition, the MGL received a feature matrix of the consonants and vowels of Turkish, which it uses to find natural classes. The results reported here were obtained by running the MGL at the 75% confidence level, which is the level that generated the results that most closely matched the human results.

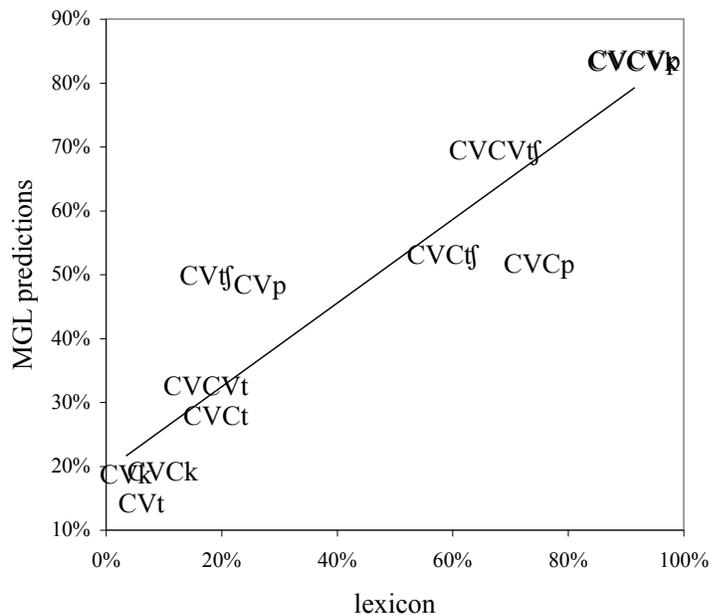
For each test item, the MGL generated alternating and non-alternating possessive forms, each form associated with a confidence score, which represents the likelihood of getting that response from a human. To calculate the proportion of alternating responses that the MGL predicts, the confidence score of each alternating response was divided by the sum of the confidence scores of the alternating and non-alternating responses. For example, given the noun *fat*, the MGL produced the form *fat-i* with a confidence of 87% and the form *fad-i* with a confidence of 23%. The predicted alternation rate for *fat* was calculated as $23\% / (23\% + 87\%) = 21\%$.¹⁰ Thus, the MGL predicted alternation rates for each of the 72 test items of the experiment.

¹⁰The MGL's confidence in *fat-i* and its confidence in *fad-i* are not guaranteed to add up to 100%, because the MGL may use different rules with different scopes for deriving the two outputs. For example, *fat-i* was derived with a rule that is limited to CVt roots, most of which do not alternate, hence the high confidence rate; whereas *fad-i* was derived with a rule that affects *t*-final stems of any size, and thus allows the relatively higher rate of alternation in CVCVt roots relative to CVt roots to boost the confidence in *fad-i*.

2.5.2 Results

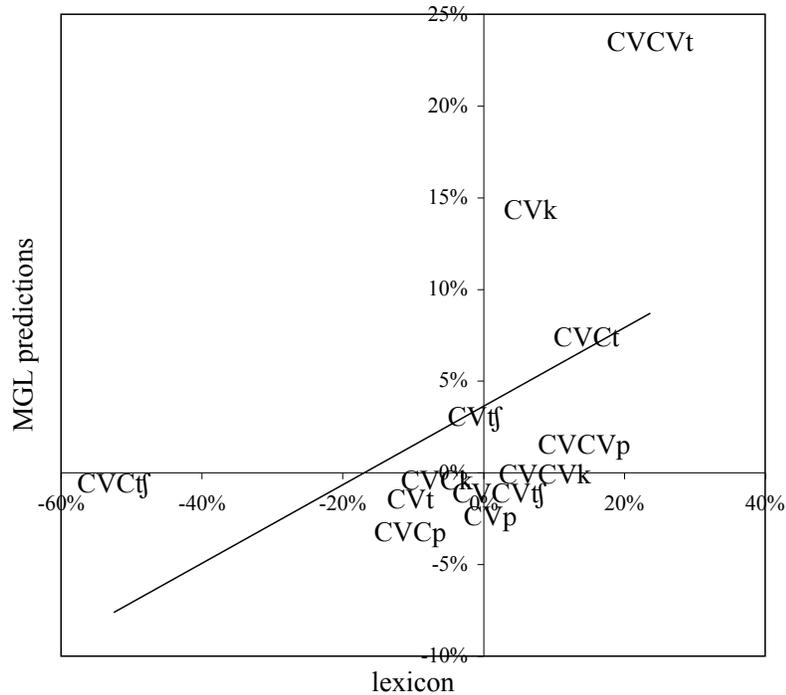
The chart in (65) shows MGL's prediction for the nonce words used in the experiment, grouped by size vs. place, plotted against the proportion of alternating words in TELL in the corresponding size and place. The MGL predictions matches the lexicon very well (Spearman's rank correlation test, $S = 18$, $\rho = .937$, $p < .001$).

(65) Rates of alternation in the lexicon, by place and size, plotted against the percentage of alternating responses predicted by the Minimal Generalization Learner.



The MGL prediction match the lexicon for the height effect as well, as shown in (66), with significant correlation (Spearman's rank correlation test, $S = 92$, $\rho = .678$, $p < .05$). This contrasts sharply with the lack of correlation between the lexical statistics and the experimental results (see 32 above).

- (66) The difference in rates of alternation between high and non-high vowels, by size and place, in the lexicon and in the MGL results.



2.5.3 Discussion

The MGL’s impressive performance in matching the lexical trends of Turkish voicing alternations were to its detriment. In out-performing the participants of the experiment described in §2.3, it failed to mimic human behavior.

The MGL is a powerful learner for phonological patterns. Given nothing but a list of paradigms and the natural classes that the segments in it form, it learned that Turkish has voicing alternations and that there are factors that are correlated with their distribution. However, since the MGL lacks a theory of possible interactions between phonological elements, it could not ignore the predictive power of vowel height and backness in determining the alternating or non-alternating status of attested nouns, and it used all the correlations it found in predicting the status of novel forms.

Humans, I argue, are biased to ignore any effect that vowel quality might have on the voicing of a neighboring consonant. This one and the same bias is observed in two

domains of linguistic investigation: In the cross-linguistical study of regular phonological phenomena, and in the language-specific study of the distribution of lexically-determined phonological processes.

The MGL results are representative of a wider range of learning algorithms, such as CART (Breiman et al. 1984) or C4.5 (Quinlan 1993), which use purely distributional properties of a lexicon to model human behavior. The MGL's advantage over these other models is that it isn't given a list of possible generalizations to explore in advance, but rather generates its own set of hypotheses. With models other than the MGL, the lack of vowel effect could be hard-wired by not supplying the model with information about vowel quality. Since these models are not specific to language and therefore don't have any information about natural phonological interactions, such an exercise would offer little insight into the problem at hand. The MGL simulation is informative specifically because it is given whole words to deal with, without additional information about which generalizations to attend to.

The MGL results show that a model that isn't equipped with a set of biases that determine the universal range of phonological interactions will be unable to successfully mimic human behavior and ignore accidental regularities in a lexicon.

2.6 UR-based approaches to final devoicing

The traditional generative analysis of Turkish voicing alternations (Lees 1961; Inkelas & Orgun 1995; Inkelas et al. 1997) attributes different underlying representations to word-final stops based on their behavior (although a different approach was suggested in Hayes 1995b). There is no explicit analysis of Turkish in terms of Optimality Theory, but an analysis in the spirit of Inkelas et al. (1997) would be something like (67). In this analysis, nouns that surface with a voiceless stop throughout the paradigm have a voiceless stop underlyingly, while stops that alternate have an underlying stop that is unspecified for [\pm voice].

- (67) a. The UR's of [at] and [tat] are /at/ and /taD/
 b. The UR of the possessive is /I/ (a high vowel)
 c. /at + I/ → [at-i] requires IDENT(voice) ≫ *VtV

at + I	IDENT(voice)	*VtV
a.  at-i		*
b. ad-i	*!	

- d. /taD + I/ → [tad-i] is consistent with IDENT(voice) ≫ *VtV

taD + I	IDENT(voice)	*VtV
a. tat-i		*!
b.  tad-i		

In this theory, IDENT(voice) dominates any relevant markedness constraints, and alternating stops have under-specified underlying representations that escape faithfulness. Underlyingly voiced stops will surface faithfully throughout their paradigm, as is observed in nouns such as *ad* ~ *ad-i* ‘name’. The deletion of dorsals can be encoded in a different representational mechanism, that of “ghost segments” (Zoll 1996), or segments whose absence from the output does not violate the regular MAX, as suggested by Joe Pater (p.c.).

This theory encodes the observed difference between alternating and non-alternating paradigms in the underlying representations, leaving the grammar consistent. Since the experiment in §2.3 shows that speakers have detailed grammatical knowledge about the propensity of final stops to alternate, it is not clear how speakers could encode this knowledge if it allowed to escape the grammar. Burying information about voicing

alternations in the lexicon would force speakers to look for generalizations directly in the lexicon, where nothing would prevent them from finding the vowel quality effects that they didn't exhibit in §2.3.

In the analysis offered in §2.4, the bare forms of nouns were assumed as their underlying representations, and it exactly this assumption that forced the speaker to find conflicting ranking arguments, and then encode lexical statistics in the grammar. The consequences of assuming surface forms as underlying forms are further explored in §4.4.

Beckman & Ringen (2004) offer a different UR-based analysis of Turkish voicing alternations. They focus on the fact that pre-vocalic voiceless stops in Turkish are aspirated, i.e. nouns like *at* ~ *at-i* are actually pronounced *at* ~ *at^h-i*. Then, they derive the three-way contrast between voiceless throughout, voiced throughout and alternating stops from an underlying contrast between aspirated, voiced, and unmarked stops, respectively. The aspirated and unmarked stops merge in the bare stem due to a constraint against final aspirated stops, and the unmarked and voiced stops merge in the possessive forms due to phonetic passive voicing of intervocalic stops.

The accuracy of Beckman & Ringen's (2004) phonetic description is not contested here¹¹. Rather, I point out that a reliance on underlying representations leaves unexplained speakers' knowledge about the distribution of voicing alternations.

2.7 Conclusions

This chapter presented a study of Turkish voicing alternations that contrasted trends found in the Turkish lexicon with the knowledge that speakers have about it, showing that speakers are biased to reproduce certain trends but not others.

¹¹In fact, assuming that Beckman & Ringen's (2004) phonetic description is accurate, then no possessive form of Turkish violates *VtV, and this constraint can no longer distinguish alternating and non-alternating forms. The speaker would have to call upon different constraints, such as IDENT(asp). Alternatively, the effect of *VtV could be observed opaquely.

Speakers chose voicing alternations when presented with novel nouns more often with poly-syllables than with mono-syllables, and with non-coronals more often than coronals, reflecting the trends in the lexicon. However, they did not choose more alternating responses when the rightmost vowel of the novel noun was high or back, ignoring the trend for more alternations in those conditions in the lexicon. The proposal made here was that lexical trends are learned in terms of typologically-responsible constraints, which are part of UG. The prediction this makes is that there is a necessary correlation between the space of regular phonological processes as observed in the world languages on one hand, and the space of irregular trends that speakers can extract from their lexicon on the other hand, since both kinds of phenomena stem from a single posited set of Universal Constraints.

A statistical analysis of the Turkish lexicon was offered, and contrasted with the results from the experiment, showing that speakers ignored a correlation between vowel quality and the voicing of a neighboring vowel. The experimental results were contrasted with the results of the MGL simulation (Albright & Hayes 2002, 2003, 2006), which over-learned the Turkish data, projecting the vowel quality effects that humans ignored.

The conclusion was that a general-purpose statistical learner could not reproduce the behavior that humans display, and that a successful theory of lexical learning must combine the ability to learn lexical trends with UG-based biases. The proposed learner identified conflicting lexical behaviors in the lexicon and resolved the conflict by cloning constraints. Once constraints are cloned, each clone keeps a list of the words it governs, assuring that existing words behave consistently. At the same time, the clones can be used in a generalized way, referring only to the proportion of words that are governed by each clone, to project the lexical trend onto novel words.

The resulting learner simulated the process of learning a lexicon without relying on general-purpose pattern matching. Rather, it used a set of Universal Constraints that were augmented by the ability to clone constraints. In the Turkish case, the simulated learner

ignored the correlation between vowel quality and consonant voicing thanks to the absence of constraints that relate the two, and thus it mimicked the behavior of the human learner.

CHAPTER 3

SURFACE-BASED LEXICAL TRENDS IN HEBREW

3.1 Introduction

In Hebrew, the plural suffix for nouns has two allomorphs: *-im* for masculine nouns and *-ot* for feminine nouns. The choice of affix is completely predictable for adjectives and loanwords, but native nouns allow exceptions both ways: some masculine nouns take *-ot*, and some feminine nouns take *-im*.

The masculine nouns that exceptionally take *-ot* are phonologically clustered. Out of the 230 *ot*-takers in a Hebrew lexicon (Bolozky & Becker 2006), 146 nouns, or 63%, have the vowel [o] in their last syllable. The results reported in §3.3 below and in Berent, Pinker & Shimron (2002, 1999) show that speakers are aware of the trend for more *-ot* in nouns that end in [o], and project this trend onto novel items. In other words, speakers' choice of plural allomorph is not determined entirely by the stem's gender or morphologically idiosyncratic properties, but also by the stem's phonological shape.

In my analysis of this case of partially phonologically determined allomorph selection, *ot*-takers with [o] in them respond to a high-ranking markedness constraint that requires an unstressed [o] to be licensed by an adjacent stressed [o] (cf. similar requirement on vowel licensing in Shona, Beckman 1997; Hayes & Wilson 2008). Markedness-based accounts of allomorph selection in OT are common in the literature, starting with Mester (1994) and continuing with Mascaró (1996), Kager (1996), Anttila (1997), and Hargus (1997), among many others. More recent work includes Paster (2006), Wolf (2008b), and Trommer (2008). Since the analysis crucially relies on the use of markedness constraints, i.e. constraints that assess output forms, regardless of the posited underlying representation,

I set out to empirically test the adequacy of accounting for lexical trends using markedness constraints.

At issue is what Albright & Hayes (2003) call source- vs. product-oriented generalizations. In the Hebrew case, one can state the correlation between a stem [o] and *-ot* in a source-oriented way, i.e. in terms of a relationship between singular and plural forms, saying that nouns that have [o] in the singular are more likely to take *-ot* in the plural. Alternatively, one can state the generalization in a product-oriented way, i.e. in terms of conditions on the plural forms only, saying that in the plural, noun stems that have [o] in them are more likely to show up with the suffix *-ot*. In Optimality Theory, generalizations that are stated in terms of markedness constraints are product-oriented, since markedness constraints only assess outputs, or products of derivations. In contrast, rule-based theories express generalizations in terms of mappings between inputs and outputs, i.e. generalizations depend on the input to the derivation, so they are source-oriented.

The source-oriented and product-oriented generalizations are almost exactly equivalent when stated over the attested lexicon of Hebrew, since each and every noun that has an [o] in the final syllable of its plural stem also has an [o] in the singular¹, and with the exception of five nouns², every noun that has an [o] in its final syllable in the singular also has an [o] in the final syllable of the plural stem.

I propose that evidence in favor of product-oriented knowledge of lexical trends can be adduced by Hebrew speakers' behavior in an artificial language setting. I present such an experiment, where speakers were taught a language that is just like Hebrew, but with two additional vowel-change rules that caused [o]'s to be present only in the singular stem or only in the plural stem, but not in both. Speakers preferred to associate the selection of

¹For nouns with the vowel pattern [o-e] in the singular, vowel deletion makes the [o] stem-final in the plural, e.g. *fomér* ~ *fomr-ím* 'guard, keeper'.

²Three nouns change the singular [o] to [u] (*xók* ~ *xukím* 'law', *tóf* ~ *tupím* 'drum' and *dóv* ~ *dubím* 'bear'), and two nouns change the singular [o] to [a] (*róf* ~ *rafím* 'head', *yóm* ~ *yamím* 'day').

–*ot* with nouns that have [o] in the plural stem rather than in the singular stem, showing that they were using surface-based, or product-oriented methods for selecting the plural allomorph.

This chapter is organized as follows: §3.2 presents the distribution of the plural allomorphs in the lexicon, and §3.3 shows that speakers project this distribution onto novel items. The analysis of these trends in terms of markedness constraints is in §3.4. Support for this analysis is presented in §3.5, with results of an artificial language experiment that shows speakers’ preference for product-oriented generalizations. The results are discussed and analyzed in §3.6. Conclusions are in §3.7.

3.2 Hebrew plurals: Lexicon study

Hebrew has two plural markers: –*im* and –*ot*. When nouns that refer to humans have an *im*-form and an *ot*-form, they invariably correspond to natural gender, as in the word for boy/girl in (68)³. At the phrase level, gender agreement on adjectives and verbs is also invariably regular.

- (68) a. yelad-ím ktan-ím jar-ím
 boy-pl little-pl sing-pl ‘little boys are singing’
 b. yelad-ót ktan-ót jar-ót
 girl-pl little-pl sing-pl ‘little girls are singing’

At the word level, native nouns can take a mismatching suffix: (69a) shows that the masculine noun *xalón* exceptionally takes –*ot* at the word level, but the accompanying adjective and verb take –*im*, revealing the true gender of the noun (Aronoff 1994). The opposite is seen with the feminine noun *nemalá* in (69b).

³When nouns that refer to humans only have one plural form, the plural affix does not necessarily conform to natural gender. For example, the native noun *fuliy-á* ~ *fuliy-ót* ‘apprentice’ can apply to either males or females. The word for ‘baby’ has gender marking in the singular (masculine *tinók* vs. feminine *tinók-et*), but the plural is *tinok-ót* for male or female babies. Not surprisingly, children often use the form *tinok-ím* to refer to male babies.

- (69) a. xalon-ót gdol-ím niftax-ím
 window-pl big-pl opening-pl ‘big windows are opening’
- b. nemal-ím ktan-ót nixnas-ót
 ant-pl small-pl entering-pl ‘small ants are coming in’

In the loanword phonology, the plural suffix selection is completely regular even at the word level: If the right edge of the singular noun is recognizable as a feminine suffix, as in *fukátf-a*, *-ot* is selected (70a), otherwise it’s *-im*, as in *blóg-im* (70b). This even applies to nouns that refer to male humans, like *koléga* (70c). Loanwords that refer to female humans but don’t have a plausible feminine suffix on them, like *madám*, mostly resist pluralization⁴ (70d).

- (70) a. fukátf-a * fukátf-im fukátf-ot ‘focaccia’
 b. blóg blóg-im * blóg-ot ‘blog’
 c. kolég-a * kolég-im kolég-ot ‘(male) colleague’
 d. madám ? madám-im ??? madám-ot ‘madam (in a brothel)’

A final factor that affects the distribution of the plural allomorphs is phonological. Masculine native nouns show a clustering of the *ot*-takers: most of the masculine nouns that exceptionally take *-ot* have [o] in their final syllable (Glinert 1989; p. 454, Aronoff 1994; p. 76). This preference for *-ot* in masculine nouns that end in [o] applies productively to novel nouns, as seen in Berent, Pinker & Shimron (1999, 2002) and in §3.3 below. The feminine native nouns are less interesting, because there are relatively few *im*-takers among them, and those few *im*-takers don’t seem to pattern in any noticeable way.

To summarize so far, there are three factors that determine plural allomorph selection without exception:

⁴Some speakers offer *madám-iy-ot* as the plural of *madám*, i.e. they add the feminine suffix *-it* to the root to make a more plausible singular feminine stem for the plural *-ot* to attach to. The change of *-it* to *-iy* before *-ot* is regular in the language (Bat-El 2008a).

- (71) a. Natural gender: Whenever a single noun stem refers to males and females, *-im* will refer to males and *-ot* will refer to females.
- b. Morpho-syntactic gender: Adjectives and verbs take *-im* with masculine nouns and *-ot* with feminine nouns. Essentially, adjectives and verbs reveal the true gender of a noun.
- c. Morpho-phonological gender: When a loan-word (i.e. a noun that keeps the stress on its stem in the plural) ends in what sounds like a feminine suffix, its plural will be in *-ot*, otherwise its plural will be in *-im*.

And there are two factors that have some power in predicting the plural allomorph selection, but these allow exceptions:

- (72) a. Morpho-syntactic gender: A native noun (i.e. a noun that loses its stress to the plural affix in the plural) usually takes *-im* if it's masculine and *-ot* if it's feminine.
- b. Phonology: The majority of native masculine nouns that take *-ot* in the plural have an [o] in their stem.

From this point on, the focus will be on native masculine nouns, and the phonological effect of a stem [o] on the selection of the plural affix. The presence of a stem [o] makes the selection of *-ot* more likely, relative to the selection of *-im* in the absence of a stem [o].

The partial predictability in the distribution of *ot*-takers is not incompatible with the existence of minimal pairs, such as those in (73), where the choice of plural affix disambiguates the meaning. Overall in the lexicon, *-ot* is more likely with a stem [o], but for any single lexical item, the selection of an affix is unpredictable.

- (73) a. himnon-ím / himnon-ót ‘national anthem’ / ‘religious hymn’
- b. tor-ím / tor-ót ‘line, queue’, ‘appointment’ / ‘turn’
- c. maamad-ím / maamad-ót ‘stand’ / ‘status’

With certain nouns, the choice of plural suffix is variable in and between speakers. Some nouns that occur variably in current usage are in (74), where the percentage indicates the proportion of *-ot* plurals out of the total plural forms found in Google⁵.

(74)	a.	fofar-ím / fofar-ót	56%	‘shofar’
	b.	dyokan-ím / dyokna-ót <i>or</i> dyokan-ót	41%	‘portrait’
	c.	kilfon-ím / kilfon-ót	11%	‘pitchfork’

For the purposes of this study, data about the distribution of *-im* and *-ot* comes from an electronic lexicon of Hebrew (Boložky & Becker 2006) that was modeled after TELL (a Turkish Electronic Living Lexicon, Inkelas et al. 2000). The lexicon lists nouns and their plurals. The nouns are mostly collected from the Even-Shoshan dictionary, and their plurals reflect the knowledge of the second author, occasionally augmented by Google searches, in an attempt to approximate an idealized native speaker. The table in (75) lists the native masculine nouns in the lexicon, arranged by the vowel in their final syllable. Recall that in this context, ‘native’ refers to unaccented nouns (Bat-El 1993; Becker 2003), i.e. nouns that surface in the plural with the stress on the plural suffix.

(75)	Final vowel	<i>n</i>	<i>ot</i> -takers	% <i>ot</i> -takers
	u	1101	6	0.5%
	i	464	8	1.7%
	a	1349	39	2.9%
	e	977	31	3.2%
	o	523	146	27.9%
	Total	4414	230	5.2%

⁵There are surely many more nouns that variably take either plural affix, but Hebrew orthography makes searching for them online a difficult task. The variable choice of the plural affix goes back to Tiberian Hebrew, where a considerable number of nouns are attested with two plural forms (Aharoni 2007), e.g. do:r-ím (Isaiah 51, verse 8) vs. do:r-ó:θ (Isaiah 41, verse 4) ‘generations’.

The data in (75) shows that *ot*-taking accounts for a fairly meager proportion (2.2%) of the native nouns that end in vowels other than [o], but almost a third of the nouns that end in [o]. The 146 *ot*-takers that end in [o] account for 63% of the 230 *ot*-takers.

There are further morpho-phonological regularities that correlate with *ot*-taking within the set of nouns that have [o] in their final syllable. For instance, *ot*-taking is completely regular for a class of tri-syllabic masculine nouns that have a stem of the shape [CiCaC–] and the suffix [–on] (e.g. *fikar-on* ‘state of drunkenness’). These nouns can be productively formed from verbs to mean ‘state of X-ness’, and with this meaning, their plural is always in *–ot*⁶. Tri-syllabic nouns in [–on] account for 54 of the 146 [o]-final *ot*-takers in (75). Of the remaining 92 [o]-final *ot*-takers, 49 end in the segments [on], but in many cases, it is hard to determine whether these segments belong to the an affix or to a stem.

Having an [o] in the root is well correlated with taking *–ot* in the plural even after allowing for the effect of the suffix [–on]. In the lexicon, this can be seen with monosyllables: Of the 70 monosyllables with [o] in them, 20 are *ot*-takers (29%), and none of these *ot*-takers end in [n]. This rate of *ot*-taking is comparable to the overall rate of *ot*-taking.

Looking at di-syllabic nouns only⁷, the effect of a root [o] is observed not only locally, but also at a distance. The table in (76) shows that having an [o] in the penultimate syllable correlates with a level of *ot*-taking that is intermediate between roots with a final [o] and roots with no [o].

⁶The etymological data in Bolozky & Becker (2006) confirms the modern productivity of *ot*-taking for [CiCaC-on] nouns. Of the 230 *ot*-takers, 216 are attested before modern Hebrew (i.e. Biblical or Mishnaic). Of the remaining 14 *ot*-takers that were created in modern times, 13 are [CiCaC-on] nouns. The remaining modern item, *dúax* ~ *dux-ót* ‘report’, is colloquially pronounced *dóx* ~ *dox-ót*, thus making every single modern *ot*-taker a noun with [o] in its stem.

⁷Bolozky & Becker (2006) list only six native nouns with an [o] in their antepenultimate syllable, and none with earlier [o]’s. All six are poly-morphemic and take *–im*. This is hardly surprising, given that few native nouns surface more than two syllables long, and all are analyzed as underlyingly disyllabic in Becker (2003).

(76)	Vowel pattern	<i>n</i>	<i>ot</i> -takers	% <i>ot</i> -takers
	a-a	589	12	2.0%
	o-a	102	12	11.8%
	a-o	163	34	20.9%

This action at a distance, however, is only observed when it is [a] that intervenes between the root's penult [o] and the plural affix:

(77)	Vowel pattern	<i>n</i>	<i>ot</i> -takers	% <i>ot</i> -takers
	o-a	102	12	11.8%
	o-e	288	0	0%
	o-i	18	0	0%
	o-u	1	0	0%

This absence of *ot*-takers in the last three rows of (77) is not necessarily entirely phonological. Nouns with an [o-e] vowel pattern often refer to male humans, in which case they always take *-im*, e.g. *torém* 'benefactor', *ƶodéd* 'robber'. Other such nouns are plausibly derived from present participles, which regularly take *-im* when masculine, e.g. *nozél* 'liquid', from *nazál* 'to flow', and *mocéc* 'pacifier', from *macác* 'to suck'. The paucity of [o-u] nouns reflects the general rarity of native nouns that combine two rounded vowels. Not much can be said about the 18 [o-i] nouns, since the expected number of *ot*-takers would be no more than two or three, and their absence could be a lexical gap. On the other hand, of the eight *ot*-takers that have [i] in the final vowel of their singular stem, only one keeps that vowel in the plural, so it's possible that having an [i] in the last vowel of the plural stem is particularly incompatible with selecting *-ot*⁸. In the following discussion, I

⁸Of the eight *ot*-takers with a final [i] in Bolozky & Becker (2006), only *kír* 'wall' has [i] in the plural, *kír-ót*. Four more are segolates that have [ye] or [ey] in the plural: *gáis* ~ *gyas-ót* 'army', *yáin* ~ *yeyn-ót* 'wine', *láil* ~ *leyl-ót* 'night', and *xáil* ~ *xeyl-ót* 'corps'. The remaining three are essentially suppletive: *reʔí* ~ *marʔ-ót* 'mirror', *arí* ~ *aray-ót* 'lion', and *prí* ~ *pe(y)r-ót* 'fruit'.

will stay agnostic about the status of the intervening vowels in (77), and assume that the presence of [o] in the stem's penultimate syllable increases the likelihood of taking *-ot* no matter what the vowel in the ultima is.

To summarize the findings: In native masculine nouns, a stem [o] is correlated with selecting the plural *-ot*. The correlation is strongest when the [o] is closest to the *-ot*, i.e. in the final syllable of the stem. A weaker correlation is observed when the [o] is in the penultimate syllable of the root, when an [a] intervenes.

3.3 Speakers' knowledge of lexical trends

To test what generalizations Hebrew speakers make about the distribution of the plural suffix, and see how these generalizations relate to the distribution of the plural suffix in the lexicon, I tested speakers' choice of plural suffix with novel words that had four vowel patterns: [a-a], [o-a], [a-o], and [i-o]. These represent words that have no [o] at all, words that have [o] in the penultimate syllable of the stem, and two kinds of words that have an [o] in their final syllable.

3.3.1 Materials and methods

For each of the four vowel patterns tested, the experiment contained 14 novel words and 6 existing words, i.e. 56 novel words and 24 existing words, or 80 in total. The 6 existing words in each vowel pattern were all native nouns of Hebrew, four of which were *ot*-takers and two were *im*-takers. All existing words were high-frequency words with frequent plural forms.

For each of the 80 words, the singular and two plurals were recorded by a male native speaker in a sound-attenuated booth onto a Macintosh computer at 44100 Hz, using Audacity. Then, for each word, two .wav sound files were created using Praat (Boersma & Weenink 2008). One file started with .5 seconds of silence, followed by the singular played twice, then the *-im* plural, and then *-ot* plural, with a second of silence following

each word. The second file was similarly constructed with the *-ot* plural first, followed by the *-im* plural. Each file was converted to .mp3 format using the LAME encoder, version 3.97 (from <http://www.mp3dev.org/>).

The experiment was conducted on a web-based interface, using Firefox. After some instructions were presented, training consisted of responding to three nouns with the vowel pattern [u-a]: an existing *ot*-taker (*sulám* ‘ladder’), an existing *im*-taker (*duxán* ‘stall’), and a novel noun (*kufár*). Feedback was given for the two existing items.

The experimental items were randomized and presented in a frame sentence that makes them masculine nouns, e.g.:

(78) ze *kamoz*, ve ze od *kamoz*.
this_{MASC} is a *kamoz*, and this_{MASC} is another *kamoz*.

beyaxad, ele fney _____
together, they’re two_{MASC} _____

The sentence appeared on the screen in Hebrew orthography, which included vowel diacritics on the target nouns. In parallel, the participants heard one of the sound files as described above, with the singular heard twice, followed by the two plural forms in random order, e.g. *kmoz-ím* and *kmoz-ót*. Using the mouse, the participants were asked to choose the form that sounded most appropriate by clicking one of two buttons.

The real words used are listed with their plurals in (79). The plural forms that were assumed to be correct are in parentheses, with the full form given if it differs from the simple concatenation of the singular root and the plural suffix. The novel words are listed in (83) below, with the experimental results.

(79) Existing words

a-a	a-o	i-o	o-a
xafáf (-ót)	makór (mekor-ót)	cinór (-ót)	olám (-ót)
zanáv (znav-ót)	xalóm (-ót)	nixóax (nixox-ót)	mosád (-ót)
mazál (-ót)	garón (gron-ót)	vilón (-ót)	ocár (-ót)
nahár (nehar-ót)	asón (-ót)	kinór (-ót)	morád (-ót)
davár (dvar-ím)	alón (-ím)	kidón (-ím)	gozál (-ím)
bacál (bcál-ím)	jaón (jeon-ím)	kiyór (-ím)	koláv (-ím)

3.3.2 Participants

The participants were 62 adult native speakers of Hebrew, students at the Hebrew University in Jerusalem. They were recruited with the generous help of Ram Frost, of the Hebrew University Psychology Department. One additional participant was excluded for making more than 60% mistakes with the actual words tested, suggesting that she misunderstood the task. A mistake was defined as a judgment that deviated from the author's knowledge of Hebrew, as given in (79), and hence from the statistics extracted from Bolozky & Becker (2006). The other 62 speakers made very few mistakes with the actual words ($M = .7$, $SD = .8$, $\max = 3$).

3.3.3 Results

The participants chose *-ot* least often with [a-a], more often with [o-a], and most often with [a-o], essentially replicating the lexical trend (80). There is a trend in the lexicon for more *-ot* after [i-o] than after [a-o], which speakers did not replicate; this is discussed in §3.4.2 below. The by-item results are in (83).

(80)

Vowel pattern	Experiment	Lexicon
a-a	26%	2%
o-a	29%	12%
a-o	32%	21%
i-o	33%	26%

The results were analyzed with a mixed-effects logistic regression in R (R Development Core Team 2007) using the *lmer* function of the LME4 package, with *participant* and *item* as random effect variables. With an unordered four-level *vowel* fixed-effect factor as a predictor and the choice of plural affix as a binary dependent variable, the vowel effect only approaches significance. With [a-a] as a baseline, [a-o] is more conducive to choosing *ot*-plurals (81), but the other two vowel patterns are not.

(81)

	Estimate	SE	<i>z</i>	<i>p</i>
(Intercept)	-1.1077	0.1431	-7.739	< 0.001
a-o	0.3425	0.1848	1.853	0.064
i-o	0.3042	0.1852	1.642	0.101
o-a	0.1678	0.1858	0.903	0.366

An inspection of the results for the individual items (in 83) raised the suspicion that some stimuli got a very high rate of *ot*-responses due to the similarity of their final syllable (or their last three segments) to the final syllable (or last three segments) of a real *ot*-taker. For example, the two stimuli that got the highest number of *ot*-responses in the [a-a] vowel pattern were *ga.rád* and *ca.gág*, and each of them shares the last syllable with the real *ot*-takers *mo.rád* ~ *morad-ót* ‘slope’ and *gág* ~ *gag-ót* ‘roof’.

To see what post-hoc effect the final syllable might have, a binary variable named *similar* was added to the analysis. The items that were given a value of 1 were *garád*, *cagág*, *kalám*, *pašáf*, *kanód*, *pacóc*, and *cikór*, due to their similarity, respectively, to *morád*

~ *morad-ót* ‘slope’, *gág* ~ *gag-ót* ‘roof’, *sulám* ~ *sulam-ót* ‘ladder’, *xafáf* ~ *xafaf-ót* ‘worry’, *nód* ~ *nod-ót* ‘flask’, *nicóc* ~ *nicoc-ót* ‘spark’, and *makór* ~ *mekor-ót* ‘source’. The other items were given a value of zero, since they did not share their final syllable with any known *ot*-taker.

The addition of *similar* as a fixed-effect variable made a highly significant improvement to the model, as determined by an ANOVA model comparison ($\chi^2(1) < .001$). Not only did *similar* come out highly significant, it allowed the effect of *vowel* to emerge (82). The adequacy of this model was verified with the *pvals.fnc* function from the *languageR* package (Baayen 2008), which left the p-values essentially unchanged.

(82)	Estimate	SE	<i>z</i>	<i>p</i>
(Intercept)	-1.3488	0.1357	-9.936	< 0.001
a-o	0.4660	0.1577	2.954	0.003
i-o	0.4977	0.1608	3.096	0.002
o-a	0.4187	0.1652	2.534	0.011
similar	0.8172	0.1698	4.814	< 0.001

With [a-a] as the baseline, each of the three vowel patterns that have [o] in them came out significantly more conducive to *ot*-responses than the baseline. An additional model that is identical to the one in (82) except for the specification of [o-a] as the baseline for *vowel* shows a significant difference between [a-a] and [o-a] ($p = .011$), but without a significant difference between [o-a] and either of [a-o] or [i-o] ($p > .1$).

Since the similarity of the final syllables of the novel items to the final syllables of existing *ot*-takers was seen to make a significant improvement, four other similarity measures were tested: (a) the initial syllable (one or two segments), (b) the initial two segments, (c) the initial three segments, and (d) the final two segments. Each of these four measures was encoded as a binary variable, following the procedure described for *similar* above. Then, each variable was added, one at a time, to the base model in (81). The first

three of these did not reach significance ($p > .1$), and their addition to the model was found unjustified by an ANOVA model comparison ($\chi^2(1) > .1$). The similarity of the final two segments did reach significance ($p < .005$) and improved the model significantly ($\chi^2(1) < .005$), but not nearly as much as the similarity of the final three segments did. I conclude that the final syllable, or last three segments, offered the best measure of similarity for the current study.

Finally, the effect of final consonants was tested by adding an unordered 13-level fixed-effect *consonant* variable to the analysis in (82). None of the levels reached significance, and overall, the addition of *consonant* did not improve the model, as determined by an ANOVA model comparison ($\chi^2(1) > .1$)

In conclusion, the vowel pattern [a-a], which has no [o] in it, produced a rate of *ot*-reponses that was significantly lower than patterns with [o] in them. The vowel pattern [o-a], with its non-final [o], did not come out significantly different from the [o]-final patterns.

