Learning Probabilistic Patterns: Influence of Homophony, L1 and Frequency

Hanbyul Song

A dissertation submitted in partial fulfilment

of the requirements for the degree of

Doctor of Philosophy

of

University College London

Division of Psychology and Language Sciences

University College London

I, Hanbyul Song, confirm that the work presented in my thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

....

Hanbyul Song

Abstract

In this thesis, I investigate whether learners' avoidance of alternation and neutralization, as well as learners' exposure to their native language (L1), affect how they learn new morpho-phonological patterns. While the effect of individual factors on morpho-phonological learning has been widely studied, whether these factors have a collective effect on learning and interact with the frequency of variants in the input has been understudied. To explore whether there are any interactive effects of these factors, I modify the type of alternations, learners' native languages, and relative frequency of variants across several repetitions of an experiment.

I exposed adult English speakers to an artificial language in which plural forms were probabilistically marked by one of two prefixes. One of the prefixes triggered either a non-neutralizing or neutralizing alternation that could create homophony. I found that English speakers generally matched the relative input frequency to their output. However, learners avoided the construction that resulted in a phonological alternation, but only when it was infrequent. This finding suggests that though there is a tendency to avoid alternations, it depends on how frequent the relative variants are in the input. Moreover, English speakers were poorer at learning the neutralizing alternation than the non-neutralizing alternation, showing their bias against neutralization that can create homophony.

Additionally, I replicated the same experiments with Korean speakers because there is abundant exposure to neutralization in their L1. I found that Korean speakers were successful at learning both neutralizing and non-neutralizing alternations, suggesting that having abundant exposure to neutralization can make new neutralization easier to learn.

Finally, I argue for a model which implements the avoidance effect as a discounting of observations that trigger homophony in the training data, rather than requiring a special constraint penalizing neutralization in the grammar. This Discount model correctly predicts the different

learning results between English and Korean speakers and provides a straightforward explanation for learners' bias against neutralization and homophony. This approach places the locus of the bias in the learning process rather than in the grammar.

Impact Statement

This thesis makes contributions in several areas of linguistics, psychology, and cognitive science. In linguistics, the findings of this thesis provide empirical evidence of learners' bias against phonological alternations and neutralization. The findings also broaden our understanding of how people learn languages by examining multiple factors (and their interactions) at the same time within one series of experiments; most experiments have examined one factor at a time. Moreover, using a probabilistic model, I explore how the experimental results can be accounted for in a general theory of learning rather than relying on a specific linguistic constraint. Such exploration can act as a bridge between linguistics and psychology/cognitive science. Beyond linguistics, this thesis may be of interest to psychologists and cognitive scientists who study how people learn and use patterns in general.

UCL Research Paper Declaration Form

referencing the doctoral candidate's own published work(s)

1. **For a research manuscript that has already been published**

   a) **What is the title of the manuscript?**

   Interaction of phonological biases and frequency in learning a probabilistic language pattern

   b) **Please include a link to or doi for the work**

   https://doi.org/10.1016/j.cognition.2022.105170

   c) **Where was the work published?**

   Cognition

   d) **Who published the work?** (e.g. OUP)

   Elsevier

   e) **When was the work published?**

   September, 2022

   f) **List the manuscript's authors in the order they appear on the publication**

   Hanbyul Song, James White

   g) **Was the work peer reviewed?**

   Yes

   h) **Have you retained the copyright?**

   Yes

   i) **Was an earlier form of the manuscript uploaded to a preprint server?** (e.g. medRxiv). If 'Yes', please give a link or doi)

   Click or tap here to enter text.
   If 'No', please seek permission from the relevant publisher and check the box next to the below statement:
   ☒
   *I acknowledge permission of the publisher named under **1d** to include in this thesis portions of the publication named as included in **1c**.*

2. **For multi-authored work, please give a statement of contribution covering all authors** (if single-author, please skip to section 4)

Hanbyul Song: Data curation, Formal analysis, Visualization, Writing – review & editing.
James White: Supervision, Conceptualization, Methodology, Formal analysis, Resources, Visualization, Writing – review & editing.

3. **In which chapter(s) of your thesis can this material be found?**

   Chapters 1 and 2

4. **e-Signatures confirming that the information above is accurate** (this form should be co-signed by the supervisor/ senior author unless this is not appropriate, e.g. if the paper was a single-author work)

   Candidate
   Hanbyul Song
   Date:
   14/03/2023

   Supervisor/ Senior Author (where appropriate)
   James White
   Date
   14/03/2023

Acknowledgement

My dissertation could not be completed without help from amazing people around me. First and foremost, I would like to express my sincere gratitude to my primary supervisor James White. I feel the most fortunate to have him as my supervisor. He is an incredible researcher, lecturer, and supervisor. For me, the process of getting this PhD has been a continuous journey of learning. Jamie taught me everything that I need to know to be a good researcher including designing experiments, running complicated statistical analyses, writing a journal paper, discussing complicated concepts, and preparing for talks at conferences, etc. I feel incredibly lucky that I had a chance to learn valuable skills and knowledge from him. I particularly appreciate him for always sparing time for me whenever I needed his advice and for his patience to answer my endless questions with great detail. There were so many times that I panicked about the problems to which I did not know the solution at the time and asked for his opinion. Jamie never hesitated to help me and redirect me to the solutions. I have not had many chances to express my gratitude towards him, but I have always appreciated his sincere and thoughtful guidance and support.

I would also like to thank my secondary supervisor Andrew Lamont for his detailed feedback and sincere advice which have brought new perspectives to this thesis. He also willingly shared his experience and thoughts, which helped me to have a positive attitude towards my research and future career. He is an amazing person and researcher to work with. I also appreciate my previous secondary supervisor John Harris from whom I learned so much on various topics of phonology. His thoughtful comments on my work gave me a chance to look at my thesis from different points of view. I must thank Andrew Nevins who invited me to events, talks and seminars and gave me chances to present my work to wider audiences. I also sincerely appreciate the academic staff of the linguistics department who had given me valuable feedback on my research and helped me to grow as a researcher.

I am indebted to my former supervisor Dongmyung Lee, at Dong-A University who guided and encouraged me to pursue further studies in phonology. It was his lecture where I first learned about phonetics and phonology, and I liked it so much that I ended up studying it for 6 more years. He never hesitated to offer me advice and guidance whenever I asked for his help. I am also sincerely grateful to my former advisor Kyuhong Hwang at Dong-A University for kindly giving me valuable advice on my academic career and guiding my research. I learned the most important concepts in morphology and syntax from his lectures. Even after I graduated, they have always spared time to listen to me, give me the most valuable insights, and help me to navigate my path. My special thanks go to Adam Albright who suggested amazing ideas for my thesis at various conferences and CreteLing. I also learned so much from Adam Chong and Yuni Kim at the ACTL 2018. I am grateful to the audience at OCP (16th), AMP (7th), and LSA (94th) for their valuable comments. The modified version of Chapter 2 was published, and I appreciate Vsevolod Kapatsinski and two other anonymous reviewers for their helpful feedback.

I also thank Andrew Clark at the Research Laboratory for Language and Speech Diversions at UCL. He miraculously solved all the technical issues I had and helped me smoothly run my experiments. I spent a lot of time in the dark basement where the lab is located, and Andrew's kindness and humour made my time there much more enjoyable. I also thank Richard Jardine who always helped solve my administrative problems. Without his help, my PhD journey could have been much slower and gloomier. I also appreciate colleagues at the CPD department, Cristina, Molly, Sonia, Veronia, and Mikael, who have supported my studies and helped me to succeed both in my research and my job.

I appreciate my fellow PhD students in the linguistics department. At the weekly writing session, Youngjin Kim has given me emotional support and courage to finish my thesis. Her kindness gave me so much energy and joy to endure stressful times. I am grateful for Florian Breit who greatly helped me to navigate my way through the linguistic world. My gratitude extends to all my fellow PhD students including Ruoying Zhao, Yara Alshaalan, Gregor Williamson, Yan Zhang, Irini Symeonidou, Jane Middleton, Emilia Molimpakis, Varvara Kus, Anna Grabovac, Diane Stoianov, Stefano Castiglione, Yiling Huo, Woraprat Manowang, Elisa Mattiauda, Erying Qin, Wenkai Tay, Shenshen Wang, and XinXin Yan and many more.

My sincere gratitude goes to my parents, Sooyoung Lee and Giyoung Song, who believed in me even when I doubted myself. They have encouraged me to pursue my studies in linguistics and supported every step of my journey. They provided me with a safe shelter whenever I needed to rest and recharge. Their unconditional love and care gave me an incredible sense of comfort and stability, which helped me to complete this PhD programme. My brother Haneol has been there for me when I needed him. His sense of humour cheered me up and gave me the courage to continue my research. Without their tremendous support and love, I could not have started or finished this thesis. I can never thank them enough for everything they have done for me. They are the best!

My sister Hangyul, words can hardly say how much I appreciate you. She has been to all of my talks at conferences, driven me to conference venues, and sat with the audience during my talks. Seeing her face in the audience gave me comfort and courage. When I was struggling, she did not think twice to come to London to emotionally support me and take care of me. She put me on priority and supported me in every way. The food she cooked for me, especially the ginger shots, kept me going for the last six months of the intense writing period I am the luckiest person to have her as my sister. I could not have finished this thesis without her.

Lastly, I am deeply thankful to Mohamed Jelali from the bottom of my heart. His love, kindness and care supported me throughout the entire PhD journey. When I was frustrated, he took me to dinner, a cup of coffee or a late-night walk around parks to give me a break from the stressful situations. We shared silly little laughs and had small sparkling moments, which melted my problems away and brought me the courage to go ahead and tackle the problems that I faced at the time. I also thank him for understanding and enduring a grumpy, frustrated, and stressed version of myself for many years. I am forever grateful for your kind heart and support.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

CHAPTER 1

Introduction

Learning a language is a task that humans generally do with ease. Most people unconsciously acquire a language and use it to communicate with one another. However, understanding how people acquire languages and what affects language learning is extremely complicated due to many complex processes involved. Upon hearing an input, language learning mechanisms process it by distinguishing contrasting sounds, recognising meaningful units and structures, and inferring meanings. Furthermore, each learning process can be influenced by various factors, such as a learner's cognitive capacity, biases, age, and previous language learning experience, as well as statistical information in the input.

Additionally, the fact that language learning occurs within learners' cognitive systems adds to the challenge of understanding the mechanism because researchers cannot directly observe it. Researchers who examine language learning can only infer it from pieces of evidence, such as existing language patterns, errors that learners make when learning a natural language, and findings from language learning experiments. Artificial language learning experiments are useful tools to study language learning, particularly when learners' output does not match the input learners were exposed to during the experiment. Such asymmetrical learning results suggest that learners have baseline assumptions about the nature of language and prefer particular patterns over others. By comparing the mismatch between the input and output, researchers can study which baseline assumptions and biases can affect the language learning mechanism.

1.1 Goals of the dissertation

This thesis examines the influence of the following four factors on learning: 1) avoidance of phonological alternations, 2) avoidance of neutralization and homophony, 3) learners' native language (L1), and 4) the frequency of variants in the input. In a natural language learning scenario, it is likely that numerous factors partially overlap and have an additive impact on the learning process than a single factor independently influences the learning. This thesis investigates multiple factors independently by controlling them to test their independent influence on language learning. Additionally, this thesis examines the additive effects of these factors on learning by simultaneously testing multiple factors at the same time within one series of experiments. Through a series of experiments, I plan to answer the following questions. The details of each question are covered in the following sections.

1) Do learners' biases against phonological alternations affect morphophonological learning?
2) Does the fact that phonological alternations often result in neutralization and homophony make it harder to learn alternations and make people even more likely to avoid them?
3) If learners have abundant exposure to neutralization in their L1, does this prior exposure facilitate learning a novel neutralizing alternation?
4) What framework can best account for neutralization and homophony avoidance during learning, including any potential influence of L1 exposure to neutralization?

The experiments implement an artificial grammar learning paradigm (AGL). In AGL experiments, learners are presented with novel language data and asked to internalize a system or capture statistical information of multiple patterns in the language. If learners are less successful at internalizing a certain pattern compared to another, it implies that they have a baseline preference or bias against the pattern. AGL is a useful method for examining a certain property of interest (e.g., a phonological process, a phonetic property, or syntactic structure, etc.) in language acquisition because unlike in natural language learning environments, researchers can fully control learning materials and learning environments in a lab. It is virtually impossible to find two natural languages that only differ in the property of interest as languages tend to have many additional differences. However, in AGL, researchers can create multiple languages that only differ in the property of interest, which allows them to explicitly test that particular property in a language learning process (Culbertson, 2023; see e.g., Hudson Kam & Newport, 2005, 2009; Wilson, 2006; Culbertson, Smolensky, & Legendre, 2012; White, 2013, 2014; White et al., 2018; Yin & White, 2018). In the current experiments, I exposed participants to a probabilistic morpho-phonological pattern in which one of the two variants triggered a phonological alternation. I tested how well participants learned the alternation and how often they chose the alternation-triggering variant in their output. In the remainder of this chapter, I cover the following concepts mentioned in the research questions: learners' bias against phonological alternation, neutralization and homophony, as well as the influence of their native language. Then, I provide an overview of each chapter of the thesis.

1.2 Influence of bias against alternation on language learning

Morphemes alternate when they occur in different forms depending on the contexts in which they appear. Phonological alternation can be triggered by adding a morpheme that can create a new context. For instance, the nasal consonant /n/ in the words *indent* [ɪ**n**dɛnt], *impossible* [ɪ**m**pɒsəbəl], and *incredible* [ɪ**ŋ**kɹɛdəbəl] alternates and specifically surfaces with the same place features as the following consonant. Phonological alternations potentially complicate language acquisition because learners not only need to recognize different forms of a morpheme but also how they surface in different environments. Previous studies show that learners disfavour phonological alternations (artificial grammar learning: Kapatsinski, 2009; Tessier, 2012; Stave, Smolek, & Kapatsinski, 2013; White, 2013, 2014; Smolek & Kapatsinski, 2018; Smolek, 2019; natural language learning:

Kerkhoff, 2007; Do, 2018). One way to explain this tendency is that learners have a preference for *paradigm uniformity* (Hayes, 1997; Steriade, 2000), which requires morphemes to have an identical surface form across morphologically related words (Hayes, 2004; White, 2017; Kapatsinski, 2018; Breiss, 2021). Words or phrases form a paradigm when they share the same root and when they are one lexical item. A unified paradigm means that a shared root has an identical surface form across members within a paradigm. However, a non-unified paradigm is found when a shared root has different surface forms across members within a paradigm.

I examine whether learners' biases against alternations affect learning a novel alternation and choosing a morphological construction. My approach to this question is to give adult English speakers a choice between two plural forms which are probabilistically marked by one of two prefixes. One of the prefixes triggers a phonological alternation whereas the other prefix never triggers alternation. A probabilistic pattern indicates that multiple forms (variants) refer to a given instance whereas a deterministic pattern means that one form is used to refer to a given incident. In probabilistic patterns, variants may have different likelihoods of being used. For instance, the word *cat* can be pronounced in at least four possible forms – [kʰæt̚], [kʰætʰ], [kʰæʔ], or [kʰæ] – and the percentage of the time that each variant is used to refer to the word *cat* depends on linguistic and social factors (Bresnan & Ford, 2010). However, some patterns in natural languages are deterministic: the plural form of the word *cat* is always *cats* ([kʰæts]), meaning no additional forms can be used for the plural of *cat*.