(83) Nonce words and the percent of *ot*-plurals chosen for them⁹

a-a		o-a		a-o		i-o	
sagáf	9%	donáf	30%	zaróf	25%	idóf	26%
takáv	23%	foláv	25%	davóv	32%	xizóv	25%
kalám	32%	sotám	38%	gafóm	32%	dimóm	21%
garád	38%	opád	26%	kanód	55%	nidód	53%
pasás	34%	xodás	19%	barós	23%	migós	25%
gaváz	9%	nokáz	21%	kamóz	38%	rizóz	49%
banác	21%	motác	38%	pacóc	40%	lixóc	43%
daláf	28%	rokáf	26%	tanóf	32%	biyóf	28%
pafáf	43%	kováf	13%	bakóf	23%	giróf	13%
zavák	17%	losák	42%	sakók	32%	fibók	11%
cagág	38%	fonág	28%	baróg	30%	ricóg	30%
bazáx	21%	sováx	21%	fadóax	47%	lifóax	40%
fanál	28%	gomál	28%	calól	25%	zihól	32%
dagár	19%	zovár	45%	galór	32%	cikór	49%

3.3.4 Discussion

Hebrew speakers productively extend the effect that a stem [o] has on the choice of the plural allomorph from their lexicon to novel nouns. In the lexicon, a stem-final [o] is more conducive to choosing *-ot* than a non-stem-final [o], which in turn is more conducive to choosing *-ot* than a stem that lacks [o] completely. In the experiment, speakers reliably reproduced the difference between the presence and absence of [o], but not the [o]'s

⁹In the plural forms, the initial [a] was deleted for [aa] and [ao] nouns, e.g. the plurals offered for *sagáf* were *sgaf-ím* and *sgaf-ót*. This was done in order to reduce the likelihood that these nouns would be interpreted as referring to humans, and thus skewing the responses towards *-im*. In the lexicon, the deletion or retention of the [a] is an idiosyncratic property of roots, but a retained [a] correlates well with animacy (Bat-El 2008b).

location. Whether speakers replicate the lexical trend of the [o]'s location is a matter for further experimentation.

The vowel effect in the experiment was only detected reliably when the similarity of the test items to actual *ot*-takers was taken into account - specifically, what mattered most was the similarity of the final syllable. In designing the stimuli in (83), I made sure that overall they didn't resemble real native nouns of Hebrew too closely. An impressionistic inspection of the results in (83), however, lead me to believe that novel nouns that share their final syllable with real *ot*-takers got a high rate of *-ot* responses, regardless of their vowel. For example, the novel noun *cagág*, which has no [o] in it, got more *-ot* responses than most nouns that do have [o], and I attribute that to the existence of the real noun *gág* ~ *gag-ót* 'roof'. The logistic regression model in (82) strongly confirmed this hypothesis. Other measures of similarity that were tested were shown to be either less useful or completely insignificant.

Berent, Pinker & Shimron (1999, 2002) report a series of experiments similar to the one I present here. They gave participants novel nouns, presented orthographically, and asked the participants to write a plural form for them. The novel nouns were chosen so as to control for their similarity to real *im*-takers and *ot*-takers, and they found that novel nouns that are similar to existing *ot*-takers elicited a higher rate of choosing *-ot*.

Berent et al. (1999, 2002) controlled for the degree of similarity of their novel items to actual items by consistently varying the number of changed features, but not by making the change in a consistent phonological position. They define three levels of similarity between novel items and real items: (a) "similar", which involves changing one feature on one segment that is not a place feature – usually a change of [voice], lateral (r vs. l), or anterior (s vs. ʃ), (b) "moderate", which involves a bigger change of one segment – usually a change of place of articulation and some other feature, and (c) "dissimilar", which involves a change in all of the consonants of the root. In the majority of cases, the "similar" and "moderate" changes altered the second syllable of the root (69% of the stimuli in in

experiment 1 of Berent et al. 1999, and 50% of the stimuli in experiments 1 and 2 in Berent et al. 2002). Yet, with a modified second syllable in more than half the stimuli, all three experiments found a significant effect of similarity to real *ot*-takers. This contrasts with the results of this study, which found the similarity effect to be strong with an unmodified second syllable.

I conclude that similarity between novel items and existing items has a clear effect on speakers' behavior, and yet the exact definition of this similarity is far from clear. For instance, the difference between the "similar" and "moderate" conditions reached significance in Berent et al. (1999) but not in Berent et al. (2002). What effects the exact degrees and locations of changes may have is still largely unanswered.

To summarize, two robust effects emerge from the current study and from Berent et al. (1999, 2002). The first is the presence of [o] in the root, which elicited a significantly higher number of *-ot* responses than roots without [o] in them. The location of the [o] in the root was not shown to have a significant effect on the speakers' responses, and it is hoped that further experimentation will be able to show this effect. The second is a similarity effect, where items that are similar to existing *ot*-takers elicited significantly more *-ot* responses than items that are not. The exact formulation of the similarity effect, however, is elusive, and would require further research.

3.4 Using markedness constraints to learn lexical trends

The lexicon study presented in §3.2 and the experimental results in §3.3 show that having [o] in the root is conducive to choosing the plural *-ot*. Additionally, in the lexicon, an [o] in the final syllable is more conducive to *-ot* than a non-final [o], although this effect was regrettably not found in the current study. In this section, I offer an analysis of this correlation in terms of markedness constraints.

The analysis is based on Optimality Theory (Prince & Smolensky 1993/2004) with the Recursive Constraint Demotion algorithm (RCD, Tesar & Smolensky 1998, 2000; Tesar

1998; Prince 2002), augmented with a mechanism of constraint cloning (Pater 2006, 2008b; Becker 2007). Cloning allows the speaker to keep track of lexical trends and build their relative strength into the grammar.

The appropriateness of using markedness constraints will be simply assumed in this section, but it is discussed and motivated empirically in §3.5, using results from an artificial language experiment.

3.4.1 Analysis

The preference of roots that have [o] for taking *-ot* is interpreted as a requirement for licensing unstressed [o]'s. In native nouns, stress shows up on the root in unsuffixed forms (e.g. *xalón* 'window'), but stress moves to the right in suffixed forms, such as the plural (e.g. *xalon-ót* 'windows'). In the plural, then, the root's [o] surfaces unstressed, where it requires licensing.

Limiting [o] (and other non-high round vowels) to prominent positions is quite common in the world languages. Many languages are known to limit [o] to the stressed syllable, as in Russian *dóm-a* ~ *dám-áx* 'at home(s)'¹⁰. Similar restrictions apply in Portuguese and elsewhere.

Other languages require [o] to be licensed by the word-initial syllable. Turkish native nouns, for instance, allow [o] only in the first syllable of the word. Shona allows [o] in the word-initial syllable, and more interestingly, an initial [o] can license an [o] later in the word (Beckman 1997; Hayes & Wilson 2008).

In the analysis proposed here, Hebrew is like Shona, but with stress: In Hebrew, [o] must be stressed, but a stressed [o] allows [o] to appear elsewhere in the word. A similar licensing effect is seen with High vowels in several Romance languages (see §3.4.4). The licensing of [o] is not a categorical restriction in Hebrew, as unstressed [o]'s are tolerated.

¹⁰In standard American English, and other dialects, [o] can be unstressed ('piano', 'fellow') word-finally, but in some dialects, especially in the South, unstressed [o] is not allowed ('piana', 'fella'). This restriction on [o] in English, however, is just a part of a wider ban on unstressed full vowels in these dialects.

The licensing effect emerges when selecting *-ot* allows its stressed [o] to license the unstressed [o] in a root via auto-segmental linking.

Regular nouns (84a) allow [o] to surface unlicensed in the plural. For *ot*-takers that have an [o] in the root-final syllable (84b), the [o] is licensed directly by stress in the singular, and by being associated with the stressed syllable in the plural. As for *ot*-takers that have a non-final [o] (84c), the [o] surfaces faithfully in the singular, just like the [o] in *alon-ím*, but it is licensed across the [a] in the plural.

(84)		Singular	Plural	
a.	Regular	a l ó n	a l o n - í m	‘oak tree’
			$\begin{array}{c} \left[\begin{array}{l} -\text{high} \\ +\text{back} \end{array} \right] \left[\begin{array}{l} +\text{high} \\ -\text{back} \end{array} \right] \\ \quad \\ \text{a} \quad \text{o} \end{array}$	
b.	Irregular	x a l ó n	x a l o n - ó t	‘window’
			$\begin{array}{c} \left[\begin{array}{l} -\text{high} \\ +\text{back} \end{array} \right] \\ \diagup \quad \\ \text{a} \quad \text{o} \end{array}$	
c.	Irregular	o l á m	o l a m - ó t	‘world’
			$\begin{array}{c} \left[\begin{array}{l} -\text{high} \\ +\text{back} \end{array} \right] \\ \diagup \quad \\ \text{o} \quad \text{a} \end{array}$	

This diagram in (84c) shows the licensing of the unstressed [o] in the root by the stressed [o] of the plural affix, skipping the intervening [a]. Alternatively, the [a] could be associated with the licensed features, and thus eliminate the skipping, since [a] is compatible with $[-\text{high}]$ and $[\text{+back}]$ specifications. Licensing a marked vowel non-locally across another vowel is attested in other languages, as discussed in Hualde (1989); Walker (2006). In the Lena Bable dialect of Spanish, the $[\text{+high}]$ feature of a word-final vowel must be licensed by the stressed vowel, skipping any intervening vowels (e.g. /trweban+u/ → [trwiban-u]). The treatment of intervening vowels, in Hebrew and cross-linguistically, is discussed in further detail in §3.4.4.

As discussed in §3.2 above, it is not clear which vowels may intervene when *-ot* is selected non-locally. The current study is not particularly committed to this question, and the analysis will go through with just minor modifications if the set of intervening vowels turns out to include just [a] or a larger set.

Among nouns that have [o] in their roots, only those that surface stressless in the plural, i.e. native nouns, could benefit from taking *-ot* in the plural. Loanwords, i.e. nouns that keep their stress on the root, would not benefit from taking *-ot*, since there is no [o] that needs licensing, and indeed loanwords do not allow exceptional *ot*-taking.

In terms of Optimality Theory (Prince & Smolensky 1993/2004), taking *-im* or *-ot* can be fruitfully understood as responding to the satisfaction of different markedness constraints.

The requirement for the masculine *-im* on masculine nouns is enforced by a morphological constraint, ϕ -MATCH, which demands gender features to match in poly-morphemic words. For an *im*-taker like *alón* (85), ϕ -MATCH outranks the constraint LOCAL(o), which requires local licensing of [o]:

(85)

$\text{alón}_{\text{MASC}} + \{\text{im}_{\text{MASC}}, \text{ot}_{\text{FEM}}\}$	ϕ -MATCH	LOCAL(o)
a. \rightarrow <i>alón-ím</i>		*
b. <i>alón-ót</i>	*!	

Conversely, an *ot*-taker like *xalón* requires a high-ranking LOCAL(o):

(86)

$xalon_{MASC} + \{im_{MASC}, ot_{FEM}\}$	LOCAL(o)	ϕ -MATCH
a. $xalon\text{-}\acute{i}m$	*!	
b. $\rightarrow xalon\text{-}\acute{o}t$		*

The constraints that enforce [o]-licensing are defined below (87-88). The constraints are modeled after Hayes & Londe (2006), who find a similar case of exceptional action at a distance in Hungarian vowel harmony. See §3.4.4 below for a discussion of other possible definitions of the constraints.

(87) LOCAL(o)

An [o] must be licensed by virtue of being stressed, or by virtue of being auto-segmentally associated to a stressed [o] in an adjacent syllable.

(88) DISTAL(o)

An [o] must be licensed by virtue of being stressed, or by virtue of being auto-segmentally associated to some stressed [o].

When the root [o] is farther away from the stressed syllable, LOCAL(o) is not satisfied with either plural affix, but DISTAL(o) prefers that the [o] be licensed across the intervening vowel. In (89), DISTAL(o) outranks ϕ -MATCH, and LOCAL(o) is unranked with respect to either of the other two constraints.

(89)

$olam_{MASC} + \{im_{MASC}, ot_{FEM}\}$	DISTAL(o)	LOCAL(o)	ϕ -MATCH
a. $olam\text{-}\acute{i}m$	*!	*	
b. $\rightarrow olam\text{-}\acute{o}t$		*	*

With an *im*-taker that has a non-final [o], it is the ranking of ϕ -MATCH over DISTAL(o) that assures the correct result (90). Just like in (89), the ranking of LOCAL(o) is immaterial.

(90)

$\text{olar}_{\text{MASC}} + \{\text{im}_{\text{MASC}}, \text{ot}_{\text{FEM}}\}$	ϕ -MATCH	DISTAL(o)	LOCAL(o)
a. \rightarrow olar-ím		*	*
b. olar-ót	*!		*

In a small number of nouns, LOCAL(o) and/or DISTAL(o) force the change of a root [o] to [u], as in (91a). Ranking ϕ -MATCH and one of LOCAL(o) or DISTAL(o) over IDENT(Hi) would give rise to the vowel alternation, as shown in (92). The number of words involved, however, is very small: It's the nouns *xok* 'law', *tof* 'drum' and *dov* 'bear', the quantifiers *kol* 'all' and *rov* 'most', and a dozen adjectives. There are only two words that display an *o* ~ *a* alternation: *rof* 'head' and *yom* 'day' (91b).

- (91) a. xók xuk-ím 'law'
 b. róʃ raf-ím 'head'

(92)

$\text{xok}_{\text{MASC}} + \{\text{im}_{\text{MASC}}, \text{ot}_{\text{FEM}}\}$	ϕ -MATCH	DISTAL(o)	LOCAL(o)	IDENT(Hi)
a. \rightarrow xuk-ím				*
b. xok-ím		*	*	
c. xok-ót	*!			

An additional effect that follows from the use of constraints that license [o] by the stressed syllable is the regularity of the plural affix selection in loanwords. In these words,

stress stays on the root¹¹, so any [o] in the stem would be equally licensed in the singular and the plural. The tableau in (93) shows the noun *blóg* ‘blog’, where the presence of the [o] cannot trigger selection of *-ot*, since LOCAL(o) is equally satisfied by either plural affix.

(93)

$\text{blóg}_{\text{MASC}} + \{\text{im}_{\text{MASC}}, \text{ot}_{\text{FEM}}\}$	ϕ -MATCH	LOCAL(o)
a. \rightarrow <i>blóg-im</i>		
b. <i>blóg-ot</i>	*!	

Similarly, if a loanword has an unstressed [o] in it, like *kétfop* ‘ketchup’, LOCAL(o) is equally unable to prefer one of the plural allomorphs over the other.

(94)

$\text{kétfop}_{\text{MASC}} + \{\text{im}_{\text{MASC}}, \text{ot}_{\text{FEM}}\}$	ϕ -MATCH	LOCAL(o)
a. \rightarrow <i>kétfop-im</i>		*
b. <i>kétfop-ot</i>	*!	**

The regular selection of *-ot* with feminine loanwords, as in *fukátfa* ~ *fukátf-ot* ‘focaccia’, does indeed introduce an unlicensed [o]. Since my analysis allows LOCAL(o) to dominate ϕ -MATCH for some nouns, one would expect that some feminine loanwords would choose *-im*, contrary to fact. However, recall that the selection of *-ot* in loanwords is not based on morpho-syntactic gender (i.e. the gender that is revealed by agreement on

¹¹If suffixation puts the stressed syllable more than three syllables away from the edge, the stress (optionally) shifts two syllables to the right (Bat-El 1993; Becker 2003), but never off the root. For example, the plural of *béybisiter* ‘male babysitter’ is either *béybisiter-im* or *beybisiter-im*, but never **beybisiter-ím*. Similarly, the plural of *béybisiter-it* ‘female babysitter’ is either *béybisiter-iy-ot* or *beybisiter-iy-ot*, but never **beybisiter-íy-ot*.

adjectives and verbs), but rather on apparent morpho-phonological gender: All and only the nouns that *appear* to be feminine by virtue of having a feminine suffix on them take *-ot*, including masculine nouns that end in *-a*, such as *kolég-a* ‘(male) colleague’. I am assuming that some other constraint enforces this pattern, a constraint that categorically outranks both LOCAL(o) and ϕ -MATCH. I call this constraint μ -MATCH, as shown in (95).

(95)

fukátfa _{FEM} + {im _{MASC} , ot _{FEM} }	μ -MATCH	ϕ -MATCH	LOCAL(o)
a. fukátf-im	*!	*	
b. \leftarrow fukátf-ot			*

Returning to native masculine nouns now, there is still the problem of selecting *-ot* for those *ot*-takers that don’t have [o] in them, such as *fém* ~ *fem-ót* ‘name’. Since neither LOCAL(o) nor DISTAL(o) can help with selecting *-ot* in the absence of a root [o], some other mechanism must be involved.

I propose that *ot*-taking can be attributed to a constraint that doesn’t refer to the root vowel, but rather penalizes some aspect of the *-im* suffix itself, e.g. * \acute{o} /HIGH, which penalizes stressed high vowels (Kenstowicz 1997; de Lacy 2004). A constraint such as *LAB would work equally well – neither constraint is otherwise clearly active in the language¹².

¹²Arguably, both constraints are relevant for Hebrew phonology in general: * \acute{o} /HIGH could be used to derive the distribution of stressed vowels in segholates, which only allow non-high stressed vowels, producing alternations like the one in *kécev* ~ *kicb-í* ‘rhythm / rhythmic’. Self-conjunction of *LAB could account for the restrictions on the distribution of labials in roots.

(96)

$fem_{MASC} + \{im_{MASC}, ot_{FEM}\}$	* \acute{o} /HIGH	ϕ -MATCH	LOCAL(o)
a. $fem\text{-}\acute{i}m$	*!		
b. $\rightarrow fem\text{-}\acute{o}t$		*	

This use of * \acute{o} /HIGH, which attributes the selection of *-ot* to marked structure that happens to appear in the suffix *-im*, makes no reference to the phonological shape of the root. This is in line with the rest of the analysis, which assumes that any vowel other than [o] is inert with respect to plural allomorph selection.

In principle, the selection of *-ot* with nouns that don't have [o] in them could be done with a purely arbitrary diacritic, with no phonological substance at all. In the analysis proposed in (96) above, however, it is hard to see why the learner would fail to notice the preference that * \acute{o} /HIGH makes, if this constraint is indeed universal and available to the learner "for free".

I leave open the possibility that in some cases, learners are left with no phonological mechanism for making the right choice in allomorph selection, and they are forced to simply list the exceptional affix-takers. Suppose that a constraint such as * \acute{o} /HIGH is unavailable to the speaker for some reason, making the observed form *fem- $\acute{o}t$* harmonically bounded, as in (97).

(97)

$fem_{MASC} + \{im_{MASC}, ot_{FEM}\}$	ϕ -MATCH	LOCAL(o)
a. $fem\text{-}\acute{i}m$		
b. $\ominus fem\text{-}\acute{o}t$	*!	

Faced with a situation as in (97), the speaker will simply list the form *fem-ót* in their lexicon (cf. a similar proposal in Tessier 2008). Once listed in the lexicon, this form will have no effect on the grammar and thus no effect on the treatment of novel nouns.

To summarize the point so far: Most masculine native nouns in Hebrew select the plural *-im* due to a high ranking morphological constraint, ϕ -MATCH. Two phonological constraints, LOCAL(o) and DISTAL(o), prefer the selection of *-ot* when there is an [o] in the final or non-final syllable of the root, respectively. Different Hebrew nouns are subject to different constraint rankings: Nouns that take *-im* are associated with a high-ranking ϕ -MATCH, while nouns with [o] in them that take *-ot* are associated with a high-ranking LOCAL(o) or DISTAL(o). Finally, *ot*-takers that don't have [o] in them are associated with a different high-ranking phonological constraint, * \acute{o} /HIGH.

3.4.2 Ranking conflicts trigger the formation of generalizations¹³

I have shown that in the lexicon, selection of *-ot* is most common with nouns that have [o] in their final syllable, less common with nouns that have [o] in their penultimate syllable, and least common with nouns that don't have [o] at all. Speakers replicated the effect that the presence of the [o] had, and it is hoped that future work will demonstrate that speakers replicate the effect of the location of the [o].

I proposed an analysis that relies on the idea that different words of the language are subject to different grammars: Masculine nouns that take *-im* are associated with a high ranking of a morphological constraint that requires the masculine affix on masculine nouns, while those masculine nouns that take *-ot* are associated with highly ranked phonological constraints, such as constraints that require a root [o] to be licensed.

The analysis must now be completed with a mechanism that allows speakers to do three things: (a) learn the correct affix to choose with existing nouns, (b) learn the relative

¹³This section introduces the basic mechanism of constraint cloning, as applied to the Hebrew data. The cloning mechanism is also described in chapter 2, and it is explored formally in chapter 4.

frequency of *ot*-taking in the lexicon relative to the presence and position of a root [o], and (c) project the frequencies of the lexicon onto novel items. Such a mechanism is outlined here, and in chapter 2; the full proposal is detailed in chapter 4.

The analysis relies on learners’ ability to identify cases where there is no single grammar that can apply successfully to all of the words of their language. The Recursive Constraint Demotion algorithm (RCD, Tesar & Smolensky 1998, 2000; Tesar 1998; Prince 2002) allows language learners to collect ranking arguments from different lexical items and find conflicting rankings.

The use of RCD is most clearly illustrated with comparative tableaux (Prince 2002), where pairs of winners and losers are compared as to how they fare on various constraints. For example, the plural form of *xalón* ‘window’ is *xalon-ót*, so the learner has to make sure that *xalon-ót* wins over the intended loser *xalon-ím*. The constraint ϕ -MATCH prefers *xalon-ím*, while the constraint LOCAL(o) prefers *xalon-ót*, so if *xalon-ót* is to win, the constraint that prefers the winner must be ranked over the constraint that prefers the loser. This situation is shown with the winner-loser pair in (98a), with LOCAL(o) assigning a W (“Winner preferring”) to it and ϕ -MATCH assigning an L (“Loser preferring”).

Similarly, the winner-loser pair in (98b) shows the *im*-taker *alón* ‘oak tree’, which requires the ranking of ϕ -MATCH over LOCAL(o).

(98)

	LOCAL(o)	ϕ -MATCH
a. <i>xalon-ót</i> \succ * <i>xalon-ím</i>	W	L
b. <i>alon-ím</i> \succ * <i>alon-ót</i>	L	W

Given a comparative tableau, the learner can extract a constraint ranking from it by finding columns that have only W’s or empty cells in them, and *installing* the constraints in those columns. Installing a constraint means that it is added to the constraint ranking

below any constraints that are already in it, and any winner-loser pairs it assigns a W to are removed from the tableau. Installing constraints continues until all winner-loser pairs are removed. In the case of (98), however, there are no constraints to install, since all the columns have both W's and L's in them.

The solution to this situation was offered by Pater (2006, 2008b), who suggested that a constraint can be *cloned* to solve the inconsistent ranking of the constraints. Cloning a constraint means that the learner makes two copies, or clones, of the constraint, and makes both clones lexically-specific. Clones are lexically-specific in the sense that they apply only to the list of lexical items that are associated with them. When a constraint is cloned, every lexical item it assigns a W to is associated with one clone, and every lexical item it assigns an L to is associated with the other clone¹⁴.

In the case at hand, suppose the learner decided to clone LOCAL(o). One clone would be associated with *xalón*, and the other would be associated with *alón* (99).

(99)

	LOCAL(o) _{xalon}	ϕ-MATCH	LOCAL(o) _{alon}
a. xalon-ót > *xalon-ím	W	L	
b. alon-ím > *alon-ót		W	L

Now there is a column that only has W's in it, and there is a constraint to install: LOCAL(o)_{xalon}. Once installed, the first winner-loser pair in (99) is removed, which leaves the column of ϕ-MATCH with only W's in it. ϕ-MATCH is installed and added to the constraint ranking below LOCAL(o)_{xalon}, and the second and last winner-loser pair in (99) is removed. The remaining constraint, LOCAL(o)_{alon} is added to the ranking below ϕ-MATCH.

The result is the grammar in (100), where there are no longer any ranking conflicts.

¹⁴This last point is a departure from Pater (2006, 2008b), see §1.1.4 for discussion.

(100) LOCAL(o)_{xalon} ≫ φ-MATCH ≫ LOCAL(o)_{alon}

As the learner encounters more nouns with [o] in their final syllable, the conflict between φ-MATCH and LOCAL(o) will cause more nouns to be associated with one of the clones of LOCAL(o). Nouns that take *-ot* will be associated with the higher ranking clone, and nouns that take *-im* will be associated with the lower ranking clone.

(101) LOCAL(o)_{xalon, makom, ...} ≫ φ-MATCH ≫ LOCAL(o)_{alon, faon, pagof, ...}

Since nouns like *fém*, which don't have [o] in them, are neither preferred nor dispreferred by LOCAL(o), they will not be assigned a W or an L by LOCAL(o), and thus will not be associated with either clone.

Of the nouns with [o] in their final syllable in Bolozky & Becker (2006), 146 are *ot*-takers and 377 are *im*-takers. A speaker who learns all of them will end up with a grammar such as the one in (102)¹⁵.

(102) LOCAL(o)_{146 items} ≫ φ-MATCH ≫ LOCAL(o)_{377 items}

The grammar in (102) achieves two goals at once: It encodes the behavior of the existing nouns of Hebrew by associating them with one of the clones of LOCAL(o), and since it has a list of *im*-takers and a list of *ot*-taker, the grammar lets the learner discover the proportion of *ot*-takers among the the nouns that have [o] in them. This information, in turn, can be used to project the relative number of *im*-takers and *ot*-takers onto novel nouns.

Once LOCAL(o) is cloned, and each clone is made lexically-specific, there is no longer a general LOCAL(o) constraint that can apply to novel items. When faced with a novel noun that has [o] in its final syllable, the speaker must decide which clone of LOCAL(o) to associate it with, and this decision will be influenced by the number of items associated with each clone. Since 27.9% of the nouns associated with clones of LOCAL(o) are associated with its higher ranking clone, the learner will have a chance of 27.9% of choosing *-ot*.

¹⁵This picture is somewhat simplified, since the set of *ot*-takers with a final [o] is not homogeneous, as described in §3.2.

(103) LOCAL(o)_{27.9%} ≫ φ-MATCH ≫ LOCAL(o)_{72.1%}

There is another, perhaps simpler way of projecting the relative strength of the two clones of LOCAL(o) onto novel items. Given a novel item, the speaker can decide that the behavior of the novel item mimics the behavior of some given noun, chosen at random from the lists of nouns associated with the clones of LOCAL(o). If such a word is chosen at random, there is a 27.9% chance of that word being associated with the higher ranking clone, thus giving the novel item a 27.9% chance of being an *ot*-taker. Either way, the result is the same: The relative strength of the trend created by the existing nouns of the language is built into the grammar, and then can be projected onto novel items.

The use of markedness constraints in this analysis builds into the grammar only those generalizations that can be expressed with plausible universal constraints, such as constraints on the licensing of [o], which is seen cross-linguistically. The lexicon may contain further generalizations that cannot be expressed in terms of plausible universal constraints, such as the fact that among the nouns that have an [o] in their final syllable, *ot*-takers with [i] in their penultimate syllable (e.g. *cinór* ‘tube’) are more common than those with [a] in their penultimate syllable (e.g. *xalón*). In the experiment presented in §3.3, speakers did not project this trend onto novel nouns, suggesting that they have never learned it. If only root [o]’s are relevant for taking *-ot*, it is expected that other vowels would be ignored. Note that the speaker cannot simply ignore any vowel that is in the penultimate syllable, since having an [o] in the penult is conducive to more *-ot*.

To summarize, this section presented a mechanism that detects inconsistent ranking arguments between lexical items, and resolves the inconsistency by cloning a constraint. Once a constraint is cloned, lexical items are associated with different clones, assuring that they surface as intended. Additionally, the difference in size between the lists of associated lexical items is available to the learner, so that the learner can project the relative strength of lexical trends onto novel items.

3.4.3 Learning specific patterns first

The previous section took on the analysis of nouns that have [o] in their final syllable, showing how speakers can learn that these nouns have two possible behaviors (*im*-taking vs. *ot*-taking), and use constraint cloning to keep track of the nouns that behave in each way. This section shows how the mechanism is applied more generally to nouns that have [o] not only in their final syllable, but anywhere in their root.

The analysis offered here has one constraint that prefers *im*-taking, ϕ -MATCH, no matter what the shape of the noun is. Three constraints prefer *ot*-taking: * \acute{o} /HIGH, which affects nouns of any shape; DISTAL(o), which affects nouns that have [o] anywhere in the stem; and LOCAL(o), which only affects nouns that have [o] in their final syllable.

This analysis organizes nouns into three sets: Nouns that have [o] in their final syllable are the most specific set, identified by LOCAL(o); nouns that have [o] in their penult are found by using DISTAL(o) to identify the set of nouns that have [o] anywhere in the stem, and taking away the nouns with final [o]; and finally nouns that don't have an [o] at all are found by taking all nouns that are affected by * \acute{o} /HIGH and removing the nouns that were found using LOCAL(o) and DISTAL(o).

This ordering that the analysis imposes on the data means that the learner has to follow it in order to discover the generalizations correctly. This can be done by ensuring that LOCAL(o) is cloned first, associating all nouns with a final [o] with its clones, and leaving other nouns unassociated. Then DISTAL(o) should be cloned, associating the nouns that have [o] in them that were left over by LOCAL(o). Finally, any nouns that would be left unassociated would be taken care of by * \acute{o} /HIGH.

To ensure that the most specific constraint is cloned first, it suffices to choose the column that has the least number of W's and L's in it, but still contains at least one of each. As seen in (104), LOCAL(o) is singled out as the most specific constraint in the comparative tableau.

(104)

	LOCAL(o)	DISTAL(o)	ϕ -MATCH
a. xalon-ót \succ *xalon-ím	W	W	L
b. alon-ím \succ *alon-ót	L	L	W
c. olam-ót \succ *olam-ím		W	L
d. olar-ím \succ *olar-ót		L	W

Simply cloning LOCAL(o), however, is not quite sufficient. As the comparative tableau in (105) shows, once LOCAL(o) is cloned, LOCAL(o)_{xalon} can be installed, removing the first winner-loser pair from the tableau. Since this again leaves the tableau with no constraints to install, another constraint will be cloned. Assuming DISTAL(o) is chosen for cloning, one of its clones will be associated with the item that DISTAL(o) assigns a W to, viz. *olám*, and the other clone will be associated with the two items that DISTAL(o) assigns a L to, viz. *alón* and *olár*.

(105)

	LOC-AL(o) _{xalon}	LOC-AL(o) _{alon}	DISTAL(o)	ϕ -MATCH
a. xalon-ót \succ *xalon-ím	W		W	L
b. alon-ím \succ *alon-ót		L	L	W
c. olam-ót \succ *olam-ím			W	L
d. olar-ím \succ *olar-ót			L	W

The problem here is that a single lexical item, *alon*, is “double-dipping”, i.e. its choice of *-ot* is listed with clones of two constraints. The grammar the learner would make from (105) is in (106).

(106) LOCAL(o)_{xalon} >> DISTAL(o)_{olam} >> ϕ -MATCH >>
 LOCAL(o)_{alon}, DISTAL(o)_{alon, olar}

While double-dipping doesn’t prevent the learner from successfully learning the real nouns of Hebrew, it makes the wrong prediction about speakers’ ability to project lexical statistics onto novel words. If DISTAL(o) has one clone that lists *ot*-takers that have a non-final [o], and another clone that lists *all* of the *im*-takers that have an [o] anywhere in the root, as in (106), speakers will underestimate the ability of non-final [o] to correlate with the selection of [o]. In the lexicon, 12 out of the 102 nouns that have the vowel pattern [a-o] are *ot*-takers, which makes their likelihood in the lexicon 11.8% (see 76 above). If these 12 *ot*-takers are weighed against all the *im*-takers that have an [o] in them, as in (106), their likelihood in the grammar would only be 5.2%. This goes contrary to the observation in §3.3 that speakers correctly reproduce the relative strength of lexical trends.

To prevent double-dipping, it is not enough to simply clone the most specific constraint available. The learner must also ignore (or “mask”) the matching W’s and L’s that are assigned by less-specific constraints once a more specific constraint is cloned. This is shown in (107), where the speaker cloned the most specific LOCAL(o) and also masked W’s and L’s that were assigned to items that are associated with the new clones.

(107)

	LOC- AL(o) _{xalon}	LOC- AL(o) _{alon}	DISTAL(o)	ϕ-MATCH
a. xalon-ót \succ *xalon-ím	W		⊗	L
b. alon-ím \succ *alon-ót		L	⊗	W
c. olam-ót \succ *olam-ím			W	L
d. olar-ím \succ *olar-ót			L	W

Recall that finding the most specific constraint to clone was done by finding the column that had the smallest number of W's and L's. After the most specific constraint is cloned, the learner searches for constraints that are more general, defined as constraints that assign a superset of the W's and L's that the cloned constraints assigns. The more general DISTAL(o) will be found this way, and W's and L's that belong to lexical items that are now associated with clones of LOCAL(o) are masked, or ignored for the purposes of cloning.

The installation of LOCAL(o)_{xalon} can be done either before or after the masking of the general W's and L's from the column of DISTAL(o). Once LOCAL(o)_{xalon} is installed, the first winner-loser pair can be removed. This leaves DISTAL(o) as the column with the least number of W's and L's, and it is cloned. Now, only *olám* and *olár* are correctly associated with clones of DISTAL(o). The resulting grammar in (108) correctly lists all and only nouns with [o] in their final syllable under clones of LOCAL(o), and all and only nouns with [o] in this non-final syllable under clones of DISTAL(o).

(108) LOCAL(o)_{xalon} \gg DISTAL(o)_{olam} \gg ϕ-MATCH \gg
 LOCAL(o)_{alon}, DISTAL(o)_{olar}

As the speaker learns the rest of the nouns of the language, the grammar in (108) will include an increasing number of lexical items, which in turn will let the speaker project their relative number onto novel items.

Nouns with no [o] in their stem are listed by * \acute{o} /HIGH once the nouns with [o] are taken care of. The comparative tableau in (109) shows all three kinds of nouns.

(109)

	LOCAL(o)	DISTAL(o)	* \acute{o} /HIGH	ϕ -MATCH
a. xalon- \acute{o} t \succ *xalon- \acute{i} m	W	W	W	L
b. alon- \acute{i} m \succ *alon- \acute{o} t	L	L	L	W
c. olam- \acute{o} t \succ *olam- \acute{i} m		W	W	L
d. olar- \acute{i} m \succ *olar- \acute{o} t		L	L	W
e. fem- \acute{o} t \succ *fed- \acute{i} m			W	L
f. fed- \acute{i} m \succ *fed- \acute{o} t			L	W

Once LOCAL(o) and DISTAL(o) are cloned, the column of * \acute{o} /HIGH will be left with only one W and one L at the bottom, due to the masking of W's and L's on general constraints. At that point, * \acute{o} /HIGH will be cloned, and its clones will be associated with nouns that don't have [o] in them. The complete grammar is the one in (110).

(110) LOCAL(o)_{xalon} \gg DISTAL(o)_{olam} \gg * \acute{o} /HIGH_{fem}
 \gg ϕ -MATCH \gg
 LOCAL(o)_{alon}, DISTAL(o)_{olar}, * \acute{o} /HIGH_{fed}

3.4.4 Vowel harmony and [o]-licensing

The analysis presented here used two constraints to enforce the licensing of [o] by a stressed syllable, locally and at a distance. This approach was inspired by Hayes & Londe (2006), who find a similar case of exceptional action at a distance in Hungarian vowel harmony. This approach, however, is not in line with most work on vowel harmony in Optimality Theory.

More commonly, vowel harmony is enforced by constraints that require features to be expressed over several segments, described in terms of auto-segmental spreading or by some other kind of structure, such as spans (McCarthy 2004) or domains (Cassimjee & Kisseberth 1998). An additional constraint, REALIZE, penalizes the expression of a feature on two non-adjacent segments, skipping a middle segment¹⁶ (Cassimjee & Kisseberth 1998). The Hebrew case can certainly be described in those terms, as in the following derivation of *olam-ót* (111).

(111)

$olam_{MASC} + \{im_{MASC}, ot_{FEM}\}$	HARMONY	REALIZE	ϕ -MATCH
a. <i>olam-ím</i>	*!		
b. <i>olam-ót</i>		*	*

The constraint HARMONY states that an [o] must be structurally associated with the stressed syllable, either by being auto-segmentally linked to a stressed [o] or by being in some other kind of structure that includes any [o] and the stressed vowel. The constraint REALIZE requires that all the elements in the domain of harmony realize the harmonic feature, i.e. it penalizes any non-[o] vowels inside the structure that imposes harmony.

¹⁶See below for further discussion of skipping.

Under this view, three kinds of Hebrew nouns can be distinguished: *ot*-takers with a non-final [o] will require HARMONY \gg REALIZE, ϕ -MATCH as in (111). Nouns with a final [o] only require HARMONY \gg ϕ -MATCH, since skipping isn't an issue when the stem [o] is adjacent to the stressed syllable. Finally, nouns with no [o] in them at all will only require *ó/HIGH \gg ϕ -MATCH, as in the other analysis. This situation is shown in (112).

(112)

	HARMONY	REALIZE	*ó/HIGH	ϕ -MATCH
a. xalon-ót \succ *xalon-ím	W		W	L
b. alon-ím \succ *alon-ót	L		L	W
c. olam-ót \succ *olam-ím	W	L	W	L
d. olar-ím \succ *olar-ót	L	W	L	W
e. jem-ót \succ *jed-ím			W	L
f. jed-ím \succ *jed-ót			L	W

In (112), the most specific constraint is REALIZE, and it singles out the nouns that have a non-final [o]. This contrasts with LOCAL(o), the most specific constraint in (109), which singled out the nouns with a final [o]. To the learner, this wouldn't matter, since either configuration allows a separation of the two kinds of nouns.

The more serious challenge in (112) is the mismatch in the preferences between REALIZE and HARMONY: REALIZE prefers *im*-taking, while HARMONY prefers *ot*-taking. This would prevent the learner from identifying HARMONY as more general than REALIZE, who would then fail to prevent double-dipping. In contrast, LOCAL(o) and

DISTAL(o) both prefer *ot*-taking, and thus make the identification of DISTAL(o) as more general a rather trivial matter.

If the definition of specific-general relationships could be extended to cover cases of constraints that make opposite choices, then the problem is solved, and the analysis in this section can proceed just like the analysis with LOCAL(o) and DISTAL(o). If this move turns out to be unwarranted, the solution will have to be found elsewhere.

The idea that vowels may be skipped by the harmonic feature is criticized by Ní Chiosáin & Padgett (2001) and Gafos (1999), among others, who claim that harmony processes never skip intervening elements. If this is right, then the long-distance licensing of [o] in Hebrew cannot be analyzed as a case of vowel harmony.

The auto-segmental and the domain/span-based approaches (111,112) assume that the harmonizing feature appears once in the output, and it associates with several segments. An alternative arises from the discussion of high vowel licensing in several dialects of Spanish (Hualde 1989), analyzed by Walker (2006) as a case of agreement by correspondence, i.e. the licensed feature appears twice in the output, not once, and thus intervening features are allowed. An analysis in terms of Walker (2006) is given in (113).

(113)

olam _{MASC} + {im _{MASC} , ot _{FEM} }	LICENSE(o)	INTEGRITY	ϕ-MATCH
a. olam-ím	*!		
b. olam-ót		*	*

In (113b), the features of the root [o] are pronounced twice, once on the root and once on the suffix. Since these two pronunciations express a single underlying set of features in two non-contiguous locations in the output, a violation of INTEGRITY is incurred. This analysis faces the same challenge that faces the analysis in (111,112): The constraint

that distinguishes local licensing from distal licensing, INTEGRITY, makes the opposite preferences with respect to the constraint that requires licensing, LICENSE(o).

Two empirical questions can weigh on the nature of the correct analysis of Hebrew. The first question is about the exact pronunciation of the interveners in words like *olam-ót*. Is the [a] that intervenes between the two [o]’s pronounced significantly differently from the [a] in *olar-ím*, where the root’s [o] is unlicensed? If the [a] is not pronounced differently, that would be evidence against the harmony-based approach (111,112).

The second empirical question is about the range of possible interveners. In the lexicon, only nouns with the vocalic pattern [o-a] are more conducive to *ot*-taking than nouns without [o] in them. It is not known how often speakers will choose *-ot* with nouns that have other interveners, e.g. [o-e], [o-i]. If [o] can be licensed across certain vowels but not others, this would be a problem for the agreement by correspondence account (113), which makes no prediction about the identity of the interveners.

Ultimately, the question is about the actual typology of vowel-vowel interactions cross-linguistically, which include vowel harmony and vowel licensing. The Hebrew case is a little different from most known cases, since it does not involve the selection of vowels only, but rather the selection of whole allomorphs that consist of active vowels and inert consonants. I conclude that the place of Hebrew in the typology of vowel-vowel interactions is not sufficiently well known to motivate a rejection of the analysis in terms of LOCAL(o) and DISTAL(o).

3.5 Product-orientedness in an artificial language¹⁷

The analysis of Hebrew plural allomorph selection proposed here relies on markedness constraints. The two allomorphs are available in the underlying representation of the plural

¹⁷This work was done in collaboration with Lena Fainleib (Tel Aviv University). We are grateful to Ram Frost, of the Hebrew University Psychology Department, for his generous help with various aspects of this work.

suffix, and they are allowed to compete in the phonology, with the assumption that choosing one of the allomorphs does not entail the deletion of the other, since only one can be chosen at a time (Mester 1994, Mascaró 1996, Anttila 1997, and many others). Simply pronouncing one of the allomorphs as it is in the UR, then, has no faithfulness cost, and therefore the choice is left to markedness constraints.

Markedness constraints only assess surface forms – in this case, the licensing of an unstressed [o] in the plural stem. These constraints have no access to the underlying representation of the root, nor to its pronunciation in the singular. It follows, then, that speakers are predicted to prefer the choice of *-ot* no matter whether the singular has an [o] in it or not.

This prediction cannot be tested with the real words of Hebrew, since every plural stem that has an [o] in it also has an [o] in the corresponding singular stem. The prediction can be tested, however, with an artificial language that is just like Hebrew, but allows plural stems that have [o] in them without a corresponding [o] in the singular. This section describes a pair of such artificial languages and how Hebrew speakers learned them.

Two languages were taught in this experiment. In both languages, singulars were plausible native nouns with an [o] or an [i] in their final syllable, and in the corresponding plural forms, [o]'s alternated with [i]'s and vice versa. The choice of the plural suffix agreed with the plural form in the “surface” language and with the singular form in the “deep” language (114). Only final vowels were varied, since they have the strongest effect on plural allomorph selection in real Hebrew.

(114)

“surface” language		“deep” language	
afív	afov-ót	afív	afov-ím
axís	axos-ót	axís	axos-ím
amíg	amog-ót	amíg	amog-ím
azíx	azox-ót	azíx	azox-ím
adíc	adoc-ót	adíc	adoc-ím
agóf	agif-ím	agóf	agif-ót
apóz	apiz-ím	apóz	apiz-ót
acók	acik-ím	acók	acik-ót
abóf	abif-ím	abóf	abif-ót
alód	alid-ím	alód	alid-ót

Note that the ten singulars are exactly identical in the two languages. The ten plural stems are also identical, but the choice of plural allomorph is different: In the “surface” language, plural stems with [o] select *-ot*, and plural stems with [i] select *-im*. In the “deep” language, it is not the plural stem, but rather the singular stem that selects *-ot* if it has [o] and *-im* if it has [i]. Another way to think about the “deep” language is to say that plural stems with [o] select *-im*, and plural stems with [i] select *-ot*.

After participants were trained and tested on one of the languages in (114), they were asked to generate plurals for the twenty nouns in (115). The responses were rated for their success in applying the vowel changes and the selection of the plural affix, where success was defined as the replacement of a singular [o] with a plural [i] and vice versa, and the selection of a plural affix according the generalization in the relevant language.

(115)	agív	amóv
	apís	adós
	axíg	afóg
	amíx	atóx
	alíc	afóc
	axíf	aróf
	aníz	ahóz
	afík	abók
	afíj	agóf
	azíd	apód

3.5.1 Materials

In the experiment, each participant was trained and tested on a language that contained 10 nouns, where each noun consisted of a random pairing of a sound and a concrete object, like a fruit or a household item. Once trained and tested, each participant was asked to generate plurals for 20 new nouns that they haven't encountered before. An additional noun was used in the beginning of the experiment for demonstration. In total, each participant encountered 31 nouns.

All the pictures of the objects used in the experiment were taken indoors, using daylight, with a Sony digital camera at 3.2 mega-pixels, then reduced to 400x300 pixels and saved as jpg files. The objects were placed on a neutral background, and positioned so as to make them as easy as possible to recognize. The objects were chosen such that their names in actual Hebrew were masculine *im*-takers. Items that were shown both in singletons and in pairs included the demonstration item, which was an almond, and the training items, which were a red onion, a potato, an apple, a persimmon, a strawberry, an artichoke, an orange, a green bell pepper, an eggplant, and a cucumber. In the plural generation phase, subjects saw the following items in pairs: pears, lemons, pomegranates, avocados, heads of

garlic, carrots, loquats, zucchinis, melons, dried apricots, uncooked steaks, beets, coconuts, prickly pears, jars of instant coffee, knives, mobile phones, power splitters, computer mouses, and bottles of olive oil. All of these were confirmed by several Israeli speakers of Hebrew to be easy to recognize and name.

The auditory materials included the singulars and plurals of the training materials shown in (114), the demonstration item, which was *axún* ~ *axun-ím*, and the plural generation items in (115). These were recorded by a male native speaker in a sound-attenuated booth onto a Macintosh computer at 44100 Hz, using Audacity. One wav file was created for each singular form, using Praat (Boersma & Weenink 2008). For each plural form, an additional file was created, which started with the singular, followed by .5 seconds of silence, followed by the singular again, another .5 seconds of silence, and then the plural form. All files were then converted to .mp3 format using the LAME encoder, version 3.97 (from <http://www.mp3dev.org/>).

3.5.2 Methods

The experiment was conducted on a web-based interface, using Firefox. Participants sat in a quiet room and wore headphones with a built-in microphone. They were recorded during the whole length of the experiment using Audacity on a single channel at 44,100 Hz. At the end of the experiment, the recording was saved as an mp3 file using the LAME encoder.

Each participant was randomly assigned to either the “surface” language or the “deep” language. Then, the training materials were generated by randomly combining the sounds from the relevant part of (114) with the ten training objects described above, to create 10 nouns that pair sound and meaning. Additionally, the twenty sounds from (115) were randomly combined with the twenty plural generation items described above, to create 20 nouns. The plural generation nouns were divided into two groups, each containing five nouns with [i] and five with [o].

Participants were told that they would learn a made-up language that is a new kind of Hebrew, and that it is written in Hebrew letters and pronounced with an Israeli accent. They were asked to memorize the words of the new language and try to figure out the regularity¹⁸ of the language.

The experiment was conducted as follows: training and testing on singulars (two rounds), training and testing on singulars and plurals (three rounds), plural generation for ten new nouns, testing on the singulars and plurals from the training phase, and plural generation for 10 additional new nouns. These phases are described more fully below.

Training started with singulars only: A picture of an object was displayed on the screen, and a sentence below it introduced the object as a masculine noun, by displaying the text in (116).

(116) Here's a nice_{MASC} _____
hine _____ nexmad_{MASC}

In parallel, the name of the object was played. The participant pressed a key to go to the next item. All 10 items were thus introduced in a random order, and then introduced again in a new random order. After each item was introduced twice, participants were tested on them. A picture of an item was displayed, along with the instructions in (117).

(117) Say in a clear voice, “this is a nice_{MASC} _____”, or “I don’t remember”
imru be-kol ram ve-barur, “ze _____ nexmad_{MASC}”, o “ani lo zoxer/et”

The whole procedure of training and testing was then repeated. Note that at this point, all participants were trained on the same materials, regardless of whether they were going to learn the “surface” language or the “deep” language.

After two rounds of training and testing on singulars, plurals were introduced. A picture of a pair of objects, e.g. two apples, was displayed, with the text in (118).

¹⁸The Hebrew word used was *xukiyút*, which depending on context, can mean ‘legality’, ‘well-formedness’, ‘regularity’, ‘pattern’, etc.

- (118) Here's one_{MASC} _____ on the right and one_{MASC} _____ on the left.
 Together, these are two_{MASC} nice_{MASC} _____.
 hine _____ exad_{MASC} mi-yamin ve _____ exad_{MASC} mi-smol.
 beyaxad, ele fney_{MASC} _____ nexmadim_{MASC}.

In parallel, the singular was played twice, followed by the plural. All 10 items were thus introduced in the singular and plural in a random order, and then introduced again in a new random order. After each item was introduced twice, participants were tested on them. A picture of a pair of items was displayed, along with the instructions in (119).

- (119) Say in a clear voice, “here there's one_{MASC} _____ on the right and one_{MASC} _____ on the left, and together these are two_{MASC} nice_{MASC} _____”.
 imru be-kol ram ve-barur, “yef po _____ exad_{MASC} mi-yamin ve _____
 exad_{MASC} mi-smol, vebeyaxad ele fney_{MASC} _____ nexmadim_{MASC}”.

The whole procedure of training and testing was repeated two more times, for a total of three rounds.

After the training and testing were over, participants were asked to generate plurals in the artificial language for nouns that they hadn't seen before, in two rounds. In the first round, five nouns with [o] and five with [i] were randomly selected from (115) and paired with meanings. A picture of one such noun was displayed with the instructions in (120), and in parallel, the noun's name was played twice.

- (120) Here's one_{MASC} _____ on the right and one_{MASC} _____ on the left. And what are they together? Say in a clear voice, “here's one_{MASC} _____ on the right and one_{MASC} _____ on the left, and together these are two_{MASC} nice_{MASC} _____”.
 Complete the sentence in a way that seems to you to be most compatible with the new kind of Hebrew you learned today.
 hine _____ exad_{MASC} mi-yamin ve _____ exad_{MASC} mi-smol. ve ma hem fney
 ele beyaxad? imru be-kol ram ve-barur “yef po _____ exad_{MASC} mi-yamin ve

_____ exad_{MASC} mi-smol, vebeyaxad ele jney_{MASC} _____ nexmadim_{MASC}”.
haflimu et ha-mifpat be-cura je-tifama laxem haxi matima la-ivrit ha-xadafa je-
lamadetem.

After the first round of plural generation, the ten nouns that speakers were trained and tested on appeared for another round of testing (no feedback was given at this point). This was done to make the participants mentally review the material they learned, reconsider any potentially unfruitful strategies, and hopefully make the next round of plural generation more consistent with the artificial language. After this round of testing, the second and last round of plural generation included the remaining ten nouns from (115), following the same procedure as in the first round of plural generation.

3.5.3 Participants

Data from a total of 60 participants was used in this study, 21 students at the Hebrew University and 39 students at the Tel Aviv University. All were born in Israel and were native speakers of Hebrew, without any self-reported hearing or vision difficulties. There were 24 males and 36 females, average age 23.4, age range 18–29¹⁹. For their time and effort, participants were either paid 20 shekels (around US\$6) or given course credit.

Four additional participants were excluded: One participant misunderstood the task, and most of the time supplied the names of objects in actual Hebrew instead of their names in the artificial language. Another participant failed to correctly repeat several of the names for novel items she had just heard, and performed badly on the other tasks, suggesting an unreported disorder of hearing or cognition. Two other participants were excluded because they did not produce any response for several items in the plural generation rounds.

¹⁹In pilots, participants over 30 were largely unable to perform minimal memorization, so 29 was chosen as a cut-off age for the current experiment.

3.5.4 Transcription and encoding

For each participant, two sections of the recording were transcribed: the testing rounds for the singulars, and the plural generations rounds. The recordings were matched up with the intended responses as they appeared on the server log, and written using a broad phonetic transcription.

For the testing rounds on the singulars, each response was given a score. A perfect score of 1 was given for a perfect recall of the expected form. Recalls with spirantized labials were also accepted, i.e. *avof* for *abof* or *afoz* for *apoz* were also given a score of 1. Pronunciations with an initial [h] (e.g. *habof* for *abof*) were also considered perfect and given a score of 1. Such pronunciations were considered to be within the normal range of variation in Hebrew, and compatible with perfect memorization. A score of .5 was given to any response that deviated from the expected form minimally, i.e. one feature on one segment (*amik* for *amig* or *apuz* for *apoz*) or by transposition of two consonants (*asix* for *axis*). A score of 0 was given to lack of recall or to any form that deviated from the expected form by more than one feature. This created a *memorization* score for each participant, on a scale of 0–20, quantifying their ability to correctly recall the singulars of the artificial languages. Since the singulars in both languages were the same, the memorization score is useful for controlling for any differences between the two groups.

The rounds of plural generation were broadly transcribed, and the plural forms were coded for their stem vowels and choice of plural affix. Most speakers produced full sentences, as indicated in (120), and a few just provided the singular and the plural without a frame sentence. No participant gave just plural forms without repeating the singulars. All participants repeated the singular forms they heard essentially perfectly, so no coding of the singulars was necessary. Speakers also had no trouble with reproducing the two consonants of the singular in the plural form, so no coding of that aspect was necessary either. Occasional initial [h]'s or the substitution of [e] for [a] in the initial syllable (*habok-ot* or *ebok-ot* for the expected *abok-ot*) were considered to be within the normal range of

variation for Hebrew, and were not taken to be errors. On each trial, a successful vowel mapping was defined as a production of an [o] in the singular and an [i] in the plural stem, or vice versa²⁰. A successful plural allomorph selection was defined as one that matches the intended generalization in the language the participant was taught, e.g. *-ot* for plurals stems with [o] in the “surface” language. A trial was categorized as successful if it had a successful vowel mapping and a successful choice of plural affix. With 20 trials each, participants were assigned a *generalization* score on a scale of 0–20.

3.5.5 Results

As expected, the “surface” language participants generalized the intended pattern better than the “deep” language participants. The table in (121) shows the proportion of trials where participants successfully changed a singular [o] to [i] and vice versa, and also selected the plural affix as expected in the language they were asked to learn. The “surface” group was equally successful in both conditions, whereas the “deep” group was worse at the change from singular [i] to plural [o] than at the change from [o] to [i].

(121)	“Surface” language	“Deep” language	difference
[o] → [i]	55%	42%	13%
[i] → [o]	54%	34%	20%
Total	54%	38%	16%

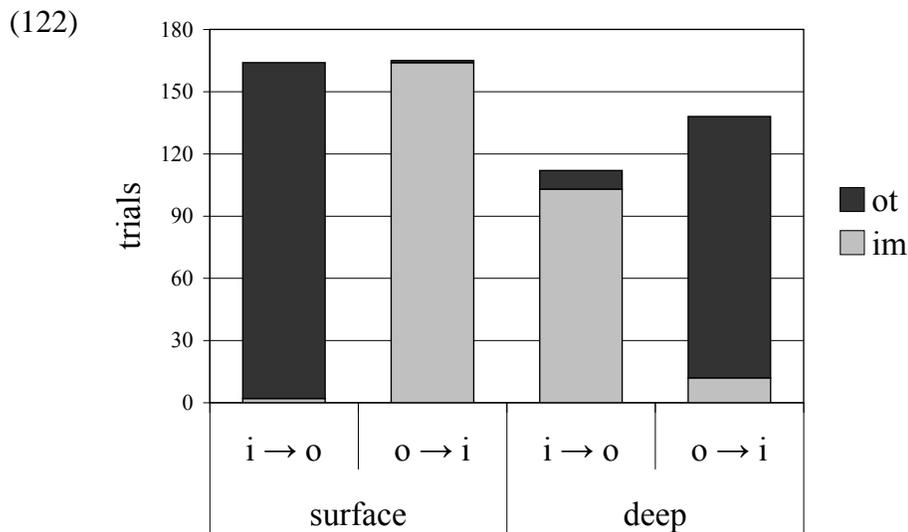
This section presents four aspects of the experimental results: (a) The participants in the “surface” language were more successful than the participants in the “deep” language, with a particular disadvantage for the “deep” group in the change from [i] to [o], shown in §3.5.5.1, (b) The two groups did not have significantly different memorization scores, and these scores correlate with the generalization scores only in the “deep” group, shown

²⁰The term “success” is used here in its statistical sense, which is judgement neutral, and simply refers to one of two possible outcomes in a binomial experiment. In this sense, a heart-attack can also be defined as a success.

in §3.5.5.2, (c) Speakers were biased towards using *-im*, proving that they were influenced by real Hebrew in the experiment, shown in §3.5.5.3, and (d) Misperception of the stimuli was marginal in both groups, and cannot account for the disadvantage of the “deep” group, shown in §3.5.5.4.

3.5.5.1 Generalization differences between the groups

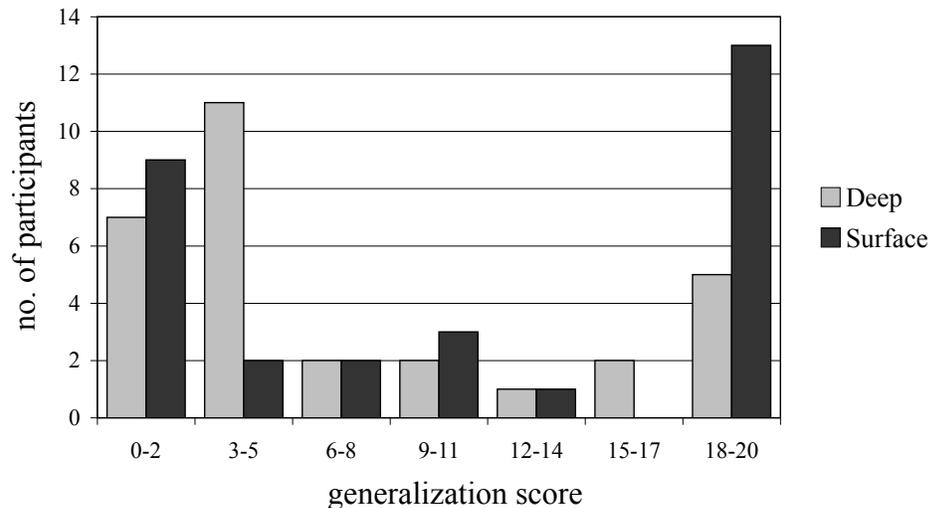
The “surface” language participants were on average more successful than the “deep” language participants at changing stem vowels from [i] to [o] and vice versa (54% vs. 38% of the trials). Given a successful stem vowel change, the “surface” language participants were better at selecting the appropriate plural affix (99% vs. 92%), as seen in (122). The “surface” language participants performed both of the required vowel changes equally well, whereas the “deep” language participants were less successful at changing [i] to [o] than [o] to [i].



A by-subject analysis shows that the generalization scores for the “surface” language participants ($n = 30$, $M = 10.9$) were on average higher than the scores for the “deep” language participants ($n = 30$, $M = 7.7$). The generalization scores were bi-modally distributed in both groups, as seen in (123), with 78% of the speakers scoring either 0–5 or 18–20. In other words, most participants either did very poorly or very well, with only

a few participants in the middle. The “surface” group is characterized by a large number of participants at the higher end of the scale, while the participants in the “deep” group are more heavily concentrated at the low end.

(123)



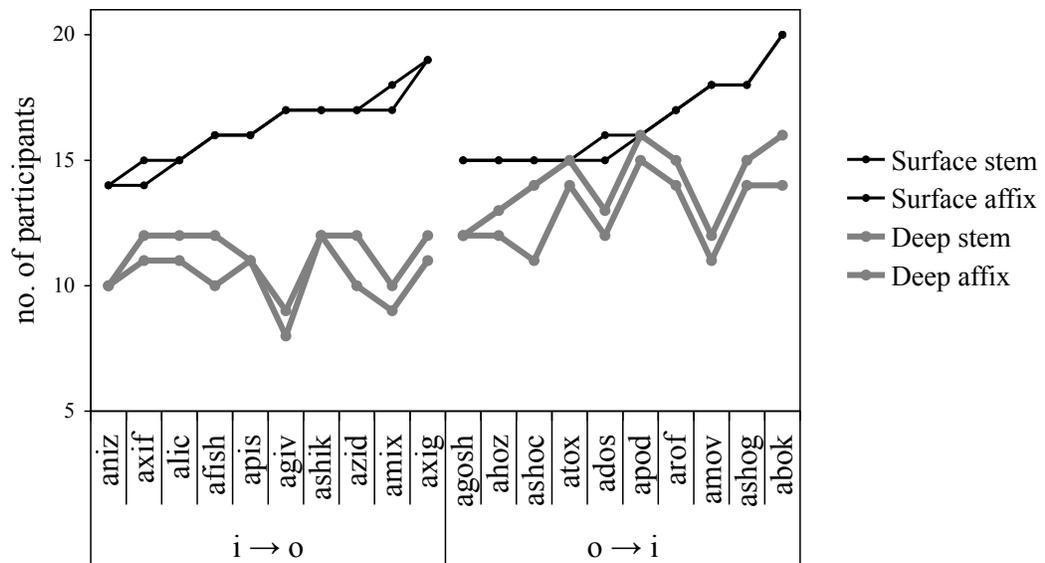
Since statistical tests that assume a normal distribution, such as the t-test, are out, the data was transformed using a cut-off point. Participants who scored above the cut-off point were given a score of 1, and the others were given a score of 0. The transformed results were compared with Fisher’s exact test. At a cut-off point of 17, the difference between the groups is significant (odds ratio 3.736, $p = .047$). The choice of 17 for the cut-off point comes from the distribution of the generalization scores in the “surface” group, where no participant scored in the 13–17 range, inclusive, suggesting that a score of 18 or above is the minimum for being considered a good generalizer.

The by-item analysis also shows a significant difference in the performance of the two groups. The chart in (124) shows the number of participants who successfully changed a stem vowel [i] in the singular to [o] in the plural and vice versa for each item, and the number of participants who successfully changed the stem vowel and also chose the expected plural affix for the language they learned. The differences between the groups are significant both for the stem vowel change only (paired t-test: $t(19) = 7.36$, $p < .001$) and

for the combined stem vowel change and affix selection (paired t-test: $t(19) = 9.25, p < .001$).

The chart in (124) also shows that given a successful stem vowel change, the “surface” language participants almost always selected the expected affix, as evidenced by the almost complete overlap of the two black lines (paired t-test: $t(19) = 1.83, p > .05$). The “deep” language participants, however, often changed the stem vowel successfully, but then failed to choose the expected affix, as evidence by the two distinct gray lines (paired t-test: $t(19) = 6.19, p < .001$).

(124) The number of participants who correctly changed stem vowels and chose appropriate plural suffixes, by item



A final thing to note about (124) is that the performance of the “surface” group participants is equally good on the items that require the change of [i] to [o] and those that require the change of [o] to [i] ($t(17.67) = .268, p > .1$), whereas the “deep” group participants performed more poorly on the items that required the change of [i] to [o] ($t(17.17) = 4.430, p < .001$).

The experimental results were analyzed with a mixed-effects logistic regression model in R (R Development Core Team 2007) using the *lmer* function of the LME4 package,

with *participant* and *item* as random effect variables. For each trial, the dependent binary variable *total success* was given a value of 1 for a successful change of stem vowel and a choice of the expected plural affix, and 0 otherwise. The predictor of interest was the unordered two-level factor *participant group* with the “surface” group as a base-line. In a simple model that had *participant group* as its only predictor, *participant group* did not reach significance. Adding another unordered two-level factor, *singular vowel*, with [i] as the baseline, and the group-vowel interaction factor, shown in (125), made a significant improvement to the model, as determined by an ANOVA model comparison ($\chi^2(1) < .01$).

(125)	Estimate	SE	<i>z</i>	<i>p</i>
(Intercept)	0.761	0.723	1.054	0.292
<i>group</i>	-1.859	1.010	-1.843	0.065
<i>vowel</i>	0.091	0.286	0.317	0.752
<i>group:vowel</i>	0.658	0.374	1.760	0.078

In (125), *participant group* has a negative coefficient, meaning that being in the “deep” group was negatively correlated with successful stem vowel change and affix selection. This effect, however, only approached the standard .05 significance level. Additionally, the interaction effect has a positive coefficient, meaning that in the “deep” group, the singular vowel [i] correlated with better success than the singular vowel [o], but this trend also only approached significance. The model stays essentially unchanged when validated with the *pvals.fnc* function from the *languageR* package (Baayen 2008). The rather modest p-values of this model are clearly due to the bi-modal distribution of the participants’ performance, as seen in (123), and evidenced in (125) by the large standard error of the *participant group* factor.

One way to bring the *participant group* variable into significance is to separate each participant’s responses to the [i] items and the [o] items, essentially nesting participants under vowels. This allows for the *participant group* effect to emerge by eliminating

the ability to observe any vowel effect. The new model, in (126), has *item* and *vowel:participant* as random effect variables and *participant group* as a fixed variable. In this model, being in the “deep” group is significantly less conducive to success than being in the “surface” group. The model stays essentially unchanged when validated with *pvals.fnc*.

(126)	Estimate	SE	<i>z</i>	<i>p</i>
(Intercept)	0.759	0.567	1.337	0.181
<i>group</i>	-1.880	0.794	-2.369	0.018

To summarize, the participants in the two groups behaved differently, with the “surface” language participants performing better than the “deep” language participants. Additionally, the “deep” language participants were less successful at changing singular [i] to [o] than vice versa. Statistical modeling of the difference between the groups with a logistic regression proved challenging, no doubt due to the bi-modal distribution of the data. While all the effects in the model in (125) were in the right direction, they only approached the .05 significance level. Finding a model that brings out the difference between the groups below the .05 level, as in (126), was done at the price of eliminating the vowel effect.

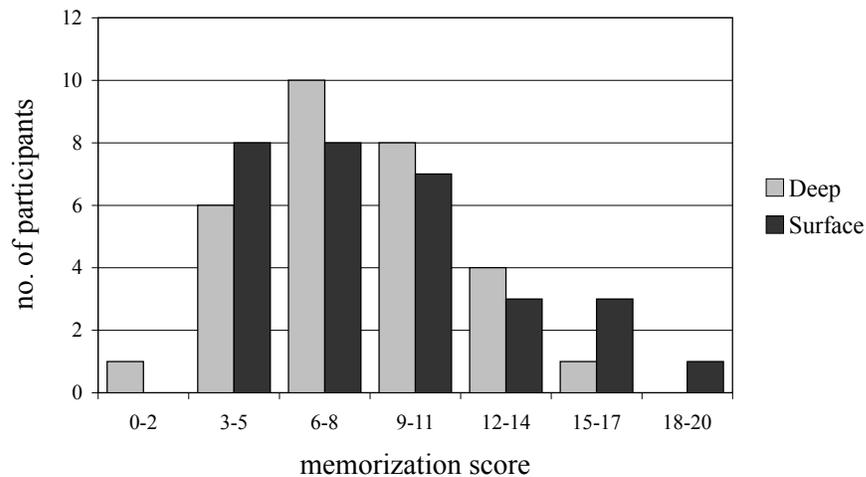
3.5.5.2 No memorization differences between the groups

Since the differences between the two languages are seen over two disjoint groups of people, it could be argued that the participants who learned the “surface” language just happened to be more alert or motivated. While participants were assigned to the two languages randomly to prevent such an effect, their memorization scores can also show that there were no clear differences between the groups in this respect.

The two groups can be compared on their ability to memorize the singular nouns in the initial part of the experiment, since participants in both groups performed the same task in that stage. As seen in (127), speakers’ scores on the memorization task are quite similar in both groups (“surface”: $n = 30$, $M = 9.12$, $SD = 4.23$; “deep”: $n = 30$, $M =$

8.48, $SD = 3.74$). The scores are approximately normally distributed in both groups²¹, and a t-test reveals that they are not significantly different ($t(57.14) = .61, p > .1$). We can safely conclude that there are no significant differences between the groups in the ability to memorize items (and by extension, in their general alertness and cognitive abilities), and that any differences between the groups in their generalization abilities, as seen in (123), mean that the two languages differ in their level of difficulty.

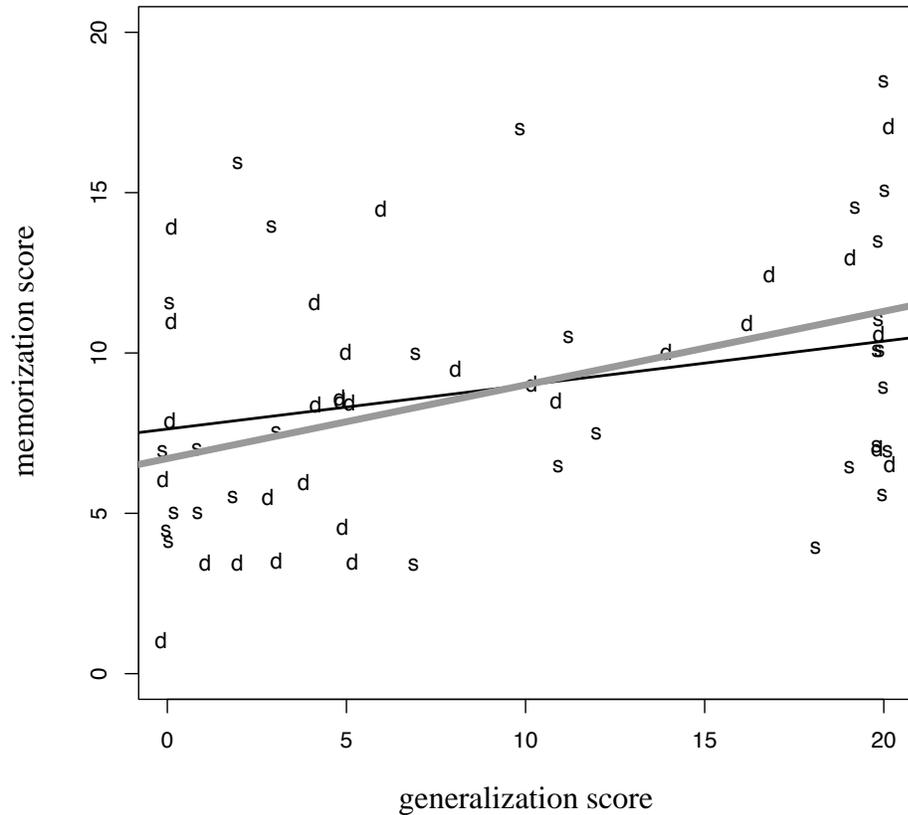
(127)



Interestingly, the correlation between the participants memorization scores and generalization scores is different in the two groups. In the chart in (128), “surface” language participants are marked with “s” and a black regression line, and the “deep” language participants are marked with “d” and a gray regression line. A little noise was added to reduce overlap between points.

²¹A Shapiro-Wilk normality test on each group reveals that the “surface” group is marginally normally distributed ($W = .92, p = .038$), and the “deep” group is solidly normally distributed ($W = .98, p > .1$).

(128)



For each group, a linear model was made using the *ols* function in R, with the generalization scores as a dependent variable and the memorization scores as a predictor. In the “surface” group, the generalization scores could not be predicted from the memorization score ($R^2 = .075$, sequential ANOVA: $F(1,28) = 6.49$, $p > .1$), but in the “deep” group, the correlation was significant ($R^2 = .188$, sequential ANOVA: $F(1,28) = 2.32$, $p < .05$).

This difference between the groups is not surprising. The “surface” language was predicted to be easy to learn, and indeed whether speakers have learned the language successfully or not had little to do with their relative alertness. The “deep” language was hard to learn, and participants had to pay attention to learn it successfully.