Previous studies examining the acquisition of probabilistic patterns show that adults generally replicate the relative frequency of variants presented in the input to their output, which refers to *frequency-matching behaviour* (a.o., Hudson Kam & Newport, 2005, 2009; Hayes, Siptár, Zuraw, & Londe, 2009; Austin, 2010; Schuler, 2017). However, learners often regularize their outputs by excessively using a more frequent pattern in their output than what they saw in the input. Regularization is more frequently found in children than in adults (artificial grammar learning experiment: Hudson Kam & Newport 2005, 2009) as well as in an emerging language learning environment (sign language learning: Ross, 2001; Senghas & Coppola, 2001; Singleton & Newport, 2004). Although adult learners tend to show frequency-matching behaviour, they also regularize their output (linguistic patterns: Hudson Kam & Newport, 2009; Culbertson et al., 2012, non-linguistic patterns: Gardner, 1957; Weir, 1964, 1972; Starling, 2013).

In experiments in Chapter 2, plural nouns are probabilistically marked by adding one of two prefix forms (*ba-* and *ni-*). While the *ni-* prefix triggers the palatalization of velar-initial stems, the *ba-* prefix never triggers an alternation. The probability of the two prefix forms in the input differs by conditions. Adult English speakers are then tested on how well they learn an alternation and how often they choose the alternation-triggering variant in their output. Through these experiments, I show that English speakers generally match the relative frequency of prefix forms presented in the input to their output; however, they shift away from the alternation-triggering prefix form when it is infrequent in the input. This finding implies that learners tend to avoid phonological alternations, but this is affected by the frequency of a variant that triggers alternations in the input.

1.3 Influence of bias against neutralization and homophony on language learning

I further investigate whether learners' aversion to a phonological alternation can be even stronger when an alternation triggers neutralization and homophony. Phonological alternations can neutralize the contrasting values between phonemes and merge distinct phonemes. This type of alternation results in neutralization at a phonological level. For example, in German, devoicing of word-final consonants can eliminate the contrasting voicing feature ([±voice]) between /d/ and /t/ (e.g., /bʁan**d**/ 'fire' - [bʁan**t**], /ʁaː**t**/ 'council/advice' - [ʁaː**t**]). Another type of neutralization is at a lexical level, in which a merge of phonemes further deletes lexical contrasts. The neutralization at a lexical level can also refer to derived homophony (Silverman, 2010): multiple distinct lexical items have an identical (or near identical) pronunciation. In German, devoicing of word-final consonants can also create derived homophony (e.g., /bun**d**/ 'league', /bun**t**/ 'colourful' - [**bunt**]).

Neutralizing alternations are particularly interesting because they impede one of the crucial functions of speech sounds, which is to distinguish sounds from one another. Previous studies on neutralization and homophony have mainly focused on diachronic change (e.g., Martinet, 1952; Hockett, 1967; King, 1967; Wedel, Kaplan, & Jackson, 2013). Martinet suggests that a pair which builds more frequently occurring minimal pairs is less likely to be neutralized in a language. Other studies compare the existing and hypothetical neutralizing alternations, showing that existing ones create less homophony than hypothetical ones (Silverman, 2010; Kaplan, 2011). Additionally, a handful of recent studies have directly examined how learners acquire patterns with neutralization and homophony (Kapatsinski, 2009; Yin & White, 2018). Yin and White's study shows that learners are poorer at learning a neutralizing alternation than a non-neutralizing alternation when the neutralizing alternation triggers homophony. Learners may disfavour neutralizing and homophony-creating patterns due to their preference for a one-to-one relationship. Learners tend to have a baseline assumption to link one label (surface form) with one category (an object, a lexical item, a morpheme, a grammatical function, etc.; Aronoff, 1976; Markman & Wachtel, 1988; Pinker, 1996). However, homophony has a many-to-one relationship because one label (surface form) can refer to multiple lexical items, which may motivate learners' baseline bias against homophony-creating patterns.

In Chapter 2, I study whether the fact that a phonological alternation triggers neutralization and homophony makes it harder to learn the alternation and whether it causes people to be even more likely to avoid the alternation-triggering prefix variant. I manipulated whether the alternation was neutralizing (and homophony-creating) or non-neutralizing, and I found that English speakers were worse at learning the alternation when it was neutralizing. However, they did not show a greater tendency to avoid the trigger of the alternation when it was neutralizing than when it was not neutralizing.

In Chapter 4, I consider two models that account for the experimental results. In one model, which I call the Discount model, I implement learners' avoidance of neutralization and homophony by discounting input data that result in derived homophony. The derived homophony can interfere with learners' ability to build a paradigmatic relationship between related words due to a bias against many-to-one relationships. This interference can hinder learners' ability to learn a phonological alternation from homophonous input. The discounting in the Discount model reflects this reduced effect of homophonous input for learning the alternation. I show that the model correctly predicts the experimental results and provides a simple and effective explanation for the avoidance of neutralization and homophony. Furthermore, I demonstrate that the Discount model outperforms an alternative model in which the avoidance of neutralization is implemented as a constraint in the grammar.

1.4 Influence of native language on language learning

In addition to learners' avoidance of phonological alternations and neutralization, I examine whether learners' existing linguistic knowledge adds to these effects. One source of existing linguistic knowledge comes from learners' exposure to their native languages (L1). Learners often *transfer* phonetic and phonological information (Rubach, 1980; Altenberg & Vago, 1983) or syntactic structure (Bhela, 1999) from their L1 when learning a natural target language (e.g., a second or a third language). L1 transfer can often result in errors in a target language. Altenberg and Vago (1983) showed that native Hungarian speakers who learn English as their second language incorrectly apply the Hungarian *regressive* voicing assimilation rule (the voice feature of the preceding obstruent is assimilated to that of the following obstruent) to English words. They often incorrectly devoice the last obstruent of the first word when the initial consonant of the following word is voiceless (e.g., *[biyon**t h**iz], for 'beyond his' [bijɑn**d h**ɪz]; Altenberg and Vago, 1983, p. 432). Although the L1 transfer effect in learning a natural target language has been widely studied, a relatively smaller number of recent studies examined the L1 effect in artificial language learning experiments. These studies showed that learners' pre-established knowledge from their L1 can transfer to learning an artificial language (statistical information: Finn & Hudson Kam, 2008; dominant patterns: White et al., 2018; Martin & Culbertson, 2020; lexical information: Tang & Baer-Henney, 2023).

In Chapter 3, I consider the additional influence of the L1 effect in morpho-phonological learning by replicating the experiment from Chapter 2 with adult native Korean speakers. Korean has numerous neutralizations, which sometimes create homophony (Silverman, 2010). Thus, native Korean speakers are more likely to be frequently exposed to neutralizing alternations in their L1 than are native English speakers. Comparing the results between Korean and English speakers would show whether having abundant neutralization in L1 can influence learning a new type of neutralization and choosing the construction that triggers neutralization. I found that Korean speakers were equally good at learning the neutralizing alternation and the non-neutralizing alternation. Compared to the results of English speakers, who were less successful at learning the neutralizing alternation than the non-neutralizing alternation, the L1 effect likely facilitated Korean speakers' acquisition of the neutralizing alternation.

1.5 Plan for the dissertation

In **Chapter 2**, I introduce learners' avoidance of phonological alternation and neutralization as well as the relative frequency of variants in the input. I then explain the experimental design that tests the potential impact of these factors on learning a morpho-phonological pattern using the artificial language learning paradigm. Next, I show that adult native English speakers generally replicated the relative frequency of prefix forms in the input to their output; however, they shifted away from the prefix form that triggered a non-neutralizing alternation only when it was infrequent in the input. Moreover, they were less successful at learning the neutralizing alternation than the non-neutralizing alternation but did not shift away from the alternation-triggering prefix.

In **Chapter 3**, I investigate an additional influence of learners' exposure to their L1 on these effects. I replicate the experiment from Chapter 2 with adult native Korean speakers. I provide an overview of neutralizations in Korean to show that Korean speakers are more frequently exposed to neutralizations in their L1 than are English speakers. I find that Korean speakers learned a neutralizing alternation as successfully as a non-neutralizing alternation, indicating the L1 effect on learning a novel neutralizing alternation. Moreover, Korean speakers matched the frequency of prefix variants presented in the input to their output.

In **Chapter 4**, I consider two frameworks, the Discount and *Neutralization models, to account for learners' avoidance of neutralization and homophony. I conclude that the Discount model, which implements the avoidance effect by causing the model to learn less effectively from homophonous input, provides a simpler and more accurate explanation for the avoidance effect than the *Neutralization model which includes a constraint within the grammar.

Finally, in **Chapter 5**, I provide a summary of the findings and future research suggestions.

CHAPTER 2

Effects of phonological alternation, neutralization, and frequency on morpho-phonological learning

2.1 Introduction

This chapter demonstrates the influence of learners' biases against alternations, neutralization, and homophony on the acquisition of novel alternations and the choice of a morphological construction. The topic of which biases can affect phonological learning is still debatable. One factor that has an impact on language learning is an *analytic bias* (Moreton, 2008)*,* which refers to learners' systematic inclination for particular patterns over others. An analytic bias, also sometimes refers to as a learning bias, can facilitate learning the preferred patterns. Previous research suggests that learning biases can also serve as a mechanism leading to language change and typological asymmetry (Hudson Kam & Newport, 2005; Wilson, 2006; Moreton, 2008; Culbertson et al., 2012; White, 2013). For instance, Moreton (2008) examined the typological asymmetries in vowel-height harmony patterns and consonant-continuancy harmony patterns across languages. Although the two patterns have a similar strength of phonetic precursors, vowel-height harmony patterns are more common than consonant-continuancy harmony patterns in languages. In an experiment, Moreton found that learners learned vowel-height harmony patterns more successfully than consonant-continuancy harmony patterns, indicating that learners are biased in favour of the vowel-height harmony patterns.

An example of a very strict view of analytic bias (a 'hard bias' view) is classical Optimality

Theory (OT) which assumes that a grammar consists of a set of universal constraints and that only patterns that are formed by different rankings of these constraints are learnable (Prince & Smolensky, 1993/2004; McCarthy & Prince, 1995). This view is often used to explain typological asymmetries found across languages. A less strict view of analytic bias (a 'soft bias' view) indicates that learners can still learn the patterns that they are biased against, but they may struggle to learn them or learn them more slowly. Language learning experiments using an artificial language can be useful tools to study learners' biases. In these experiments, learners are required to learn specific patterns, and they often result in outputs that are different from what they saw in the input or are less successful at learning a certain pattern compared to another. Asymmetric results imply that certain biases affected the learning and resulted in the discrepancy between input and output.

In the following sections, I first introduce learners' bias against phonological alternations and neutralization. I then cover the effect of the relative frequency of variants in the input on choosing the variants in the output and provide an overview of the experiments.

2.2 Learners' avoidance of phonological alternation

Morphemes often alternate and are realized as different surface forms according to the contexts in which they occur. Phonological alternations can be triggered by adding a morpheme that creates a new context. For example, the word-final /k/ in *electric* [ilɛktɹɪ**k**] alternates to [s] when a morpheme '*-ity*' is attached to the word *electric+ity* [ilɛktɹɪ**s**əti]. Learners often show their bias against phonological alternations (artificial language learning: Kapatsinski, 2009, 2013; Tessier, 2012; Stave, Smolek, & Kapatsinski, 2013; White, 2014, 2017; natural language learning: Kerkhoff, 2007; Do, 2018).

One explanation for learners' aversion to phonological alternations is that they prefer to maintain the surface form of a shared root across morphologically related words. In language change, alternations among morphologically related forms are often deleted, which is referred to as *paradigm* or *analogical levelling* (Bybee & Brewer, 1980; Tiersma, 1982; Hayes, 1997; McCarthy, 1998). Paradigm levelling results in consistency among inflectional forms (e.g., pre-classical Latin, hono:**s** → hono**r,** has changed to maintain consistency in hono:**r**is, hono:**r**i:, hono:**r**em; Hock, 1991; Albright, 2005). From this observation, learners' bias in favour of *paradigm uniformity* was proposed (Steriade, 2000). Paradigm uniformity indicates that learners have a baseline assumption that a shared root has an identical surface form across the members in a paradigm. Words or phrases are in a paradigm when they share a root (e.g., {*fung-us, fung-i, fung-al* …}) A unified paradigm means that a shared root has an identical surface form across the members of a paradigm (Steriade, 2000). To illustrate, in (1) below, (1a) shows a unified paradigm, as the shared root /dɛns/ has an identical surface realization ([dɛns]) across words within a paradigm. (1b) shows a non-unified paradigm, as the shared root has multiple surface realizations, [ɪlɛktɹɪk], [ɪlɛktɹɪ**s**], and [ɪlɛktɹɪ**ʃ**] across the members in a paradigm. The multiple surface realizations of the shared root in a non-unified paradigm are triggered by phonological alternations (e.g., k→s and k→ʃ in (1b)).

(1) Examples of a unified and non-unified paradigm.

a. unified form: {*dense*: dɛns, dɛnsəti, dɛnsnəs, dɛnsɪfaɪ, dɛnsli …}

b. non-unified form: {*electric*: ɪlɛktɹɪk, ɪlɛktɹɪ**s**əti, ɪlɛktɹɪ**ʃ**ən, ɪlɛktɹɪkəl, ɪlɛktɹɪks …}

Previous studies show evidence that children learning their native languages are biased in favour of paradigm uniformity. In modern Greek, the legal plural form (the form that is used by adults) for [e**x**o] 'I have' is [ˈe**ç**ete] 'you-pl. have'. Kazazis (1969) found that a 4-year-old child learning modern Greek produced an illegal form, *[ˈe**x**ete], instead of the legal form, [ˈe**ç**ete]. The illegal form, *[ˈe**x**ete], produced by children implies that their baseline assumption is to preserve the surface form of the shared morpheme ([ex-]) across morphologically related words, indicating that children have an aversion to the phonological alternation that triggers a non-unified paradigm.

In addition, further evidence of children's preference for paradigm uniformity when learning their native languages comes from a recent study showing that Korean children avoid alternations (Do, 2018). Korean children (ages 4-7) were asked to orally produce inflected forms of verbs in a sentence completion task. Younger children (ages 4-5) mostly maintained the surface form of the base stem in inflected forms, but adults alternated the base stem (e.g., base stem: [jəɾ] 'open', correct form (adult's form): [**jə**-n-da] 'open-PRES-DECL', younger children's forms: *[**jəɾ**-ɨ-n-da]). Children often used the morphologically correct inflectional form, but they did not apply the phonological alternation by incorrectly adding a filler /-ɨ/ to maintain the base stem([jəɾ]). [1] Older children (ages 6-7) used the correct form more frequently than younger children, but they also used extended morphological structures that are not preferred by adults more often than younger children, presumably as a way to avoid the alternation. Older children attached additional morphemes, such as an auxiliary that slightly alters the meaning but maintains the surface realization of the stem (e.g., older children's forms: [**jəɾ**-ə-bo-n-da] 'open COMP-try-PRES-DECL'). Crucially, when the alternation of a verbal stem was syntactically obligatory, older children mostly produced correct forms, but younger children still made errors. This finding suggests that although children know the correct alternation, their preference for a unified paradigm still leads them to avoid the alternation and affects their choice of morphological structure. These findings also demonstrate that children choose alternative morphological constructions and select alternative stems to avoid non-unified paradigms.