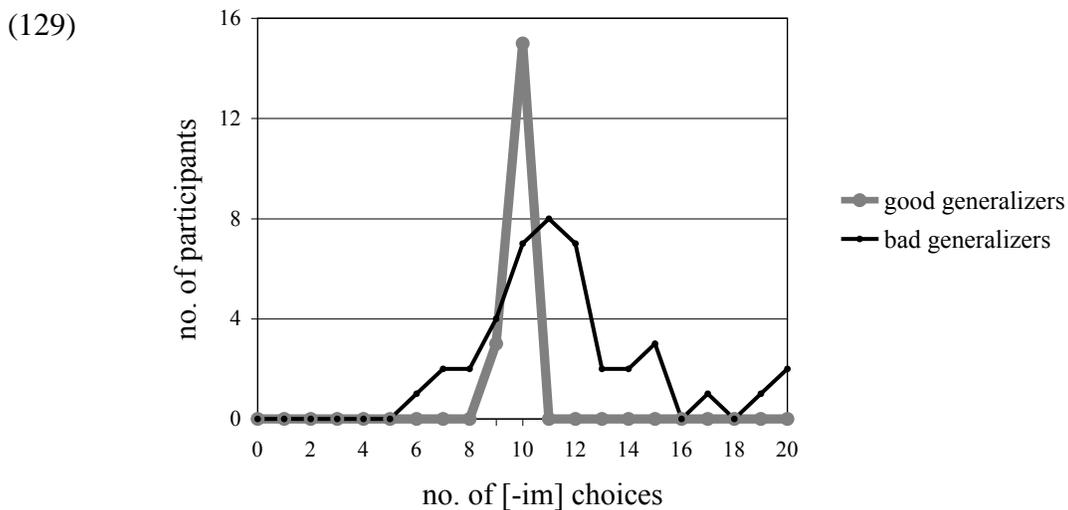
3.5.5.3 Bias towards *-im*

There is good reason to believe that participants in this experiment were influenced by their knowledge of real Hebrew in dealing with the two artificial languages.

The experimental stimuli were balanced between [-im] and [-ot], and indeed in order to get a perfect generalization score of 20, participants had to choose [-im] exactly 10 times, and thus show no preference for [-im] over [-ot].

However, the words of the artificial languages were presented as masculine nouns, as indicated by the adjectives and numerals that agreed with them in the various frame sentences. Since masculine nouns in real Hebrew are heavily biased towards [-im], the influence of real Hebrew would bias speakers towards [-im].

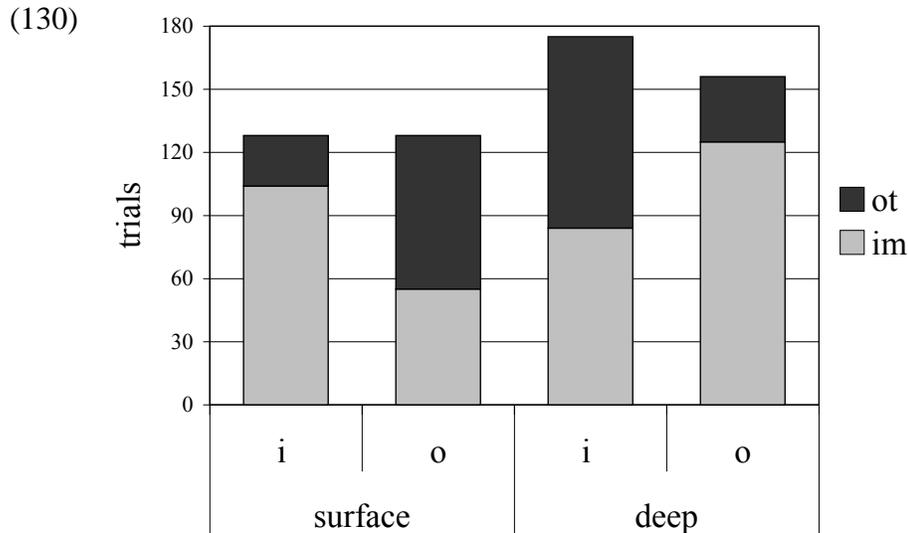
Indeed, the good generalizers (i.e. those who scored 18 and above) have their choices of [-im] concentrated at 10, while the bad generalizers (i.e. those who scored 17 or less) have their choices of [-im] concentrated above 10, as seen in (129).



The number of [-im] choices for the good generalizers was not significantly different from 10 ($n = 18$, $M = 9.83$, Wilcoxon test with $\mu = 10$, $V < 100$, $p > .1$). The bad generalizers chose the masculine [-im] significantly more often than the feminine [-ot], showing that they treated the new words as masculine Hebrew nouns, and extended the preference for [-im] from real Hebrew to the artificial nouns ($n = 42$, $M = 11.64$, Wilcoxon test with $\mu = 10$, $V > 670$, $p < .01$). The choice of [-im] comes out as significantly greater than 10 even when all participants are included ($n = 60$, $M = 11.10$, Wilcoxon test with $\mu = 10$, $V > 1200$, $p < .05$).

3.5.5.4 Errors and vowel perception

Speakers who failed to change stem vowels correctly from [i] to [o] or vice versa usually left the stem vowel unchanged. The distribution of trials with unchanged stem vowels is shown in (130), where each column indicates the number of responses with *-im* and the number of responses with *-ot* for each unchanged stem vowel.



Mirroring the finding in (122) above, the “surface” group is seen to be more successful, with only 43% of the trials leaving the stem vowel unchanged, compared to 55% of the trials in the “deep” group. Again, the “surface” group is equally successful with either stem vowel, but the “deep” group leaves more [i]’s than [o]’s unchanged.

It is instructive that the vast majority of unsuccessful trials, in both groups, leaves the stem vowel unchanged (94% and 95% of the unsuccessful trials, in the “surface” group and “deep” group, respectively). This means that speakers had virtually no difficulty in perceiving the stem vowels correctly in the singular and in the plural, leading them to choose either [i] or [o] in the plural stem, but no other vowel.

In 34 trials (2.8% of the total number of trials), speakers made a spurious vowel change, i.e. the speakers realized that some vowel change must be applied, but didn’t change an [i] to [o] or vice versa. At this rate, these are no more than experimental noise. Of the 60 participants, only 12 made spurious vowel changes (six from each group), and only six

participants made a spurious vowel change in more than one trial (three from each group). The most common spurious changes were to [u], which is the vowel that [o] is most likely to be misperceived as, with 12 trials changing [i] to [u] and 7 trials changing [o] to [u], for a total of 19 trials, or a mere 1.6% of the total number of trials.

3.5.5.5 Summary of the experimental results

In conclusion, we see that Hebrew speakers responded to the two languages in very different ways: The “surface” language was significantly easier to generalize. Generalization scores in both languages were bi-modally distributed, with speakers who were good generalizers and speakers who were bad generalizers. A significantly larger proportion of the speakers of the “surface” language were good generalizers relative to the speakers of the “deep” language.

Speakers of the “surface” language were equally successful in changing [i] to [o] and [o] to [i], while the “deep” language speakers were less successful with the [i] to [o] change relative to the [o] to [i] change. In both groups, speakers perceived stem vowels correctly in the vast majority of the time, as evidenced by the small number of trials with spurious vowel changes. The influence of real Hebrew on the artificial languages was seen in the bias that speakers had towards selection of [-im].

3.6 Discussion and analysis

The experimental results show that in selecting plural allomorphs in Hebrew, speakers make their decisions based on the surface form of plural nouns, not based on their underlying form or their singular form. This section shows how the greater success of the “surface” language participants follows naturally from the Optimality Theoretic analysis I offered for Hebrew in §3.4.

3.6.1 The role of Universal Grammar in learning alternations

The participants in both languages had to learn the same two new vowel mappings, from [o] to [i] and vice versa, with the difference being only in the selection of the plural affix that accompanies the change. Without a proper theory of affix selection, it might be surprising that a difference in affix selection between two languages is causing a difference in the ability to perform stem vowel changes between the two languages.

In the “surface” language, the introduction of an [o] into a plural stem was always accompanied by the selection of *-ot*, so no violations of LOCAL(o) were introduced. Nouns with [o] in the singular were expected to change it to [i] and to select *-im*, in which case leaving the singular [o] intact would have created a violation of LOCAL(o). Thus, in the “surface” language, LOCAL(o) allows the smooth alternation of [i] with [o] due to the selection of *-ot*, and encourages the alternations of [o] to [i] with the selection of *-im*. The plurals in the “surface” language never violate LOCAL(o), making the changes from [i] to [o] and from [o] to [i] equally good from the markedness point of view, and indeed speakers were equally successful with both changes.

In the “deep” language, the introduction of an [o] in a plural stem was accompanied by the selection of *-im*, thus introducing a violation of LOCAL(o). Singular [o]’s were expected to change to [i], thus eliminating the potential for a violation of LOCAL(o). Thus, in the “deep” language, only plurals that change [i] to [o] introduce a violation of LOCAL(o), and indeed speakers were less successful in changing [i] to [o] relative to changing [o] to [i].

Under my analysis of Hebrew, then, the greater success of the “surface” speakers at vowel alternations in the stem follows naturally from the distribution of the plural affixes in the two language. Choosing *-ot* is compatible with changing a stem vowel to [o] and with retaining a singular [o], while choosing *-im* is compatible with neither retaining a singular [o] nor with introducing a plural [o].

As for finding a constraint ranking for the two languages, it again emerges that the “surface” language is easier to analyze, and is thus expected to be easier to learn: In the “surface” language, nouns that have an [o] in their plural stem always select *-ot*, so LOCAL(o) can be uniformly ranked over ϕ -MATCH. Nouns that have [i] in their plural stem always select *-im*, which is compatible with a uniform ranking of ϕ -MATCH over * $\acute{\sigma}$ /HIGH. Under this view, the “surface” language is just a simpler, more extreme expression of actual Hebrew. The single constraint ranking in (131) can be successfully used to provide the correct choice of plural affix for the “surface” language.

(131) LOCAL(o) \gg ϕ -MATCH \gg * $\acute{\sigma}$ /HIGH

In the “deep” language, speakers cannot find a single constraint ranking for the language that uses the markedness constraints that are active in the plural allomorph selection of actual Hebrew. Since nouns with [i] in their plural stems always take *-ot*, a speaker could rank * $\acute{\sigma}$ /HIGH over ϕ -MATCH, but that would entail selection of *-ot* for *all* nouns, contrary to overt evidence. Nouns with [o] in their plural stems always take *-im* in the “deep” language, which would imply ranking ϕ -MATCH over LOCAL(o). This ranking leaves LOCAL(o) completely inactive in the artificial language, and attributes all of the *ot*-selection of the language to * $\acute{\sigma}$ /HIGH, contrary to the situation in real Hebrew, where most *ot*-selection is due to LOCAL(o). Finding a grammar for the “deep” language would require constraint cloning, as shown in (132). The nouns that have a known plural will be divided between the two clones of * $\acute{\sigma}$ /HIGH.

(132) * $\acute{\sigma}$ /HIGH_{afiv, axis, amig, azix, adic} \gg ϕ -MATCH \gg * $\acute{\sigma}$ /HIGH_{agof, apoz, acok, abof, alod},
 LOCAL(o)

While the grammar in (132) allows the participant to correctly select a plural affix once they have heard the correct plural form, it does not allow them to generalize correctly to forms that were only given in the singular. While the nouns with [i] and the nouns with [o] are neatly divided between the clones of * $\acute{\sigma}$ /HIGH, they are listed under a constraint

that is indifferent to the vowel of the stem, and hence this neat division cannot be reliably extended to novel items.

Another possibility that might be available to the participants in the “deep” language is to use the OCP (Obligatory Contour Principle, Goldsmith 1976) to choose the plural allomorph that has a vowel that is not identical to the last vowel of the root. An OCP effect on vowels is observed in actual Hebrew, where the combination of two [o]’s inside a root is quite rare, and the combination of two [i]’s is even rarer. Extending the effect of the OCP from roots to whole words would give the participant a single grammar to derive the “deep” language. Using the OCP this way still makes the “deep” language more different from actual Hebrew than the “surface” language: In the “surface” language, OCP is only active inside roots, like real Hebrew, while in the “deep” language, the OCP needs to apply across morpheme boundaries, unlike real Hebrew. Even with the OCP, then, the “deep” language is predicted to be harder to learn than the “surface” language.

3.6.2 Stem changes and allomorph selection

A question remains about the mechanism(s) that participants have used to apply vowel changes to the noun stems. Vowel changes in paradigmatic relations are ubiquitous in Hebrew. In making verbs and deverbal nouns, speakers of Hebrew are able to impose vowel mappings on words regardless of the words’ input vowels. For example, the loanword *lúp* ‘loop’ can give rise to the verb *liplép* ‘to loop’, with nothing left of the input’s [u]. For an OT-based account of Hebrew vowel changes in verbs, see Ussishkin (2000). In nouns, however, it’s less clear that Hebrew allows arbitrary vowel changes.

The most common vowel change in nouns involves an alternation between [e] and [a], as in *mélex* ~ *melaxím* ‘king’. Other vowel alternations are much less common, such as the change from [o] to [u] or from [o] to [a], as in (91) above. All vowel changes, then, are limited to plausible phonologically-driven changes, with mid vowels either rising to their corresponding high vowels or lowering to [a], both of which can be construed as vowel

reduction. Excluding the changes that go from various vowels to [a], no nouns involve a change of vowel backness or vowel rounding.

In the artificial languages, vowel changes involve backness and rounding that don't map onto [a], and thus represent a qualitative departure from real Hebrew. Since seemingly arbitrary vowel mappings are allowed in verbs, however, there is reason to believe that speakers did not go outside their grammatical system to learn the mappings, but only outside their nominal system.

Another perspective on the difference between the two artificial languages is offered by the phonological cycle (Chomsky & Halle 1968; Kiparsky 2000). If the theory allows the vowel change to apply independently of the addition of the plural affix, then the "surface" language applies the vowel change first and then chooses the plural affix to go with the changed vowel, while the "deep" language selects the plural affix first, and then changes the stem vowel. The "deep" language, under this view, renders the effect of LOCAL(o) opaque, since the vowels it operates over are no longer in the surface representation. In a version of Optimality Theory where morphological and phonological operations apply one at a time, as in Wolf (2008b), both languages respect LOCAL(o), but the "deep" language does so opaquely. Are opaque languages inherently more difficult to learn than transparent languages? The answer to that is not known. Most known cases of opacity in the world languages, if not all, are historically innovative, suggesting that even if speakers might be biased against opacity, this bias can certainly be overcome. Additionally, children innovate opaque interactions that don't exist in the adult language they're learning (Jesney 2007). If the only difference between the two artificial languages is the transparency of the pattern, it's not clear that the difference in difficulty that participants had is predicted.

There is reason to believe, however, that Hebrew speakers would not allow the vowel change to apply independently of the affix selection. Semantically, the vowel changes and plural affixes were associated with a single unit of meaning, namely, plurality. Even if a single morpheme is expressed in two different ways, it's hard to see how the two

changes could apply in two different levels of the cycle. Furthermore, vowel changes alone never mark plurality in actual Hebrew. Each and every plural noun in real Hebrew is marked with either *-im* or *-ot*, regardless of any vowel change. This is different from the situation in Arabic, where vowel changes in the stem and concatenated plural suffixes are in complementary distribution, and each mark plurality separately²².

If it is agreed that both the vowel change and the plural affix selection must happen at the same level in the cycle, then the theory of allomorph selection in Paster (2006) makes the peculiar prediction that it's the "deep" language that would be the more natural one for speakers. In this theory, allomorph selection is only allowed to refer to the shape that a stem has in the input to the current level in the cycle. In the "deep" language, then, the plural allomorphs harmonize with the vowel of the singular, while in the "surface" language, the plural allomorphs are chosen to go against the phonologically preferred pattern.

3.6.3 The limited role of phonotactics

My analysis of the experimental results relies on the activity of two markedness constraints that are quite specific and typologically-supported: LOCAL(o), which penalizes unstressed [o]'s unless followed by a stressed [o], and *ó/HIGH, which penalizes stressed high vowels. My analysis predicts that the "surface" language would be easier to learn than the "deep" language. One could argue, however, that the preference for the "surface" language could also be stated in much more general terms, as a simple reflection of Hebrew phonotactics. In this section I show that a simple projection of Hebrew phonotactics predicts that the "surface" language is actually harder than the "deep" language.

Looking at the attested vowel combinations in the singular forms of Hebrew shows a preference for non-identical vowels. The table in (133) shows counts from Bolozky &

²²In Arabic paradigms like *wazir* ~ *wuzara:ʔ* 'minister', it is plausible that *-a:ʔ* is a suffix, but it never marks the plural on its own; it always accompanies a vowel change that marks the plural. In contrast, the plural suffixes *-una* and *-at*, as in *katib* ~ *katib-una* 'writer', always mark the plural on their own, and are never accompanied by a vowel change.

Becker (2006) for all singular native nouns that contain the relevant vowel sequences and counts for native masculine di-syllabic nouns only. Both counts show that disharmonic vowel sequences are more frequent than harmonic ones.

(133)

Vowel combination	All singulars	Di-syllabic masculines
i-o	286	107
o-i	132	8
i-i	126	2
o-o	21	8

Perhaps counts of vowel combinations in plural nouns are more relevant for comparing preferences that speakers make in the plurals of the artificial languages. The table in (134) gives the counts for plurals by the final vowel of their stem, broken down by gender.

(134)

Stem-affix combination	Masculine	Feminine	Total
...i-ot	6	1070	1076
...o-im	527	5	532
...i-im	437	7	444
...o-ot	147	178	325

The totals in (134) again show a preference for disharmonic vowel sequences over harmonic ones, so if speakers are thought to select plural suffixes based on phonotactic considerations, the “deep” language is predicted to be easier than the “surface” language, contrary to fact. Even considering the masculine nouns alone makes the same wrong prediction: Since [-im] is the most frequently used affix with either stem vowel, participants would be predicted to prefer the selection of [-im] after any stem vowel, whereas in fact, speaker preferred [-im] only with a stem [i].

The experimental results cannot be reduced, then, to a mere preference for frequent vowel patterns, since speakers actively prefer patterns that are less frequent phonotactically. In my interpretation of the results, speakers analyze the artificial languages in terms of constraints that are active in real Hebrew. A simple projection of the phonotactics of real Hebrew onto the artificial languages, without the mediation of a grammar, makes the wrong prediction.

3.6.4 Learning alternations without Universal Grammar

The two languages taught in this experiment were formally equally complex. The singulars and the plural stems were identical in both, and the choice of plural suffix was completely predictable from the shape of either the singular stem or the plural stem. A learner who uses a simple information-theoretic approach should find the two languages equally hard to learn, unlike the human subjects, who found the “surface” language significantly easier.

The results are challenging for a source-oriented model of phonology, such as the Minimal Generalization Learner (MGL, Albright & Hayes 2003). In the MGL, the selection of the affixes is relativized to observed changes between paradigmatically related forms. In the case of Hebrew, the MGL would identify two changes: going from nothing to [im] and going from nothing to [ot]. These changes compete for the real words of Hebrew, so the addition of [im] would mis-fire with an *ot*-taker, and vice versa. This is why each change is associated with a success rate, which is the number of words it derives correctly divided by the number of words it can apply to. Simplifying the MGL results greatly, its analysis of Hebrew is seen in (135)²³. The addition of [im] at the end of the word has a

²³The actual output of the MGL contains hundreds of rules, and requires some interpretation. For instance, the MGL rules don't abstract over the root-final consonants directly, as shown simplistically in (135). Rather, the MGL creates rules that refer to each individual segment, and then gradually abstracts from them using natural classes. The picture in (135) also abstracts away from cases of vowel deletion, which cause the MGL to identify a change that is wider than the simple addition of [im] or [ot]: For example, in *zanáv* ~ *znavót* 'tail', the change is from [anáv] to [navót], and the suffix [ot] is not analyzed separately from the deletion of the root vowel.

high success rate, since most masculine nouns are *im*-takers. The addition of [ot] at the end of just any word would have a low success rate, but the addition of [ot] to a word that ends in [o] followed by a consonant would have a reasonably high success rate.

(135)	change	environment	success rate
	$\emptyset \rightarrow [im]$	/ ___ #	~97%
	$\emptyset \rightarrow [ot]$	/ ___ #	~3%
	$\emptyset \rightarrow [ot]$	/ o C ___ #	~30%

The MGL result is impressive in that it manages to extract a set of generalizations from the rather complex raw data: It identifies the suffixes, and it identifies the kind of nouns that take them. In this model, however, the similarity between the suffixes and their environment is accidental: It learns nothing about vowel harmony, and could equally well learn a language, Hebrew', where choosing *-ot* is correlated with any other phonological property of the root.

When the MGL is applied to the two artificial languages, it identifies two changes in each language, as shown in (136). The two changes have a success rate of 100% in the two languages, since the plural allomorph selection is completely regular. Crucially, these four changes are not attested in real Hebrew at all, so the two languages are equally different from real Hebrew, and are thus predicted to be equally easy or equally hard for native speakers. Due to the vowel change in the stem, the MGL can no longer separate the suffixes [im] and [ot] from the stem.

(136)	“surface” language	“deep” language
	$o C \rightarrow [i C im]$	$o C \rightarrow [i C ot]$
	$i C \rightarrow [o C ot]$	$i C \rightarrow [o C im]$

Albright & Hayes (2003) recognized this aspect of the MGL in its treatment of the vowel changes in the English past tense. English speakers use the vowel [o] (as in *drove*,

rode) to form the past tense of novel verbs, regardless of the vowel in the present tense. In real English, only the four vowels [aɪ, eɪ, iː, uː] change to [o] in the past²⁴, but speakers identify [o] as a good marker of the past tense with little regard for what the present tense vowel is, and extend the use of [o] to unattested vowel mappings (while still preferring mappings that resemble existing mappings). Albright & Hayes (2003) point out that a model of human behavior must include the ability to state generalizations about derived forms separately from the bases they are derived from. I claim that the use of markedness constraints, as proposed here, is suitable for doing just that.

3.6.5 The role of the grammar of real Hebrew

The participants' responses in the experiment make it clear that they identified the plural affixes of the artificial language with the plural affixes of actual Hebrew. All the plural forms that participants produced contained a well-formed plural affix, either *-im* or *-ot*²⁵. Furthermore, speakers were quite successful at recognizing that the choice of affix depends on the vowels of the root, but except for one speaker, they never selected the vowels of the plural suffix independently of its consonants, but rather treated them as two whole units, *-im* and *-ot*, just like in real Hebrew.

Whenever the participants produced plural forms, either repeating forms they have heard or generating plurals that they haven't heard, they pronounced them all with final stress without fail. This indicates that the nouns of the artificial languages were not accepted as just any nouns of Hebrew, but more specifically as native nouns of Hebrew.

²⁴Examples: *drive* ~ *drove*, *break* ~ *broke*, *freeze* ~ *froze*, and *choose* ~ *chose*.

²⁵A single participant offered the following four paradigms: *amov* ~ *amivit*, *agiv* ~ *agivit*, *atox* ~ *atixit*, and *afoc* ~ *axifoc*. The rest of this participant's responses were unremarkable, with either *-im* or *-ot* in them.

With loanwords, plurals are formed without moving the stress away from the root, so a pluralized loanword will never surface with final stress²⁶.

Finally, the preference for *-im* over *-ot* in the experiment, as discussed in §3.5, is the clearest indication that participants accepted the artificial nouns as nouns of Hebrew. In the artificial languages, *-im* and *-ot* were equally represented, so the higher frequency of *-im* responses must be attributed to the influence of real Hebrew. It is very likely that speakers accepted the artificial nouns as masculine, especially given the numerals and adjectives that agreed in gender with those nouns in the various frame sentences. However, *-im* is more frequent than *-ot* in real Hebrew overall (since masculine nouns are more than twice as common as feminine nouns), so speakers can show a bias for *-im* even if they ignore the cues for masculine gender in the experiment.

3.6.6 Source-oriented generalizations?

The aim of this chapter is to highlight the importance of product-oriented generalizations in phonology, yet it is obviously still the case the languages have source-oriented generalizations. Even the Hebrew plural affix, which I have shown to be subject to a product-oriented generalization, is also subject to a source-oriented generalization: Loanwords that end in [-a] in the singular invariably take the plural [-ot], regardless of their gender, as noted in (70) and (71). In other words, the choice of plural affix must also be sensitive to some aspect of the input to the derivation.

In Optimality Theory, there are two ways in which an output can be sensitive to the input: The activity of faithfulness can force identity between an input and an output, or some mechanism of opacity can give rise to structure that depends phonologically on some aspect of the input, e.g. in the Tiberian Hebrew /defʔ/ → [defe], the second [e] in the

²⁶Some nouns that are etymologically borrowed were fully nativized and now get final stress in the plural, e.g. *balon-ím* ‘balloon’. These nouns are all di-syllabic, just like the majority of native Hebrew nouns (Becker 2003).

output is not present due to faithfulness, but its presence depends on the presence of the glottal stop in the input (McCarthy 2007a).

Faust (2008) offers an analysis of Hebrew in which the plural affix [-ot] phonologically contains the feminine suffix [-a]. In terms of OT, this would mean that nouns that end in [-a] select [-ot] via input-output faithfulness to a [-high] feature. An alternative analysis would attribute the selection of [-ot] to output-output faithfulness (Benua 1997) to the [-high] feature in [-a]. I leave the exact solution of this issue to future work.

3.7 Conclusions

This chapter examined the distribution of the two plural suffixes *-im* and *-ot* on Hebrew nouns. The lexicon study showed a connection between having [o] in the root and a preference for selecting *-ot*, with the preference being stronger when the [o] is final, and thus adjacent to the suffix, and weaker when the [o] is non-final in the root. In a novel word test, speakers replicated the effect that [o] had in the lexicon, choosing *-ot* as a plural suffix most often with novel roots that have an [o] in their final syllable, and least often with roots that don't have [o] at all.

I offered an OT-based analysis of plural allomorph selection in Hebrew, which relied on a mechanism of constraint cloning to build lexical trends into the grammar, and project those trends onto novel nouns. In the analysis, allomorph selection was understood to be without faithfulness cost, and therefore only markedness constraints were involved in the analysis.

Since markedness constraints only assess output forms, they have no access to underlying representations or to paradigmatically related forms. In deriving Hebrew plurals, the selection of *-ot* is predicted to correlate with the presence of [o] in the plural stem, regardless of the vowels of the singular. Since in real Hebrew, the presence of [o] in a plural stem always corresponds to the presence of [o] in the singular, the prediction cannot be tested on the real words of the language.

To test whether the selection of the plural affix is sensitive to the vowels of the input or the vowels of the output, I created a pair of artificial languages, where a singular [i] alternates with a plural [o] and vice versa. In one language, the selection of *-ot* correlated with the presence of [o] in the plural stem, and in the other language, the selection of *-ot* correlated with the presence of [o] in the singular stem. As predicted, speakers were significantly more successful at generalizing the language where the selection of *-ot* correlated with the presence of [o] in the plural stem.

The artificial languages were designed and presented as languages that are just like real Hebrew, with the only difference being the vowel changes from [o] to [i] and vice versa, which don't occur in real Hebrew. To insure that singulars and plurals are correctly paired, participants never heard or produced a plural form without hearing or producing its singular in the same trial. Indeed, the experimental results show that the participants accepted the artificial nouns as native nouns of Hebrew, evidenced by their generation of plural forms with final stress and a bias towards *-im*.

The prediction of the markedness-based analysis, which favors the language that pairs *-ot* with plural [o]'s, was contrasted with an MGL-based analysis (Albright & Hayes 2003), which predicts that the two languages would be equally different from Hebrew, and thus equally difficult for Hebrew speakers. The point is applicable more generally to any analysis that relies on general pattern-finding mechanisms that don't have any expectations about what a possible human language is. Since the two artificial languages are formally equally complex, with the exact same amount of information in them, there is no a priori reason to prefer generalizations about output forms over generalizations about input forms. Additionally, I have shown that the experimental results cannot be reduced to a mere phonotactic preference, since the phonotactics of real Hebrew prefer the pairing of non-identical vowels over identical vowels.

In real Hebrew, the connection between [o] in the stem and the selection of *-ot* is equally reliable when stated over singulars or over plurals: One can say that singulars with

[o] often choose *-ot*, or one can say that plural stems with [o] often choose *-ot*. And yet, the results of the artificial language experiment show that speakers are biased to choose the plural-based interpretation over the singular-based interpretation. This bias follows naturally from the analysis I offer, which attributes allomorph selection to the activity of universal markedness constraints, as is standardly assumed in the OT literature.

CHAPTER 4

LEXICAL TRENDS AS OPTIMALITY THEORETIC GRAMMARS

4.1 Introduction

The results presented in chapters 2 and 3 were used to motivate a framework, based on Optimality Theory (Prince & Smolensky 1993/2004), that learns lexical trends and projects them onto novel items. The mechanism for learning a lexical trend from an ambient language relied on the Recursive Constraint Demotion algorithm (RCD, Tesar & Smolensky 1998, 2000; Tesar 1998; Prince 2002), augmented with a mechanism of constraint cloning (Pater 2006, 2008b).

This chapter goes on to develop this version of OT in greater detail and in greater generality. It starts with a discussion of the cloning mechanism in §4.2, with a focus on the question of identifying the constraint to clone. Then, the learning algorithm is fleshed out formally in §4.3. The learning algorithm assumes that when learning paradigms, the surface form of the base of the paradigm is always taken to be its underlying form, and non-surface-true underlying representations are limited to affixes only. This assumption is explored and motivated in §4.4. The use of OT constraints to account for lexical trends makes predictions about the typology of lexical trends, and §4.5 explores this typology. Conclusions are offered in §4.6.

4.2 Choosing the constraint to clone

The cloning algorithm proposed here is designed to achieve two goals: (a) resolve inconsistent ranking arguments, allowing the learner to use RCD and find a grammar even when faced with an inconsistent lexicon, and (b) learn a grammar that reflects statistical

trends in the lexicon and allows the learner to project these trends onto novel items. This section shows how these goals are achieved, first by identifying the situations in which cloning helps the learner find a consistent grammar, and then by showing how the choice of constraint to clone bears on the lexical statistics that get encoded in the grammar.

Constraint cloning allows the learner to accommodate inconsistent patterns in the language they're exposed to, and learn the relative strength of each pattern. When a language presents multiple inconsistent patterns, each with its own relative strength, as seen in chapters 2 and 3, multiple constraints will be cloned. In such a case, the learner will need a mechanism that allows them to list their lexical items with the various clones in a way that replicates the relative prevalence of each pattern in the data.

This section provides a formal mechanism for achieving this goal by answering two main questions: Firstly, in what situations does constraint cloning help with finding a consistent grammar for the language? And secondly, in what situations is the choice of constraint to clone crucial? It will turn out that depending on the data that is available to the learner, the choice of constraint to clone can become crucial or cease to be crucial. This in turn will mean that cloning is always relative to available language data, and that as more data becomes available, decisions about cloning will be reconsidered.

4.2.1 Minimal conflict

Constraint cloning is a solution for inconsistency. Recall that inconsistency is found by the Recursive Constraint Demotion algorithm (RCD, Tesar & Smolensky 1998, 2000; Tesar 1998; Prince 2002), which takes a Support, i.e. a set of winner-loser pairs, and tries to use it to discover a grammar. The RCD operates by finding columns that contain at least one W and no L's in them, and "installing" them, meaning that any winner-loser pairs that get a W from the installed constraints are removed from the Support. The constraints are then added beneath any previously installed constraints. When all the winner-loser pairs are thus removed, any remaining constraints are added to the grammar, and RCD concludes.

There is no guarantee, of course, that RCD will manage to install all the constraints and empty out the Support. When there is no column available that has no L's in it, RCD will give up, or stall. In some cases, such as the trivial (hypothetical) example in (137), cloning will not help. The intended winner is harmonically bounded, i.e. no constraint prefers it over the loser, indicating that something else went wrong: The learner made a wrong assumption about some underlying representation, for instance.¹

(137)

	C
a. winner \succ loser	L

Cloning the constraint in (137) wouldn't help, since making two clones of the constraint would still leave the intended winner without any constraint that prefers it over the intended loser. Having both W's and L's in a column won't help either, as in the minimal situation in (138).

(138)

	C
a. winner1 \succ loser1	W
b. winner2 \succ loser2	L

Cloning the constraint in (138), listing winner1 with one clone and winner2 with another clone, would allow the installation of one clone, removing the first winner-loser

¹A harmonically bounded winner can also be unbounded by adding a constraint that prefers the winner to the loser. Here I assume that a fixed, Universal set of constraints is always available to the learner, so there is no mechanism for adding constraints as needed beyond cloning. See, however, §4.5.3 for an example of subcategorizing constraints to affixes.

pair from the Support, but leaving winner2 in the same situation as in (137). Winner2 has no constraint that prefers it to the intended loser, i.e. it is harmonically bounded, so no grammar can help it.

Just one constraint, then, in and by itself, can never lead to fruitful constraint cloning. The minimal inconsistent scenario that can be helped by cloning involves two conflicting constraints, as in (139).

(139)

	C1	C2
a. winner1 \succ loser1	W	L
b. winner2 \succ loser2	L	W

From this minimal scenario, cloning either constraint will solve the inconsistency. The result of cloning C1 is in (140). One clone of C1 is listed with all the items that it assigns a W to, in this case winner1, and the other clone is listed with all the items that C1 assigns an L to, in this case winner2.

(140)

	$C1_{winner1}$	C2	$C1_{winner2}$
a. winner1 \succ loser1	W	L	
b. winner2 \succ loser2		W	L

When RCD is applied to (140), $C1_{winner1}$ gets installed first, and the first winner-loser pair is removed from the Support. The column of C2 is left without any L's in it, so C2 is installed, and the second winner-loser pair is removed. The remaining constraint, $C1_{winner2}$, is added at the bottom, and the resulting grammar is $C1_{winner1} \gg C2 \gg C1_{winner2}$.

Equivalently, C2 could have been chosen for cloning, with the resulting grammar being $C2_{winner2} \gg C1 \gg C2_{winner1}$. These two grammars are both fully consistent, and both successfully resolve the inconsistency by putting winner1 and winner2 in two different “bins”. Assuming that each of winner1 and winner2 represent a number of lexical items, successfully separating them and making their relative numbers accessible to the learner will make the lexical trend available, no matter which of C1 or C2 is chosen for cloning.

4.2.2 Two independent conflicts

More complex situations arise when the language has two or more lexical trends in it, which leads to two or more conflicts that need to be resolved by cloning. I examine these situations below.

Completely independent trends, as in (141), present no challenge to the learner. They are simply two instances of a minimal conflict, as in (139). Cloning any of the constraints will solve one conflict, which in turn will only leave two constraints available for cloning, and cloning either of those will solve the other conflict. This is shown below.

(141)

	C1	C2	C3	C4
a. winner1 \succ loser1	W	L		
b. winner2 \succ loser2	L	W		
c. winner3 \succ loser3			W	L
d. winner4 \succ loser4			L	W

If C1 is chosen for cloning first, $C1_{winner1}$ will be installed, which will then allow C2 to be installed. The first two winner-loser pairs will be removed from the Support, which leaves $C1_{winner2}$ ready for installation. Now the situation with C3 and C4 is reduced to a

simple minimal conflict, which can be resolved by cloning either constraint. If C3 is then chosen for cloning, the resulting grammar would be the one in (142).

$$(142) \quad C1_{winner1} \gg C2 \gg C1_{winner2} \\ \gg C3_{winner3} \gg C4 \gg C3_{winner4}$$

If C3 is chosen for cloning first, and then C1 is chosen, the grammar is slightly different:

$$(143) \quad C3_{winner3} \gg C4 \gg C3_{winner4} \\ \gg C1_{winner1} \gg C2 \gg C1_{winner2}$$

Since the learner has no evidence for assuming that C1 and C2 interact with C3 and C4, they should be equally happy with the grammars in (142) and (143). If evidence comes along later about an interaction between the pairs of constraints, that might have an effect on the choice of constraints to clone the next time RCD is run.

When dealing with minimal conflicts, as seen in (139) and (141) above, the choice of the constraint to clone is free. Such a situation was seen in §2.4.6, where in Turkish post-vocalic dorsals and post-sonorant dorsals were forming two separate trends, governed by separate constraints. Each conflict is defined by a pair of constraints, and there is no overlap in the constraints.

4.2.3 Overlapping conflicts

In real languages, conflicting ranking arguments can overlap: Two different lexical trends can be defined using just three constraints, with one constraint serving as the pivot for both trends. This is the situation in Turkish, where stem-final coronal and palatal stops both have a trend of voicing intervocalically, i.e. IDENT(voice) is serving as the pivot for both *VtV and *VtʃV (see §2.4.2). A situation like this is shown in (144), where the specific forms and constraints are abstracted from.

(144)

	C1	C2	C3
a. winner1 \succ loser1	W	L	
b. winner2 \succ loser2	L	W	
c. winner3 \succ loser3		L	W
d. winner4 \succ loser4		W	L

In this situation, the choice of constraint to clone becomes crucial for extracting lexical trends from the data correctly. To see this, consider what happens if the speaker wrongly chooses to clone C2, as show in (145).

(145)

	C1	$C2_{winner1, winner3}$	$C2_{winner2, winner4}$	C3
a. winner1 \succ loser1	W	L		
b. winner2 \succ loser2	L		W	
c. winner3 \succ loser3		L		W
d. winner4 \succ loser4			W	L

The Support in (145) allows the speaker to install $C2_{winner2, winner4}$, and then remove the second and fourth winner-loser pairs. Then, C1 and C3 will be installed, removing the other two winner-loser pairs. The resulting grammar is the one in (146).

(146) $C2_{winner2, winner4} \gg C1, C3 \gg C2_{winner1, winner3}$

While this grammar will correctly derive listed words, it only extracted one lexical trend where the data presented two. The ranking arguments in (144) show that winner1 and winner2 form one trend, while winner3 and winner4 form a different trend. The grammar in (146) collapsed the two trends, putting winner1 and winner3 in one bin, and winner2 and winner4 in another bin.

In Turkish, for instance, cloning IDENT(voice) instead of either *VtV or *VtʃV would put *t*-final and *tʃ*-final nouns in the same bin, causing the speaker to assign the same likelihood of voicing to a novel item with a final [t] and a novel item with a final [tʃ]. Actual speakers don't do that, but rather prefer alternations with *tʃ*-final novel nouns, reflecting the lexical statistics (§2.3).

When there are multiple constraints to clone, as in (144), the learner must choose the constraint that has the smallest number of W's and L's in its column. Choosing the column with the minimal number of W's and L's is not an arbitrary choice; it is the way to ensure that a minimal number of lexical items are identified as a part of a lexical trend, leaving other lexical items to the care of other trends or to the regular grammar.

In (144), C1 and C3 are each equally eligible for cloning, with 2 non-empty cells each in their respective columns, compared with the 4 non-empty cells of C2. Choosing either C1 or C3 for cloning would produce the intended result, where the speaker identifies the two lexical trends that are in the data. If C1 is chosen, the learner can install $C1_{winner1}$ and remove the first winner-loser pair from the Support. The new Support, with the first winner-loser pair crossed out, is shown in (147).

(147)

	$C1_{winner1}$	$C1_{winner2}$	C2	C3
a. winner1 \succ loser1	W		L	
b. winner2 \succ loser2		L	W	
c. winner3 \succ loser3			L	W
d. winner4 \succ loser4			W	L

Cloning C1 left a Support that is still inconsistent, as the columns for C2 and C3 still have both W's and L's in them. Looking again for the constraint that has the fewest non-empty cells in its column, C3 is chosen for cloning, since it has fewer non-empty cells than C2. Once C3 is cloned, $C3_{winner3}$ is installed, winner3's winner-loser pair is removed from the Support, and this allows C2 to be installed. The winner-loser pairs of winner2 and winner4 are removed, leaving the Support empty, which in turn lets the remaining $C1_{winner2}$ and $C3_{winner4}$ be installed, leading to the grammar in (148).

(148) $C1_{winner1} \gg C3_{winner3} \gg C2 \gg C1_{winner2}, C3_{winner4}$

The two trends are successfully captured by the clones of C1 and C3, with C2 serving as a pivot for both. Cloning C3 first would have resulted in almost exactly the same grammar, just with $C3_{winner3} \gg C1_{winner1}$. Since C1 and C3 don't interact directly, their relative ranking doesn't matter.

4.2.4 Interim summary: Choosing the least populated column

To summarize so far: The minimal situation where cloning constraints is a useful tool for resolving inconsistencies involves two constraints, each with both W's and L's in their columns. When the W's and L's that the two constraints assign are exactly opposite, as in

(139) or (141), the choice of constraint to clone is inconsequential. When conflicts involve unequal numbers of W's and L's, as in (144), lexical trends are correctly identified only if the least populated column (i.e. the column with the minimal number of W's and L's) is chosen for cloning first.

Choosing the least populated column guarantees that the minimal number of lexical items is listed with clones, which in turn guarantees that the learner makes the finest distinctions that their Universal constraint set can express.²

Choosing the least populated column to clone is beneficial for identifying lexical trends even when only one trend is involved. Consider the situation in (149), where C1 and C2 make exactly opposite demands on winner1 and winner2, but C2 is neutral with respect to winner3.

(149)

	C1	C2
a. winner1 \succ loser1	W	L
b. winner2 \succ loser2	L	W
c. winner3 \succ loser3	W	

If C1 is wrongly chosen for cloning, winner1 and winner3 will be listed with one clone of C1, and winner2 will be listed with the other clone. The resulting grammar would be $C1_{winner1, winner3} \gg C2 \gg C1_{winner2}$, where winner1 and winner3 end up in the same “bin”, and thus wrongly skew the lexical trend in favor of winner1. The problem here is that only winner1 and winner2 are part of the minimal conflict. Winner3 is not a part of the conflict,

²This aspect of learning is analogous to the way the Minimal Generalization Learner (Albright & Hayes 2002, 2003, 2006) starts with the finest generalization it can make, i.e. over single words, and then gradually expands the scope of generalization. In the learning method proposed here, there is only one level of granularity, which is dictated by the constraints in CON.

and shouldn't be made part of it by mis-cloning. If the least populated column, that of C2, is chosen for cloning, only winner1 and winner2 will be listed with its clones, and the resulting grammar would be $C2_{winner2} \gg C1 \gg C2_{winner1}$. This gives the intended result, where only winner1 and winner2 are listed with clones, and winner3 remains a nameless player in the regular grammar.

Correctly choosing a constraint to clone, then, can be crucial in some cases but not others. In (139), either constraint can be correctly cloned, but only one correct option is available in (149). It is the addition of another winner-loser pair that makes the choice crucial in (149). This means that as the learner is exposed to more data about their language, the choice of constraint to clone can change from being free to being crucial; therefore, the learner could make decisions about cloning that will turn out to be wrong as more data is discovered. To avoid such problems, where an early decision causes a mistake down the road, cloning must not be permanent. Constraints are cloned as necessary when RCD is run, and a grammar is reached, but when a new winner-loser pair is added to the Support, RCD makes a fresh start with all the constraints in their pristine, pre-cloned state.

Choosing the least populated column is a necessary condition on cloning, but one more move is needed to clone correctly in cases of trends that are in general-specific relationship. This additional move is explained below.

4.2.5 General-specific relations between trends; masking

When a language presents two lexical trends to the learner, the two trends can be completely independent, as seen in (141), or they can overlap, as seen in (144). A third kind of relationship between trends involves one trend that is governed by a constraint that assesses a subset of the W's and L's that another constraint assesses, as seen in (150).

(150)

	C1	C2	C3
a. winner1 \succ loser1	W	W	L
b. winner2 \succ loser2	L	L	W
c. winner3 \succ loser3		W	L
d. winner4 \succ loser4		L	W

Cases like this were encountered in Turkish (§2.4.3) and in Hebrew (§3.4.3). In Hebrew, speakers learn that [o] in the final syllable of a noun is most conducive to selecting the plural suffix [-ot], and that an [o] in the penultimate syllable is less so. That is, they learn two separate trends. The two trends can be captured by a specific constraint that prefers a plural suffix with [o] in it when adjacent to a stem [o], and a more general constraint that prefers a plural suffix with [o] no matter how far it is from the [o] of the stem. The more specific constraint, which demands adjacency, can be used to list the nouns with an [o] in the final syllable of their stems, leaving the nouns with a non-final [o] to the care of the more general constraint. There is no need for the theoretically undesirable constraint that prefers a plural suffix with [o] only when the stem has an [o] that is *not* adjacent to the plural suffix.

A simple inspection of (150) reveals that C1 is more specific than C2, since C1 assigns a proper subset of the W's and L's that C2 assigns. The least populated column in (150) that contains both W's and L's is that of C1, so C1 is chosen for cloning.

However, simply cloning C1 will not allow the learner to correctly learn the lexical trends of the language. To see this, consider the result of cloning C1, shown in (151), with the first clone of C1 installed, and the first winner-loser pair crossed out.

(151)

	$C1_{winner1}$	$C1_{winner2}$	C2	C3
a. winner1 \succ loser1	W		W	L
b. winner2 \succ loser2		L	L	W
c. winner3 \succ loser3			W	L
d. winner4 \succ loser4			L	W

At this point, C2 and C3 are equally eligible for cloning, since they each have a total of three W's and L's. If C2 is chosen for cloning, one of its clones will be listed with winner3, and the other clone will be listed with winner2 and winner4. The resulting grammar would be the one in (152).

(152) $C1_{winner1} \gg C2_{winner3} \gg C3 \gg C1_{winner2}, C2_{\{winner2, winner4\}}$

The grammar in (152) is not quite right: While it correctly puts winner1 and winner2 in two separate bins, it also incorrectly puts winner2 in the same bin with winner4, in effect allowing winner2 to “double dip” and skew the lexical statistics in its favor. Recall that each of the winners in (151) represents a class of lexical items. If winner3 and winner4 each represent a relatively small number of items, and winner2 represents a large number of items, the learner would learn a trend that is quite different from the actual trend in the lexicon.

In the Hebrew case, double-dipping means that nouns with an [o] in their final syllable are learned correctly (pitting 34 *ot*-takers against 129 *im*-takers), but nouns with a non-final [o] are not. The more general constraint that prefers a plural with [o] no matter where the stem [o] is will pit 12 *ot*-takers with a non-final [o] against *all* 219 of the *im*-takers that have [o] in them, not only against the 90 *im*-takers that have a non-final [o]. This means

that the likelihood of *ot*-taking in the presence of a non-final [o] would be predicted to be $12/(12+219) = 5\%$, whereas the lexical statistics predict a likelihood of $12/(12+90) = 12\%$. In other words, double-dipping is reducing the likelihood of *ot*-taking by more than half. The experimental results presented in §3.3 are not as conclusive as one could hope for, but they suggest that lexical statistics are learned correctly, without the skewing created by double-dipping.

To learn lexical statistics correctly, the learner has to prevent lexical items from double-dipping. This is achieved by “masking” the extra W’s and L’s from any general constraints, where masking a W or an L means that it is ignored for the purposes of cloning. Formally, what the learner does is first clone a constraint and list lexical items with it; then, the learner identifies constraints that assign a superset of the W’s and L’s of the cloned constraint, and remove W’s and L’s from the superset constraints, such that lexical items that were just listed with the specific constraint are protected from another listing. This is shown in (153), where the L that C2 assigns to winner2 is masked in the Support. The W that C2 assigns to winner1 is also masked, even though that W will be gone anyway when $C1_{winner1}$ is installed and the winner-loser pair is removed.

(153)

	$C1_{winner1}$	$C1_{winner2}$	C2	C3
a. winner1 \succ loser1	W		W	L
b. winner2 \succ loser2		L	L	W
c. winner3 \succ loser3			W	L
d. winner4 \succ loser4			L	W

After general W's and L's were masked from the Support, C2 now has the least populated column. When C2 is cloned, winner3 and winner4 are listed with its clones, leading to the grammar in (154).

$$(154) \quad C1_{winner1} \gg C2_{winner3} \gg C3 \gg C1_{winner2}, C2_{winner4}$$

This grammar successfully captures the two trends in the data, with each class of lexical items listed with only one clone of one constraint.

4.2.6 Remaining questions

The least populated column metric, augmented by the masking mechanism for preventing double-dipping, were shown to be sufficient for correctly learning lexical trends. The examples shown so far involved conflicts between pairs of constraints. These abstract examples correspond to the scenarios seen in actual languages in chapters 2 and 3 and in the rest of this chapter. It is possible, however, that a single lexical trend could involve a conflict between more than two constraints. While such cases are not currently known in human languages, they are explored below for the sake of completeness.

The simplest form of constraint conflict involves two constraints, but a single conflict can involve any number of constraints. The tableau in (155) illustrates a conflict that involves four constraints (cf. Pater 2008a for a similar scenario). While no constraint can be installed in this scenario without cloning, cloning any one of the constraints will solve the inconsistency.

(155)

	C1	C2	C3	C4
a. winner1 \succ loser1	W	L		
b. winner2 \succ loser2		W	L	
c. winner3 \succ loser3			W	L
d. winner4 \succ loser4	L			W

Cloning C3, for instance, and listing winner2 and winner3 with its clones, will lead to the grammar $C3_{winner3} \gg C4 \gg C1 \gg C2, C3_{winner2}$. If a different constraint is chosen for cloning, different lexical items will be made part of a lexical trend. For example, if C1 is cloned, winner1 and winner4 will be made part of the lexical trend. In other words, cloning any one of the constraints in (155) will resolve the conflict, but different predictions are made about the lexical trend involved.

A scenario similar to the one in (155) is in (156), where a single conflict involves three constraints, and cloning any of the three would solve the inconsistency.

(156)

	C1	C2	C3
a. winner1 \succ loser1	W	L	L
b. winner2 \succ loser2	L	W	W
c. winner3 \succ loser3	W	W	L

Cloning C_1 , for example, would make $winner_1$ and $winner_3$ listed with one of its clones, and installing that clone would leave only the second winner-loser pair in the Support. This would allow C_2 and C_3 to be installed. The resulting grammar would be $C_{1_{winner_1, winner_3}} \gg C_2, C_3 \gg C_{1_{winner_2}}$. If, however, C_2 is cloned, $winner_2$ and $winner_3$ are now forming a class of items that gets listed with a clone, eventually leading to the grammar $C_{2_{winner_2, winner_3}} \gg C_1 \gg C_3, C_{winner_1}$. We see again that choosing any one of the constraints to clone solves the inconsistency, but the resulting lexical trends are different: Cloning C_1 puts $winner_1$ and $winner_3$ in the same bin, while cloning C_2 puts $winner_2$ and $winner_3$ in the same bin. Since it is not known whether natural languages produce situations such as the one in (155) or the one in (156), it is not known whether this is a problem.

4.3 The cloning algorithm

The previous section presented the basic mechanics of cloning, focusing on the choice of constraint to clone. This section adds in the details, presenting an algorithm for learning an OT grammar that incorporates cloned constraints. The algorithm is based on the Recursive Constraint Demotion Algorithm (RCD, Tesar & Smolensky 1998, 2000; Tesar 1998; Prince 2002), augmented with a mechanism for Inconsistency Resolution that is based on Pater (2006, 2008b).

The learner's goal in the proposed model is to discover the phonological realization of the morphological categories in their language. The morphological structure, including its meaning and any associated hierarchical structure is taken here to be given. The learner needs to discover the phonological underlying representation of the various morphemes and the phonological processes that take place as these morphemes are combined to make words, even if these phonological processes apply to some morphemes and not others.

This section starts by presenting the original RCD in §4.3.1, and then adding the cloning mechanism for resolving inconsistency in §4.3.2. The properties of the new object

introduced, the clone, are discussed and formalized in §4.3.3. The application of a grammar that has cloned constraints in it is discussed in §4.3.4.

The Cloning RCD operates on a Support, which contains three kinds of linguistic objects: winners, underlying forms, and losers. Of these, only the winners are directly observable surface forms. Underlying forms and losers must be provided by a separate mechanism, and §§4.3.1-4.3.4 presuppose that the underlying forms and losers are given. In the remaining two sections, some ideas for creating underlying forms and losers are explored. The search of underlying forms is taken on in §4.3.5, and the generation of losers is touched on in §4.3.6.

4.3.1 Background: RCD

RCD is an algorithm for learning a grammar, given a set of universal constraints and a prepared table of winner-loser pairs.³ This table is also called the Support in Tesar & Prince (2006). In each winner-loser pair, the winner is a surface form of the adult language that the learner is exposed to, and the loser is some other form, provided by the learner or the analyst, that the winner is compared to. In each pair, it is assumed that the winner and the loser are derived from a single underlying representation, also provided by the learner or the analyst.

A winner-loser pair, then, is prepared by taking each output form of the language, assigning an underlying form and a loser to it, and comparing how the winner and the loser fare on the set of universal constraints. A sample winner and loser are shown in (157), where the winner is the surface form [atʃ-i], and the analyst provided the underlying form /atʃ + i/ and the loser *[aɕ-i].

³Tesar & Smolensky (1998, 2000) define their version of RCD to operate on mark-data pairs. These were later replaced by winner-loser pairs, which abstract from the number of violation marks to a simple binary distinction (Prince 2002 et seq.)

(157)

/atʃ + i/	IDENT(voice) _{σ1}	*VʃV	ONSET
a.  atʃ-i		*	*
b. aɔʃ-i	*!		*

When the winner and the loser in (157) are made into a winner-loser pair, as in (158), the number of violation marks in each column is compared (Prince 2002). A “winner-favoring” constraint, or a constraint that assigns more violation marks to the loser than to the winner, assigns a W to the pair. Similarly, a “loser-favoring” constraint is one that assigns more violations to the winner than it does to the loser, and this is marked by an L. A constraint that assigns the same number of violations to the winner and to the loser, like ONSET in this example, leaves an empty cell in (158).

(158)

	IDENT(voice) _{σ1}	*VʃV	ONSET
a. atʃ-i > aɔʃ-i	W	L	

Once the Support is ready, even with just one winner-loser pair, as in (158), RCD can run on it. RCD produces a stratified hierarchy of the constraints by finding constraints that have at least one W and no L’s in their column, and “installing” them. Installing constraints means that they are added to the constraint hierarchy below any previously installed constraints, and any winner-loser pairs they assign W’s to are removed from the Support. RCD is done when the Support is emptied out, and any constraints that were left over are installed at the bottom of the hierarchy. In (158), RCD first identifies IDENT(voice)_{σ1} as a constraint that has at least one W and no L’s in its column, and installs it. This removes the single winner-loser pair in the Support, so RCD can finish by installing

*VfV and ONSET below IDENT(voice)_{σ1}. The resulting grammar is IDENT(voice)_{σ1} ≫ *VfV, ONSET.

RCD is described formally in (159). It starts with a Support, and finds all the constraints that are ready to install (159a). It finds the winner-loser pairs that these constraints assign a W to, removes these winner-loser pairs from the Support (159b-i), adds these constraints to the developing constraint hierarchy (159b-ii), and removes these constraints from the Support (159b-iii). Once the Support is empty, any remaining constraints are added to the hierarchy, and RCD is done.

(159) **RCD Algorithm** (after Tesar & Smolensky 1998, 2000)

Given a Support S ,

Given a set of constraints in S , *not-yet-ranked constraints*,

H := a new constraint hierarchy.

While S is not empty, repeat:

- a. *current-stratum* := all the constraints in *not-yet-ranked constraints* that have at least one W and no L's in their column in S
- b. If *current-stratum* $\neq \emptyset$,
 - i. remove winner-loser pairs that are assigned a W by any constraint in *current-stratum*.
 - ii. put *current-stratum* as the next stratum in H , and
 - iii. remove *current-stratum* from *not-yet-ranked constraints*

Put *not-yet-ranked constraints* as the next stratum in H .

Return H .

RCD is guaranteed to find a ranking of the constraints in a given Support if the data in the Support was created from some ranking of the constraints (Tesar & Smolensky 2000, p. 109). If, however, the language data does not come from a single ranking, RCD is not guaranteed to find a ranking. This is shown with the fragment of Turkish in (160), where

the first winner-loser pair was created by the grammar $\text{IDENT(voice)}_{\sigma 1} \gg *V\text{t}fV$, and the second winner-loser pair was created by the opposite ranking.

(160)

	$*V\text{t}fV$	$\text{IDENT(voice)}_{\sigma 1}$
a. $\text{at}f\text{-i} \succ \text{at}\zeta\text{-i}$	L	W
b. $\text{ta}\zeta\text{-i} \succ \text{ta}f\text{-i}$	W	L

Given the Support in (160), RCD will not be able to find a constraint that has at least one W and no L's in its column. With no constraints to install, the Support cannot be emptied out, and RCD stalls. In situations like these, constraint cloning can potentially let RCD find a grammar, as explained in the next section.

4.3.2 Cloning RCD

Constraint cloning (Pater 2006, 2008b) is a mechanism for finding a grammar given inconsistent language data. Cloning attempts to find a grammar by duplicating an existing constraint, and making each copy of the original constraint applicable to a subset of the lexical items that appear in the Support. In the simplest case, each winner-loser pair in the Support contains a unique lexical item, but this assumption is not necessary for successful cloning.

The result of cloning $*V\text{t}fV$ in (160) is shown in (161). There are now two clones of $*V\text{t}fV$, and each one has a limited domain: One clone has the lexical item *taf* in its domain, and other clone has *atf*.⁴ Additionally, each item in the domain of a clone is annotated for the constraints that are the source of the conflict, in this case, $\text{IDENT(voice)}_{\sigma 1}$. For more about the need to annotate the domains with conflictors, see §2.4.4.

⁴The lexical item both winners share, the possessive suffix, is dealt with separately, see below.

(161)

	$*V_{\text{tf}}V_{\langle \text{taf}, \text{IDENT}_{\sigma_1} \rangle}$	$\text{IDENT}(\text{voice})_{\sigma_1}$	$*V_{\text{tf}}V_{\langle \text{af}, \text{IDENT}_{\sigma_1} \rangle}$
a. $\text{ta}\zeta\text{-i} \succ \text{taf-i}$	W	L	
b. $\text{af-i} \succ \text{a}\zeta\text{-i}$		W	L

Once $*V_{\text{tf}}V$ is cloned, RCD can apply to the Support exactly as described in (159): First, $*V_{\text{tf}}V_{\langle \text{taf}, \text{IDENT}(\text{voice})_{\sigma_1} \rangle}$ is installed, and the first winner-loser pair is removed from the Support. Then, $\text{IDENT}(\text{voice})_{\sigma_1}$ is installed, and the second winner-loser pair is removed. With the Support emptied out, RCD is done, and the constraint that was left over, $*V_{\text{tf}}V_{\langle \text{af}, \text{IDENT}(\text{voice})_{\sigma_1} \rangle}$, is added at the bottom of the hierarchy. The resulting stratified constraint hierarchy, or the grammar, is $*V_{\text{tf}}V_{\langle \text{taf}, \text{IDENT}(\text{voice})_{\sigma_1} \rangle} \gg \text{IDENT}(\text{voice})_{\sigma_1} \gg *V_{\text{tf}}V_{\langle \text{af}, \text{IDENT}(\text{voice})_{\sigma_1} \rangle}$.

The Cloning RCD Algorithm takes a Support and returns a stratified constraint hierarchy, just like the original RCD as given in (159). The Cloning RCD differs by potentially returning a hierarchy in which some constraints are cloned.

The cloning RCD is described formally in (162). It is identical to the original RCD in its installation procedure (162a,b). Cloning is triggered by a non-empty Support that has no constraints available for installing (162c). The algorithm chooses a constraint to clone by considering relevant candidates. Candidates for cloning are constraints that have at least one W in their column (162c-i). Naturally, candidates for cloning also have at least one L in their column, since if there were any constraints that had at least one W and no L's in their column, they could have been installed directly, without cloning. Of the candidates for cloning, constraints that have the smallest total of W's and L's in their column are preferred (162c-ii). If multiple constraints tie for the fewest W's and L's, one of them is chosen at random (162c-iii). Then, the cloning subroutine (described in 163 and 167 below) is called, which takes the current Support and the constraint to clone, and

returns a new Support (162c-iv). Once a constraint is cloned, RCD continues its attempt to install constraints and empty out the Support.

(162) **Cloning RCD Algorithm**

Given a Support S ,

not-yet-ranked constraints := a set of constraints in S .

H := a new constraint hierarchy.

While S is not empty, repeat:

- a. *current-stratum* := all the constraints in *not-yet-ranked constraints* that have at least one W and no L's in their column in S
- b. If *current-stratum* $\neq \emptyset$,
 - i. remove winner-loser pairs that are assigned a W from any constraint in *current-stratum*.
 - ii. put *current-stratum* as the next stratum in H , and
 - iii. remove *current-stratum* from *not-yet-ranked constraints*
- c. If *current-stratum* = \emptyset ,
 - i. *cloning-candidates* := the constraints in *not-yet-ranked constraints* that have at least one W in their column
 - ii. *cloning-candidates* := $\min(\text{W+L}, \text{cloning-candidates})$ ⁵
 - iii. *cloning-candidate* := some constraint $C \in \text{cloning-candidates}$
 - iv. $S := \text{clone}(S, C)$

Put *not-yet-ranked constraints* as the next stratum in H .

Return H .

⁵The function *min* takes a set of constraints in a Support and a type of object to count, and returns the subset of constraints that have the smallest number of the object to count. In this case, *min* counts non-empty cells (i.e. W's and L's).

The cloning subroutine is described formally in (163). It starts by identifying constraints that are more general than the constraint to clone (163a). This is done because W's and L's from general constraints will have to be masked, as described in §4.2.5. Two clones are made, one to collect winners that are assigned a W by the constraint to clone, and one for winners that are assigned an L (163b). Each clone is simply a copy of the original constraint, i.e. it is the same function from linguistic objects to violation marks. Once copied, the clones are given the empty set as their domain (163c), which means that they no longer assign violation marks to any linguistic object. The clones are added to the Support (163d), and since their domains are empty, their columns don't have any W's or L's in them. Now, winners that get W's or L's from the original constraints are divided between the clones. As the algorithm is stated here, the whole winner is put in the domain of a clone, rather than some morpheme(s) inside it. The issue of finding the morphemes that are responsible for the conflicting ranking arguments is discussed in §4.3.3 below. Starting with winners that the original constraint assigns a W to (163e), each winner is added to the domain of the W-collecting clone (163e-i), which causes the W-collecting clone to assign a W to the winner. Each winner is also annotated with a reference to the constraint(s) that caused the conflict, i.e. the constraint(s) that assign an L to the winner (163e-ii). Finally, if there are more general constraints that assign W's to the winner, those W's are masked from the Support, as explained in §4.2.5. The same procedure applies to the winners that the original constraint assigns an L to (163f), but with the W's and L's switched around. After the clones are properly created, the original constraint is removed from the Support (163g).

(163) **Cloning subroutine** (preliminary version, see final version in (167))

Given a support S and a constraint to clone $C \in S$,

- a. *general constraints* := constraints that assign a superset of the W's and L's that C assigns.

- b. Create two constraints, C_w and C_L , such that for any linguistic object x , $C(x) = C_w(x) = C_L(x)$.
- c. Make \emptyset the domain of C_w and C_L .
- d. Add C_w and C_L to S .
- e. For each winner opt that C assigns a W to,
 - i. $conflictors :=$ the constraints in S that assign an L to opt .
 - ii. Add $\langle opt, conflictors \rangle$ to the domain of C_w
 - iii. Mask any W's that *general constraints* assign to opt .
- f. For each winner opt that C assigns an L to,
 - i. $conflictors :=$ the constraints in S that assign a W to opt .
 - ii. Add $\langle opt, conflictors \rangle$ to the domain of C_L
 - iii. Mask any L's that *general constraints* assign to opt .
- g. Delete C .
- h. Return S .

Like the original RCD, the Cloning RCD is not guaranteed to empty out the Support and produce a stratified constraint hierarchy. For example, the presence of a harmonically bounded winner will prevent the algorithm from finding a grammar, and no cloning will help with that, as seen in (137) and (138). Tesar & Smolensky (2000) prove that the original RCD is guaranteed to find a grammar given data that was produced by a consistent grammar. It is likely that the cloning RCD has the same condition for success, but a general formulation of the kinds of Supports that the Cloning RCD will be able to process completely is a matter for future research.

4.3.3 The domain of cloned constraints

The Cloning RCD was defined in (162) to apply to any Support, but it was designed with a more specific goal in mind. The case studies in chapters 2 and 3 explored speakers'

ability to learn a morphological category (the plural in Hebrew, the possessive in Turkish) whose phonological expression involved partially unpredictable behavior, and also project the partial predictability onto novel items. For example, a Turkish *tʃ*-final noun can keep the voiceless [tʃ] in the possessive, or it can alternate with the voiced [dʒ]. The choice between the voiceless and the alternating stop is partially predictable given the size of the noun: Among the existing *tʃ*-final nouns of Turkish, the alternators are a minority among the mono-syllabic nouns, and a majority among the poly-syllabic nouns. Speakers replicate this difference in novel nouns, choosing alternating stops more often with poly-syllables than with mono-syllables.

To achieve speakers' ability to replicate lexical trends, lexical items are added to the domain of clones, based on each item's behavior with respect to the clone. Since the clones assess the morpho-phonological properties of lexical items, it follows from (162) that the domains of clones contain lexical items that share morpho-phonological properties. Once these domains are set up, they give speakers access to the relative prevalence of each pattern in the lexicon, allowing them to project this relative prevalence onto novel items.

The point to develop here is the exact nature of the domain of cloned constraints. Given two winners that require opposite constraint rankings, and hence are put in the domains of two different clones, it is not a logical necessity to add the entire winner to the domain of the clone. It could be that some part of the winner, e.g. its root, is put in the domain of the clone. A related question is about the ability of a clone to assess violations: If a clone of *VtʃV has the bi-morphemic form [taɖʒ-i] 'crown.POSSESSIVE' in its domain, how does it treat a form that has just one of the two morphemes, such as the homophonous [taɖʒ-i] 'crown.ACCUSATIVE'? And what happens if an additional morpheme intervenes between the root and the possessive suffix, e.g. [taɖʒ-i-ni] 'crown.POSSESSIVE.ACCUSATIVE'? These questions are addressed in this section.

The answer I offer is that when the Cloning RCD adds a poly-morphemic word to the domain of a clone, it separates the word into its immediate morphological constituents,

i.e. the morphologically outermost affix and its stem. Then, the Cloning RCD adds an ordered triplet to the domain of the clone, which consists of the stem, the category of the outermost affix, and the conflicting constraints. For example, given the form *tağ-i* ‘crown.POSSESSIVE’, the ordered triplet will consist of the root /taɟ/, the morphological category “POSSESSIVE”, and any relevant constraints.

Effectively, this decomposition of the form allows the speaker to learn two things about the grammar of their language, simultaneously: The speaker learns a fact about the behavior of the root /taɟ/, and a fact about the possessive affix. Each of these facts can influence the speaker’s treatment of novel words. To see how, consider the fragment of Turkish in (164), taken from TELL (Inkelas et al. 2000), where *amaɟ* and *anaɟ* behave consistently in the possessive and in the accusative, but the final stop of *avuɟ* is voiced in the possessive and voiceless in the accusative.

(164)	Bare noun	Possessive	Accusative	
a.	amaɟ	amaɟ-i	amaɟ-i	‘goal’
b.	anaɟ	anaɟ-i	anaɟ-i	‘female cub’
c.	avuɟ	avuɟ-u	avutɟ-u	‘fist’

Making a Support from (164) yields (165), and running the Cloning RCD on it yields the grammar in (166).

more likely than the accusative to cause inter-vocalic voicing, or vice versa: In addition to nouns like *avutf*, which voice in the possessive but not in the accusative, there are also roots that voice in the accusative but not in the possessive, such as *kukutf* ‘stone (of a fruit)’, possessive *kukutf-u*, accusative *kukuɕ-u*.

In the current proposal, the learner is free to learn the behavior of each root and affix combination separately if they have observed this behavior in the ambient language. They are not forced to assign a single behavior to each root. The learner is biased, however, to assign consistent behavior to nouns across affixes, as discussed in §4.3.4. If some noun has been observed with a voiceless stop with one or more affixes, it is likely to have a voiceless stop in forms of the noun that the learner hasn’t observed yet. If a noun has been observed to alternate with some affixes and not to alternate with others, the speaker is free to choose either behavior with forms of the noun that they haven’t observed yet.

The final version of the cloning subroutine of the Cloning RCD Algorithm is given in (167) below. It differs from (163) in the kind of object that is added to the domain of a clone. Rather than adding an ordered pair of a winner and conflicting constraints, (167) defines an ordered triplet of a stem, an affix and a set of conflicting constraints. If the winner is mono-morphemic, it is defined as the stem, and the affix slot remains empty (167e). If the winner is poly-morphemic, it is decomposed into its immediate constituents, i.e. the outermost affix and the stem that it attaches to. The outermost affix refers to the highest affix in a morphological tree structure, or in a derivational model, the last affix in a derivation.

(167) **Cloning subroutine** (final version)

Given a support S and a constraint to clone $C \in S$,

- a. *general constraints* := constraints that assign a superset of the W’s and L’s that C assigns.
- b. $C_w := C_L :=$ a constraint such that for any x , $C(x) = C_w(x) = C_L(x)$.

- c. Make \emptyset the domain of C_W and C_L .
- d. Add C_W and C_L to S .
- e. For each winner opt that C assigns a W to,
 - i. If opt is morphologically complex,
 - $opt_{AFF} :=$ the outermost affix in opt .
 - $opt_{STEM} :=$ the stem of opt_{AFF} .
 Else,
 - $opt_{AFF} :=$ null.
 - $opt_{STEM} := opt$.
 - ii. $conflictors :=$ the constraints in S that assign an L to opt .
 - iii. Add $\langle opt_{STEM}, opt_{AFF}, conflictors \rangle$ to the domain of C_W
 - iv. Mask any W's that *general constraints* assign to opt .
- f. For each winner opt that C assigns an L to,
 - i. If opt is morphologically complex,
 - $opt_{AFF} :=$ the outermost affix in opt .
 - $opt_{STEM} :=$ the stem of opt_{AFF} .
 Else,
 - $opt_{AFF} :=$ null.
 - $opt_{STEM} := opt$.
 - ii. $conflictors :=$ the constraints in S that assign a W to opt .
 - iii. Add $\langle opt_{STEM}, opt_{AFF}, conflictors \rangle$ to the domain of C_L
 - iv. Mask any L's that *general constraints* assign to opt .
- g. Delete C .
- h. Return S .

The decomposition of winners into their immediate constituents gives the learner access to lexical information about stems, affixes, and the constraint conflicts that they involve. This allows the learner to project their grammar onto combinations of stems and affixes that they haven't seen before, such as a known stem and a known affix that were previously only observed separately, or a novel stem with a known affix, etc.

4.3.4 Applying a grammar with cloned constraints

The grammar in (168) below is repeated from (166), with the addition of the dative form of *anaŋf*. It allows $*V\text{ŋ}V$ to rank either above or below $\text{IDENT}(\text{voice})$. In any given derivation that uses (168), only one ranking is chosen, so the grammar is categorical for any given derivation. The choice of ranking, however, depends on the input to the derivation and how well it matches the items listed in the grammar, so the choice of ranking can be probabilistic in some cases and categorical in others.⁶

$$(168) \quad *V\text{ŋ}V \begin{matrix} \langle \text{ama}\text{ŋf}, \text{POSS}, \text{IDENT}(\text{voice}) \rangle \\ \langle \text{ama}\text{ŋf}, \text{ACC}, \text{IDENT}(\text{voice}) \rangle \\ \langle \text{avu}\text{ŋf}, \text{POSS}, \text{IDENT}(\text{voice}) \rangle \end{matrix} \gg \text{IDENT}(\text{voice}) \gg *V\text{ŋ}V \begin{matrix} \langle \text{ana}\text{ŋf}, \text{POSS}, \text{IDENT}(\text{voice}) \rangle \\ \langle \text{ana}\text{ŋf}, \text{ACC}, \text{IDENT}(\text{voice}) \rangle \\ \langle \text{ana}\text{ŋf}, \text{DAT}, \text{IDENT}(\text{voice}) \rangle \\ \langle \text{avu}\text{ŋf}, \text{ACC}, \text{IDENT}(\text{voice}) \rangle \end{matrix}$$

Given the grammar in (168), the ranking between $*V\text{ŋ}V$ and $\text{IDENT}(\text{voice})$ in any single derivation depends on how well the input to the derivation matches the domains of the clones of $*V\text{ŋ}V$. If the matching is complete, the choice of ranking is categorical. If the matching is partial, the choice is potentially stochastic.

If the speaker wishes to reuse a form that they have heard before, such as the possessive form of *amaŋf*, they will find an exact match for it in the high-ranking clone of $*V\text{ŋ}V$. Using the grammar $*V\text{ŋ}V \gg \text{IDENT}(\text{voice})$, the outcome can only be *amaŋŋ-i*. In this case, then, the choice of ranking is categorical.

Given a novel *ŋ*-final root, however, and asked to derive its possessive form, there is no single listing in the grammar that matches the outcome perfectly. There are, however,

⁶For a comparison with other probabilistic approaches in OT, see §4.3.7.

three listed possessive forms. Since two of the listed possessives are in the high-ranking clone, and one is in the low-ranking clone, the speaker is twice as likely to derive the novel root using the high-ranking clone, i.e. the effect of the grammar in (168) is stochastic when deriving a possessive form of a novel root. Deriving the dative form of the same novel root would be categorical, with a single listing of a dative on the low-ranking clone.