Furthermore, a study using an artificial grammar learning paradigm (AGL) showed that children's tendency to avoid a phonological alternation to maintain a unified paradigm has been

---

[1] Children seemed to incorrectly add /-ɨ/ because /jəɾ/ only appears before a vowel. This is an incorrect stem allomorph but they presumably used it to keep the surface form of the base stem ([jəɾ]).

found among children as young as 4 years old (Tessier, 2012). In an experiment using AGL, children were exposed to consonant sequences (e.g., [...bd…] and […ʧd…]) that were presented in two different contexts. When sequences were in morphologically complex count nouns (e.g., stem + plural suffix), the first consonant was in a stem-final position and the second consonant was in a suffix-initial consonant (e.g., one po**b** and two po**b**+**d**eɫ, one wa**ʧ** and two wa**ʧ+d**eɫ). When the same sequences were in monomorphemic mass nouns, the same sequences were within a stem (e.g., some gɪ**db**it, some zɪ**ʧd**ɪn). The results showed that children made more errors and changed the sequences more often when the sequences were in morpheme-medial positions than when they were in stem-suffix boundaries (e.g., where the stem would be analysed as part of a paradigm). These findings suggest that children prefer to keep the same surface realization of a shared stem across a paradigm, showing their preference for paradigm uniformity.

In the current experiments, learners' preference for paradigm uniformity is tested using an artificial language learning experiment. During the experiments, learners were required to learn a probabilistic pattern in which one of the morphological constructions triggered an alternation of stems. I examined how often learners chose the construction that triggered an alternation in their output and how well they learned the alternation.

2.3 Learners' avoidance of neutralization and homophony

The second factor I investigate is whether people find it harder to learn phonological alternations when they trigger neutralization and homophony and choose the alternation-triggering pattern less often in their output. A phonological system has contrasting units that distinguish one sound from another, and these units can build phonemic categories. Phonological alternations often remove and neutralize contrasting values of phonemes in certain phonological contexts. For example, in English, the contrasting voicing feature distinguish /t/ and /d/, which further differentiates the words ***pat*** from ***pad***. When the word-final /t/ and /d/ are followed by the *-ing* suffix, they both alternate to [ɾ] (e.g., pat-*ing* – [pæ**ɾ**iŋ], pad-*ing* – [pæ**ɾ**iŋ]), meaning that the alternation neutralizes the contrasting voicing feature and phonemes are no longer distinguishable (/t/, /d/ - [ɾ]). This observation raises a further question: if learners already disfavour phonological alternations, would they show an even stronger bias against phonological alternations that trigger neutralization?

Neutralization appears at two levels: a phonological and lexical level. Neutralization at the phonological level occurs when the contrasting values that distinguish phonemes are merged, and the phonemes are no longer distinct. For instance, in a number of languages, the voicing feature of word- or syllable-final obstruents are often neutralized to [-voice] (Trubetzkoy, 1939; Moulton, 1962; Dinnsen, 1985; Slowiaczek & Dinnsen,1985), meaning that some underlying voiced and voiceless obstruents are no longer distinguishable in surface forms (e.g., German: /anoʁɐk/ 'jacket' – [anoʁɐ**k**], /klu:g/ intelligent' – [klu:**k**]). In these examples, the neutralization of the voicing feature eliminates the distinction between phonemes /k/ and /g/[2].

Additionally, neutralization appears at a lexical level when the merge of contrasting features further results in the identical (or near identical) pronunciation of distinct lexical items, creating derived homophony (Silverman, 2010). In Korean, the laryngeal features of the obstruents, /ʧ, ʧʰ/, are neutralized to the unreleased [t˺] in word- or syllable-final positions, which results in the identical surface form for the distinct lexical items (e.g., /piʧ / 'a debt', /piʧʰ/ 'light' - [pit˺]). Additionally, in German, the neutralization of contrasting voicing feature of word-final obstruents often results in homophony (e.g., *rad* /ʁa:**d**/ 'wheel', *rat* /ʁa:**t**/ 'council' – [ʁa:**t**]). However, neutralizations do not always create derived homophony. In Korean, /t'ok/ 'exactly/decisively' is realized as [t'ok˺] due to the neutralization of laryngeal features of a word-final obstruent, but the neutralization does not create homophony because there are no words like /t'ok'/ or / t'okʰ/ that can be realized as [t'ok˺].

---

[2] Previous studies have shown that the neutralization of the voicing feature of domain final obstruents can be incomplete (see Roettger, Winter, Grawunder, Kirby, and Grice, 2014 for detailed reviews).

Hence, neutralization can potentially create derived homophony, but it does not always create homophony.

Assuming that the primary goal of languages is clear and efficient communication, languages should have developed to maximize the accuracy of delivering and perceiving distinct sounds and words. However, neutralization (especially at a lexical level) often jeopardizes this goal by merging distinct sounds. Neutralization can particularly obstruct clear communication when listeners cannot use external context to distinguish phonemes or lexical items (Blevins & Wedel, 2009). External context generally refers to information that can be obtained other than phonetic and phonological contexts (e.g., grammatical categories of words, syntactic structures, or pragmatic information).

Several studies have focused on how sound changes can be blocked when they create homophony, particularly when this causes ambiguity in communication. For instance, Martinet (1952) suggests that functional load[3] plays a pivotal role in the application of a neutralizing alternation. Although the definition of the functional load varies, Martinet (1952, p.8) defined it as "the number of lexical pairs which could be complete [homonyms] if it were not that one word of the pair presents one member A of the opposition where the other shows the other member B". The functional load of an opposition is high[4] when the opposition builds a number of frequently occurring minimal pairs. It can be also high when an opposition forms minimal pairs that play a crucial part in distinguishing and understanding sentences. An opposition with a high functional load is less likely to be neutralized because doing so may result in many homophonies and potentially jeopardize communication.

[3] Martinet expanded the *functional load* mentioned earlier by Jakobson (1931), Mathesius (1931), and Trubetzkoy (1939). The term functional load is interchangeable with 'functional yield' or 'functional burden'.

[4] In this paper, the terms 'high' or 'low' are used to indicate the functional importance of maintaining the phonemic oppositions; however, the terms can be replaced by numerical ratings if needed (Martinet, 1952).

Martinet (1952) showed that some long and short vowels in Parisian French were likely to be merged as their functional load is low. Old French used to have contrasting lengths of the vowels /i/-/i:/, /ü/-/ü:/, /u/-/u:/, /e/-/e:/, and the contrasting lengths were primarily used to mark genders of words. However, the genders of words are also marked by obligatory gender markers adjacent to pronouns or articles, indicating that there was a low chance that the contrasting vowel lengths built minimal pairs and acted as a sole cue to determine the genders of words. Hence, the functional load of the vowel lengths was low, which could have caused the merge of vowel lengths.

In addition, Blevins and Wedel (2009) proposed that expected sound changes cannot occur when they result in homophony that causes ambiguity. Neutralization can particularly obstruct clear communication when listeners cannot use external context to distinguish phonemes or lexical items (Blevins & Wedel, 2009). In modern Banoni, a Western Oceanic language, the historic distinction between long and short vowels has mostly been neutralized (e.g., /man-**a**:/ 'give-1SGO" and /man-**a**/ "give -3SGO" – [man**a**]; Lincoln, 1976a). However, the neutralization of vowel lengths is inhibited in the words /tam**a**/ 'father' and /tam**a**:/ 'my father'. This is because the homophony [mana] is likely to appear in a context where there are external cues to correctly categorize the lexical items and interpret sentences (e.g., [mana vai] "give me it" or [mana i] "give him/her it'). However, the words /tam**a**/ 'father' and /tam**a**:/ 'my father' are likely to appear in the same or similar syntactic, semantic, and pragmatic contexts (no other external information can assist in categorizing the lexical items). Accordingly, eliminating the contrasting vowel lengths of word-final positions can trigger ambiguity, and the neutralization of vowel lengths is likely to be inhibited.

Learners might also be biased against neutralization and homophony because they have a preference for a one-to-one relationship. Learners may have a baseline preference to exclusively assign one label (surface form) to one category (Aronoff, 1976; Markman & Wachtel, 1988; Pinker, 1996). For example, if learners assign the label 'trees' to refer to a group of plants that grow tall and have branches, bark, and leaves, they tend to exclude other possible labels (e.g., 'birds', 'cars', etc.)

to refer to the same group of plants. However, homophony has a many-to-one relationship because one label (surface form) refers to multiple categories (e.g., one form [pæɾiŋ] can be either *pat-ing* or *pad-ing* in North American English), which could cause learners to be biased against homophony.

Markman and Wachtel (1988) proposed that children may initially assume that labels that are used to refer to categories of items are mutually exclusive. In a series of experiments, Markman and Wachtel showed that, upon hearing a novel nonce label, children tended to assign the novel label to refer to an unfamiliar item (e.g., a cherry pitter or a radish rosette maker) of which they might not know the label. When they already knew the labels referring to familiar items (e.g., a cup or a banana), they frequently rejected the new label to refer to the familiar item. Instead, children often used the novel label to refer to a specific part (or a sub-category) of the familiar items.

In addition, the *uniqueness entry principle* suggests that children have a baseline assumption to link one structure (morpheme) to one grammatical function unless they receive direct evidence that multiple structures can refer to the same function (Wexler & Culicover, 1980; Pinker, 1996). Pinker suggested that children are likely to need more pieces of evidence to establish a many-to-one relationship between multiple structures and a grammatical function compared to a one-to-one relationship. Furthermore, children are more prone to ignore morphemes (and to not use them to refer to a grammatical function) when the morphemes do not have phonetic properties or are homophonous with another morpheme within a paradigm. For instance, if children already use the morpheme, *-ed,* to indicate a past tense of verbs (e.g., *walk–walked*) in English, they may particularly need more evidence to add a zero-inflected morpheme (e.g., *cut–cut*, *put–put*) that does not have phonetic properties or a morpheme that triggers homophony between past and past participles of verbs (e.g., *spring–sprung–sprung, loose–lost–lost*) to indicate the past tense of verbs as well as the existing *-ed* morpheme. Furthermore, Aronoff (1976) suggested that new forms are less likely to be derived when there is already an existing form with the same meaning in the lexicon. For instance, *+ity* is often attached to an adjective ending with *-ous* to form a nominative (e.g., curious –

curiosity). However, when the nominal form of an adjective is already occupied by another form, attaching *+ity* is blocked (e.g., glorious – glory, *gloriosity; Aronoff, 1976, p. 44). These studies indicate that learners may have an *a priori* bias that causes them to consider patterns that obey a one-to-one relationship to be more likely. Therefore, learners may need more observations to successfully learn patterns that create neutralization and homophony, due to them representing a many-to-one relationship.

A few recent studies directly examine the effect of neutralization and homophony on phonological alternations. Yin and White (2018) examined the effect through an artificial language learning experiment. During the experiment, participants were required to learn singular nonce words (CVCVC) and matching plural forms which were created by adding the suffix /-i/ (CVCVC-i). The plural suffix /-i/ triggered the palatalization of stem-final alveolar consonants (/t, d, s, z/→[ʧ, ʤ, ʃ, ʒ]/__i), which could be non-neutralizing or neutralizing (e.g., t, ʧ→[ʧ], d, ʤ→[ʤ]). In addition, learners were exposed to both neutralizing and non-neutralizing alternations during the experiment, but the number of homophonous words created by the neutralizing alternation varied depending on the condition. The experimental results showed that learners were poorer at learning the neutralizing alternation compared to the non-neutralizing alternation when the neutralizing alternation triggered homophony. Thus, Yin and White showed empirical evidence that learners are biased against homophony-creating neutralization.

Kapatsinski (2009, 2012, 2013) also showed that learning a novel phonological alternation can be affected by neutralization and homophony. In a series of experiments using an artificial language learning paradigm, Kapatsinski exposed adults to plural words that were formed by adding either a plural suffix /-i/ or /-a/ to singular stems. The velar-final stems were palatalized when followed by the suffix /-i/ (/-k/, /-g/→[-ʧi], [-ʤi]). In one of the conditions, additional stems ending with palatal consonants, /-ʧ/ or /-ʤ/, were introduced to trigger neutralization and create homophony. The experiments also compared two learning tasks. One task was designed to encourage participants to acquire source-oriented generalizations by emphasizing pairs of singular and plural forms (e.g., the plural form was created by adding a suffix to the singular stem). The other task intended to favour product-oriented generalizations where ideal plural forms were highlighted (e.g., plural forms should end with -*Ci* or -*Ca*; Bybee, 2001) but a connection between a singular and plural form was weakened.

The extra cases of neutralization resulted in different learning outcomes in the two learning tasks. In the source-oriented task, the additional neutralization cases decreased the percentage of the time that learners palatalized velar-final stems. This was not affected by the neutralization and homophony themselves but by the fact that the additional neutralization cases strengthened the particular source-oriented generalization that a plural was formed by adding a suffix to a stem with no stem change (e.g., [-k] → [-ki]). Accordingly, the additional neutralization cases were likely to encourage learners to add a plural suffix to velar-final stems without palatalizing them. In contrast, the additional neutralization cases in the product-oriented task enhanced the application of palatalization because they supported the product-oriented generalization that the plural forms should end with [-ʧi] or [-ʤi]. Thus, learners were more likely to palatalize velar-final stems to have the plural forms ending with [-ʧi] or [-ʤi].

The experiments in this study will take a similar approach to that of Kapatsinski. Both experiments test how often learners alternate stems given two variable plural affixes and examine the neutralization and homophony avoidance effect on learning phonological alternations. However, one major difference between the two experiments is that, in Kapatsinski's study, additional neutralization cases were introduced, which caused greater support for product-oriented generalizations and therefore encouraged palatalization. In the current experiments, I manipulated neutralization across conditions without changing the total number of trials or the amount of support for product-oriented generalizations. In this way, I could test the effect of neutralization and homophony without introducing a confound in terms of support for product-oriented generalizations.

2.4 Learning a probabilistic pattern: frequency matching and regularization

In the current experiments, participants were required to learn an alternation that was triggered by one of the variants in a probabilistic pattern. I exposed learners to a probabilistic pattern instead of a deterministic pattern because natural languages often have probabilistic patterns. Probabilistic patterns indicate that multiple variants (i.e., surface forms, words, phrases, patterns, etc.) can be used to refer to one instance. For example, an English suffix, *-ing*, can be realized as either [-ɪŋ] (e.g., *working*) or [-ɪn] (e.g., *workin'*) in England and Northern America, and the likelihood of each variant being used for *-ing* may vary depending on several factors, such as a grammatical category, perception, a social class, and sex (Fischer, 1958; Woods, 1979; Houston, 1985; Roberts, 1994). Conversely, deterministic patterns mean that which form appears is fully predictable: the past tense of the word *work* always has the same surface form (i.e., [wɝːkt]).

Learning a probabilistic pattern seems more challenging than learning a deterministic pattern because learners need to track the likelihood of each variant and choose which variant to use in their output. In some cases, the usage of variants in a probabilistic pattern can be arbitrary and inconsistent. The unpredictable use of variants can be found among learners who learn a second language (L2) later in their lives. They often arbitrarily omit or use tense, aspect, and agree markings in L2 (Johnson & Newport, 1989; Birdsong, 1999). Unpredictable use of variants can also be found in native language cases. In Russian, some prepositions can be either realized as C or CV forms before certain lexical items (e.g., [**s** mnˈoʐ̥ətvəm] or [**se** mnˈoʐ̥ətvəm] 'with a large amount (mathematical) set'; Linzen et al., 2013, p.455). As the two realizations ([se] and [s]) of the preposition inconsistently appear with the same stem, phonological (contextual) factors alone cannot predict the likelihood of each variant.

One of the factors that can influence learners' choice of a variant in probabilistic patterns is the relative frequency of variants in the pattern. There are two major kinds of frequency that can affect learning probabilistic linguistic patterns: *token* and *type frequency*. Token frequency refers to how often a unit (e.g., a word, an item, or a pattern, etc.) occurs in a language. For instance, the token frequency of the verb *have* is likely to be higher than that of the verb *obnubilate*. Bybee (2001) suggested that there are two possible effects of token frequency. First, items with higher token frequency tend to be resistant to grammatical changes and paradigm levelling, indicating that frequent words are prone to having different surface forms among morphologically related words (Bybee, 2001; Phillips, 2001). For example, the irregular past forms of frequent verbs (e.g., *keep–kept* and *sleep–slept*) are likely to keep their irregular forms whereas the irregular past forms of less frequent verbs (e.g., *weep–wept, creep–crept*, and *leap–leapt*) are more likely to alternate to regular past forms (e.g., *weeped, creeped,* and *leaped*) that maintain the surface form of the shared morpheme (Hooper, 1976; Bybee, 2001, p. 12). Second, other alternations such as reduction are more quickly applied to items with higher token frequency. For instance, phrases and words that occur frequently tend to be reduced (e.g., shortening phrases: 'be going to' – 'gonna' [gʌɾ̃ə]; Krug 1998; the reduction of syllabicity in the ə+[+sononant] sequence: memory [memᵊɹi], camera [kæmᵊɹə]; Hooper, 1976; Bybee, 2001; reduction of word-final /t/ and /d/: just [ʤʌs(t)], and [æn(d)]; Bybee, 2000b, 2001).

By contrast, type frequency indicates how often a certain pattern (or a phenomenon) is found or used among the full set of patterns in a language. This is observable with plural forms in English, where adding the /-s/ plural suffix (e.g., a dog- dog**s**) is more likely to occur than all other plural forms (e.g., an octopus-octop**i**, a phenomenon-phenomen**a**, a fish-**fish**), suggesting that type frequency of plural forms using the /-s/ suffix is higher than other plural forms. Type frequency is closely related to the productivity of a pattern, which implies that a pattern with high type frequency is more likely to be applied to new words (e.g., borrowed words or new forms) than a pattern with low type frequency (Bybee, 2001). As a result, when English speakers hear a new noun, they may add the /-s/ plural suffix to the new noun to make its plural form.

When learning a probabilistic linguistic pattern, adult learners generally show *frequency-matching* (or probability-matching) behaviour, meaning that they track the relative frequency of variants[8] presented in the input and replicate the frequency in their output (Hudson Kam & Newport, 2005, 2009; Hayes, Siptár, Zuraw, & Londe, 2009; Austin, 2010; Schuler, 2017). Ernestus and Baayen (2003) tested whether learners replicate the frequency of voicing patterns in their L1 to novel nonce words. They exposed Dutch speakers to nonce words that ended with various obstruents and studied how often they chose the underlying voicing features for the given word-final obstruents. In Dutch, the voicing feature in word-final obstruents is mostly neutralized. The percentage of the time that word-final obstruents are underlyingly [+voice] is highest in word-final velar fricatives and lowest in word-final bilabial stops (Baayen, Piepenbrock, & Gulikers,1995). The results showed that learners closely applied the proportion of the underlying voicing features in Dutch to the nonce words. For example, 74.1% of the participants chose the underlying voice feature of the word-final obstruent in a nonce word [te:s] to be voiced (/te:**z**/), which closely matched the percentage of the time (76.5%) that word-final alveolar fricative preceded by a long vowel was underlying [+voice] in the corpus. Additionally, Hayes et al. (2009) tested whether adult Hungarian speakers reflect probabilistic patterns in their native language to nonce words. They found that certain stem types exhibit variation in what type of suffix they take (front or back) and participants matched these frequencies in their responses.

Though learners often show probability-matching behaviour, they sometimes regularize their output with the frequent variant by using it even more often than what they saw in the input (and using the infrequent variant less often). Regularizing can lead to more deterministic patterns by reducing the amount of variability in a pattern. Learners' tendency to regularize their output can be found in learning non-linguistic probabilistic patterns (Gardner, 1957; Weir, 1964; Starling, 2013).

[8] It could be type or token frequency depending on the study and what is being learned.

Gardner (1957) showed that adult learners regularized their output with the most frequent light pattern presented in the input when the input was complicated. During the experiment, adults saw that one light was lit at a time. The number of lights being lit and the proportion of the time that each light was lit differed by conditions. When two lights were lit during the exposure (e.g., light A flashed 70% of the time and light B flashed 30% of the time), learners matched the proportion of each light being lit in their responses, showing frequency-matching behaviour; however, when three lights were lit (e.g., light A flashed 70% of the time, light B and C each flashed 15% of the time), learners started to overly use the most frequent light (light A) and regularized their responses.

In addition to regularizing non-linguistic probabilistic patterns, children are more likely to regularize their output than adults in language learning (Ross & Newport, 1996; Ross, 2001; Singleton & Newport, 2004; Hudson Kam & Newport, 2005, 2009). It has been shown in previous studies that deaf children often produced more systematic sign languages by regularizing their output when they were exposed to inconsistent sign languages from their previous generation, (Nicaraguan sign language: Senghas, Coppola, Newport, & Supasilla, 1997; Senghas & Coppola 2001, American sign language: Ross & Newport, 1996; Ross, 2001; Singleton & Newport, 2004). For example, Singleton and Newport (2004) studied a deaf child, Simon (age 7), who learned a sign language from his parents, who were themselves non-native speakers of the sign language (as they learned it in their late teens). The parents' sign language showed inconsistent use of morphemes and grammar; however, Simon's sign language was more consistent and systematic by using the frequent variant even more often in the language (e.g., the movement morphemes were used approximately 70% of the time in the parents' speech whereas they were used 88% of the time in Simon's speech).

Previous artificial language learning experiments have also shown that children tend to regularize their output. Hudson Kam and Newport (2005) showed that, when learning probabilistic noun-determiner patterns, children generally regularized their output by always using or omitting the determiners in their output. During the experiment, children were exposed to nonce sentences (V-S-

O word order) where determiners probabilistically occurred with nouns (e.g., in a noun phrase, /rʊŋmawt po/ 'bowling-ball DETERMINER', the determiner /po/ was presented 45%, 60%, 75%, and 100% of times in the input conditions). Some responses from children showed that they even generated their own system in their output (that was not presented in the input), which made their outputs more deterministic. Adults, however, matched the input frequency to their output.

In the follow-up study, Hudson Kam and Newport (2009) showed that adults also regularized their output when the input was more complicated. During the experiment, Hudson Kam and Newport (2009) replicated the experiments with adults and children with a modification. Learners were exposed to the input where the main determiner was presented with nouns 60% of the time (e.g., in a noun phrase /mæwznɛr ka/ 'boat DETERMINER', the determiner /ka/ was presented 60% of the time in the input). This time, they were introduced to additional filler determiners in the remaining 40%. Depending on the input conditions, a number of additional filler determiners (2, 4, 8, and 16 additional filler determiners) were presented with each noun in the remaining 40% of the input. The main and filler determiners were inconsistently and probabilistically presented in the input, meaning that the determiners were evenly distributed across lexical items and syntactic positions (e.g., intransitive subjects, transitive subjects, and transitive objects). For instance, in a condition with 16 filler determiners, participants heard the main determiner (e.g., /ka/) 60% of the time across nouns and, in the remaining 40% of the time, they heard 16 additional filler determiners (e.g., /bɪp/, /fu/, /zlæ/…) with nouns, indicating that each filler determiner was inconsistently presented 2.5% of the time in the input.

The results showed that children used the main determiner in the output more frequently than what they had seen in the input, making their output more systematic. Adults also regularized their output with the main determiner when the input had filler determiners compared to when the input did not. Especially, when there were 16 filler determiners in the input, adults regularized their output by overly using the main determiner 90% of the time in their output (the input frequency of the main determiner was 60% of the time). However, when the filler determiners were presented with specific nouns (hence presented in consistent patterns) in the second experiment, adults did not regularize their output even when presented with 16 filler determiners. This result suggests that adults can conditionally regularize their output by excessively using the more frequent pattern in the input to their output. Namely, several elements such as the number, the relative frequency, and the consistency of variants influence whether or not adults regularize their output.

Another study showed that learners regularize their output when they disfavour one of the variants. Culbertson, Smolensky, and Legendre (2012) examined existing typological asymmetries of noun and modifier orders. Across languages, patterns with consistent (regular) noun and modifier orders (Orders 1 and 2 in Table 1) are more frequently observed than inconsistent orders (Orders 3 and 4). Among the inconsistent orders, Order 4 is the least common across languages, raising the possibility that learners might be biased against Order 4 (Greenberg, 1963; Hawkins, 1983; Rijkhoff, 1998; Hurford, 2003).

Culbertson et al. conducted an experiment using a mixture-shift paradigm to examine whether participants' learning results were parallel to the typological ranking of the four orders. During the experiment, adults were exposed to probabilistic word order (the majority pattern was presented in 70% of the trials, and the minority pattern was presented in 30% of trials). Depending on the condition, the majority pattern matched one of the orders in Table 1 (e.g., adjectives and numerals precede nouns 70% of the time and they follow nouns 30% of the time). The results showed that adult learners generally regularized their output by boosting the frequency of the majority pattern in their output; however, when Order 4 was the majority pattern, they did not regularize their output. Instead, they used the majority and minority patterns similar to what they had seen in the input (70% and 30% of the time). This result implies that learners do not regularize their output with the frequent variant when they are biased against it, suggesting the effect of biases in learning probabilistic patterns.

Table 1. Four patterns of ordering determiners (adjectives and numerals) and nouns across languages. (Cited from Culbertson et al., 2012, p. 309)

Order 1. Adjective-Noun & Numeral-Noun
Order 2. Noun-Adjective & Noun-Numeral
Order 3. Noun-Adjective & Numeral-Noun
Order 4. Adjective-Noun & Noun-Numeral

In the current study, I adopt the mixture-shift experimental paradigm used by Culbertson et al. (2012) and test the interactive effect of the frequency of variants in the input and learners' biases against phonological alternation and neutralization in morpho-phonological learning. During the experiments, learners are exposed to two prefix forms that are presented in probabilistic patterns (a frequent form appeared 66.7% of the time and an infrequent form appeared 33.3% of the time in the input). Additionally, one of the prefix forms triggers alternations of stems. I further examine whether learners' choice of the prefix form that triggers alternations and the learnability of the alternations are affected by the relative frequency of the prefix form in the input. In the next section, I provide an overview of the experiments, which is followed by detailed experimental methods.

2.5 Experiment

2.5.1 Experiment Overview

During the experiment, adult participants were exposed to an artificial language and asked to learn 10 singular stems, a phonological rule, and plural forms which were probabilistically marked by one of two prefixes. The experiment consisted of four phases: stem learning, stem test, prefix learning, and prefix test. In the stem learning phase, participants learned 10 singular nouns

(CVCV). They were then tested on how well they remembered the nouns that they had learned. During the prefix learning phase, participants learned the plural forms of the nouns, which were formed by adding either of two prefixes, *ni-* or *ba-* (ni-CVCV or ba-CVCV). Participants were exposed to both prefix forms for each stem, but the proportion of times that each plural form appeared in the input was manipulated. This variable (Prefix Frequency) had three levels: a Frequent *ba-*, a Frequent *ni-*, and a 50-50 condition. In the Frequent *ba-* condition, the *ba-*stem was presented in two-thirds of trials, and the *ni-*stem was presented in one-third of trials; in the frequent *ni-* condition, the proportion of the prefix forms was reversed; in the 50-50 condition, both prefix forms were presented equally (50% of the trials each).

Furthermore, in the prefix learning phase, participants learned a phonological alternation triggered by the *ni-* prefix. The types of alternation were manipulated (Alternation Type). In the No Alternation condition, the *ni-* prefix did not trigger alternation (e.g., singular: [kimu], plural: [nikimu]). This condition was used as the baseline to examine learners' choice of prefix form without the influence of phonological alternation. In the Non-Neutralizing condition, the *ni-* prefix triggered the palatalization of velar-initial stems (e.g., singular: [kimu], plural: [niʧimu]). In this condition, I intended to test whether learners' choice of the *ni-* prefix was affected by the fact that it triggered phonological alternation. Finally, in the Neutralizing condition, palatalization triggered by *ni-* was neutralizing and resulted in homophony half of the time. Two singular stems were changed to make the palatalizing alternation phonologically neutralizing and homophony-creating (singular words: [ʧimu] and [kimu], plural form: [niʧimu]). This condition allowed me to test whether the fact that the phonological alternation was neutralizing and homophony-creating affected learners' acquisition of the rule and their choice of the prefix forms. Finally, in the prefix test phase, participants were asked to verbally produce the plural forms to test what they had learned. The experiment was a 3 (Prefix Frequency) x 3 (Alternation Type) between-subjects design, and participants were assigned to one of the nine conditions.

2.5.2 Method

2.5.2.1 Participants

Two hundred and twenty-five adult native English speakers (162 females, 63 males; mean age = 24, age range = 18-58) completed this experiment, and each participant was assigned to one of the nine conditions.[9] An additional 50 participants were recruited but excluded from the analysis because they failed to reach the 75% accuracy criterion in the stem test phase or did not complete the experiment. Most participants were Southern British English speakers, but some participants spoke another variety of English (e.g., North American English, Singapore English, Malaysian English, etc.). The experiment was conducted at University College London. Participants were recruited using the UCL Psychology Subject Pool, and they received credit or monetary compensation.

2.5.2.2 Materials

For the stem learning and stem test phases, 10 CVCV nonce words were used for singular noun stems. Each singular noun stem was presented with a singular object (e.g., a strawberry, a bucket, etc.; see Figure 1). The singular noun stems were identical in the No Alternation and Non-Neutralizing conditions. In the two conditions, the consonants for singular noun stems were chosen from 8 consonants that exist in both English and Korean (/p, b, t, d, k, g, m, n/). Crucially, in the Neutralizing condition, 2 consonants, /t/ and /d/, were correspondingly exchanged for /ʧ/ and /ʤ/, meaning that, in the Neutralizing condition, the consonants for singular noun stems were chosen from the following 8 consonants: (/p, b, **ʧ**, **ʤ**, k, g, m, n/). The production of voiceless stops was aspirated ([p^h, t^h, $ʧ^h$, k^h]) in all positions to help English speakers distinguish them from the

[9] Each condition had 25 participants except for the Frequent *ni-* Non-Neutralizing condition (26 participants) and the 50-50 Non-Neutralizing condition (24 participants).

corresponding voiced stops. The six consonants, /p, b, t, d, m, n/, in the No Alternation and Non-Neutralizing conditions or those, /p, b, ʧ, ʤ, m, n/, in the Neutralizing condition, were placed in the stem-initial position ($\underline{\mathbf{C_1}}V_1C_2V_2$) once. The critical velar stops, /k, g/, were targeted by the palatalization rule in the Non-Neutralizing and Neutralizing conditions and were placed stem-initially twice. Thus, there were six singular noun stems each starting with one of /p, b, t, d, m, n/ in the No Alternation and Non-Neutralizing conditions, and six singular noun stems each starting with one of /p, b, ʧ, ʤ, m, n/ in the Neutralizing condition. In all conditions, there were two stems starting with /k/, and two stems starting with /g/. The consonants were used a similar number of times in the stem-medial position ($C_1V_1\underline{\mathbf{C_2}}V_2$).