The effect of the grammar in (168) is not necessarily categorical with a new combination of known morphemes. If the speaker wished to derive the dative form of *anatf*, they will find two matches for the root *anatf* in the low-ranking clone of *VtʃV, and one match for the dative in the same low-ranking clone. So the dative form of *anatf* is guaranteed to be derived using the low-ranking clone. The dative form of *amatf*, however, presents a conflict: There are two listings for the root *amatf* with the high-ranking clone, and one listing for the dative with the low-ranking clone. The speaker will have to weigh both factors in making their decision. It is not necessarily the case that roots and affixes have the same weight in determining the outcome of the grammar, since for any given combination of root and affix, it is likely that there will be many more listings for the affix than for the root, but it is not clear that in real languages, the affix generally prevails in such cases. The current proposal limits itself to pointing out that a grammar like the one in (168) can potentially generate a stochastic outcome given a new combination of two known morphemes.

A separate question about the application of a grammar with cloned constraints has to do with the scope of the clone over a phonological form that has multiple morphemes in it. The final voiceless stop of the root *avutʃ*, for instance, becomes voiced in the possessive, but it surfaces faithfully in the accusative (164, 168). This root can combine with both affixes to make the form *avuç-~~u~~-nu* ‘fist.POSS.ACC’,⁷ with the possessive followed by the accusative

⁷The morphological affiliation of the *n* that appears between the affixes is unclear. An *n* appears in Turkish whenever a third person possessive suffix is followed by a case suffix. Since this *n* also appears before consonant-initial case suffixes, it is not there to repair a hiatus.

(the opposite order is ill-formed in Turkish). Unsurprisingly, the possessive affix, which is closer to the root, prevails.

The local effect of lexically-specific behavior is discussed by Pater (2008b), who suggests that a locality condition be built into the definition of a lexically-specific constraint: A markedness constraint assesses a violation only if the marked structure it specifies contains a phonological exponent of an exceptional morpheme that's in the domain of the constraint. This is a representational approach to locality. The alternative that I would like to suggest here is the derivational approach to locality, as suggested by Wolf (2008b), based on a derivational model of Optimality Theory (OT-CC, McCarthy 2007a). If the form *avuç-u-nu* is derived by first combining *avuç* with the possessive, then the final stop will become voiced, following the specification in (168). In the next step of the derivation, the addition of the accusative no longer creates a derived environment for the markedness constraint *VɬV, so despite the fact that the combination of the root *avuç* and the accusative suffix is specified as one that blocks inter-vocalic voicing, the root-final stop cannot be turned voiceless again. However, as pointed out by Pater (in preparation), it is not yet known how to make derivational models of OT compatible with constraint demotion algorithms of the type used here. A full integration of the derivational approach to locality will require additional research.

4.3.5 Searching for the UR of affixes

The discussion in §4.3 has so far presupposed the existence of a Support that contained observed forms of the language as winners, and in addition, underlying representations and losers that were supplied by the analyst. The language learner will have to provide their own underlying representations and losers, of course. This section offers a mechanism for discovering the UR's that the learner needs, while still assuming that losers are provided by the analyst.

A fully general mechanism for finding underlying representations algorithmically is yet to be proposed, although significant headway was made by Tesar (2006), Merchant (2008), and in parallel lines of work, also by Boersma (2001) and Apoussidou (2007) and by Jarosz (2006). A central component of the current proposal is the assumption that roots are always surface-true, so the search for non-surface-true forms is limited to affixes. Since cross-linguistically, affixes are small in size and in number compared to roots, the search for their UR's is likely to produce manageable results in realistic cases.

The algorithm starts with a given affix, such as the possessive affix in Turkish, and a set of stems that combine with it. In this situation, there might be a lexically-specific phonological process involved, also known as a lexical trend. Each affix defines a set of paradigms, or a set of pairs of output forms, where each pair consists of a base and a derived form. A prerequisite for discovering the lexical trend is to assume the surface form of the base as its underlying form. The reasons for this prerequisite are discussed in detail in §4.4, but in a nutshell, the problem is that assigning non-surface-true information to the base could prevent the learner from cloning constraints and listing roots in their domains, making lexical trends unavailable to the grammar.

In the cases presented below, the base is a simple bare root. In some languages, however, bare roots do not surface, and the bases of affixation already have some obligatory inflection on them, such as a third person marker or a nominative marker. To learn a trend in such a situation, the learner will have to identify the presence of this affix and strip it off. This extra step is abstracted from in the present discussion, and the assessment of its impact on the process is left for future work.

In the Turkish possessive, assuming the surface form of the base and the surface form of the possessive suffix as their respective UR is all the speaker needs to learn the lexical trend. These surface-true underlying forms will allow the speaker to discover conflicting evidence about the ranking of, e.g., *V₁V and IDENT(voice), as discussed in §4.4 and in chapter 2.

In other cases, more work is needed: In the Dutch past tense, for example, the past tense suffix can show up as either [-tə] or [-də] (169). The underlying form of the suffix lies in a fairly large space of plausible hypotheses: It could be identical to just one of the surface forms, i.e. /-tə/ or /-də/, or it could be both forms (where they are allowed to compete as allomorphs), or it could be some non-surface-true form, such as / [+voice] -də/ with a floating [+voice] feature, or it could be a combination of surface-true form(s) with non-surface-true form(s).

(169)

Imperative	Past tense	
stɔp	stɔp-tə	‘stop’
tɔp	tɔb-də	‘worry’

Given the assumption that the UR’s of [stɔp] and [tɔp] are /stɔp/ and /tɔp/, the learner can start their search for the UR of the past tense suffix by testing each of its surface forms as a hypothesis. This is a good place to start, since with *n* surface forms of the suffix, there are exactly *n* hypotheses to test. In (170), for example, both roots are tested with the hypothesis that the UR of the suffix is /tə/. This hypothesis must be rejected, since it generates a harmonically bounded winner, as seen in the winner-loser pair that has no W’s in its row (170b).

(170)

/... p/ + /tə/	IDENT(voice) _{ROOT}	IDENT(voice) _{ONSET}
a. stɔp-tə > stɔb-də	W	W
b. tɔb-də > tɔp-tə	L	L

The hypothesis that the UR of the affix is /də/ is tested in (171). This hypothesis generates an inconsistent grammar, but it is a grammar that can be rendered consistent by

cloning. Subjecting the Support in (171) to the Cloning RCD would return a consistent grammar with one of the constraints in it cloned.

(171)

$/\dots p/ + /d\partial/$	IDENT(voice) _{ROOT}	IDENT(voice) _{ONSET}
a. $st\partial p-t\partial \succ st\partial b-d\partial$	W	L
b. $t\partial b-d\partial \succ t\partial p-t\partial$	L	W

Once a consistent grammar is found, the speaker can declare the search for the UR of the affix successful. With the the UR of the suffix in place, the learner has UR's for all of their morphemes, since roots are taken to always have surface-true UR's.

It is worth pointing out that a more permissive hypothesis, which assumes the two surface forms of the past tense affix as underlying representations of two allomorphs that are allowed to compete, as in (172), actually fares worse than the simple hypothesis in (171). With both forms of the affix to choose from, the winner in (172b) is harmonically bounded. Cloning IDENT(voice)_{ROOT} can't help, because once IDENT(voice)_{ROOT}(*stop,PAST, \emptyset)* is installed, and (172a) is removed from the Support, there is no W in the Support to empty it out.

(172)

$/\dots p/ + \{ /d\partial/, /t\partial/ \}$	IDENT(voice) _{ROOT}	IDENT(voice) _{ONSET}
a. $st\partial p-t\partial \succ st\partial b-d\partial$	W	
b. $t\partial b-d\partial \succ t\partial p-t\partial$	L	

In the case of the Dutch past tense, then, the existence of multiple surface forms of the suffix did not cause an explosion of the search space for the UR. Trying one surface form

at a time (which grows linearly with the number of surface forms) is sufficient for finding one form that can serve as the UR.

In some cases, however, the UR of the affix cannot be simply one of its surface forms, as in the Korean accusative (173), taken from Albright (2008), and discussed in fuller detail in §4.4.2. Word-finally, the only coronal obstruent that Korean allows is an unreleased voiceless unaspirate dental. Upon the addition of a vowel-initial suffix, if an aspirated stop emerges, it can either be dental or pre-palatal.

(173)	Bare noun	Accusative	
	nat ^ɾ	nat ^h -il	‘piece’
	nat ^ɾ	nat ^h -il	‘face’

Given the assumption that the two roots in (173) are underlyingly identical to their base form, i.e. /nat^ɾ/, taking the surface form of the accusative suffix as its UR cannot derive the different observed forms, as shown in (174), where the winner [nat^h-il] is harmonically bounded.

(174)	/nat ^ɾ / + /il/	IDENT(asp)	IDENT(anterior)	IDENT(voice)
a.	nat ^h -il > nad-il	L		W
b.	nat ^h -il > nat ^ɾ -il		L	

To find out what needs to be changed about the UR of the affix, the learner can compare the intended winner [nat^h-il] to the current winner [nat^ɾ-il], given their current hypothesis about the UR of [nat^h-il]. This is shown in (174b), and it reveals that the accusative involves a change of the feature [anterior], and prompts the speaker to add [anterior] as a floating feature to the affix.

Adding the feature that causes an unfaithful mapping to the UR of the suffix will now rescue [nat^h-il] from its predicament (175).

(175)

/nat [?] / + /[+ant] il/	IDENT(anterior)	MAX(float)
a. nat ^h -il > nat ^h -il	W	L
b. nat ^h -il > nat ^h -il	L	W

Running the Cloning RCD on (175) can produce a consistent grammar by cloning either of the two available constraints. If IDENT(anterior) is cloned, the resulting grammar would be the one in (176).

(176) IDENT(anterior)_{<nat[?] (piece),ACC,MAX(float)>} ≫ MAX(float)
 ≫ IDENT(anterior)_{<nat[?] (face),ACC,MAX(float)>}

The addition of a floating [anterior] feature to the accusative suffix resolved the harmonic bounding in (174) and allowed the speaker to reach the grammar in (176).

So far, the learner was shown to be able to deal with cases of multiple allomorphs of a suffix, as in the Dutch past tense, and with cases of a single surface form of the suffix that required floating structure, as in Korean accusative. If the language presents both allomorphy and the need for floating features in the context of a single suffix, the learner will need to consider both of these aspects of the phonology in their search for the UR.

The learner will have to balance two strategies: Trying out combinations of surface forms as competing allomorphs, and trying out adding floating features to (any of) the surface forms. Since combining surface forms makes the hypothesis space grow exponentially with the number of forms involved,⁸ while adding floating features only

⁸Two surface forms give rise to three combinations, three forms give rise to 7 combinations, and n forms give rise to $2^n - 1$ combinations.

doubles the number of hypotheses,⁹ it makes sense not to exhaust the combinations of surface forms before floating features are tried out. The learner should interleave the two strategies: Start with each surface form of the affix as the UR, try adding floating features if necessary, and only if a consistent grammar could not be found, go on to try combinations of surface forms.¹⁰

The search for the UR of affixes is given formally in (177). It starts with a set of bases, and some morphological category that they can serve as stems for. The notation $[b-a]$ refers to the surface form that results from combining b with a , but does not presuppose linear order. Before the search starts, the underlying form of the bases are fixed as their surface forms (177c). Then, the learner starts collecting hypotheses about the underlying form of the affix. A complete hypothesis about the underlying form of an affix can be a single string of phonological elements, or a set of strings, like the set $\{/eɪ, æn, ə, ŋ/\}$ for the English indefinite article. Looking at one paradigm at a time, though, as in (177e), each hypothesis will be just one string.¹¹

Next, these strings are combined to form sets of strings (177f). The set of hypotheses starts with single strings, then goes on to pairs of strings, as so forth. This ordering is meant to favor hypotheses that minimize the number of strings in the UR of the suffix (as is standardly assumed in generative linguistics, e.g. in chapter 6 of Kenstowicz & Kisseberth 1979), since the first hypothesis that is tested and found to work is also the last

⁹The addition of floating features only doubles the number of hypotheses if two things are true: (a) All of the features that distinguish the intended winner from the most similar available winner, as determined by faithfulness violations, are added as floating features to the affix, and (b) these floating features are added to all of the allomorphs of the affix. If either of these assumptions is too strong, then the space of hypotheses will not just double, but grow even bigger.

¹⁰The Korean accusative is not free of allomorphy, since it surfaces as $[-riɭ]$ when attached to vowel-final stems (e.g. $pori \sim pori-riɭ$ ‘barley’). Tracing both $[-iɭ]$ and $[-riɭ]$ to the same underlying form was proposed in Odden (1993), but this analysis is not pursued here due to concern about the plausibility of deleting the $[r]$ after consonant-final roots, in light of McCarthy (2007b). In this case, then, the learner will have to try out each of $[-iɭ]$ and $[-riɭ]$, with and without a floating $[anterior]$, before they decide that both allomorphs are listed in the UR of the accusative, both with the floating $[anterior]$.

¹¹The term “string” is used here loosely to refer to an autosegmental phonological structure that can include floating features.

one tested. Each hypothesis in turn is tested (177g) by letting the hypothesized UR's map to the observed forms, augmented by losers that are supplied by the analyst. If the hypothesis is not successful, presumably because it gave rise to harmonically bounded winners, the learner tries to enrich the current hypothesis with floating features before abandoning it (177g-v).¹²

(177) Support Preparation, **the search for UR's**

- a. Given B , a set of well-formed surface forms, or bases,
- b. and given an affix a that can combine with any form in B to make a well-formed surface form $[b-a]$,
- c. For every $b \in B$, $/b/ := [b]$.
- d. $A := \emptyset$ (a set of hypotheses about $/a/$)
- e. For every b , find all the segments that are in $[b-a]$ but not in $[b]$. Add these segments as an element of A .
- f. $P :=$ a stratified hierarchy of hypotheses about $/a/$, such that the n^{th} stratum in P , $P_n = \{p \in \mathcal{P}(A) : |p| = n\}$
- g. For each stratum $P_n \in P$, starting with $n = 1$,
For each element $/a/$ in P_n ,
 - i. Make a Support S ,
 - ii. For each element in $b \in B$, designate $/b-a/$ as the UR of $[b-a]$
 - iii. Supply loser(s) as necessary, and add winner-loser pair(s) to S
 - iv. Run the Cloning RCD on S .
 - v. If RCD finds a consistent grammar, adopt $/a/$ and stop. Otherwise, find the harmonically bounded winners in S , and if they are assessed L's by

¹²It is not known what the learner should do if a multi-string hypothesis needs to be enriched with floating features. Are the floating features added to each of the strings, or only to some strings? Further research is required on this point.

faithfulness constraints, add the distinctive features that they refer to as floating features in /*a*/, and repeat steps i–iv.

The algorithm as it is formulated here does not guarantee that the learner will be protected from exploring an exponentially large number of hypotheses about the UR of the suffix; it simply biases the speaker to find the simplest successful hypothesis as early as possible. Since cross-linguistically, affixes are small in size and in number compared to roots, fixing the UR's of roots as necessarily surface-true and allowing non-surface true UR's only for affixes is likely to produce very manageable results in realistic cases.

While the procedure in (177) will find a grammar for the cases discussed in this chapter, a general characterization of the range of cases where (177) will succeed is a matter for future research.

4.3.6 Supplying losers

Recall that the Cloning RCD applies to a Support, which is a set of winner-loser pairs, where the winner and the loser in each pair are derived from a single underlying representation. The winners are given to the learner by the ambient language, since these are the surface forms that the learner hears. The underlying representations can be found given the method described in §4.3.5 above. This section now goes on to show how the learner gets the final piece of the puzzle, the losers.

In Error-driven learning, as proposed by Tesar & Smolensky (2000) et seq., the speaker starts with a grammar that potentially differs from the adult grammar. A discrepancy between the learner's current grammar and the target grammar is discovered when the learner passes an adult form through their grammar, and notices that the output of their own grammar is different from the adult form. In this situation, the learner's own output is marked as a loser, and it is paired with the adult form to make a winner-loser pair.

For instance, a child who is learning a language that allows codas, like Turkish, might produce the adult form [pak] as [pa], deleting the coda consonant. When the adult form

and the learner's form are different, i.e. an error is made, the learner pairs the adult form with their own form, to make a winner-loser pair (178).

(178)

	MAX	NOCODA
a. pak \succ pa	W	L

Applying RCD to the Support in (178) will give rise to the ranking MAX \gg NOCODA. This ranking in turn allows the learner to produce codas faithfully. Learning is error-driven in the sense that learning continues for as long as the learner generates forms that are different from the observed adult forms, and therefore the learner's grammar is not yet identical to the adult's grammar. If, as I suggest, learning lexical trends requires an adjustment to the grammar each time a new word is learned, then error-driven learning will continue for as long as new words are learned.

Initially, the child will assume the simple case, where one constraint ranking will account for all of the phonology of the language, no constraints are cloned, and no constraints list lexical items. A learner of Turkish will be able to maintain this hypothesis until they are confronted with positive evidence for inconsistency. For *t*-final words, this will happen when the learner discovers at least one *t*-final noun that alternates (e.g. *tat* \sim *tad-i* 'taste') and at least one *t*-final noun that doesn't (e.g. *at* \sim *at-i* 'horse'). When the first alternating noun is discovered (e.g. *tat* \sim *tad-i*), the learner will demote IDENT(voice) to rank below *VtV.¹³ The winner-loser pair *tad-i* \succ *tat-i* will be kept as evidence for the new ranking. Then, when the learner encounters the non-alternating *at*, their grammar will wrongly produce the alternating possessive form **ad-i*. If the learner observes that the

¹³Recall that is ranking is only necessary when the learner discovers the existence of morphological paradigms. In unanalyzed forms of the language, intervocalic [t]'s are allowed to stay voiceless (e.g. *ata* 'father' vs. *ada* 'island'), so the learner has previously learned that IDENT(voice) \gg *VtV as part of learning the phonotactics of the language.

adult form is actually *at-i*, they will form the winner-loser pair *at-i* \succ *ad-i*, which directly conflicts with *tad-i* \succ *tat-i*, as shown in (179).

(179)

	*VtV	IDENT(voice)
a. <i>tad-i</i> \succ <i>tat-i</i>	W	L
b. <i>at-i</i> \succ <i>ad-i</i>	L	W

At this point, inconsistency is detected, and *VtV is cloned. The resulting grammar is $*VtV_{tat} \ggg \text{IDENT(voice)} \ggg *VtV_{at}$. From this point on, the learner is actually juggling two conflicting grammars, not just one, since there are two grammatical ways for *VtV to be ranked relative to IDENT(voice). If the learner encounters a new *t*-final noun, such as *kat* ‘floor’, with the possessive form *kat-i*, they will need to decide whether *kat* belongs to the grammar $*VtV \ggg \text{IDENT(voice)}$ or to the grammar $\text{IDENT(voice)} \ggg *VtV$, or in other words, which clone of *VtV should list the new item.

To find out, the learner can simply try both grammars by temporarily listing *kat* with each clone of *VtV. The temporary listing will yield two different results: The observed adult form *kat-i* is produced by the lower clone of *VtV, and the ungrammatical **kad-i* is produced by the higher clone. Since the two grammars yielded different results, one of them being the adult form, the learner can pair the adult form with the other form, and add them as a new winner-loser pair to their Support. Running the Cloning RCD again yields the new grammar $*VtV_{tat} \ggg \text{IDENT(voice)} \ggg *VtV_{\{at, kat\}}$, where *kat* is permanently listed with the lower-ranking clone of *VtV.

Trying out both grammars also helps with weeding out words that aren’t affected by *VtV, such as the *l*-final *yil* ‘year’. A temporary listing of *yil* with either clone of *VtV generates *yil-i* as the possessive form, which is identical to the adult possessive form. Since both grammars agree on the winner, the learner can conclude that the ranking of *VtV is

irrelevant for the derivation of *yil*, and there is no need to update the grammar with this lexical item.

This process goes on with every new word the learner encounters, with non-*t*-final words going unlisted, and *t*-final words being listed with one of the clones of *VtV. The resulting grammar contains a list of alternating *t*-final nouns and a list of non-alternating *t*-final nouns. Now, when a learner encounters a novel *t*-final noun, and they don't know what the possessive form of it is, they can make an estimate that is based on the words they have learned. If the list of non-alternating *t*-final nouns has 102 items in it, and the list of alternating *t*-final nouns has 18 items in it (as in TELL, Inkelas et al. 2000), then the chance of the novel noun to be alternating is 18 out of (18+102), which is 15%.

This method for generating losers and using them to feed the Cloning RCD is essentially identical to the original proposal of Tesar & Smolensky (2000), with the added assumption that error-driven learning continues as long as the speaker continues to learn new lexical items. The need to run a new form through more than one grammar, however, raises a concern about the number of those grammars. If a learner has cloned n constraints, that means that they are potentially dealing with 2^n grammars, which in turn means that every new form they encounter must be run through each of these 2^n grammars, thus greatly increasing the computational load for the learner. This worry is almost certainly overstated here. Given that lexical trends can be independent of each other, as seen in §4.2.2, trying out all of their combinations will be wasteful, since it will suffice to test two grammars for any set of independent trends. Because the learner can find out whether trends are dependent on each other by inspecting the Support, they will be able to use this information to reduce the number of grammar to try out. The details of this mechanism are left for future work.

4.3.7 Exceptionality and variation

The Cloning RCD algorithm offered here presupposes the existence of only two kinds of phonological processes: Regular processes, which apply to all available lexical

items, and lexically-specific processes, which always apply to one list of lexical items and never apply to a second list of items. This is an oversimplification, of course. While lexically-specific processes typically do not involve variation for most of the items involved, variation is not completely absent. Of the 3002 stop-final nouns in TELL (Inkelas et al. 2000), for instance, the vast majority behave consistently, but 103 items (3%) show the voicing alternation optionally. Note that the data in TELL represents inter-speaker variation, since it records the knowledge of a single native speaker. A variable grammar is needed for the representation of a single speakers' grammar, not just for the grammar of the speech community.

In the Cloning RCD, the variable behavior of a lexical item can be represented in two ways: Either the lexical item is listed with both clones of a constraint, in which case it is predicted to undergo the relevant lexically-specific process 50% of the time, or the lexical item resists listing, in which case it is predicted to undergo the lexically-specific process as often as novel items do (cf. a similar suggestion in Pater 2008b). If the learner hears an item behaving inconsistently in the ambient language, it seems plausible that they will refrain from listing the item, or that they will list it twice. This approach predicts that lexical items that undergo a lexically-specific process optionally will show one of the two behaviors mentioned above; unfortunately, it is not known whether this prediction is correct or not.¹⁴

A different approach to variability in Optimality Theory is stochastic grammar (Boersma 1997, Boersma & Hayes 2001, et seq.), where constraints are arranged on a numerical scale, and each constraint defines a normal distribution somewhere on the scale. Each time the grammar is used in a derivation, a ranking value for each constraint is assigned by sampling

¹⁴Some suggestive, possibly promising, numbers come from Google searches on Hebrew *ot*-takers. Most Hebrew nouns take one of the plural suffixes, *-im* or *-ot*, categorically. Searching for the two plural forms for each item and comparing the number of hits, this categorical behavior is reflected in a rate of *ot*-taking that is close to 0% or to 100% for any given item. A small number of items have *ot*-taking rates in the 40–50% range, and smaller number still have rates in the 10–20% range. Interestingly, no items were found in the 50–97% range.

from its distribution, and if two constraints have non-negligibly overlapping distributions, their ranking relative to each other can change between derivations. Stochastic grammars are usually learned with the Gradual Learning Algorithm (GLA, Boersma 1997; Boersma & Hayes 2001), which approaches the target grammar gradually by incrementally adjusting the relative scaling of the constraints in response to errors. Stochastic grammar was designed to deal with cases of regular variability, where a phonological process is variable with little relation to any lexical item involved. In lexical trends, however, each known lexical item usually behaves categorically, and the trend created by the aggregation of lexical items causes stochastic behavior with novel items.

Zuraw (2000) offers an analysis of Tagalog's lexical trend of nasal substitution that combines the GLA with a constraint called USELISTED. The GLA learns a stochastic grammar that affects novel words, while USELISTED protects stored forms from variation. A similar analysis of exceptions to vowel harmony in Hungarian is offered by Hayes & Londe (2006).

Recently discovered problems with the GLA cast doubt on its usefulness in analyzing lexical trends. One such problem is raised in Hayes & Londe (2006): When the GLA notices a winner that needs high-ranking faithfulness, it promotes all the faithfulness constraints that prefer that winner. Since general faithfulness constraints, by definition, prefer more winners than specific faithfulness constraints, general faithfulness will be promoted too fast, causing the learner to learn a superset language. See also Tessier (2007) for a discussion of the same problem arising in learning the regular phonology of a language. Additionally, a rather serious problem with the GLA is that it is not guaranteed to converge on a ranking in certain situations, as discovered by Pater (2008a). It should be noted, however, that the USELISTED mechanism is conceptually separate from the GLA, and could potentially be used in conjunction with a more successful theory for learning stochastic grammar.

The cloning approach offered here and the USELISTED approach share a core property: They both incorporate lexical listing into an OT grammar, thus allowing a single grammar to apply categorically to known items and stochastically to novel items. Arguably, the cloning approach is more appealing on theoretical grounds, since it is more parsimonious: It relies on the familiar markedness and faithfulness constraints of OT, and does not introduce a new kind of constraint that directly accesses the lexicon. Additionally, the cloning approach makes an unmediated connection between lexical listing and the projection of trends, as both follow from the association of lexical items with clones. In contrast, the USELISTED approach relies on a separate learning mechanism to ensure that the stochastic grammar is synchronized with stored lexical entries.

4.4 Moving hidden structure into the grammar

The model proposed here builds speakers' knowledge of lexical trends into a constraint ranking, augmented with cloned constraints. If the language has an irregular phonological process, and the irregularity can be expressed in phonological terms, then the speaker uses cloned constraints to list the lexical items involved, and the resulting constraint ranking is used to project the lexical trend onto novel items.

One consequence of this approach is that information about inconsistent patterns in lexical items is built into the grammar rather than being stored in the lexicon. In Turkish, for instance, my analysis attributes the difference between alternating stops (e.g. *tat* ~ *tad-i* 'taste') and non-alternating stops (e.g. *at* ~ *at-i* 'horse') to lexically-specific rankings of faithfulness and markedness constraints. Both kinds of words have a voiceless stop in their UR's (i.e. /tat/, /at/), but the voiceless stop doesn't always surface faithfully. In contrast, the traditional generative analysis of Turkish (Inkelas & Orgun 1995; Inkelas et al. 1997; Petrova et al. 2006) attributes the difference to the underlying representations: Non-alternating stops are underlyingly voiceless (or aspirated in Petrova et al. 2006), and alternating stops are underlyingly unspecified for voice (i.e. /taD/, /at/).

My approach is in keeping with a central principle of generative linguistics, which seeks to identify predictable patterns in lexical items and use the grammar to derive them. My approach is not in keeping, however, with a tradition of attributing hidden structure to underlying representations. In Turkish, the alternating or non-alternating nature of a stem-final stop is hidden in the bare form of the noun, and it is discovered by examining the noun's suffixed form. In the traditional generative approach, the hidden structure is encoded in the roots, while my approach attributes the hidden structure to the grammar via listing of roots with clones. See, however, Hayes (1995b, 1999) for arguments against the use of underlying representations to encode hidden structure, including an analysis of Turkish along the lines I propose here in Hayes (1995b).

In this section, I examine the mechanism of attributing hidden structure to various parts of the linguistic apparatus and how it relates to learning lexical trends. I will show that lexical trends can be discovered only if the learner is biased to attribute hidden structure to the grammar first, or to a combination of the grammar and the underlying representations of affixes. When hidden structure is forced into underlying representations of the roots, it is “lost” to the grammar, and speakers are predicted not to learn lexical trends in such cases.

4.4.1 Hidden structure in the grammar: Turkish

The distribution of voicing alternation in Turkish is available to speakers: They know how many words have alternating stops and how many have non-alternating stops, and they keep this information separately for the stops in the different places of articulation, and within each place, for mono-syllabic nouns separately from poly-syllabic nouns.

The first step in making this information available to the grammar is assuming that the bare form of the noun is also its underlying representation. This will force the learner to attribute the behavior of the stem-final stop to the grammar, as seen in (180). The derivations of *at-i* and *tad-i* require different grammars because they both have a voiceless stop underlyingly.

- (180) a. The UR's of [at] and [tat] are /at/ and /tat/
 b. The UR of the possessive is /I/ (a high vowel)
 c. /at + I/ → [at-i] requires IDENT(voice) ≫ *VtV
 /tat + I/ → [tad-i] requires *VtV ≫ IDENT(voice)

The inconsistent ranking requirements in (180) trigger constraint cloning, and then a listing of words under the two clones, as discussed above. In contrast, the classic generative analysis of Turkish (Inkelas & Orgun 1995; Inkelas et al. 1997) assumes that the stem-final stops in *at-i* and *tad-i* differ in the underlying representation, as in (181).

- (181) a. The UR's of [at] and [tat] are /at/ and /taD/
 b. The UR of the possessive is /I/ (a high vowel)
 c. /at + I/ → [at-i] requires IDENT(voice) ≫ *VtV

at + I	IDENT(voice)	*VtV
a.  at-i		*
b. ad-i	*!	

- d. /taD + I/ → [tad-i] is consistent with IDENT(voice) ≫ *VtV

taD + I	IDENT(voice)	*VtV
a. tat-i		*!
b.  tad-i		

In the UR-based analysis, the grammar is consistent for all the words of the language (i.e. IDENT(voice) \gg *VtV), and therefore the learner is left without a way to build lexical statistics into their grammar.

In principle, speakers can find the relevant lexical statistics by going directly to the lexicon and extracting the relevant information from it, as is practiced in Analogical Modeling of Language (AML, Skousen 1989, 1993) and in Spread Activation models (Schreuder & Baayen 1995; Krott et al. 2001). When going to the lexicon directly, however, the speaker will not be biased by UG to find only grammatically-principled generalizations. Any kind of regularity in the lexicon could be discovered and projected onto novel items, contrary to fact: In the Turkish lexicon, there is a trend for more voicing alternations after high vowels than after low vowels, yet speakers show no sign of having learned this trend. Since cross-linguistically, vowel height cannot affect the voicing of a following stop, this is the expected result. To learn all and only the phonologically plausible generalizations about their lexicon, language speakers must encode these generalizations in their grammar, where they can benefit from the biases imposed by UG.

Assuming the base form of a noun as its underlying representation means that any additional aspects of the noun's behavior that are not directly observable in the base form will have to be attributed to other aspects of the linguistic system. Given the standard OT framework that uses underlying representations of roots and affixes and a constraint ranking, if hidden properties of roots are blocked from being attributed to those roots, hidden properties can only be attributed to the underlying representations of affixes or to the grammar. In the Turkish case, the difference between *at* and *tat* could logically be attributed to the allomorph of the possessive suffix that they take: *at* would take a simple high vowel, while *tat* would take an affix that consists of a high vowel and a floating [+voice] feature, as in (182).

- (182) a. The UR's of [at] and [tat] are /at/ and /tat/
b. The possessive has two allomorphs: /I/ and /[+voice] I/

- c. /at + I/ → [at-i]
 /tat + [+voice] I/ → [tad-i]
- d. Each allomorph of the possessive lists the roots it takes:
 /I/ takes /at/, /ot/, /sepet/, ...
 /[+voice] I/ takes /tat/, /kanat/, ...

Assuming that the floating [+voice] is protected by MAX(float), as in Wolf (2007), then the single constraint ranking MAX(float) ≫ IDENT(voice) will derive all the words of the language. In this scenario, either each root would be marked for the affix it takes, or equivalently, each affix will be marked for the roots it takes. The grammar would be consistent: Faithfulness to underlying [voice] specification would outrank *VtV, and faithfulness to floating features will be ranked higher than simple faithfulness. This scenario makes a slight improvement over the attribution of voicing information to roots: Since roots will be listed with two different affixes, the learner will have information about how many roots there are of each kind, and thus learn a lexical trend. However, roots of all sizes and of all final stops will be listed by the same two allomorphs of the possessive suffix, preventing the Turkish learner from identifying the trends for each place and size separately. Encoding hidden structure by proliferating affix allomorphs, then, does not allow the learner to discover the full range of trends in their language. In principle, the learner could assign allomorphs of the possessive suffix for nouns of different sizes and final stops, but there would be no reason for them to do that, since simply stipulating two allomorphs would be enough to make the grammar consistent.

Since encoding the hidden behavior of lexical items in the underlying representations of either roots or suffixes leaves the learner with no way or reason to identify lexical trends, encoding such behavior in the grammar is left as the only logical option. Capturing hidden behavior in terms of cloned constraints ensures that lexical trends are identified in terms of constraints, i.e. it ensures that trends are captured in phonological terms, using the variety

of phonological primitives that constraints are sensitive to, such as marked combinations of features, preferred alignments of phonological elements, positional faithfulness, etc.

Contrasted with traditional generative analyses, the proposal made here “reverses” the effect of the phonology. Instead of assigning the hidden aspects of bases to their underlying representation, and then neutralizing them in the unaffixed form, as is done traditionally, I propose that the surface forms of bases are assumed as their underlying form, and any properties of the base that emerge only in suffixed forms are achieved by constraint interaction. In the simple case of Turkish, where the only hidden property of nominal roots is the voicing of their final stop, the analysis in terms of cloned constraints is not only clearly feasible, it is also the only analysis that allows speakers to capture the variety of lexical trends that the language has.

Assuming the base form as the underlying representation has the added benefit of obviating the search for non-surface-true underlying representations. This search requires a significant amount of computation, as shown by Tesar (2006) and Merchant (2008), and in parallel lines of work, also by Boersma (2001) and Apoussidou (2007) and by Jarosz (2006), who specifically look at “final-devoicing languages”, i.e. languages like Turkish, where the behavior of root-final stops is hidden in the bare form of the root. In the proposals mentioned above, the search for the optimal lexicon not only involves a rather large search space, it is also done in parallel with a search for a constraint ranking for the language. In my proposal, the learner is only trying to learn a constraint ranking, which is shown in Tesar & Smolensky (1998) to be quite efficient, and probably more efficient than searching for a ranking and a lexicon. An explicit proof that my approach requires a lighter computational load, however, is left for future work.

4.4.2 Hidden structure in affixes: Korean

In the discussion of Turkish above, attributing hidden structure of roots to the grammar was shown to be the only way to make the full range of lexical trends

available to the speaker. The principle of attributing predictable (or in this case, semi-predictable) information to the grammar is well-established in linguistics. Attributing predictable information to underlying representations prevents the learner from discovering generalizations.

In some cases, however, there is no way to attribute the full range of alternations that are observed in a language to rankings of plausible universal constraints. One such case is the final neutralization of obstruents in Korean, discussed briefly in §4.3.5 above, where not only laryngeal features (aspiration and voicing) but also manner (stop vs. fricative) and coronal place (dental vs. post-alveolar) are neutralized. All these contrasts appear before the accusative suffix, as in (183), taken from Albright (2008).

(183)	Bare noun	Accusative		
	nat̚	nasil	‘sickle’	375
	nat̚	nat̚ʰil	‘face’	160
	nat̚	nat̚ʰil	‘piece’	113
	nat̚	naɕil	‘daytime’	17
	nat̚	nadil	‘grain’	1

The rightmost column in (183) shows the number of words in a dictionary of Korean that end with each of the coronal obstruents in their spelling, indicating the historical pronunciation of these nouns. The fricative [s] is the most common coronal root-finally in the accusative, and the aspirated [t̚ʰ] and [t̚ʰ] are quite common as well. The voiced [ɕ] and [d] are much less common, and the glottalized/tense coronals of the language are absent completely.

Albright (2008) discusses recent innovations in Korean, where speakers extend the common [t̚] ~ [s] and [t̚] ~ [t̚ʰ] alternations of the accusative at the expense of [t̚] ~ [t̚ʰ], [t̚] ~ [d], and [t̚] ~ [ɕ]. In other words, speakers extend the most frequent mappings and remove the less frequent ones (“the rich get richer”). While Albright analyzes this

preference for frequent mappings with a general-purpose learner, i.e. a learner that doesn't incorporate substantive Universal Grammar principles, I suggest that an analysis in terms of plausible markedness constraints is within reach.

First, if the language learner assumes the base form /nat⁷/ as the underlying representation of all the roots in (183), and assumes /il/ as the underlying representation of the accusative suffix, they can learn several facts about Korean.

Korean does not allow voiceless unaspirated stops intervocalically – intervocalic stops must be either voiced or aspirated. Since the base has a voiceless unaspirated stop, this stop will not surface faithfully. Stops that surface aspirated in the accusative are faithful to the voicelessness of the base (184), while stops that surface voiced are faithful to the lack of aspiration in the base (185). A sample derivation is shown in (186).

(184) /nat⁷+ il/ → [nat^hil], [nat^hʰil]
 requires *VTV,¹⁵ IDENT(voice) ≫ IDENT(asp)

(185) /nat⁷+ il/ → [nadil], [naɕil]
 requires *VTV, IDENT(asp) ≫ IDENT(voice)

(186)

/nat ⁷ + il/	*VTV	IDENT(voice)	IDENT(asp)
a. natil	*!		
b. nadil		*!	
c. nat ^h il			*

¹⁵I am taking *VTV to be a constraint that penalizes intervocalic voiceless unaspirates. One can imagine a different analysis, where markedness penalizes any intervocalic voiceless stop, either aspirated or not. This will change the details, but not the main point, which is that the appearance of different stem-final obstruents in the accusative is due to constraint interaction, not to faithfulness to a non-surface-true UR.

Given the assumption of /nat̚/ as the UR of the root, the learner gets conflicting evidence for the ranking of IDENT(voice) relative to IDENT(asp). Constraint cloning will follow, resulting in a learner that keeps track of the number of root-final aspirated coronals and voiced coronals:

$$(187) \text{ IDENT(voice)}_{\{113+160 \text{ items}\}} \gg \text{IDENT(asp)} \gg \text{IDENT(voice)}_{\{1+17 \text{ items}\}}$$

The lexical trend that is created by the existing nouns of Korean predicts that speakers will prefer coronals that become aspirated in the accusative 94% of the time, and coronals that become voiced only 6% of the time.

The mapping of /t̚/ to [s] can also be attributed to the ranking of plausible markedness constraints. Assibilation, a process that turns stops into fricatives, is widely attested cross-linguistically before high vowels (Kim 2001). I use the constraint *TI, which penalizes stops before high vowels. Roots that surface with a stop of any kind in the accusative rank faithfulness to the continuancy of the base over *TI (188), while *TI outranks faithfulness in nouns that map the /t̚/ to [s] (189).

$$(188) \text{ /nat}^+ \text{ il/} \rightarrow [\text{nat}^{\text{h}}\text{il}], [\text{nat}^{\text{f}}\text{il}], [\text{nadil}], [\text{na}\text{ɕ}\text{il}]$$

requires IDENT(cont) \gg *TI

$$(189) \text{ /nat}^+ \text{ il/} \rightarrow [\text{nasil}]$$

requires *TI \gg IDENT(cont)

The conflicting ranking conditions cause the cloning of IDENT(cont), which allows the speaker to learn that the mapping of /t̚/ to [s] affects 56% of the *t*-final nouns in the language.

$$(190) \text{ IDENT(cont)}_{\{113+160+1+17 \text{ items}\}} \gg *TI \gg \text{IDENT(cont)}_{\{375 \text{ items}\}}$$

The learner's work is not quite done. In a fair number of nouns, a final /t̚/ maps to [ɕ] or [t̚^h]. Are there plausible constraints that will map /nat̚+ il/ to [naɕil] or [nat̚^hil]? Note

that the vowel of the accusative suffix is not a front vowel. Palatalization of [t] to [tʃ] is quite common before a front vowel or glide, but not common at all in their absence.

Essentially, the learner is in a situation where they want /nat^ɹ+iɪ/ to map to [natʃ^hiɪ], but the closest they can get is [nat^hiɪ]. There is no constraint that prefers the intended winner [natʃ^hiɪ] to the loser [nat^hiɪ], and as seen in (191), there is at least one faithfulness constraint that prefers the loser, IDENT(anterior). The intended winner is harmonically bounded.

(191)

/nat ^ɹ +iɪ/	*TI	IDENT(asp)	IDENT(ant)
a. natʃ ^h iɪ > nat ^h iɪ			L

When an intended winner is harmonically bounded, no reranking or cloning can help unbound it. What must change is the underlying representation. In this case, since the faithfulness constraint IDENT(anterior) is responsible for the harmonic bounding, the learner will take the feature that this constraint refers to, i.e. [−anterior], and add it as a floating feature to the accusative suffix.¹⁶ This is an instance of a more general strategy: The learner will find features that are missing in the harmonically bounded intended winner, as identified by faithfulness violations, and attribute them as floating features to the underlying representation of the relevant affix. As will be shown shortly, attributing hidden structure to affixes expands the range of lexical trends that the speaker can account for.

Once the missing [−anterior] feature is floating in the UR of the accusative affix, the mapping of /t/ to [tʃ^h] or [tʃ] is possible, and simply involves faithfulness to floating features with MAX(float). Stops that stay [+anterior] in the accusative are faithful to the [+anterior] root's [t^ɹ] rather than to the floating feature.

¹⁶The learner will also try adding [+anterior] to the accusative suffix, but they will quickly find out that this move does no good.

- (192) a. /nat^ɾ+ [−ant] il/ → [nat^hil], [naɕil]
 requires MAX(float) ≫ IDENT(ant)
- b. /nat^ɾ+ [−ant] il/ → [nat^hil], [nadil]
 requires IDENT(ant) ≫ MAX(float)

When *TI is highly ranked, and a coronal fricative surfaces before the accusative suffix, Korean won't allow the floating [−anterior] to surface faithfully, because the language as a whole is not faithful to [anterior] on fricatives. This is ensured by the high-ranking constraint *f, which in turn is dominated by *si, making [f] surface before [i] and [s] surface elsewhere. Since the high-ranking *f makes either ranking of IDENT(ant) and MAX(float) compatible with the winner, no items that surface with [s] in the accusative will be listed with clones of IDENT(ant).

- (193) /nat^ɾ+ [−ant] il/ → [nasil]
 requires *f ≫ IDENT(ant), MAX(float)

Since the learner has conflicting evidence about the ranking of IDENT(ant), they will clone it. Among the nouns that surface with a stop in the accusative, 61% are predicted to surface as [t^h] or [ɕ] rather than as [t^h] or [d].

- (194) *f ≫ IDENT(ant)_{113+1 items} ≫ MAX(float) ≫ IDENT(ant)_{160+17 items}

After the addition of the floating [−anterior] to the UR of the accusative suffix, the learner can account for all the mappings that they observe, and they can correctly learn the proportion of each of the five stem-final coronals in the language. The preferences that the grammar makes are given in (195), showing that the grammar successfully replicates the lexical counts given in (183).

(195)	IDENT(cont)	IDENT(voice)	IDENT(ant)
[s]	56%		= 56%
[tʰ]		94%	61% = 25%
[tʰ]	44%		39% = 16%
[ɕ]		6%	61% = 2%
[d]			39% = 1%

This analysis of Korean attributes hidden marked structure to the underlying representation of the accusative affix. Once this underlying representation is set up, forms that lack the marked structure of the affix are listed with high ranking faithfulness or with other markedness constraints. Since different nouns will require different rankings, lexical trends will be learned. To summarize the result so far: assuming bases as underlying forms, and attributing marked structure that appears in derived forms to the relevant affix, leads the learner to assume different rankings for different words, which in turn leads to learning of trends.

This analysis of the lexical trends that govern accusative forms depends on the shape of the accusative affix. For instance, the high vowel in the accusative form allowed the learner to attribute the mapping of /tʰ/ to [s] to the constraint *TI, which penalizes stops before high vowels. It is expected, then, that each suffix of the language will be treated separately.

The nominative paradigms of (196), from Albright (2008), show that when an affix begins in a front high vowel, stops and fricatives are regularly palatalized. Since the pattern is regular, the markedness constraint that demands palatalization ranks over faithfulness, and therefore, the derivation of the nominative forms will not involve faithfulness to the feature [anterior] at all, and no instances of [anterior] will be attributed to the underlying representation of the nominative suffix.

(196)

Bare noun	Accusative	Nominative	
nat ^ɾ	nasil	nafi	‘sickle’
nat ^ɾ	nat ^h il	nat ^h i	‘face’
nat ^ɾ	nat ^h il	nat ^h i	‘piece’
nat ^ɾ	naɕil	naɕi	‘daytime’
nat ^ɾ	nadil	naɕi	‘grain’

Korean also has two suffixes that surface as [-e]: the locative and the dative. Kang (2002) shows that the mapping of /t^ɾ/ to [s] is more frequent for the nominative and accusative than it is for the locative and dative. In other words, the affixes that don't have a high vowel in them are less conducive to assibilation. In my analysis, the suffixes that have a high vowel can cause assibilation simply by virtue of having a high segment in their surface form. The suffixes that have a non-high vowel, in contrast, can only cause assibilation if the learner adds a floating [-continuant] feature to their underlying representation. While this difference doesn't necessarily have to correspond to frequency data, since assibilation is equally possible with any suffix once a floating feature is added to suffixes that lack a high vowel, it is instructive that speakers are more reluctant to assibilate in the environment where assibilation requires an extra learning step of adding a floating feature to the UR.

In addition, Albright (2008) reports that while speakers most commonly innovate the mapping of /t^ɾ/ to [s] and [t^h] in the accusative, they prefer the mapping to [t^h] in the locative. Albright suggests that the preference for [t^h] is a result of the accidentally high number of nouns that historically had [t^h] and that are frequently used in the locative, such as the words for field and corner. Since in my analysis, lexical trends are computed for each affix individually, accidentally skewed distributions can be learned: If more items happen to require the ranking of *TI over IDENT(cont) in the accusative than in the locative, then the probability of mapping /t^ɾ/ to [s] will be higher in the accusative.

To summarize, the Korean case shows that assuming the bare form of the noun as the underlying representation of the root and assuming only segments as the underlying form of the affix might not be enough in every situation. When learning the lexical trends for the accusative forms, for instance, the Korean learner will discover that their language requires paradigms that change the feature [anterior]. If they proceed to add a floating [–anterior] to the underlying representation of the accusative suffix, they can learn the full range of behaviors seen in the accusative. The learner will have to make a similar move with the dative suffix, which requires assibilation in the absence of a high vowel; the learner can derive the full range of observed paradigms and also learn the lexical trends involved by adding a floating [–continuant] feature to the underlying form of the suffix.

Speakers can learn lexical trends so long as hidden structure is not buried in the underlying representation of roots. Adding hidden structure to the underlying representation of affixes does not present a danger so long as the affixes themselves are not proliferated. Compare the single representation of the accusative suffix in (192), which allows the learner to identify the full range of lexical trends, with the unfortunate proliferation of affixes in (182), which leaves the learner with an incomplete account of the trends in their language.

4.4.3 Interim summary: Generalizing across affixes

In the approach to linguistic analysis that I present here, learners find lexical trends in their language, and build those trends into their grammar. In order to find lexical trends, learners must assume the bare forms of roots as their underlying representations and assume that affixes are only composed of segments. If the paradigms involved contain hidden structure, it will not be trapped in the underlying representations of the roots and affixes, and will therefore become available to the grammar.

If the speaker discovers that they cannot account for all the derived forms that they are exposed to, because some intended winners are harmonically bounded, they will try to make any required features float in the underlying representation of the relevant affix.

These floating features can be identified by comparing the intended winner and the current output of the grammar that most resembles it (as determined by faithfulness constraints), and examining the features that are referenced in faithfulness constraints that distinguish the two forms. With the enriched underlying representation of the affix, the speaker can go on to discover any lexical trends that are lurking in the data.

A recurrent theme in this approach is the separate treatment of different affixes: The Korean learner, for instance, learns a separate grammar for each of the affixes of their language. The palatalization of [tʰ] to [tʃʰ] is a lexical trend with the dative suffix [-e], but the same trend is weaker with the homophonous locative suffix [-e]. Similarly, a lexical trend that involves the assibilation of root-final stops is seen in the accusative affix [-il], but the same trend is weaker with the nearly homophonous topic suffix [-in] (Kang 2002).

The same phenomenon is reported in Tagalog (Zuraw 2000; p. 33), where a stem can be subject to nasal substitution with some affixes but not others. Indeed, Zuraw shows that Tagalog has different lexical trends for different affixes of the language.

Similarly, in Turkish, the difference between the alternating stop of *tat* and the non-alternating stop of *at* is attributed to the grammar of the possessive suffix, and nothing prevents these two roots from behaving differently with other suffixes. This prediction is borne out. TELL (Inkelas et al. 2000) lists the possessive and the accusative forms of nouns. Both of these suffixes are homophonous with stop-final nouns, consisting simply of a high vowel. While most final stops are either voiced or voiceless in both forms, some nouns have a voiced stop in the possessive and not the accusative, and other nouns have a voiced stop in the accusative and not the possessive (see §4.3.4).

In the traditional generative analysis, the hidden structure of the root is attributed to its underlying representation, and then its behavior is predicted to be the same with any affix that allows the hidden structure to surface. In Turkish, assuming /taD/ as the underlying representation of *tat* predicts that the final stop will surface voiced with any vowel-initial suffix, contrary to the observed facts.

In the approach that attributes hidden structure to the grammar, roots are not required to behave uniformly with different affixes. There is a bias, however, for assigning consistent behavior to roots, as discussed in §4.3.4. In Turkish, for instance, once a root is observed to alternate in the possessive, the grammar will record this fact by connecting three pieces: the root, the possessive affix, and a conflict between constraints. When the speaker wishes to generate the same root with a different suffix, say the accusative, and the same constraint conflict is involved, the root's possessive entry will match the root in the accusative, and bias the speaker to assign the same behavior to the root with both affixes.

4.4.4 Hidden structure in roots: English

In the various lexical trends that were discussed in this chapter, it was always the case that a relatively simple concatenation of a root and affix, together with some lexically-specific rankings, allowed the speaker to map one form onto a morphologically related form. Quite clearly, this is not always the case. Extreme examples of phonologically intractable mappings are usually described as suppletion, like the English *go* ~ *went*. In cases like these, the learner has no choice but to store the form *went* as an unanalyzed whole, and nothing about this form becomes available to the grammar of the past tense.

Other cases might not be as clear as *go* ~ *went*. The English past tense includes seven verbs that end in [ɔt]: *teach* ~ *taught*, *catch* ~ *caught*, *think* ~ *thought*, *bring* ~ *brought*, *seek* ~ *sought*, *fight* ~ *fought*, and *buy* ~ *bought*. Can these verbs be mapped onto their past tense using phonological machinery?

While mapping a verb like [fart] to [fɔt] is relatively faithful, involving only the replacement of the vowel, verbs like [brɪŋ] and [sɪk] keep nothing but their onset in the past. One can imagine that for those verbs, an allomorph of the past tense suffix that consists of a pair of floating segments, /^{ɔt}/, can dock correctly and replace the root segments. In such an analysis, MAX(float) would ensure that both segments dock at the cost of faithfulness to the root.

(197)

/ sik + ^{ot} /	MAX(float)	MAX(root)
a.  sɔt		**
b. sɔk	*	*
c. sik	**	

With the vast majority of English verbs giving evidence for the ranking MAX(root) \gg MAX(float), and seven verbs giving evidence for the opposite ranking, the learner can clone one of these two constraints, and thus give a small probability to ot -taking. However, these two constraints don't refer to any phonological aspect of the root (other than the existence of segments in it), and therefore cloning them will give the learner no information about the possible shapes of ot -takers.

This seems to be the right outcome: The ot -takers in English are not phonologically patterned in any way beyond being monosyllabic, so any kind of monosyllable would be a candidate for ot -taking. Since the ot -takers represent such a small minority of the monosyllabic verbs of English, speakers are predicted to be reluctant to project ot -taking onto novel roots.

Another consideration with the derivation of ot -takers is the availability of the regular past suffix, /-d/. When deriving the past tense of [sik], the candidate [sikt] is quite appealing: It is completely faithful to the root and to the past suffix (modulo the completely regular voicing assimilation), and even the worst aspect of it, the final [kt] cluster, is quite widely attested in English. The appeal of the regular [sikt] might cue the learner to the possibility that something non-phonological is going on, and prompt them to simply store [sɔt] as an unanalyzed whole.

Both ways of dealing with [sɔt] – cloning MAX(float) or storing it as a whole – are equally bad for finding out what kind of roots are ɔt-takers. Indeed, English speakers are reluctant to generalize ɔt-taking, or to do so in any phonologically principled way. In other words, the speaker doesn't necessarily always have to decide whether a certain pattern is suppletive or not. They may treat what's essentially a suppletive pattern with their grammatical machinery, but if the grammar tells them nothing about the shape of the relevant lexical items (e.g. due to the lack of involvement of markedness constraints, as in (197)), then no damage is done, since the pattern cannot be extended usefully.

4.4.5 The naturalness of lexical trends: Dutch

Dutch exhibits voicing alternations between bare roots (which in the case of verbs can be heard in the imperative) and affixed forms, as in (198). In the lexicon, the proportion of alternating consonants depends on the identity of the consonant, and speakers project these proportions unto novel items, as shown by Ernestus & Baayen (2003). The phonology of Dutch raises two questions that relate to the naturalness of lexical trends: (a) the issue of natural relationships between lexical trends, and (b) the functional grounding, or naturalness of each lexical trend.

(198)	Imperative	Infinitive	Past tense	
	tɔp	tɔb-ən	tɔb-də	'worry'
	stɔp	stɔp-ən	stɔp-tə	'stop'

Ernestus & Baayen (2003) report that in the lexicon, the proportion of alternating labial stops is smaller than the proportion of alternating coronal stops, and speakers replicate this preference in their treatment of novel words. In the model I propose, Dutch speakers will clone IDENT(voice) relative to constraints on voiced codas,¹⁷ and collect the stop-final

¹⁷In Ernestus & Baayen (2003), speakers' knowledge was tested with novel past tense forms, where the stem-final stop is in coda position. In the infinitive, and before other vowel-initial suffixes, the stem-final stop

words of their language, like Turkish speakers. A portion of the adult grammar of Dutch is given in (199).

$$(199) \quad \text{IDENT(voice)}_{\langle *b]_{\sigma}, 210 \text{ items} \rangle} \gg *b]_{\sigma}, *d]_{\sigma} \gg \text{IDENT(voice)}_{\langle *b]_{\sigma}, 20 \text{ items} \rangle} \\ \langle *d]_{\sigma}, 542 \text{ items} \rangle \quad \langle *d]_{\sigma}, 177 \text{ items} \rangle$$

The 210 words of Dutch that have a non-alternating [p] are collected by the clone of IDENT(voice) that ranks above $*b]_{\sigma}$, and the 20 words that have an alternating [p] are collected by the lower ranking clone. This makes the proportion of alternating [p]'s, which is 9%, available to the speaker. The *t*-final words of Dutch are similarly collected by the clones of IDENT(voice), allowing the speaker to discover that 25% of final [t]'s alternate.

The surprising aspect here is that universally, speakers are expected to prefer voicing in labials over voicing in coronals or dorsals. For example, among the languages that have a voicing contrast in stops in at least one place of articulation, [p] is more likely to be absent than [t] or [k], and [b] is more likely to be present than [d] or [g] (Maddieson 1984; pp. 35–36). The speakers of Dutch have a grammar that makes the opposite preference, giving a higher probability to [p] than to [t].

The ability of Dutch speakers to learn an unnatural relationship between lexical trends is not surprising given my approach. Different lexical trends are controlled by different pairs of constraints, and the strength of one trend is not expected to interact with the strength of another. Dutch speakers use the clones of IDENT(voice) to keep track of *p*-final nouns by listing them with $*b]_{\sigma}$, and keep track of *t*-final nouns by listing them with $*d]_{\sigma}$. The number of words listed under clones of one constraint does not affect the number of words listed under clones of another constraint. The prediction that the relationship between lexical trends need not be natural is borne out by the Dutch data.

A second intriguing aspect of Dutch voicing alternations is the effect of the vowel that precedes the stem-final consonant. In the lexicon, alternation are on average most common

could be argued to be in coda position as well, if it is taken to be ambi-syllabic, as proposed by van der Hulst (1985) et seq.

following a (non-high) long vowel or a diphthong and least common after the short non-high vowels. The high vowels, which in Dutch are phonetically short and don't have long counterparts, give rise to a rate of voicing alternation that is intermediate between the long vowels and the non-high short vowels. Vowel length, however, is a rather poor predictor of consonant voicing in the lexicon: In the GLM statistical analysis that Ernestus & Baayen (2003) report, vowel length has a very modest effect on the voicing of the following obstruent ($p = .053$). A comparison of long vowels and high vowels only shows a more robust effect ($p = .017$).

In the experimental results, the vowel effect was solid ($p = .004$). Long vowels were significantly more conducive to voicing of stem-final obstruents than short vowels of any height. There was no significant difference between the short high and short non-high vowels.

It is instructive that Dutch speakers imposed a natural trend on the data: The different vowel qualities of Dutch were abstracted away from, since universally, vowel quality (height, backness, tenseness, roundness) has no power to affect the voicing of a following consonant. Only vowel length is universally correlated with voicing, with long vowels (either pure or diphthongal) being conducive to following voiced codas and short vowels to following voiceless codas (Lisker & Abramson 1964; Ohala 1983; Volatis & Miller 1992).

Given a family of universal constraints such as $\{*V:p\}_\sigma$, $*V:t\}_\sigma$, $*V:k\}_\sigma$, which penalizes voiceless coda obstruents after a long vowel, and the more general family $\{*b\}_\sigma$, $*d\}_\sigma$, $*g\}_\sigma$, speakers will be able to keep track of alternation rates of obstruents that follow long vowels separately from the alternation rates of obstruents after short vowels. In the experiment that Ernestus & Baayen (2003) report, speakers were given bare verbal roots (e.g. $de:p$), and were asked to add the past tense suffix, which is $[-d\text{ə}]$ or $[-t\text{ə}]$. When choosing between the two possible outcomes, $[de:p-t\text{ə}]$ and $[de:b-d\text{ə}]$, the root-final consonant is guaranteed to be in the coda, and thus its voicing is expected to interact with the length of the preceding vowel.

In conclusion, the Dutch facts highlight two aspects of the theory: Firstly, they show that while natural constraints are used to keep track of lexical trends, there is no necessary connection between separate trends. If a language gives a higher probability to voicing coronals than to voicing labials in its lexicon, speakers will be able to learn these trends and project them onto novel nouns. Secondly, speakers are only able to learn lexical trends that are stated in terms of natural constraints. When the lexicon gives a higher probability to a voicing alternation after high vowels (or other vowel qualities), speakers will fail to replicate this effect in their treatment of novel words. Speakers can only replicate relationships like the ones between voicing and vowel length, since vowel length is naturally correlated with consonant voicing cross-linguistically, unlike vowel height, backness, tenseness, or rounding.

4.5 Cloning and the typology of lexical trends

Using an OT-based model to account for lexical trends makes predictions about the range of possible lexical trends and their relationship to the regular phonology of the language. In this section, the predicted typology is explored, and its correspondence with the observed range of trends is assessed.

4.5.1 Contrast neutralization and creation

Lexical trends, as I define them here, are observed in derived, or affixed forms. When a morphological category is expressed overtly by affixation, affixation may cause some phonological process to take place, or block an otherwise regular process (see Pater 2006 for a related discussion). If the phonological process does not regularly apply to all eligible affixed forms, or if the process is not blocked in all eligible affixed forms, a lexical trend arises. The two kinds of interactions are schematized in (200) and (201).

(200) *Affixation neutralizes a contrast that exists in roots*

In roots: $F \gg M$

In affixed forms: some roots require $F \gg M$, some $M \gg F$ ¹⁸

(201) *Affixation creates a contrast that doesn't exist in roots*

In roots: $M \gg F$

In affixed forms: some roots require $F \gg M$, some $M \gg F$

The Turkish example previously discussed is of the neutralizing type, as in (200): Generally in Turkish, voiced and voiceless stops contrast intervocalically, as in the minimal pair *ata* 'father' vs. *ada* 'island', showing that IDENT(voice) ranks above *VtV, i.e. $F \gg M$. In nouns like *tat* ~ *tad-i*, the voiceless [t] of the bare noun becomes [d] when intervocalic, showing that *tat* requires *VtV to rank higher than IDENT(voice), i.e. $M \gg F$.

Another lexical trend of the neutralizing type is nasal substitution in Tagalog, studied by Zuraw (2000). In Tagalog, nasals can be followed by stops inside roots (e.g. *gindáj* 'unsteadiness on feet'), but when certain nasal-final suffixes are attached to certain stop-initial stops, the nasal-stop cluster does not surface faithfully, and a single nasal stop is pronounced instead (e.g. /maŋ**-bigáj**/ → *ma-migáj* 'to distribute'). Zuraw attributes nasal substitution to the markedness constraint NASSUB (although she does not commit to its functional grounding), i.e. a markedness constraint that is freely violable inside roots due to high-ranking faithfulness, but the same constraint neutralizes the nasal/oral distinction in some affixed forms.

Lexical trends that create phonological contrasts, as in (201), are attested in a number of Celtic languages.¹⁹ In these languages, consonant mutation often creates consonants or consonant clusters that are only attested in mutated forms, never in underived forms

¹⁸Following Wolf (2008b), I am assuming that the effect of a markedness constraint M can be limited to derived environments using principles of OT-CC (McCarthy 2007a), and without having to hard-wire the limitation to derived environments into the definition of the constraint.

¹⁹The following discussion of Irish benefitted from the wisdom of Matt Wolf and Emily Elfner.

(Ní Chiosáin 1991; Wolf 2008a). In Irish, for example, a word-initial [m] usually mutates into a nasal glide, [w̃], but the mutation is blocked in some words. Since [w̃] is generally banned in Irish, we can conclude that *w̃ outranks faithfulness constraints such as IDENT(cont). In derived environments, [w̃] is usually allowed, but some words exceptionally block mutation, such as *meid* ‘amount’, which does not turn into **w̃eid*. Assuming that mutation is due to faithfulness to a floating feature, as proposed in Wolf (2007), MAX(float) must outrank *w̃ for most words of Irish,²⁰ while the exceptions require *w̃ to rank above MAX(float), leading to an inconsistent grammar that must be resolved by cloning.

4.5.2 Competing repairs

In addition to the trends that follow the schemata in (200) and (201), a third kind of lexical trend can be caused by exceptional ranking of two faithfulness constraints, as schematized in (202).

(202) *Affixation respects markedness by deploying two different repairs*

In roots: M ≫ F1, F2

In affixed forms: some roots require M ≫ F1 ≫ F2, some M ≫ F2 ≫ F1

A case that can be described in terms of (202) is the zero-marked past tense of English verbs, as discussed in §1.1.2. In English, final clusters of alveolar stops (t, d) are not allowed, so the constraint that bans these clusters, *DD, is undominated in the language. There is no evidence that can bear on how these clusters are repaired inside roots: A hypothetical root such as *[pɛdd] could surface as [pɛd], [pɛnd], [pɛdɪd], or several other options. In the past tense, however, comparing *t*-final and *d*-final roots and their past tense forms reveals that most verbs repair the alveolar stop clusters by epenthesis (e.g. /gard +

²⁰Note that MAX(float) is not active in roots, since a hypothetical root with a floating [-cont] in it could give rise to [w̃], contrary to fact. So generally in Irish, *w̃ ≫ MAX(float), and the effect of MAX(float) must be limited to derived environments.

d/ → gardɪd), while some verbs repair the cluster by deletion (e.g. /sprɛd + d/ → sprɛd). Verbs like *guide* require *DD ≫ MAX ≫ DEP, while verbs like *spread* require the opposite ranking of the faithfulness constraints, i.e. *DD ≫ DEP ≫ MAX.

4.5.3 Exceptional emergence of the unmarked

The fourth and last kind of lexical trend involves a faithfulness constraint that dominates two conflicting markedness constraints. In roots, the effect of the markedness constraints is not felt, due to the overriding faithfulness. In affixed forms, however, allomorph selection allows the markedness effect to emerge without a faithfulness cost. This kind of lexical trend is schematized in (203).

(203) *Allomorph selection responds to competing markedness effects*

In roots: F ≫ M1, M2

In affixed forms: some roots require F ≫ M1 ≫ M2, some F ≫ M2 ≫ M1

Trends that are structured as in (203), where there is no faithfulness cost to the irregular behavior, are expected in irregular allomorph selection. Since allomorphs are selected with no faithfulness cost (Mascaró 1996 et seq.), the effect of different markedness constraints can emerge.

One case that is described in the terms of (203) is plural allomorph selection in Hebrew nouns (see chapter 3 for a full discussion). Masculine nouns usually take the masculine plural affix *-im*, but some masculine nouns exceptionally select the feminine plural affix *-ot*. Most of those exceptional nouns have [o] in them, which I suggest is done to satisfy LICENSE(o), a markedness constraint that requires unstressed [o] to be licensed by a stressed [o]. Since Hebrew roots allow unstressed [o] in them freely, faithfulness outranks LICENSE(o) generally in the language. In affixed forms, regular nouns take *-im* due to MATCH(gender), a morphological markedness constraint that requires the masculine suffix on masculine stems, so for those nouns, MATCH(gender) ≫ LICENSE(o). Masculine nouns with [o] in their root that select the feminine *-ot* require LICENSE(o) ≫ MATCH(gender).

An example that does not rely on morphologically-oriented constraints, only on phonological ones, comes from the Turkish aorist²¹ (Lees 1961; Napikoğlu & Ketrez 2006). This verbal suffix shows up in three forms, shown in (204). The distribution is regular in all but CVC roots that end in {r, l, n}. The aorist suffix is simply [r] after vowel-final stems of any length; it is [-Ir]²² after poly-syllables that end in a consonant; and [-Er] after mono-syllables that end in an obstruent or glide. For mono-syllabic nouns that end in {r, l, n}, some roots take [-Ir], and others take [-Er].

(204)	Shape of stem	Affix	Examples ²³
	V-final	-r	de-r, ye-r, uyu-r, bafla-r
	C-final poly-syllables	-Ir	gerek-ir, tjalif-ir
	Obstruent-final mono-syllables	-Er	bit-er,öp-er
	{r, l, n}-final mono-syllables	-Ir	kal-ir, gör-ür
		-Er	dal-ar, ör-er

The analysis in terms of markedness is fairly straightforward once some simple assumptions about Turkish stress are made. In line with Hayes (1995a), I assume that stress in Turkish, which by default falls on the word-final syllable, is trochaic, meaning that the stressed final syllable is in a foot by itself. Little is reported about secondary stress in Turkish, but assuming it shows up on every other syllable from the ultima, a mono-syllabic stem like [bit] shows up in the aorist with an unparsed syllable: bi(t-ér). Longer stems will have another foot before the stressed one: (gère)(k-ír). In other words, both [-Er] and [-Ir] show up inside the strong foot of the word (the main stressed foot), but [-Er] additionally demands to be in the initial, or leftmost foot of the word. To ensure that

²¹I am indebted to Matt Wolf and John Kingston for their help in the following analysis.

²²The capital I represents a high vowel that gets its backness and roundness from the preceding vowel. The capital E represents a non-high unrounded vowel that gets its backness from the preceding vowel.

²³Glosses: say, eat, sleep, begin / need, work / finish, kiss / stay, see / dive, knit.

[-Er] only appears when it's inside the leftmost foot of the word, it is subcategorized to the categorical alignment constraint ALIGNL-BY-FOOT,²⁴ which requires that no foot be preceded by another foot in the word (McCarthy 2003). In mono-syllables, ALIGNL-BY-FOOT is equally satisfied by [-Er] and [-Ir], and the decision is passed down to * $\acute{\sigma}$ /HIGH, a constraint that penalizes stressed high vowels (205-206).

(205)

/gerek + {-Er, -Ir}/	ALIGNL-BY-F _{-Er}	* $\acute{\sigma}$ /HIGH
a. \Rightarrow (gère)(k-ír)		*
b. (gère)(k-ér)	*!	

(206)

/bit + {-Er, -Ir}/	ALIGNL-BY-F _{-Er}	* $\acute{\sigma}$ /HIGH
a. bi(t-ír)		*!
b. \Rightarrow bi(t-ér)		

In mono-syllables that end in {r, l, n}, the constraint that penalizes stressed high vowels conflicts with a constraint that penalizes non-high vowels between sonorants, *RER. The sonorants {r, l, n} have a high first formant, like low vowels, so *RER enforces dissimilation in the height of the first formant, penalizing the lack of contour created by a sequence of sounds with a high first formant.²⁵

²⁴This subcategorization of an affix to a markedness constraint is distinct from constraint cloning, and belongs to the realm of prosodic morphology. For a more famous example, compare the subcategorization of the Tagalog infix -um- to ALIGN-L-BY- σ in McCarthy (2003).

²⁵The interaction between sonorants and vowel height is active elsewhere in Turkish: Coda {r, l, n} lower a preceding [ɛ] to [æ] across the board — presumably an assimilation effect. As expected when the same

(207)

/kal + {-Er, -Ir}/	ALIGNL- BY-F _{-Er}	*RER	*ó/HIGH
a. \rightarrow ka(l-ír)			*
b. ka(l-ár)		*!	

CVR roots that take [-Ir], like *kal*, require *RER \gg *ó/HIGH, while CVR roots that take [-Er], like *dal*, require the opposite ranking. This in turn will lead to the cloning of *RER. The ranking arguments are summarized in (208).

(208)

	ALIGNL- BY-F _{-Er}	*ó/HIGH	*RER
a. gerek-ir \succ gerek-er	W	L	
b. bit-er \succ bit-ir		W	
c. dal-ar \succ dal-ir		W	L
d. kal-ir \succ kal-ar		L	W

Once ALIGNL-BY-F_{-Er} is installed in (208), and the first winner-loser pair is removed from the Support, the conflict between *ó/HIGH and *RER is apparent. Note that no faithfulness cost is associated with the selection of the allomorphs of the aorist, and all the work is done by ranking general and lexically-specific markedness constraints.

phonetic factor causes both dissimilation and assimilation in the same language, the dissimilation affect is more restricted: Dissimilation is operative only in allomorph selection in verbs.

It might be worth noting that the distribution of the Turkish aorist is irregular only in those cases where one sonorant from the aorist suffix and one sonorant from the root flank a vowel. In other words, the irregular pattern is not phonologically arbitrary. My UG-based analysis expresses this non-accidental nature of the distribution by the use of the markedness constraint *RER.

The Turkish case is parallel to the analysis of the English verbs offered above, which crucially relies on the fact that the past tense consists of an alveolar stop and that the verbs that exceptionally don't take it end in an alveolar stop. The distribution of the lexical exceptions is not phonologically arbitrary, but rather follows from a constraint against clusters of alveolar stops.

4.6 Conclusions

This chapter presents a theory of speakers' knowledge of irregular morphology. I claim that speakers use an Optimality Theoretic grammar to identify irregular patterns in their lexicon and extract partial phonological regularities from it. The theory relies on the Recursive Constraint Demotion algorithm (Tesar & Smolensky 1998, 2000; Tesar 1998; Prince 2002), augmented with a mechanism of constraint cloning (Pater 2006, 2008b).

Once it is discovered that different lexical items require different constraint rankings, a constraint is cloned, and each clone lists lexical items with it. As the speaker learns the words of their language, lexical statistics are gradually built into the grammar. The resulting grammar is able to give consistent behavior to listed items, and also project the trend that is created by the listed items stochastically onto novel items.

I offer a formal theory of cloning, which involves the "least populated column" metric for identifying constraints to clone, augmented with "masking", which is a measure for preventing double-dipping, ensuring that lexical trends are represented correctly in the grammar. I formalize the learning algorithm as a variant of RCD with error-driven learning, including a method for finding underlying representations. In order to make

lexical statistics available to the grammar, the learner must make sure that these statistics are not buried in the lexicon via the assignment of abstract underlying representations to roots. I present an algorithm for minimizing the information in the lexicon by assuming the surface form of the base as the underlying representation, and by minimizing the number of allomorphs that affixes have. Minimizing the information in underlying representations has as a necessary consequence the attribution of more information to the grammar.

The use of the constraints of Optimality Theory to express lexical trends predicts a typology of trends. I explore this typology and show that all of its predictions correspond to observed lexical trends.