For vowels, I used /i, a, u/ as these vowels are cross-linguistically common, very distinct from one another and exist in both English and Korean. Each vowel was used a similar number of times in the V_1 and V_2 positions. The velar stops (/k, g/) did not appear in the C_2 position when preceded by the vowel /i/. Stress was always placed on the first syllable of stems. Words that sound similar to existing English and Korean words were avoided.

Figure 1. Examples of a picture of a singular object presented in the stem learning and test phases, and a picture of multiple objects presented in the prefix learning and test phases.

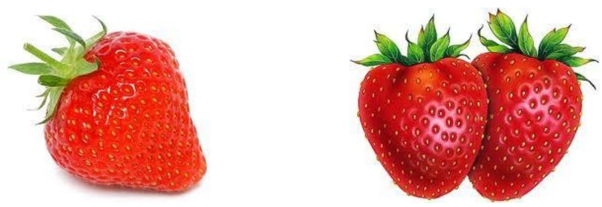

niba (singular stem)   baniba/niniba (plural form)

For the prefix learning phase, I took the 10 singular noun stems that participants had already learned in the stem learning phase and created plural forms by adding a prefix, *ba-* or *ni-*, to each stem (*ba*-CVCV or *ni*-CVCV). In the No Alternation condition, none of the prefixes triggered the alternation of stems (e.g., singular: [kimu], plural: [nikimu]). In the Non-Neutralizing and Neutralizing conditions, the *ni-* prefix triggered the palatalization of velar-initial stems (e.g., singular: [kimu], plural: [niʧimu]; see (2)).[10] Each plural form appeared with a picture of multiple objects (see Figure 1). I chose a prefix instead of a suffix to present progressive palatalization, which is less common than regressive palatalization across languages (Bateman, 2007), so that the palatalization rule was less familiar to adult native English speakers and the learning of the rule was less affected by learners' native language. English has a few regressive palatalizations (k→s/__i: *electric–electricity*, t→ʃ/__i: *note–notion*), but velar consonants are not palatalized when preceded by /i-/.

(2) Palatalization of velar-initial stems in the Non-Neutralizing and the Neutralizing conditions.

/k/→[ʧ] /i___
/g/→[ʤ] /i___

In the Neutralizing condition, I changed the two singular noun stems so that the palatalization that was triggered by *ni-* became neutralizing and homophony-creating. See Table 2 for examples of crucial singular noun stems and their plural forms in each Alternation Type. Two alveolar-initial stems were exchanged for palatal-initial stems: [timu] and [dapi] were correspondingly exchanged for [ʧimu] and [ʤapi] in the Neutralizing condition. The four stems ([kimu], [gapi], [kuta], and [gaku]) were palatalized when preceded by *ni-*, and two of them ([kimu], [gapi]) became homophonous with the distinct stems ([ʧimu] and [ʤapi]) in ni+stem

[10] During the experiment, the alternation-triggering prefix was not counterbalanced between *ni-* and *ba-*. Instead, the palatalization was only triggered by *ni-* (but never by *ba-*) because palatalization frequently occurs next to high front vowels across languages to ease articulatory effort (Keating & Lahiri, 1993), but not next to low vowels. Due to this point, counterbalancing which prefix triggered the palatalizing rule would introduce the additional factor of whether the rule was phonetically natural or not. Moreover, the results of the No Alternation condition suggest that learners did not have an inherent preference for the *ni-* or *ba-* prefix (see section 2.6.2).

plural forms (e.g., singular: [kimu], [ʧimu]; ni+stem plural forms: [niʧimu]; singular: [gapi], [ʤapi]; ni+stem plural forms: [niʤapi]). For the other two velar-initial stems ([kuta] and [gaku]), the *ni-* prefix triggered phonological neutralization (distinct sets of phonemes /k, ʧ/ and /g, ʤ/ respectively merge to [ʧ] and [ʤ]) but did not create homophony. Except for the two stems, the rest of the stems in the Neutralizing condition were identical to the other two conditions.

The stimuli used in this experiment were recorded in a soundproof booth by a phonetically trained female native Cypriot-Greek speaker so that the stimuli would be distinct from both of the languages spoken by the participants (i.e., English and Korean). The RODE-NT1-A large diaphragm condenser microphone and the audio interface recorder (RME Finface UC) were used for the recording. The sampling rate was 44,100 Hz and 16 bits. The stimuli were rescaled for intensity and the volume was adjusted to a comfortable listening volume.

Table 2. Examples of crucial singular noun stems and their plural forms by the Alternation Types.

Condition	Stem (Singular)	Prefix-Stem (Plural)	
No Alternation Condition	ˈkimu	baˈkimu	niˈkimu
	ˈtimu	baˈtimu	niˈtimu
	ˈgapi	baˈgapi	niˈgapi
	ˈdapi	baˈdapi	niˈdapi
	ˈkuta	baˈkuta	niˈkuta
	ˈgaku	baˈgaku	niˈgaku
Non-Neutralizing Condition	ˈkimu	baˈkimu	niˈʧimu
	ˈtimu	baˈtimu	niˈtimu
	ˈgapi	baˈgapi	niˈʤapi
	ˈdapi	baˈdapi	niˈdapi
	ˈkuta	baˈkuta	niˈʧuta
	ˈgaku	baˈgaku	niˈʤaku
Neutralizing Condition	ˈkimu	baˈkimu	niˈ**ʧ**imu
	ˈʧimu	baˈʧimu	
	ˈgapi	baˈgapi	niˈ**ʤ**api
	ˈʤapi	baˈʤapi	
	ˈkuta	baˈkuta	niˈʧuta
	ˈgaku	baˈgaku	niˈʤaku

2.5.2.3 Procedure

The experiment consisted of four phases: a stem learning, a stem test, a prefix learning, and a prefix test phase. The stem learning and test phases were designed to introduce singular words, and the prefix learning and test phases were designed to introduce plural forms of the words to participants. The experiment was coded in Python using Pygame by James White. Before starting the experiment, participants were informed that their goal was to fit into an alien society by successfully communicating with other aliens using their language; however, participants were unaware of the purpose of this experiment. They were also told to remember the words and matching pictures as accurately as they could by paying close attention to the picture that appeared on the computer screen and the sound that they heard from an alien instructor. The alien instructor appeared as a drawing of an alien on the right side of the screen. When the alien instructor said a word, a bubble appeared on the screen, and then participants heard the sound files. For the stem test and prefix test phases, I used a production test: participants were asked to verbally produce the answer into a microphone.

Phase 1: Stem learning

During the stem learning phase, participants were introduced to a set of 10 singular CVCV nonce words (noun stems). Participants were exposed to 100 self-paced trials divided into 10 blocks. Ten singular stems were randomly allocated in each block, and each block was repeated 10 times (10 stems x 10 blocks), meaning that each stem occurred once within each block in a random order. Each trial began with an alien instructor on the right side of the screen,[11] and then the picture of a single item appeared on the left side of the screen. Once the picture of a single item appeared, a

---

[11] I appreciate Van de Vijver and Baer-Henney (2014) for sharing the image of the alien for this study.

bubble appeared above the alien instructor (to indicate the alien is speaking), and participants heard the matching word through headphones. Once participants heard the stem, they were asked to repeat what they heard out loud. They could move to the next trial by pressing the space bar. The stems were only presented with audio; hence, no orthography was provided at any point in the experiment.

Phase 2: Stem test

Once participants finished the stem learning phase, they were tested on how well they remembered the stems that they had just learned. They were instructed to orally produce the word into a microphone once they saw the picture of a single item on the screen. They were also informed that they would get feedback (to indicate the correct answer) from the alien instructor after they produced an answer. When their answer matched that of the alien instructor without an error, they gained 5 points, but if the participants made an error in their answer, they did not get any points. Furthermore, an experimenter informed them that they needed at least 190 points (75% of trials) to pass the stem test phase, and they could see their score at the bottom of the screen. The stem test phase had 50 trials, and the items presented during stem test were identical to those in the stem learning phase, meaning that there were no novel items. At first, participants saw the alien instructor on the right side of the screen, and then the picture of a single item appeared on the left side of the screen. Participants produced the answer into the microphone, and the answer was recorded. Once they produced an answer, an experimenter outside of the recording room (who monitored the experiment over the headphones and the screen) coded it as correct or incorrect. After the experimenter coded the answer, a bubble appeared above the alien instructor, and the participants heard feedback (the correct form) and saw the score changed at the bottom of the screen (if their response was correct). Figure 2 shows the overview of the process in the stem test and prefix test phases. Participants who achieved at least 75% accuracy (190 points) moved to the prefix learning phase, but participants who did not reach the accuracy criterion repeated the same test one more time; however, if they failed the second time, the experiment terminated, and they were excluded from the analysis.

Figure 2. Overview of the procedure of each trial for the stem and prefix test phases.

 Current Score: 10	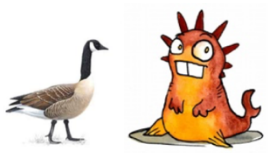 Current Score: 10		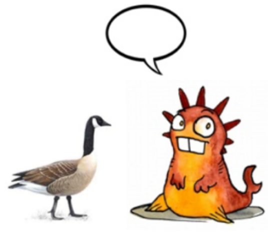 Current Score: 15
• Alien appears.	• Picture of object appears.	• Participant produces response. • Response recorded. • Experimenter codes response.	• Participant hears feedback. • Score updated.

Phase 3: Prefix learning

The prefix learning phase introduced the plural forms of the singular noun stems that participants had already learned in the stem learning phase. The plural forms were made by adding either of the prefixes *ni-* or *ba-* to the stems. Participants were exposed to 60 self-paced trials in which 10 plural words were randomly allocated into each block, which was repeated 6 times (10 plural forms x 6). Before starting the prefix learning phase, participants were told that there was more than one way to form plural words in this alien language. Participants were exposed to both ba+stem and ni+stem plural forms of each stem during the prefix learning phase; however, the frequency of each plural form presented in the input differed depending on the prefix frequency condition (see Table 3).

Table 3. Frequency of plural forms in the prefix learning and test phases.

<table>
<tr><th></th><th>Frequent ba- condition</th><th>Frequent ni- condition</th><th>50-50 condition</th></tr>
<tr><td>Frequent prefix form</td><td>ba+stem
40 trials
(66.7%)
(4 per stem)</td><td>ni+stem
40 trials
(66.7%)
(4 per stem)</td><td rowspan="2">ba+stem
30 trials
(50%)
ni+stem
30 trials
(50%)
(3 per stem)</td></tr>
<tr><td>Infrequent prefix form</td><td>ni+stem
20 trials
(33.3%)
(2 per stem)</td><td>ba+stem
20 trials
(33.3%)
(2 per stem)</td></tr>
</table>

Table 4. Number of trails showing the stem alternation in the prefix learning and test phase by conditions. Note: half of the trials in the Neutralizing condition triggered homophony.

<table>
<tr><th></th><th>Frequent ba- condition</th><th>Frequent ni- condition</th><th>50-50 condition</th></tr>
<tr><td>No Alternation</td><td>0 /60 trials (0%)</td><td>0 /60 trials (0%)</td><td>0 /60 trials (0%)</td></tr>
<tr><td>Non-Neutralizing</td><td rowspan="2">8/60 trials (13%)</td><td rowspan="2">16/60 trials (27%)</td><td rowspan="2">12/60 trials (20%)</td></tr>
<tr><td>Neutralizing*</td></tr>
</table>

In the Frequent *ba-* condition, ba+stem was the frequent pattern, occurring in two-thirds (66.7%) of trials, whereas ni+stem was the infrequent pattern, occurring in one-third (33.3%) of trials. For example, during the prefix learning phase, participants assigned to the Frequent *ba-* Neutralizing condition heard [bakuta] four times and [niʧuta] two times for the plural form of [kuta] ‘a cat’. In contrast, in the Frequent *ni-* condition, ni+stem was the frequent pattern (occurring in 66.7% of the trials) whereas ba+stem was an infrequent pattern (occurring in 33.3% of the trials). In the 50-50 control condition, both ba+stem and ni+stem were presented equally in one-half (50%) of trials. The number of trials demonstrating the palatalization of the velar-initial stems in each condition can be seen in Table 4. The stimuli of the prefix test phase were identical to those of the prefix learning phase.

Phase 4: Prefix test

In the prefix test phase, learners were tested on the same plural forms that they had already learned in the prefix learning phase. They were exposed to 60 trials in total. The procedure of the prefix test phase was identical to that of the stem test phase; however, the coding and scoring systems were different. Once participants produced an answer, the experimenter coded it depending on which prefix was used, whether the stem was correct or incorrect, whether or not the initial stem consonant was correctly altered, and whether the stem consonant was alternated with some other sound or not changed. This more elaborate coding system ensured that all aspects of the responses could be analysed at a later point.

Once the experimenter coded the response, the speech bubble appeared, and participants heard one of the two plural forms (ba+stem or ni+stem) as the feedback from the alien instructor, and saw their score changed. The proportion of ni+stem and ba+stem presented during the prefix test phase matched with the proportion that participants were exposed to in the prefix learning phase. For instance, if a participant was assigned to the Frequent *ba-* condition, the participant heard ba+stem 40 times and ni+stem 20 times in both the prefix learning and prefix test phases. Next, participants saw their scores changed at the bottom of the screen. Participants gained 10 points if they produced the correct plural form that matched what the alien instructor produced (e.g., the participant said [niʤapi], and the alien also said [niʤapi]). They received 5 points if the stem form in their response matched that of the alien, but the prefix form differed from that of the alien (e.g., the participant said [niʤapi], and the alien said [baʤapi]). Otherwise, they did not receive any points (e.g., the participant said [bagapi], and the alien said [niʤapi]).[12] If the

[12] The scoring system had a coding error that was only found in the later analysis process. The correct (intended) scoring system is the following: when learners' responses were correct and matched with the alien's, they got 10 points (e.g., participant's response: [niʧimu], alien's response: [niʧimu]), and when learners' responses were phonologically correct but did not match the alien's (i.e., they picked a different prefix), they got 5 points (e.g., participant's response: [bakimu], alien's response: [niʧimu]). In other cases, they did not gain any points. The potential effects of the erroneous scoring system on the experimental results can be found in section 2.7.6.

participant did not respond within 10 seconds, the trial was coded as an error, participants heard the correct form from the alien, and automatically moved to the next trial.

Following Culbertson et al. (2012), I provided the score in the prefix test phase, even though there was no score threshold in this phase, to encourage participants to answer probabilistically. Since both prefix forms were correct answers, participants could potentially use only one prefix and ignore the other in the prefix test phase. However, to obtain the highest score, the best strategy was to match the frequency distribution of the prefix forms produced by the alien in the prefix test phase. Once participants completed the prefix test phase, I debriefed them by asking them questions about their thought processes during the experiment and their response strategies in the prefix test phase.

2.6 Analysis and results

2.6.1 Analysis plan and predictions

I first tested the hypothesis that learners avoid using the morphological construction that triggers phonological alternation due to their bias against phonological alternation. To test this hypothesis, I examined how often participants chose the frequent and infrequent prefix forms in their responses during the prefix test phase. I predicted that participants would choose to use the alternation-triggering prefix (*ni-*) less often than what they had seen in the input due to their preference for a unified paradigm. They were also more likely to shift away from the *ni-* prefix when it was infrequent in the input.

Additionally, I tested the hypothesis that learners find it more difficult to learn neutralizing alternations that create homophony compared to non-neutralizing alternations. I compared how often participants applied the non-neutralizing and neutralizing alternations after *ni-* (correct cases) and *ba-* (over-generalization error cases). I predicted that learners would be poorer at learning the neutralizing alternation compared to the non-neutralizing alternation due to their tendency to avoid homophony. This means that they would apply the neutralizing alternation less often than the non-neutralizing alternation after *ni-* and apply the neutralizing alternation more often than the non-neutralizing alternation after *ba-* (e.g., make more over-generalization errors).

2.6.2 Selection of frequent and infrequent prefix forms

I first compared the proportion of the time that participants used the frequent prefix form in their output during the prefix test phase. Table 5 and Figure 3 show the percentage of the time that participants used the frequent variant in their output according to the Prefix Frequency (Frequent *ba-*, Frequent *ni-*, and 50-50) and Alternation Type (No Alternation, Non-Neutralizing, and Neutralizing). In the Frequent *ba-* and Frequent *ni-* conditions, the frequent variant was ba+stem and ni+stem, respectively; however, there was no frequent variant in the 50-50 condition, and ba+stem was chosen arbitrarily to be coded as the frequent variant. Participants' responses with the wrong stem and prefix (neither *ba-* or *ni-*) as well as those that were not produced within the 10-second time limit were excluded from the analysis.

Table 5. Mean percentage of times in which participants chose the frequent prefix option in the prefix test phase.

Prefix Frequency	Alternation Type		
	No Alternation	Non-Neutralizing	Neutralizing
50-50	47%	51%	48%
Frequent *ba-*	59%	76%	67%
Frequent *ni-*	64%	66%	66%

Figure 3. Proportion of trials in which participants chose the frequent prefix option in the prefix test phase. Note: Dashed lines show the relative frequency of the frequent prefix form presented in the prefix learning phase. Error bars show 95% confidence intervals.

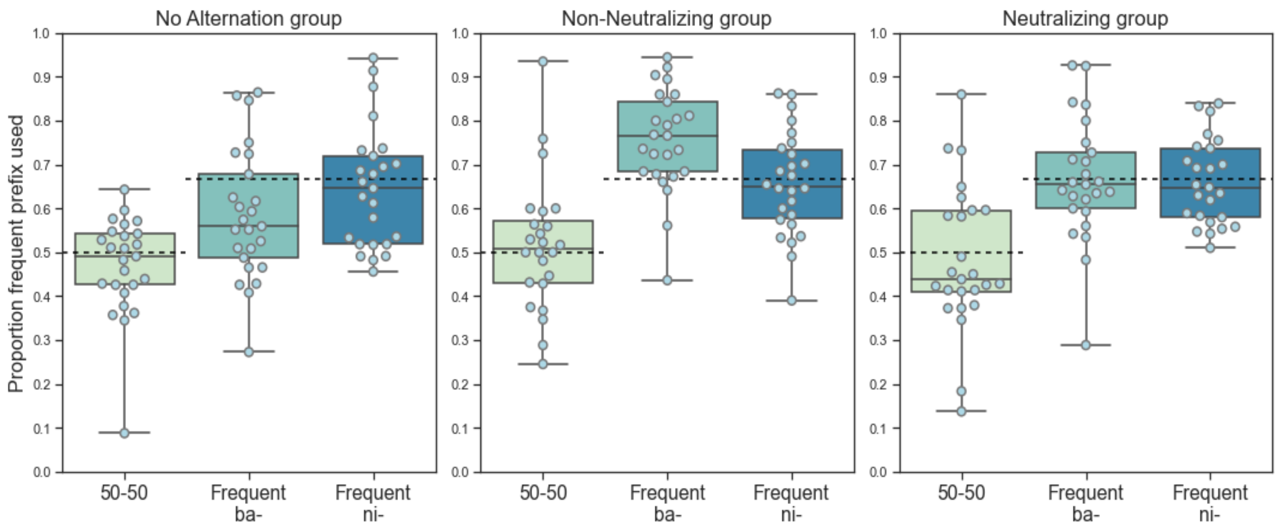

The aggregate results in Table 5 and Figure 3 show that participants generally matched the frequency of the frequent variant presented in the input to their output. The dashed line shows the percentage of the time that the frequent prefix form was presented in the input. The percentage of the time that participants used the frequent variant (arbitrarily chosen to be ba+stem) in the 50-50 condition was close to 50% in all Alternation Types, which matched the frequency presented in the input. In the Frequent *ba-* and Frequent *ni-* conditions, the proportion of the time that the frequent variant was used in learners' output generally matched the proportion of the time in the input (66.7%) except for one condition. In the Frequent *ba-* Non-Neutralizing condition, the proportion of the time that learners chose the frequent variant in their output was higher (75.7%) than what they had seen in the input (66.7%).

The results were analysed using a mixed effects logistic regression model (Jaeger, 2008) using *lme4* packages (Bates et al., 2015) implemented in R (R Core Team, 2018). The model had fixed effects of Prefix Frequency (Frequent *ba-*, Frequent *ni-*, and 50-50), Alternation Type (No Alternation, Non-Neutralizing, and Neutralizing). The model also included the interaction of the fixed effects and had a random intercept of Stem and Subject (random slopes were not warranted in the model; Barr, Levy, Scheepers, & Tily, 2013). First, I conducted a likelihood ratio test with a backward stepwise comparison using the *anova()* function in R (see Barr et al., 2013) to compare the full model to a subset model with the interaction effects removed. The likelihood ratio test indicated that the interaction effects significantly improved the model's fit ($\chi^2(4) = 11.26$, p = .02); hence, the full model was reported as the final model. Table 6 shows the summary of the fixed effects of the final model.

Table 6. Summary of the fixed effects in the final model predicting selection of the frequent prefix for English speakers.

Fixed effect	**Estimate**	**Standard error**	**Wald *z***	***p*-value**
Intercept	–0.14	0.12	–1.17	.24
Non-Neutralizing	0.21	0.17	1.19	.24
Neutralizing	0.07	0.17	0.42	.68
Frequent *ba-*	**0.51**	**0.17**	**2.99**	**< .01****
Frequent *ni-*	**0.79**	**0.17**	**4.57**	**< .001*****
Non-Neutralizing & Frequent *ba-*	**0.63**	**0.25**	**2.55**	**.01***
Non-Neutralizing & Frequent *ni-*	–0.16	0.24	–0.65	.52
Neutralizing & Frequent *ba-*	0.30	0.24	1.23	.22
Neutralizing & Frequent *ni-*	–0.02	0.24	–0.09	.93

The reference group (intercept) is the 50-50 No Alternation condition, **R code for the final model:** glmer(UsedFrequentPrefix ~ PrefixFrequency*AlternationType + (1|Subject) + (1|Stem), data=data, family=binomial), **R code for the full model:** glmer(UsedFrequentPrefix ~ PrefixFrequency*AlternationType + (1|Subject) + (1|Stem), data=data, family=binomial).

The intercept (the 50-50 No Alternation condition) of the model and the simple effect Alternation Type (Non-Neutralizing and Neutralizing) as not significant. The fixed effect of Prefix Frequency (Frequent *ba-* and Frequent *ni-*) was positive and significant, indicating that when the input had a relatively more frequent variant, participants used the frequent prefix more than 50% of the time in their output. Furthermore, the interaction effects of Prefix Frequency (Frequent *ba-*) and Alternation Type (Non-Neutralizing) reached significance, suggesting that when *ba-* was frequent in the input, participants used the frequent variant significantly more often in the Non-Neutralizing condition than other conditions. Specifically, they shifted away from the *ni-* prefix that triggered a non-neutralizing alternation when it was infrequent in the input. No other interaction effects were significant in the model, suggesting that the shift away from *ni-* was limited to the Frequent *ba-* Non-Neutralizing condition.

In addition, I examined the percentage of the time that learners chose *ni-* with velar-initial and non-velar-initial stems to test if learners' choice of prefix depended on whether or not the stem could alternate. The *ni-* prefix triggered the palatalization of velar-initial stems whereas the prefix never triggered the palatalization of non-velar-initial stems.

Table 7. Mean usage of *ni-* for velar-initial and non-velar-initial stems, by condition.

Alternation Type	Stem-initial consonant	Prefix Frequency		
		50-50	Frequent *ba-*	Frequent *ni-*
No Alternation	Velar	50%	43%	62%
	Non-velar	55%	40%	66%
Non-Neutralizing	Velar	50%	25%	67%
	Non-velar	48%	24%	66%
Neutralizing	Velar	49%	30%	61%
	Non-velar	53%	35%	70%

Table 7 shows the percentage of time that participants chose each prefix form for the velar-initial and non-velar-initial stems according to the Prefix Frequency and Alternation Type. The aggregate results indicate that participants chose to use the *ni-* prefix with velar-initial and non-velar-initial stems a similar amount of time in their output. A mixed effects logistic regression was implemented to test the proportion of the time that participants chose to use the *ni-* prefix with velar-initial and non-velar-initial stems. The model had the fixed effect of Stem Type (velar-initial and non-velar-initial stems), the random intercepts for Stem and Subject, and the random slope of Stem Type by Subject.[13] The model only included the Non-Neutralizing and Neutralizing conditions because none of the stems were palatalized in the No Alternation condition. The intercept of the model was non-velar-initial stems. The fixed effect of Stem Type (velar-initial stems) did not significantly improve the model ($\chi^2(1) = 0.27$, $p = .60$), indicating that the percentage of the time that participants used the *ni-* prefix was comparable across the Stem Type. This finding suggests that although the *ni-* prefix only triggered the palatalization of the velar-initial stems, learners' choice of prefix seems independent of which stem followed.

2.6.3 Frequency of variants in the input vs. the output

In addition to analysing the aggregate mean results, I turned my attention to individual participants to see whether they regularized their output with the *ba-* prefix. If there was a tendency to shift away from *ni-* that triggered alternation, learners would have over-used *ba-* in their output more frequently than what they had seen in the input. Following Culbertson et al., I conducted a one-sample sign test for each condition to evaluate the number of participants who shifted away from *ni-* (and over-used *ba-* instead; see Table 8). Bonferroni-adjusted alpha level of .0056 (.05/9) was used to correct for multiple comparisons.

---

[13] The full model: glmer(UsedNi~Stem Type+(1+StemType|Subject)+(1|Stem), data=AlternationData, family=binomial)

Each group had 25 participants,[14] and if there was no tendency to shift away from the *ni*- prefix, approximately half of the group (12.5) was expected to regularize their output. The most interesting result was found in the Frequent *ba*- Non-Neutralizing condition, where 21 out of 25 participants used the frequent variant more frequently in their output compared to the input ($p < .0056$). In other conditions, participants did not show any tendency to shift away from the *ni*- prefix, and they instead showed probability-matching behaviour. Interestingly, they did not shift away from *ni*- when it triggered a neutralizing alternation that could create homophony in the Neutralizing condition. The result of the one-sample sign test parallels the results of the mixed effects logistic regression models presented in the previous section; participants mostly showed frequency-matching behaviour and they regularized their output with the frequent variant only in the Frequent *ba*- Non-Neutralizing condition.

To summarize, adult English speakers generally matched the frequency of variants presented in the input to their output. When there was a frequent variant in the input, learners successfully chose to use the frequent variant more frequently than chance level (50%) in their output. However, when *ni*- was infrequent and triggered a non-neutralizing alternation, learners shifted away from it and regularized their output with *ba*-. When *ni*- triggered a neutralizing alternation, they did not shift away from *ni*- even when it was infrequent. The aggregate results were supported by the individual results showing that individual participants did not shift away from the *ni*- prefix except when it was infrequent and triggered the non-neutralizing alternation. Moreover, they used the alternation-triggering prefix (*ni*-) similarly often with the velar-initial stems (that alternated after *ni*-) and non-velar-initial stems (that did not alternate), indicating that the prefix choice was a global decision, made regardless of which stem followed. Thus, these results indicate that adult English speakers

[14] The Frequent *ni*- Non-Neutralizing condition had 26 participants, and the 50-50 Non-Neutralizing condition had 24 participants.

generally showed probability-matching behaviour except for one condition where they regularized their output with the frequent variant.

Table 8. Number of participants who shifted towards *ba-* and *ni-*.

Condition		Shifted towards *ba-*	Shifted towards *ni-*	*p*-value (sign test)
No Alternation	50-50	12 / 25 (48%)	13 / 25 (52%)	1
	Frequent *ba-*	7 / 25 (28%)[15]	18 / 25 (72%)	.04
	Frequent *ni-*	14 / 25 (56%)	11 / 25 (44%)	.69
Non-Neutralizing	50-50	12 / 24 (50%)	9 / 24 (38%)	1
	Frequent *ba-*	**21 / 25 (84%)**	**4 / 25(16%)**	**< .0056**
	Frequent *ni-*	14 / 26 (54%)	12 / 26 (46%)	.85
Neutralizing	50-50	9 / 25 (36%)	16 / 25 (64%)	.23
	Frequent *ba-*	10 / 25 (40%)	15 / 25 (60%)	.42
	Frequent *ni-*	14 / 25 (56%)	11 / 25 (44%)	.69

2.6.4 Application of the palatalizing rule

In addition to the proportion of times that participants chose prefix forms in the output, I examined how successfully participants applied the palatalizing rule to velar-initial stems. Recall that participants were required to learn that the *ni-* prefix triggered the palatalization of the velar-initial stem (/k/→[ʧ]/i__, /g/→[ʤ]/i__). In the Neutralizing condition, the same palatalizing rule was neutralizing (/k, ʧ/→[ʧ]/i__, /g, ʤ/→[ʤ]/i__) and homophony-creating. If learners

[15] In the Frequent *ba-* No Alternation condition, only 7 participants out of 25 used the frequent variant more frequently in their output than in the input. The condition does not reach significance, but it suggests that participants seemed to show a very slight preference for *ni-* not *ba-*, when there were no alternations.

successfully learned the palatalizing rule presented during the experiment, they should palatalize velar-initial stems a high percentage of the time after *ni-* but not palatalize velar-initial stems after *ba-*. Figure 4 presents the proportion of cases in which participants applied the palatalization to velar-initial stems after *ni-* (where the rule should apply) and *ba-* (where the application of the rule was an over-generalization error) according to the Prefix Frequency and Alternation Type (the No Alternation condition was excluded from the analysis because prefixes did not trigger alternations in the condition).

Looking at the aggregate results, participants applied the palatalizing rule to the stems more frequently when using the *ni-* prefix than when using the *ba-* prefix in all conditions. Participants also incorrectly palatalized the velar-initial stems after *ba-* more often when the *ni-* prefix was frequent in the input. When using *ni-*, participants applied the palatalization to velar-initial stems more frequently in the Non-Neutralizing condition than in the Neutralizing condition. Conversely, the proportion of the time that velar-initial stems were palatalized after *ba-* was higher in the Neutralizing condition than in the Non-Neutralizing condition.

Figure 4. Proportion of times that participants palatalized velar-initial stems when using *ni-* (left) and using *ba-* (right).