CHAPTER 5

CONCLUSIONS AND FUTURE DIRECTIONS

5.1 Summary of the dissertation

This dissertation started with two empirical observations about two biases that humans have in their treatment of their lexicon: They ignore unnatural interactions between phonological elements (chapter 2), and they state generalizations based on the surface properties of lexical items (chapter 3). These observations were taken as evidence for a model of grammar that has built-in expectations about the naturalness of phonological operations, and that states phonological generalizations in terms of constraints on surface forms. As it happens, Optimality Theory (Prince & Smolensky 1993/2004) is such a model, and this work developed an OT-based model for learning lexically-specific phonology and for projecting the learned statistics onto novel items (chapter 4).

In Turkish, voicing alternations affect stem-final stops in some nouns (e.g. *tat* ~ *tad-i* ‘taste’), but not in others (e.g. *at* ~ *ad-i* ‘horse’). While it is not predictable whether any given lexical item will voice or not, voicing alternations are tightly correlated with the phonological shape of nouns when averaged over the lexicon. Specifically, voicing alternations are correlated with the size of nouns, with the identity of the final stops, and with the height and backness of the noun’s last vowel. When learning their language, Turkish speakers don’t content themselves with learning the behavior of individual items; they also learn about correlations between the shapes of nouns and the likelihood that they will display voicing alternations, and when given a novel noun, they match its likelihood of alternation to the likelihood of alternation of similar nouns. The question was what nouns count as being similar to the given novel noun. It turned out that the size of the noun and

the identity of the final stop were used in assessing similarity, but the quality of the noun's last vowel was ignored.

The notion of similarity that humans use, then, is biased to notice some aspects of phonological structure and ignore others. I claimed that it is not a coincidence that universally, vowel quality never affects the voicing of a neighboring consonant, but rather that this is due to Universal Grammar. Since Universal Grammar doesn't have a mechanism that correlates vowel quality with obstruent voicing, this correlation is absent both from regular phonological processes cross-linguistically and from irregular phonological patterns of exceptionality in individual languages. In Optimality Theory, the observed array of phonological processes follows from the structure of CON, the set of universal constraints. By deriving irregular patterns of exceptions from this same set of constraints, the generalization about the natural patterning of exceptions is predicted.

In Hebrew, the plural marker on nouns has two allomorphs, *-im* and *-ot*. While in some contexts the choice of allomorph is morphological, with *-im* being masculine and *-ot* feminine, the choice is also phonological. Masculine nouns with [o] in their stem are more likely to select *-ot* than masculine nouns that don't have [o]. This irregular pattern was captured in OT in terms of lexically-specific rankings of markedness constraints. Since markedness constraints assess output forms only, the OT account predicted that the choice of allomorph depends on the presence of [o] in the plural stem, without any regard to the vowels of the singular stem. Because nouns that have [o] in their plural stem also have [o] in their singular stem, Hebrew doesn't offer speakers evidence about which stem matters, and speakers could learn Hebrew equally well by generalizing over vowels of plural stems or over vowels of singular stems.

To see which stem speakers look to in their generalizations, Hebrew speakers were taught one of two languages in an artificial language experiment: One language paired *-ot* with plural stem [o], and another paired *-ot* with singular stem [o]. In both languages, vowel changes that are absent from real Hebrew restricted [o] to appear only in the singular

stem or only in the plural stem for any given paradigm. Speakers were more successful learning the language that paired *-ot* with stems that have [o] in the plural, as predicted by the analysis that uses markedness constraints.

The formal properties of the proposed OT-based model were explored and motivated in chapter 4. In this model, the inconsistent behavior of lexical items under affixation gives rise to conflicting rankings of universal constraints. These rankings in turn are used to classify the lexical items involved by cloning constraints and listing lexical items with clones. The resulting grammar captures the behavior of known items, so they can be derived to correctly produce adult forms, and it also uses the relative numbers of the recorded items to apply probabilistically to novel items, as humans do.

The analysis of Turkish in chapter 2 had to proceed in what Hayes (1999) calls “inside-out” fashion, i.e. assuming that the base is identical to its surface form, without using properties of derived forms to enrich the underlying form of the base. This move was generalized to a claim that universally, the underlying form of the root is identical to the surface form of the base, and that abstract underlying forms are limited to affixes. The implications for Turkish and a variety of other languages were explored. Finally, the range of exceptionality that was predicted from the use of markedness and faithfulness constraints was explored and shown to be fully instantiated.

5.2 Future directions

This final section explores some of the broader ramifications of the proposals made in this dissertation, specifically with regard to the predicted naturalness of lexical organization and the concomitant revised view of morpho-phonological analysis.

5.2.1 Unnatural phonology

It was seen that Turkish speakers do not project the effect that vowel quality has on stop voicing in their lexicon onto novel items, and I have claimed that this is due to the

unnaturalness of the correlation. I have also shown that the results in Ernestus & Baayen (2003) are instructively similar: Dutch speakers project the effect of vowel length on stop voicing, but not the effect of vowel quality. Looking at regular phonological phenomena in the languages of the world, it is seen that vowel length correlates with stop voicing, but vowel quality does not. Naturalness, it is claimed, determines the range of possible phonological interactions, and this in turn predicts the range of regular and irregular phonology.

The claim that all phonology is natural, however, is controversial. Pierrehumbert (2006) shows that English velar softening (e.g. *electri[k]* ~ *electri[s]ity*) is extended by speakers to novel items, yet this process is unnatural, given that it has never been observed as a regular process in any language. The view that phonology is not necessarily natural is taken by Evolutionary Phonology (Blevins 2004 et seq.), where naturalness only affects diachronic change, but not synchronic grammar. A more nuanced view is offered in a study of Hungarian vowel harmony by Hayes et al. (to appear), who show that Hungarian speakers project both natural and unnatural trends from their lexicon, but that the unnatural trends are projected more weakly than the natural ones. In an artificial language experiment, Moreton (2008) finds that speakers are biased to learn natural generalizations more successfully, but unnatural generalizations are learned as well. Similarly, Kawahara (2008) argues for a model of synchronic grammar that combines natural and unnatural constraints.

Ultimately, the question is an empirical one: In what situations does naturalness bias the behavior of speakers, and to what degree? The answer offered in this work, namely that naturalness can prevent any learning of some aspect of the lexicon, may turn out, with the accumulation of more evidence, to be too strong to be fully general.

5.2.2 The search for underlying representations

A necessary component of making lexical trends available to the grammar, I have shown, is assuming that roots always have surface-true underlying representations. This approach was taken in Hayes (1999), who went as far as to suggest doing away with underlying representations altogether, based on evidence that speakers of Yidiñ do not use derived forms to build consistent underlying representations for roots. Similar claims about the role of the surface forms of bases were made in Albright (2008), mostly based on evidence from historical change that suggests the restructuring of the grammar after the loss of phonological material from roots.

This approach contrasts sharply with the tradition in generative linguistics, which looks to bases and derived forms to glean information about underlying representations of roots, with the stated goal of making the grammar as regular and as general as possible (see e.g. chapter 6 of Kenstowicz & Kisseberth 1979, and more recently in Odden 2005). This model of the grammar has been explored formally under the rubrics of surgery or contrast analysis, using paradigmatic information to piece together an abstract underlying representation (Tesar et al. 2003; Tesar 2004; Alderete et al. 2005; Tesar 2006; Merchant 2008). The goal of reaching a consistent grammar also informs the approach taken in Boersma (2001), Apoussidou (2007), and Jarosz (2006).

The evidence, it seems to me, is squarely on the side of those who don't allow abstract underlying representations for roots. Speakers use grammatical tools to predict derived forms from the surface forms of bases, and the (partially) predictable information that speakers have should be made available to the grammar, and not be relegated to the lexicon via abstract underlying representations. This is not to say, however, that the issue is closed. Specifically, two thorny issues remain: The role of underlying representations in the proper treatment of opacity, and their role in the treatment of sentence phonology.

Opaque generalizations are ones that depend on some property of the UR, not on the surface form. For example, Beduin Arabic allows [a] in open syllables only in syllables

that were opened by epenthesis, and not in syllables that are open via a faithful parsing of the input (McCarthy 2007a). The learning mechanism offered in this dissertation would not be able to learn such a generalization. There is hope, however, that a mechanism along the lines of the “free ride” algorithm (McCarthy 2005) could be incorporated to give the learner access to such hidden generalizations. Moreover, little is known about speakers’ behavior when faced with the need to learn both irregular phonology and opaque phonology in the same language, and hence any attempt to reconcile these two aspects of phonology should be accompanied by an attempt to collect the relevant empirical evidence.

Another challenge for a theory that rejects the possibility of non-surface-true underlying representations for roots comes from the range of phenomena known as sentence phonology. In Chizigula (Kenstowicz & Kisseberth 1990), for instance, some words that have a Low tone throughout in isolation will appear with a High tone after the copula *ni*, and some other words will appear with a falling tone after the same copula. Kenstowicz & Kisseberth (1990) use these alternations to motivate abstract underlying representations that include tones that never get realized in their underlying position. The challenge to the learner and to the analyst is the need to attribute the change in the surface forms of words to some phonological element of the phrase, and since the size of phrases is unbounded, the range of hypotheses to entertain is also, at least on first sight, unbounded.

It is instructive, perhaps, that word-level phenomena often recapitulate the phrase-level phenomena: In Chizigula, the appearance of a contrast between high and falling tone is also seen word-internally under prefixation. This means that the speaker can first learn a certain amount of word-level phonology from the prefixes and suffixes of their language, and if they can generalize these lessons to inform their hypotheses about the phrasal level, then perhaps most of the work will be done. Additionally, the range of non-local phonological interactions between words at the phrase level is essentially limited to tone; all other phonological features can only cross word boundaries in local interactions via assimilation. These facts suggest that the space of hypotheses that the speaker has to search is not, in

fact, unbounded at the phrasal level, and that the space can be limited by language-specific learning of the word-level phonology and by universal expectations about the range of phenomena that are accessible to the phrasal phonology.

5.2.3 Issues in lexical organization

The phonological analyses offered in this work incorporate a great deal of lexical information into the Optimality Theoretic grammar, in the form of constraint clones that are associated with lists of stems. One wonders, then, what is the full range of interactions that should be admitted between lexical items and the grammar, and how these are learned.

Widely used and essentially uncontroversial are constraints that refer to lexical classes such as nouns (see Smith 2001 for a review). The need for affix-specific grammars has also been widely recognized in the literature, starting with the analysis of Tagalog infixation in terms of affix-specific alignment constraints (Prince & Smolensky 1993/2004; McCarthy 2003), and expanding to other domains of prosodic morphology, as in, e.g. Flack (2007b), Gouskova (2007), and §4.5.3 above. In these cases, the grammar is enriched with reference to morphological categories such as “noun” or “benefactive” that are needed elsewhere in the grammar, and are thus not assumed to add much of a burden to the learner. However, a formal mechanism for learning these constraint indexations is yet to be proposed.

Making a connection between the grammar and an arbitrary list of lexical items, however, has also been proposed under the name of lexical stratification (Itô & Mester 1995, 1999, 2003; Kawahara et al. 2003; Féry 2003; Becker 2003; Gelbart 2005; Rice 2006; Jurgec 2009, among others). The association of grammars with arbitrarily defined lists of items is conceptually akin to the treatment of lexical exceptions offered in this dissertation, and perhaps these two areas of phonology should be handled with the same theoretical machinery. Much of the work on lexical stratification is interested in the clustering of phonological properties, such as the characterization of Yamato Japanese by several different phonotactic restrictions, whereas lexical exceptions as defined in this

dissertation involve just one phonological process. If the lists of items that are associated with different clones are biased to be similar to each other, then maybe the clustering of phonological properties could be derived: Being exceptional in one way will bias towards being exceptional in some other way, thus creating phonologically-defined clusters in the lexicon.

5.2.4 Lexically-specific grammars and phonotactics

This dissertation focuses on paradigmatic relations between words, using them to learn a grammar that derives one morphological category from another; this learning happens separately from what the speakers learn about the static phonotactic generalizations about their language. This is possibly a shortcoming of the theory, since morpho-phonological alternations have been claimed to recapitulate the phonotactics of the language (“the duplication problem”, Clayton 1976; Kenstowicz & Kisseberth 1977), and Optimality Theory is expected to be able to unify these two aspects of the phonology (McCarthy 2002; pp. 71–75).

An interesting idea in this direction comes from Coetzee (2008), who suggests that phonotactics are learned by promoting word-specific clones of faithfulness constraints one by one, instead of promoting lexically-neutral constraints, as is generally practiced. It is possible that this approach can be shown to produce the attested kinds of knowledge that speakers have of the the phonotactics of their language, but this work is yet to be done.

BIBLIOGRAPHY

- Aharoni, Moti (2007). Different plural forms (curot rabim shonot). http://tora.us.fm/tnk1/messages/ljon_cwrt_yxyd_wcwrt_rbym_0.html.
- Albright, Adam (2008). Explaining universal tendencies and language particulars in analogical change. In Jeff Good (ed.) *Language Universals and Language Change*, Oxford University Press. 36 pp.
- Albright, Adam, Argelia Andrade & Bruce Hayes (2001). Segmental environments of Spanish diphthongization. In Adam Albright & Taehong Cho (eds.) *UCLA Working Papers in Linguistics 7 (Papers in Phonology 5)*, UCLA. 117–151.
- Albright, Adam & Bruce Hayes (2002). Modeling English past tense intuitions with minimal generalization. In Michael Maxwell (ed.) *Proceedings of the sixth meeting of the ACL special interest group in computational phonology*. Philadelphia: ACL, 58–69.
- Albright, Adam & Bruce Hayes (2003). Rules vs. Analogy in English past tenses: a computational experimental study. *Cognition* **90**. 119–161.
- Albright, Adam & Bruce Hayes (2006). Modeling productivity with the gradual learning algorithm: The problem of accidentally exceptionless generalizations. In Gisbert Fanselow, Caroline Féry, Matthias Schlesewsky & Ralf Vogel (eds.) *Gradience in Grammar*, Oxford University Press. 185–204.
- Alderete, John, Adrian Brasoveanu, Nazarré Merchant, Alan Prince & Bruce Tesar (2005). Contrast analysis aids the learning of phonological underlying forms. In John Alderete, Chung hye Han & Alexei Kochetov (eds.) *Proceedings of WCCFL 24*, Cascadilla Press. 34–42.
- Anttila, Arto (1997). *Variation in Finnish phonology and morphology*. Ph.D. dissertation, Stanford University, Stanford, CA.
- Anttila, Arto (2002). Morphologically conditioned phonological alternations. *Natural Language and Linguistic Theory* **20**. 1–42.
- Apoussidou, Diana (2007). *The Learnability of Metrical Phonology*. Ph.D. dissertation, University of Amsterdam.
- Aronoff, Mark (1994). *Morphology by itself*. MIT Press.
- Baayen, R. Harald (2008). *Analyzing Linguistic Data: A practical introduction to statistics*. Cambridge University Press.

- Bat-El, Outi (1993). Parasitic metrification in the modern hebrew stress system. *The Linguistic Review* **10**. 189–210.
- Bat-El, Outi (2008a). A gap in the feminine paradigm of Hebrew: A consequence of identity avoidance in the suffix domain. In Curt Rice (ed.) *When Nothing Wins: Modelling Ungrammaticality in OT*, London: Equinox.
- Bat-El, Outi (2008b). Morphologically conditioned V-∅ alternation in Hebrew: Distinction among nouns, adjectives & participles, and verbs. In Sharon Armon-Lotem, Gabi Danon & Susan Rothstein (eds.) *Generative Approaches to Hebrew Linguistics*.
- Becker, Michael (2003). Lexical statification of Hebrew: The disyllabic maximum. In Yehuda N. Falk (ed.) *Proceedings of the Israel Association for Theoretical Linguistics 19*.
- Becker, Michael (2007). From the lexicon to a stochastic grammar. Workshop on Variation, Gradience and Frequency in Phonology, Stanford Univesity.
- Becker, Michael, Nihan Ketrez & Andrew Nevins (2007). The surfeit of the stimulus: Analytic biases filter lexical statistics in Turkish devoicing neutralization. Ms. UMass Amherst.
- Beckman, Jill (1997). Positional faithfulness, positional neutralisation and Shona vowel harmony. *Phonology* **14**. 1–46.
- Beckman, Jill (1998). *Positional Faithfulness*. Ph.D. dissertation, UMass, Amherst.
- Beckman, Jill N. & Catherine O. Ringen (2004). Contrast and redundancy in OT. In Vineeta Chand, Ann Kelleher, Angelo J. Rodríguez & Benjamin Schmeiser (eds.) *Proceedings of WCCFL 23*, Cascadilla Press. 85–98.
- Benua, Laura (1997). *Transderivational Identity: Phonological Relations between Words*. Ph.D. dissertation, University of Massachusetts, Amherst.
- Berent, Iris, Steven Pinker & Joseph Shimron (1999). Default nominal inflection in Hebrew: Evidence for mental variables. *Cognition* **72**. 1–44.
- Berent, Iris, Steven Pinker & Joseph Shimron (2002). The nature of regularity and irregularity: Evidence from Hebrew nominal inflection. *Journal of Psycholinguistic Research* **31**. 459–502.
- Berko, Jean (1958). The child's learning of english morphology. *Word* **14**. 150–177.
- Blevins, Juliette (2004). *Evolutionary phonology: the emergence of sound patterns*. Cambridge: Cambridge University Press.
- Boersma, Paul (1997). How we learn variation, optionality, and probability. In *IFA Proceedings*.

- Boersma, Paul (2001). Phonology-semantics interaction in ot, and its acquisition. In Robert Kirchner, Wolf Wikeley & Joe Pater (eds.) *Papers in Experimental and Theoretical Linguistics*, University of Alberta, vol. 6. 24–35.
- Boersma, Paul & Bruce Hayes (2001). Empirical tests of the gradual learning algorithm. *Linguistic Inquiry* **32**. 45–86.
- Boersma, Paul & David Weenink (2008). Praat: Doing phonetics by computer (version 5.0.24). [Computer program].
- Bolozky, Shmuel & Michael Becker (2006). Living Lexicon of Hebrew Nouns. Ms. UMass Amherst.
- Borowsky, Toni (1987). Antigemination in English phonology. *Linguistic Inquiry* **18**. 671–678.
- Breiman, Leo, Jerome Friedman, R. A. Olshen & Charles J. Stone (1984). *Classification and regression trees*. Belmont, CA: Wadsworth International Group.
- Cassimjee, Farida & Charles Kisseberth (1998). Optimal domains theory and bantu tonology: A case study from isixhosa and shingazidja. In Larry Hyman & Charles Kisseberth (eds.) *Theoretical Aspects of Bantu Tone*, Stanford: CSLI. 33–132.
- Chomsky, Noam & Morris Halle (1968). *The Sound Pattern of English*. the MIT press.
- Clayton, Mary L. (1976). The redundance of underlying morpheme-structure conditions. *Language* **52**. 295–313.
- Coetzee, Andries W. (2008). Grammaticality and ungrammaticality in phonology. *Language* **84**. 218–257.
- Daelemans, Walter, Jakub Zavrel, Ko Van Der Sloot & Antal Van Den Bosch (2002). Timbl: Tilburg memory based learner reference guide, version 4.2. Tilburg: Computational Linguistics, Tilburg University.
- de Lacy, Paul (2004). Markedness conflation in optimality theory. *Phonology* **21**. 145–199.
- Ernestus, Miriam & R. Harald Baayen (2003). Predicting the Unpredictable: Interpreting Neutralized Segments in Dutch. *Language* **79**. 5–38.
- Faust, Noam (2008). In a position to spell out: The syntactic structure of vocalic patterns in Hebrew. Ms. University of Paris VII.
- Féry, Caroline (2003). Final devoicing and the stratification of the lexicon in german. In Jeroen van de Weijer, Vincent J. van Heuven & Harry van der Hulst (eds.) *The Phonological Spectrum 1: Segmental Structure*, John Benjamins. 145–169.
- Flack, Kathryn (2007a). *The sources of phonological markedness*. Ph.D. dissertation, UMass Amherst.

- Flack, Kathryn (2007b). Templatic morphology and indexed markedness constraints. *Linguistic Inquiry* **38**.
- Gafos, Adamantios (1999). *The Articulatory Basis of Locality in Phonology*. Garland, New York.
- Gelbart, Ben (2005). *Perception of foreignness*. Ph.D. dissertation, UMass Amherst.
- Glinert, Lewis (1989). *The grammar of Modern Hebrew*. Cambridge University Press.
- Goldsmith, John (1976). *Autosegmental Phonology*. Ph.D. dissertation, MIT.
- Gouskova, Maria (2007). The reduplicative template in tonkawa. *Phonology* **24**. 367–396.
- Hargus, Sharon (1997). Augmentation as affixation in athabaskan languages. *Phonology* **14**. 177–220.
- Hayes, Bruce (1995a). *Metrical Stress Theory: Principles and Case Studies*. University of Chicago Press.
- Hayes, Bruce (1995b). On what to teach the undergraduates: Some changing orthodoxies in phonological theory. In Ik-Hwan Lee (ed.) *Linguistics in the Morning Calm 3*, Seoul: Hanshin. 59–77.
- Hayes, Bruce (1999). Phonological Restructuring in Yidj and its Theoretical Consequences. In Ben Hermans & Marc van Oostendorp (eds.) *The derivational residue in phonology*, Amsterdam: Benjamins. 175–205.
- Hayes, Bruce (2004). Phonological acquisition in optimality theory: the early stages. In René Kager, Joe Pater & Wim Zonneveld (eds.) *Constraints in Phonological Acquisition*, Cambridge: Cambridge University Press. 158–203.
- Hayes, Bruce & Zsuzsa Londe (2006). Stochastic phonological knowledge: the case of Hungarian vowel harmony. *Phonology* **23**. 59–104.
- Hayes, Bruce & Colin Wilson (2008). A maximum entropy model of phonotactics and phonotactic learning. *Linguistic Inquiry* **39**. 379–440.
- Hayes, Bruce, Kie Zuraw, Péter Siptár & Zsuzsa Londe (to appear). Natural and unnatural constraints in hungarian vowel harmony. *Language* .
- Hoard, James (1978). Obstruent voicing in Gitksan: Some implications for distinctive feature theory. In Eung-Do Cook & Jonathan Kaye (eds.) *Linguistic Studies of Native Canada*, Vancouver: UBC. 111–119.
- Hualde, Jose (1989). Autosegmental and metrical spreading in the vowel-harmony systems of northwestern Spain. *Linguistics* **27**. 773–805.
- Inkelas, Sharon, Aylin Kuntay, John Lowe, Orhan Orgun & Ronald Sprouse (2000). Turkish electronic living lexicon (TELL). Website, <http://socrates.berkeley.edu:7037/>.

- Inkelas, Sharon & Cemil Orhan Orgun (1995). Level ordering and economy in the lexical phonology of turkish. *Language* **71**. 763–793.
- Inkelas, Sharon, Cemil Orhan Orgun & Cheryl Zoll (1997). The implications of lexical exceptions for the nature of the grammar. In Iggy Roca (ed.) *Derivations and Constraints in Phonology*, Oxford: Clarendon. 393–418.
- Itô, Junko & Armin Mester (1995). The core-periphery structure in the lexicon and constraints on re-ranking. In Jill Beckman, Laura Walsh Dickey & Suzanne Urbanczyk (eds.) *Papers in Optimality Theory*, GLSA. 181–210.
- Itô, Junko & Armin Mester (1999). The phonological lexicon. In Natsuko Tsujimura (ed.) *The Handbook of Japanese Linguistics*, Blackwell. 62–100.
- Itô, Junko & Armin Mester (2003). *Japanese Morphophonemics: Markedness and Word Structure*. MIT Press.
- Jarosz, Gaja (2006). *Rich Lexicons and Restrictive Grammars - Maximum Likelihood Learning in Optimality Theory*. Ph.D. dissertation, Johns Hopkins University.
- Jesney, Karen (2007). Child chain shifts as faithfulness to input prominence. In Alyona Belikova, Luisa Meroni & Mari Umeda (eds.) *Proceedings of the 2nd Conference on Generative Approaches to Language Acquisition – North America*. Somerville, MA: Cascadilla Proceedings Project, 188–199.
- Jurjec, Peter (2009). Disjunctive lexical stratification.
- Kager, René (1996). On affix allomorphy and syllable counting. In U. Kleinhenz (ed.) *interfaces in phonology*, Berlin: Akademie Verlag. 155–171.
- Kang, Yoonjung (2002). Sound changes affecting noun-final coronal obstruents in Korean. Paper presented at the 12th Japanese/Korean Linguistics Conference, CUNY.
- Kawahara, Shigeto (2008). Phonetic naturalness and unnaturalness in japanese loanword phonology. *Journal of East Asian Linguistics* **17**. 317–330.
- Kawahara, Shigeto, Nishimura Kohei & Hajime Ono (2003). Sino-japanese as the most unmarked stratum in japanese. *Japanese/Korean linguistics* **12**. 140–151.
- Kenstowicz, Michael (1997). Quality-Sensitive Stress. *Rivista di Linguistica* **9.1**. 157–187.
- Kenstowicz, Michael & Charles Kisseberth (1977). *Topics in Generative Phonology*. Academic Press, New York.
- Kenstowicz, Michael & Charles Kisseberth (1979). *Generative Phonology: Description and Theory*. Academic Press, New York.
- Kenstowicz, Michael & Charles Kisseberth (1990). Chizigula tonology: the word and beyond. In Sharon Inkelas & Draga Zec (eds.) *The Phonology-Syntax Connection*, Chicago: The University of Chicago Press. 163–194.

- Kim, Hyunsoon (2001). A phonetically based account of phonological stop assibilation. *Phonology* **18**. 81–108.
- Kiparsky, Paul (2000). Opacity and cyclicity. *The Linguistic Review* **17**. 351–367.
- Kirchner, Robert (1998). *An Effort-Based Approach to Consonant Lenition*. Ph.D. dissertation, UCLA.
- Kisseberth, Charles (1970). On the functional unity of phonological rules. *Linguistic Inquiry* **1**. 291–306.
- Koehn, Edward & Sally Koehn (1986). Apalai. In Desmond C. Derbyshire & Geoffrey K. Pullum (eds.) *Handbook of Amazonian languages*, Berlin: Mouton de Gruyter, vol. 1. 33–127.
- Krott, Andrea, R. Harald Baayen & Robert Schreuder (2001). Analogy in morphology: Modeling the choice of linking morphemes in dutch. *Linguistics* **39**. 51–93.
- Lees, Robert (1961). *The Phonology of Modern Standard Turkish*. Bloomington: Indiana University Press.
- Lisker, Leigh & Arthur Abramson (1964). A cross-language study of voicing in initial stops: Acoustical measurements. *Word* **20**. 384–422.
- Maddieson, Ian (1984). *Patterns of Sounds*. Cambridge University Press.
- Mahanta, Shakuntala (2007). *Directionality and Locality in Vowel Harmony*. Ph.D. dissertation, Utrecht University.
- Mascaró, Joan (1996). External allomorphy as emergence of the unmarked. In Jacques Durand & Bernard Laks (eds.) *Current Trends in Phonology: Models and Methods*, European Studies Research Institute, University of Salford. 473–483.
- McCarthy, John J. (2002). *A Thematic Guide to Optimality Theory*. Cambridge University Press.
- McCarthy, John J. (2003). OT constraints are categorical. *Phonology* **20**. 75–138.
- McCarthy, John J. (2004). Headed spans and autosegmental spreading. Ms., University of Massachusetts, Amherst. [Available as ROA-685].
- McCarthy, John J. (2005). Taking a Free Ride in Morphophonemic Learning. *Catalan Journal of Linguistics* **4**. 19–55.
- McCarthy, John J. (2007a). *Hidden Generalizations: Phonological Opacity in Optimality Theory*. London: Equinox Publishing Company.
- McCarthy, John J. (2007b). Slouching toward optimality. *Phonological Studies (Journal of the Phonological Society of Japan)* **7**.

- Merchant, Nazarré (2008). *Discovering Underlying Forms: Contrast Pairs and Ranking*. Ph.D. dissertation, Rutgers University.
- Mester, Armin (1994). The quantitative trochee in Latin. *Natural Language and Linguistic Theory* **12**. 1–61.
- Moreton, Elliott (2008). Analytic bias and phonological typology. *Phonology* **25**. 83–127.
- Moreton, Elliott & Erik Thomas (2007). Origins of Canadian Raising in voiceless-coda effects: a case study in phonologization. In Jennifer Cole & José I. Hualde (eds.) *Papers in Laboratory Phonology 9*, Mouton de Gruyter. 37–64.
- Napikoglu, Mine & Nihan Ketrez (2006). Children's overregularizations and irregularizations of the Turkish aorist. In *BUCLD 30: Proceedings of the Boston University Conference on Language Development, Volume 2*, Somerville, MA: Cascadilla Press. 399–410.
- Ní Chiosáin, Maire (1991). *Topics in the Phonology of Irish*. Ph.D. dissertation, University of Massachusetts.
- Ní Chiosáin, Maire & Jaye Padgett (2001). Markedness, segment realization, and locality in spreading. In *Segmental phonology in Optimality Theory*, Cambridge University Press. 118–156.
- Odden, David (1993). Interaction between modules in lexical phonology. In Sharon Hargus & Ellen Kaisse (eds.) *Phonetics and Phonology 4: Studies in Lexical Phonology*, San Diego: Academic Press. 111–144.
- Odden, David (2005). *Introducing Phonology*. Cambridge University Press.
- Ohala, John (1983). The origin of sound patterns in vocal tract constraints. In *The production of speech*, New York: Springer-Verlag. 189–216.
- Paster, Mary (2006). *Phonological Conditions on Affixation*. Ph.D. dissertation, University of California, Berkeley.
- Pater, Joe (2000). *NC̥. In K. Kusumoto (ed.) *Proceedings of NELS*. Amherst, MA: GLSA, vol. 26, 227–239.
- Pater, Joe (2005). Learning a stratified grammar. In Manuella R. Clark-Cotton Alejna Brugos & Seungwan Ha (eds.) *Proceedings of the 29th Boston University Conference on Language Development*, Somerville, MA: Cascadilla Press. 482–492.
- Pater, Joe (2006). The locus of exceptionality: Morpheme-specific phonology as constraint indexation. In Leah Bateman & Adam Werle (eds.) *UMOP: Papers in Optimality Theory III*, Amherst, MA: GLSA. 1–36.
- Pater, Joe (2008a). Gradual learning and convergence. *Linguistic Inquiry* **39**. 334–345.

- Pater, Joe (2008b). Morpheme-specific phonology: Constraint indexation and inconsistency resolution. In Steve Parker (ed.) *Phonological Argumentation: Essays on Evidence and Motivation*, Equinox. 1–33.
- Petrova, Olga, Rosemary Plapp, Catherine Ringen & Szilárd Szentgyörgyi (2006). Voice and aspiration: Evidence from Russian, Hungarian, German, Swedish, and Turkish. *The Linguistic Review* **23**. 1–35.
- Pierrehumbert, Janet B. (2006). The statistical basis of an unnatural alternation. In Louis Goldstein, D.H. Whalen & Catherine T. Best (eds.) *Laboratory Phonology VIII, Varieties of Phonological Competence*, Berlin: Mouton de Gruyter. 81–107.
- Pinker, Steven & Alan Prince (1988). On language and connectionism: Analysis of a parallel distributed processing model of language acquisition. *Cognition* **28**. 73–193.
- Prince, Alan (2002). Arguing optimality. ROA 562-1102.
- Prince, Alan & Paul Smolensky (1993/2004). *Optimality Theory: Constraint Interaction in Generative Grammar*. Oxford: Blackwell. [ROA-537].
- Prince, Alan & Bruce Tesar (1999). Learning phonotactic distributions. Technical Report RuCCS-TR-54, Rutgers Center for Cognitive Science, Rutgers University, New Brunswick [Available as ROA-353].
- Quinlan, J. R. (1993). *C4.5: Programs for Machine Learning*. San Mateo, CA: Morgan Kaufman.
- R Development Core Team (2007). R: A language and environment for statistical computing. Vienna, Austria. (<http://www.R-project.org>).
- Rice, Curt (2006). Norwegian stress and quantity: The implications of loanwords. *Lingua* **116**. 1171–1194.
- Schreuder, Robert & R. Harald Baayen (1995). Modeling morphological processing. In Laurie B. Feldman (ed.) *Morphological aspects of language processing*, Lawrence Erlbaum, Hillsdale, NJ. 131–54.
- Skousen, Royal (1989). *Analogical Modeling of Language*. Dordrecht: Kluwer.
- Skousen, Royal (1993). *Analogy and structure*. Dordrecht: Kluwer.
- Smith, Jennifer L. (2001). Lexical category and phonological contrast. In Robert Kirchner, Joe Pater & Wolf Wikely (eds.) *Papers in Experimental and Theoretical Linguistics 6: Workshop on the Lexicon in Phonetics and Phonology*, Edmonton: University of Alberta. 61–72.
- Tesar, Bruce (1995). *Computational Optimality Theory*. Ph.D. dissertation, University of Colorado at Boulder. ROA 90-0000.

- Tesar, Bruce (1998). Using the mutual inconsistency of structural descriptions to overcome ambiguity in language learning. In P. Tamanji & K. Kusumoto (eds.) *Proceedings of NELS 28*. Amherst, MA: GLSA, 469–483.
- Tesar, Bruce (2004). Contrast analysis in phonological learning. ROA 695.
- Tesar, Bruce (2006). Learning from paradigmatic information. In *Proceedings of NELS 36*.
- Tesar, Bruce, John Alderete, Graham Horwood, Nazarré Merchant, Koichi Nishitani & Alan Prince (2003). Surgery in language learning. In G. Garding & M. Tsujimura (eds.) *Proceedings of WCCFL 22*, Somerville, MA: Cascadilla Press. 477–490.
- Tesar, Bruce & Alan Prince (2006). Using phonotactics to learn phonological alternations. In *CLS 39*. Available as ROA-620.
- Tesar, Bruce & Paul Smolensky (1998). Learnability in optimality theory. *Linguistic Inquiry* **29**. 229–268.
- Tesar, Bruce & Paul Smolensky (2000). *Learnability in Optimality Theory*. Cambridge, MA: MIT Press.
- Tessier, Anne-Michelle (2007). *Biases and stages in phonological acquisition*. Ph.D. dissertation, University of Massachusetts, Amherst.
- Tessier, Anne-Michelle (2008). Children’s exceptional words: Lexical challenges from the time-course of L1 phonological acquisition. Talk given at the University of Western Ontario.
- Trommer, Jochen (2008). Syllable-counting allomorphy by indexed constraints. Talk given at OCP 5.
- Ussishkin, Adam (2000). *The Emergence of Fixed Prosody*. Ph.D. dissertation, UC Santa Cruz.
- van der Hulst, Harry (1985). Ambisyllabicity in dutch. In Hans Bennis & Frits Beukema (eds.) *Linguistics in the Netherlands 1985*, Dordrecht: Foris. 57–66.
- Volatis, Lydia & Joanne Miller (1992). Phonetic prototypes: Influence of place of articulation and speaking rate on the internal structure of voicing categories. *J. Acoustic Soc. Am.* **92.2**. 723–735.
- Walker, Rachel (2006). Long-distance metaphony: A generalized licensing proposal. Talk given at the PhonologyFest Workshop, Indiana University.
- Wells, John Christopher (1982). *Accents of English*, vol. 2. Cambridge University Press.
- Wolf, Matthew (2007). For an autosegmental theory of mutation. In Leah Bateman, Michael O’Keefe, Ehren Reilly & Adam Werle (eds.) *UMOP 32: Papers in Optimality Theory III*, Amherst, MA: GLSA. 315–404.

- Wolf, Matthew (2008a). Mutation and learnability in Optimality Theory. In Martin Wolkow (ed.) *Proceedings of NELS 38*. GLSA, 1–25.
- Wolf, Matthew (2008b). *Optimal Interleaving: Serial Phonology-Morphology Interaction in a Constraint-Based Model*. Ph.D. dissertation, UMass Amherst.
- Zoll, Cheryl (1996). *Parsing below the segment in a constraint-based framework*. Ph.D. dissertation, UC Berkeley.
- Zuraw, Kie (2000). *Patterned Exceptions in Phonology*. Ph.D. dissertation, UCLA.