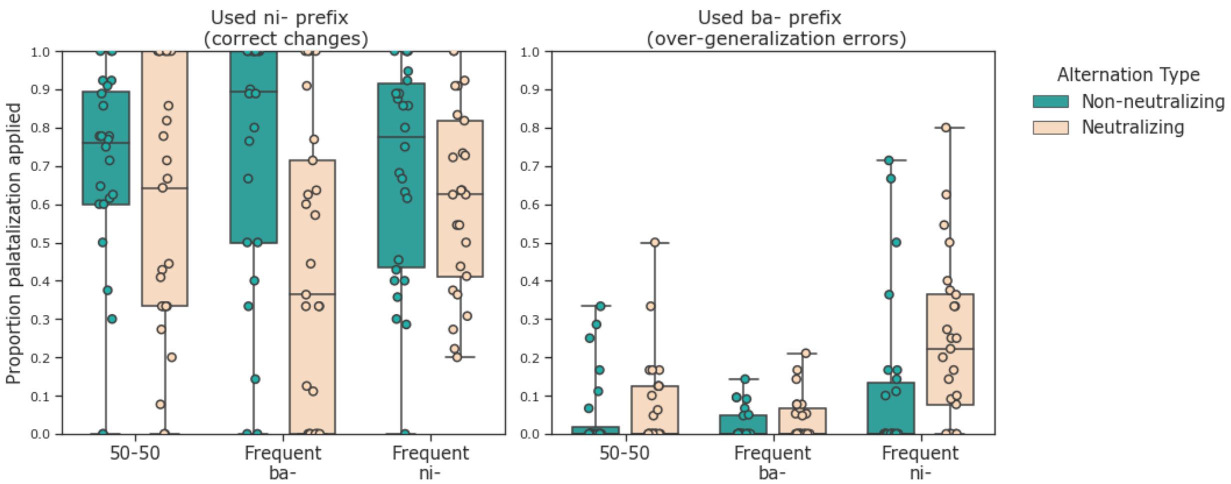

A mixed effects logistic regression model was used to test the proportion of times that palatalization was applied to stems after *ba*- and *ni*-. The model followed the general method reported in section 2.6.2. The results were analysed in two separate models. One model contained the trials in which participants palatalized velar-initial stems after the *ni*- prefix (the correct cases), and the other model included the trials in which participants palatalized velar-initial stems after the *ba*- prefix (the incorrect over-generalization cases). Both models contained the fixed effect of the Prefix Frequency (Frequent *ba*-, Frequent *ni*-, and 50-50), Alternation Type (Non-Neutralizing and Neutralizing) and their interaction effects. The model predicting the application after *ni*- had the random intercepts of Subject and Stem, and the model predicting the application after *ba*- only had the random intercepts for Subject (but not for Stem) because the model failed to converge. In each model, the backwards stepwise comparison was implemented using *anova( )* to compare a model to a subset model with one effect removed at a time.

Table 9. Summary of the fixed effects in the final model predicting palatalization after *ni*- (correct application) for English speakers.

Fixed effect	Estimate	Standard error	Wald z	p-value
Intercept	**1.30**	**0.22**	**5.80**	**< .001*** **
Neutralizing	**–0.98**	**0.30**	**–3.30**	**< .001*** **

The reference group (intercept) is the 50-50 Non-Neutralizing condition, **R code for the final model:** glmer(AppliedPalatalization ~ AlternationType + (1|Subject) + (1|Stem), data=ni-data, family=binominal), **R code for the full model:** glmer(AppliedPalatalization ~ AlternationType*PrefixFrequency + (1|Subject) + (1|Stem), data=ni-data, family=binominal)

Table 9 shows the summary of the fixed effects from the final model predicting the application of palatalization after *ni*-. A likelihood ratio test indicated that Alternation Type by Prefix Frequency interaction effects ($\chi^2(2) = 4.02$, p = .13) and Prefix Frequency (Frequent *ba*- and

Frequent *ni*-) effect ($\chi^2(2) = 1.54$, p = .46) did not significantly improve the model's fit. Hence, these effects were removed from the model. The effect of Alternation Type (Neutralizing) reached significance ($\chi^2(1) = 10.48$, p = .001) and was retained in the model. The intercept (representing the 50-50 Non-Neutralizing condition) was significant and positive. The simple effect of Alternation Type (Neutralizing) was significant and negative. These results suggest that learners frequently palatalized velar-initial stems after *ni*- but they palatalized the stems significantly less often in the Neutralizing condition compared to the Non-Neutralizing condition (aggregate results in the Non-Neutralizing condition; 50-50: 74.1%, Frequent *ba*-: 73.7%, Frequent *ni*-: 71.4% vs. in the Neutralizing condition; 50-50: 58.6%, Frequent *ba*-: 45.5%, Frequent *ni*-: 62%). Overall, adult English speakers palatalized velar-initial stems after *ni*- less successfully when the *ni*- prefix triggered the neutralizing alternation than when it triggered the non-neutralizing alternation.

Table 10. Summary of the fixed effects in the final model predicting palatalization after *ba*- (over-generalization error) for English speakers.

Fixed effect	Estimate	Standard error	Wald *z*	*p*-value
Intercept	**–3.74**	**0.36**	**–10.29**	**< .001*****
Neutralizing	**0.90**	**0.30**	**2.96**	**< .01****
Frequent *ba*-	–0.72	0.39	–1.86	.06
Frequent *ni*-	**1.36**	**0.35**	**3.85**	**< .001*****

The reference group (intercept) is the 50-50 Non-Neutralizing condition, **R code for the model:** glmer(AppliedPalatalization ~ AlternationType + PrefixFrequency + (1|Subject), data=ba-data, family=binominal) **R code for the full model:** glmer(AppliedPalatalization ~ AlternationType*PrefixFrequency + (1|Subject) + (1|Stem), data=ni-data, family=binominal)

Table 10 shows the summary of the fixed effects of the final models predicting the application of palatalization after *ba*- (over-generalization error cases). A likelihood ratio test indicated that the Alternation Type by Prefix Frequency interaction effects did not significantly improve the model's fit ($\chi^2(2) = 0.71$, $p = .70$); hence, they were removed from the model. The fixed effects of Alternation Type (Neutralizing) ($\chi^2(1) = 8.84$, $p < .01$) and Prefix Frequency (Frequent *ba-* and Frequent *ni-*) ($\chi^2(2) = 31.19$, $p < .001$) significantly improved the model fit, and they remained in the model. The significant positive simple effect of Alternation Type (Neutralizing) and the lack of significant interaction effects involving Alternation Type indicate that learners made more over-generalization errors when the alternation was neutralizing compared to when it was non-neutralizing. Also, the simple effect of Prefix Frequency (Frequent *ni-*) was significant and positive, suggesting that learners made more over-generalization errors when *ni-* was frequent in the input than when it was not (aggregate results in the Non-Neutralizing condition; 50-50: 3.7%, Frequent *ba-*: 2.2%, Frequent *ni-*: 11% vs. in the Neutralizing condition; 50-50: 8.2%, Frequent *ba-*: 4%, Frequent *ni-*: 24.5%). Thus, learners made more over-generalization errors when the alternation was neutralizing than when it was non-neutralizing. They also made more errors when the *ni-* prefix was the frequent variant in the input.

Table 11. Percentage of times that participants incorrectly palatalized non-velar-initial stems after *ba-* and *ni-*.

Condition	Palatalization after *ba-* (%)	Palatalization after *ni-* (%)
50-50 Non-Neutralizing	0	0.3
Frequent *ba-* Non-Neutralizing	0	0.5
Frequent *ni-* Non-Neutralizing	0	0.4
50-50 Neutralizing	0.3	1.1
Frequent *ba-* Neutralizing	0	0
Frequent *ni-* Neutralizing	0.8	0.2

In addition, I investigated the percentage of times that learners palatalized non-velar-initial stems to examine whether learners over-generalized the palatalization to non-velar-initial stems. Recall that participants never observed non-velar-initial stems being palatalized after either *ba-* or *ni-*. Table 11 shows the percentage of times that non-velar-initial stems were palatalized after each prefix according to the conditions. The over-generalization to non-velar-initial stems occurred less than 1.2% of the time across the conditions when using either the *ba-* or *ni-* prefix, indicating that learners rarely over-generalized the palatalization to non-velar-initial stems in any of the conditions.

Finally, I examined how often learners palatalized stems when the *ni-* prefix did and did not trigger homophony to test if application of the palatalizing rule would vary depending on whether it created homophony in a specific instance. To review, in the Neutralizing condition, the *ni-* prefix triggered neutralizing alternations, and half of them created homophony while the other half did not. In the prefix learning phase, four velar-initial stems ([kimu], [gapi], [kuta], [gaku]) were palatalized when preceded by the *ni-* prefix. In the Neutralizing condition, two velar-initial stems ([kimu], [gapi]) and two palatal-initial stems ([ʧimu], [ʤapi]) became homophonous in their *ni-*stem plural forms after the palatalization (e.g., [kimu], [ʧimu] → [niʧimu]; [gapi], [ʤapi] → [niʤapi]). The *ni-* prefix also triggered palatalization of the other two velar-initial stems ([kuta], [gaku]), which merged distinct phonemes (/k, ʧ/ → [ʧ], /g, ʤ/→ [ʤ]) but these stems did not create homophony. Hence, I compared the proportion of the time that participants applied palatalization to these two groups of stems, which I refer to as Homophony stems and No Homophony stems, respectively (see Figure 5). The data from the Non-Neutralizing and No Alternation conditions were excluded from the analysis because the palatalization never triggered homophony in those conditions.

A mixed effects logistic regression model was used to test the proportion of times that Homophony and No Homophony stems were palatalized after *ba-* and *ni-*. The proportion was tested in two separate models. The first model included the trials where the Homophony and No Homophony stems were palatalized after *ni-* and the second model included the trials in which the stems were palatalized after *ba-*. Both models had a fixed effect of the Homophony (Homophony and No Homophony stem) and Prefix Frequency (Frequent *ba-*, Frequent *ni-*, and 50-50), and their interaction. The model predicting the palatalization after *ni-* had random intercepts of Subject and

Stem, and the model predicting the palatalization after *ba-* had a random intercept of Subject as the model failed to converge. The backwards stepwise comparison was implemented in each model using *anova( )* to compare the full model to a subset model.

Figure 5. Proportion of times that participants palatalized velar-initial stems that became homophony and that did not after *ni-* (left) and *ba-* (right) in the Neutralizing condition.

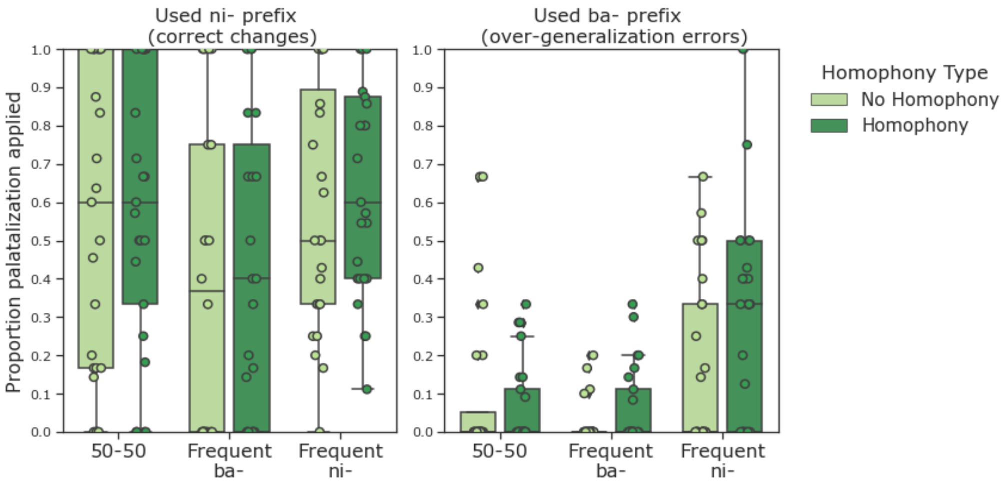

Table 12. Summary of the fixed effects in the final model predicting palatalization of Homophony and No Homophony stems after *ni-* (correct cases) for English speakers.

Fixed effect	**Estimate**	**Standard error**	**Wald *z***	***p*-value**
Intercept	0.29	0.29	0.99	.32

The reference group (intercept) is No Homophony stems in the 50-50 condition. **R code for the final model:** glmer(AppliedPalatalization ~ +(1 |Subject) + (1|Stem), data=ni_homophony, family=binominal) **R code for the full model:** glmer(AppliedPalatalization ~ +(1 |Subject) + Homophony*PrefixFrequency +(1 |Subject) + (1|Stem), data=ni_homophony, family=binominal)

Table 12 shows the summary of the fixed effects from the final model predicting palatalization of Homophony and No Homophony stems after the *ni-* prefix. The interaction effect of Homophony by Prefix Frequency ($\chi^2(2) = 0.09$, p =.95), the simple effect of Homophony ($\chi^2(1) =$

0.66, p =.41) and Prefix Frequency (Frequent *ba*- and Frequent *ni*-) ($\chi^2(2) = 5.18$, p =.07) did not improve the model fit and they were removed from the model. The intercept (No Homophony stem in the 50-50 condition) did not reach significance, suggesting that learners palatalized No Homophony stems after *ni*- near chance level in the 50-50 condition. Overall, this finding suggests that the percentage of times that learners applied the neutralizing alternation to velar-initial stems was similar across the stems that created homophony and those that did not.

Table 13. Summary of the fixed effects in the final model predicting palatalization of Homophony and No Homophony stems after *ba*- (over-generalization errors) for English speakers.

Fixed effect	**Estimate**	**Standard error**	**Wald *z***	***p*-value**
Intercept	**–3.1**	**0.37**	**–8.31**	**< .001*****
Homophony stem	**0.62**	**0.26**	**2.43**	**< .05***
Frequent *ba*-	-0.77	0.45	-1.7	.09
Frequent *ni*-	**1.51**	**0.41**	**3.72**	**<.001*****

The reference group (intercept) is No Homophony stems in the 50-50 condition. **R code for the final model:** glmer(AppliedPalatalization ~ Homophony+PrefixFrequency + (1 |Subject) , data=ba_homophony, family=binominal) **R code for the full model:** glmer(AppliedPalatalization ~ Homophony*PrefixFrequency + (1 |Subject) , data=ba_homophony, family=binominal)

Table 13 shows the summary of the fixed effects in the final model predicting the palatalization of Homophony and No Homophony stems after the *ba*- prefix. The interaction effects of the Homophony by Prefix Frequency did not significantly improve the model fit ($\chi^2(2) = 5.55$, p=.06); hence, these were removed from the model. The simple effects of Homophony ($\chi^2(1) = 6.02$, p =.01) and Prefix Frequency ($\chi^2(2) = 27.76$, p <.001) had a significant effect on the model's fit and they remained in the model. The significant and negative intercept suggests that learners palatalized No Homophony stems after *ba*- less frequently than chance in the 50-50 condition. The significant positive simple effect of Homophony (Homophony stem) and the lack of significant interaction effects involving Homophony indicate that learners made more over-generalization errors for stems that created homophony in all conditions. The significant and positive simple effect of Prefix Frequency (Frequent *ni-*) suggests that they also made more over-generalization errors when *ni-* was frequent in the input.

To summarize this section, adult English speakers displayed better acquisition of the palatalization rule (i.e., more correct applications and fewer over-generalization errors) when it was non-neutralizing than when it was neutralizing. Moreover, when comparing the stems that created homophony to stems that did not, there was no difference in the rate of correct application, but there were more over-generalization errors for the stems that created homophony.

2.6.5 Debriefing analysis

In this section, I take a different approach to examine the results by looking into participants' comments gathered during the debriefing. These comments provide more detailed information about which strategies participants used in the experiment and what they learned. Once participants completed the experiment, an experimenter asked participants to describe their thought process throughout the experiment, how they decided to answer in the prefix test phase, and their response strategies in the prefix test phase. The experimenter recorded the comments, and I subsequently coded the responses according to how much detail the participants were able to express about the phonological rule that they had learned using an ordinal scale with values from 1 (the least detailed) to 5 (the most detailed). The results for Neutralizing and Non-Neutralizing conditions are presented in Table 14, along with the criteria used to code the responses into each of the five levels. Responses from the No Alternation condition were not coded since no phonological rule was presented in that condition.

Table 14. Breakdown of participants' comments during debriefing along with coding criteria. *Note that four comments in the Neutralizing condition that were coded as Category 3 explicitly mentioned homophony creation, but the comments did not include the specific sound change.

Scale value	Coding Criteria	Non-Neutralizing Condition				Neutralizing Condition			
		50-50	Frequent *ba-*	Frequent *ni-*	**Total**	50-50	Frequent *ba-*	Frequent *ni-*	**Total**
5.	Explicitly mentioned at least one of the phonological changes (k→ʧ or g→ʤ).	3	4	2	9	2	0	4	**6**
4.	Mentioned that the *ba-* prefix does not trigger an alternation but the *ni-* prefix triggers an alternation without describing the explicit phonological changes.	9	10	11	**30**	2	4	1	7
3.	Mentioned that noun stems are changed, but nothing about the specific sounds or triggers involved.	4	1	4	9	3	6	9	18
2.	Mentioned two different options for plural forms but did not specifically mention that a change occurred.	5	7	6	**18**	8	10	7	25
1.	Did not mention any phonological changes or two options for plural forms.	3	3	3	9	10	5	3	18

I first compared the percentage of comments allocated in Categories 4 and 5 to the percentage of comments in Category 3 and below between the Neutralizing and Non-Neutralizing conditions. Category 4 was chosen as the reference point for comparison because the comments in Category 4 start mentioning some specific details on the alternation. In the Non-Neutralizing condition, 52% of the comments (39 out of 75) were allocated to Category 4 and above, and in the Neutralizing condition, 18% of the comments (13 out of 74) were in Category 4 and above, indicating that a larger number of participants in the Non-Neutralizing condition were able to provide some specific details about the alternation than were those in the Neutralizing condition. To compare the amount of detail participants were able to verbalize, I used a Mann-Whitney U test to examine the distribution of comments across the five categories in both conditions. The results showed that the median level of detail of the comments was significantly greater in the Non-Neutralizing condition (Mdn = 4) than in the Neutralizing condition (Mdn = 2; $U = 3672$, $p < 0.01$), which indicates that the significantly larger number of comments in the Non-Neutralizing condition provided more elaborate and explicit details on the alternation than in the Neutralizing condition. This finding suggests that participants in the Neutralizing condition internalized fewer details about the alternation compared to those in the Non-Neutralizing condition, or at least they were less able to verbalize these details. Although these self-reporting comments should be considered carefully and with scepticism, they support the aggregate results showing that English speakers were less successful at learning the neutralizing alternation compared to the non-neutralizing alternation.

2.7 Discussion

The primary goal of this experiment was to determine whether learners’ aversion to a phonological alternation and neutralization, along with the frequency of variants in the input, influence learning a probabilistic morpho-phonological pattern. I found that adult English speakers generally matched the relative frequency of variants presented in the input to their output, showing probability-matching behaviour. However, they shifted away from the alternation-triggering variant (*ni-*) only when it was infrequent in the input and triggered the non-neutralizing alternation. I also showed that English speakers were less successful at learning the neutralizing alternation that can create homophony compared to the non-neutralizing alternation. They made more over-generalization errors and fewer correct cases when the alternation was neutralizing than when it was non-neutralizing. They particularly made more over-generalization errors when the neutralizing alternation triggered homophony compared to when it did not and when the alternation-triggering prefix (*ni-*) was frequent in the input.

2.7.1 Frequency matching

The findings of this experiment provide a piece of evidence that adult learners have a strong tendency to match the frequency of variants in the input to their output when learning probabilistic patterns in an artificial language learning experiment. These findings parallel those of previous studies indicating adult learners' tendency to show probability-matching behaviour (probabilistic artificial languages: Hudson Kam & Newport, 2005, 2009; Austin, 2010; Schuler, 2017; probabilistic native language patterns: Ernestus & Baayen, 2003; Hayes et al., 2009; probabilistic non-linguistic patterns: Gardner, 1957; Weir, 1972).

In the current study, we expected learners to be biased against and shift away from the *ni-* prefix that triggered the alternation of velar-initial stems. However, learners tended to match the frequency of variants presented in the input to their output instead of shifting away from *ni-* (except for one condition). This frequency-matching behaviour is supported by the individual results showing that only half of the learners in each condition shifted away from one of the variants in their output. This finding is consistent with previous studies that show adults are less likely to regularize their output than children (Hudson Kam & Newport, 2005, 2009; Austin, 2010; Schuler, 2017). The experimental design of the current study could have encouraged learners to use both variants in their output, which could have led to probability-matching behaviour.

2.7.2 Regularization

I predicted that when *ni-* triggered a phonological alternation, learners would shift away from it to avoid an alternation. The experimental results showed that they shifted away from the *ni-* prefix but only when *ni-* triggered the non-neutralizing alternation and was infrequent in the input. The individual results showed that 21 out of 25 (85%) participants regularized their output with the *ba-* prefix in the Frequent *ba-* Non-Neutralizing condition. When the *ni-* prefix was presented 50% of the time or was frequent in the input (66.7%), they did not shift away from it and matched the input frequency to their output. Learners were successful at applying the non-neutralizing alternation (with an average accuracy of 74%); nonetheless, they shifted away from *ni-* when it triggered the non-neutralizing alternation and when it was infrequent in the input. These findings imply that learners avoided using the *ni-* prefix to maintain the surface form of the stem in the singular and plural forms. This result is also in line with previous studies showing learners' preference for paradigm uniformity in artificial language experiments (Kapatsinski, 2009, 2010, 2013; Tessier, 2012; Stave et al., 2013; White, 2013, 2014; Smolek & Kapatsinski, 2018; McMullin & Hansson, 2019; Smolek, 2019). A number of previous studies on paradigm uniformity mainly focused on how well learners apply an alternation. In the current study, I further show that the paradigm uniformity effect influences learners' choice of morphological structures, which was also found in a handful of recent studies (Stave et al., 2013; Smolek, 2019, Ch. 5; Do, 2018). For example, Do (2018) found that Korean children chose an alternative morphological structure or added an extra morphological structure to avoid making a non-unified paradigm even when they knew the correct morphological structure (see section 2.2).

Conversely, the findings of the current study are not in line with a previous study that used a similar learning paradigm. In Culbertson et al.'s (2012) study, adult learners were asked to learn probabilistic word order patterns, and they generally regularized their output with the frequent variant unless they were biased against the frequent word order. In contrast, in the current study, learners generally did not regularize their output with the frequent variant. The discrepant results of the two studies are particularly interesting because the experimental design of the current study was adopted from the paradigm used in Culbertson et al.'s study (2012).

Among various factors that could trigger probability-matching behaviour or regularization of the output, the type of input that participants are required to learn during the experiment, such as word or allomorph selections (as in the current study and Hudson Kam & Newport, 2005, 2009) or word orders (as in Culbertson et al., 2012), could influence the learning results. Another factor that could affect the learning results is the relative frequency of the variants in the input. Although the probability of the frequent variant in the current study (66.7%) was close to the one in Culbertson et al.'s study (70%), learners may start regularizing their output when the frequent variant is presented closer to 70% of the time in the input (Culbertson p.c.).

The complexity of the input can also affect learning a probabilistic pattern. Hudson Kam and Newport (2009) showed that adults regularized their output with the frequent variant that was presented 60% of the time in the input, which is even lower than the probability of the frequent variant presented in the current study. In their study, there were additional filler items in the infrequent patterns, and the items were inconsistently presented in the input (revisit section 2.4 for details), which possibly made the input more complicated. Similarly, Schuler (2017) found that learners regularized their output when they were required to learn a larger number of lexical items or asked to produce the answer within a very short time window (1.5 seconds), which can challenge adults' cognitive systems. The complexity of input can also trigger regularization in learning non-linguistic patterns. Gardner (1957) showed that adults matched the relative frequency when only two lights were presented, but they began to boost the frequency of the most frequent light at the expense of the infrequent light when three lights were presented.

The complexity of the input may affect learning probabilistic patterns because it can induce cognitive challenges. Learners may need longer working memory capacity or complicated processes to acquire a more complicated input (Hudson Kam & Newport 2005, 2009; Schuler 2017). Hudson Kam and Newport suggested that adults are likely to have a stronger working memory capacity to store and process complicated probabilistic patterns than children; hence, they can track the relative frequency of variants in the input and reflect it in their output. In contrast, children may have a weaker memory capacity to process complicated input (Less-is-More theory: Newport, 1990; Elman, 1993; Schuler, 2017; retrieving forms: Gathercole, 1998; Cowan, 1997; Weinart & Schneider, 1995), which can lead them to the strategy of regularizing their output with the frequent variant instead of tracking the relative frequency of variants.

Similarly, Schuler (2017) proposed that the different learning results between adults and children can be motivated by their memory and cognitive control systems. Children specifically have a stronger procedural memory system that stores and processes regular forms (patterns) whereas they have a weaker declarative memory system that manages to learn and store arbitrary and irregular forms (Ullman et al., 1997; Ullman, 2001b). The procedural memory system often functions the highest during childhood and tends to decline as children grow (Walton et al., 1992; Wolansky et al., 1999; Fredriksson et al., 2000; Schlaug, 2001; Janacsek et al., 2012). In other words, children can have a less developed cognitive control system that may encourage them to focus on the most regular and dominant form in the input and use it in their output (Diamond & Doar, 1989; Thompson-Schill et al., 2009). Adults, on the other hand, may have a more developed cognitive control system that enables them to process more complicated and arbitrary patterns in the input and reflect them in their output. When an input is too complicated and challenges adults' cognitive capacity, they can regularize their output (and produce more child-like responses).

The participants in the current study may not have been as cognitively challenged as those in Schuler's (2017) or Hudson Kam and Newport's (2009) studies. The participants in the current study were required to learn a relatively small number (10) of singular stems and asked to produce the plural forms within a 10-second time window (participants in Schuler's were required to learn 36 nonce words or asked to answer within a 1.5-second time window). They were also not introduced to additional filler items in the infrequent pattern as in Hudson Kam and Newport's study (2009). However, in the current study, the input that learners were required to learn may have imposed cognitive challenges in different ways. The complexity of alternation gradually increased in the Alternation Types. In the simplest case, the No Alternation condition, the learners were required to learn the probabilistic pattern where the plural form was marked by either *ba-* or *ni-*, and neither prefix triggered the alternation of stems. The participants in this condition matched the input frequency of prefix variants to their output and showed no tendency to regularize. The Non-Neutralizing condition was slightly more complicated because participants had to acquire that the *ni-* prefix triggered the alternation in addition to tracking the relative frequency of the prefix variants. The participants in this condition also generally frequency matched, but they regularized their output with the frequent variant particularly when the infrequent variant triggered an alternation. The Neutralizing condition was the most complicated case because there was an added factor of homophony. In this condition, learners had to acquire phonological neutralization that created homophony 50% of the time in addition to tracking the relative frequency of prefix forms.

Nonetheless, in the Neutralizing condition, English speakers did not shift away from the alternation-triggering prefix *ni-* even when it was the infrequent variant in the input. There are two possible explanations for this result. First, English speakers were less successful at learning the neutralizing alternation, meaning there was less pressure to avoid the trigger (*ni-*) of the neutralizing alternation. To recall, English speakers made fewer correct applications and more over-generalization errors in the Neutralizing condition compared to the Non-Neutralizing condition, implying that they were less certain about when the neutralizing alternation should be applied. As they had a weaker motivation to avoid the neutralizing-triggering prefix (*ni-*) in their output, they were less likely to regularize their output with the frequent prefix that did not trigger an alternation. By contrast, English speakers in the Non-Neutralizing condition successfully learned the non-neutralizing alternation, potentially resulting in more pressure to avoid the alternation-triggering prefix in their output. Another possible explanation is that the regularization found in the Frequent *ba-* Non-Neutralizing condition could have been a sampling error. Although the aggregate results, the individual results (sign test), and the debriefing comments all pointed to a significant and robust regularization effect, the effect was only found in one out of 9 conditions. The fact that this was limited to just one condition means that we should be mindful of the possibility that it could be a Type 1 error.

The finding that the participants in the current study did not regularize the output with the most complicated input (i.e., in the Neutralizing condition) raises a few questions: which elements of the input can add complexity and impose challenges to adults' cognitive capacity, and which elements can motivate regularization? In the current study, adding an alternation on top of the probabilistically variable patterns potentially led learners towards regularization (if the Frequent *ba-* effect was not a sampling error) but adding an extra factor of homophony did not. On the other hand, we cannot fully confirm that either element directly influenced the complexity of the input per se and burdened adults' cognitive capacity, in the same way that going from two to three lights as Gardner (1957) did or modifying the number of items or a response time in Schuler (2017). Nonetheless, both factors in the current experiment added complexity in terms of what participants were required to track and learn overall. These observations highlight the need for further research to better understand which elements of the input can lead adult learners to regularize their output.

2.7.3 Choice of prefix according to the stem type

Finally, I did not find a difference in using *ni-* with velar-initial stems (where palatalization should have been applied after *ni-*) or with other stems (where no alternation was applied). When learners shifted away from the *ni-* prefix in the Frequent *ba-* Non-Neutralizing condition, they used *ni-* less often with all stem types. Kapatsinski (2010) suggested two models that can account for the choice of suffix and palatalization from Russian loanword adaptation. In the two-stage model, the choice of an affix occurs first, and then the alternation of the base to which the affix is attached occurs (if needed). Because the modification of the base happens after the choice of the affix, whether the base form alternates or not is less likely to affect the choice of the affix. In the one-stage model, the choice of the affix and the modification of the base occur simultaneously, indicating that the modification of the base form is likely to influence the choice of the affix. The results of the current study are more consistent with a two-stage model where affix selection precedes phonology rather than a one-stage model where both occur together. Recall that the *ni-* prefix triggered the non-neutralizing or the neutralizing alternations of velar-initial stems, but the *ba-* prefix never triggered alternations. In eight out of nine conditions in the Non-Neutralizing and Neutralizing conditions, learners matched the input frequency of the prefixes to their output with all stem types (velar-initial or non-velar-initial stems). Particularly, in the Frequent *ba-* Non-Neutralizing condition where learners avoided using *ni-*, there was no difference between the percentage of times that they used *ni-* with velar-initial stems and with non-velar initial stems. This finding suggests that learners were not taking into account whether or not the stem would be alternated when choosing the prefix form. Hence, the two-stage model suggested by Kapatsinski (2010) can provide a more suitable explanation for the experimental results than the one-stage model. However, the fact that learners were asked to learn prefix forms instead of suffix forms may have caused the experimental results to more closely resemble the two-stage model by default because learners may have a tendency to plan words from left to right (Roelofs, 1996).

2.7.4 Homophony avoidance

In the present study, I expected to see a greater difficulty in learning a neutralizing alternation than a non-neutralizing alternation. The results show that adult learners were less successful at learning the neutralizing alternation compared to the non-neutralizing alternation. They made more over-generalization errors by incorrectly palatalizing the velar-initial stems after *ba-* and made fewer correct cases by palatalizing the velar-initial stems after *ni-* in all Prefix Frequency conditions, indicating that they were poorer at internalizing the neutralizing alternation itself as well as the trigger of the alternation. Indeed, the debriefing statements bolster the finding that learners were less able to articulate details on the neutralizing alternation than the non-neutralizing alternation. This finding suggests that lexical learning and phonological learning occur in parallel and affect each other. If these are two unrelated learning processes, having one surface form that refers to two distinct lexical items would not have affected learning the neutralizing alternation in the current experiments.

The less successful learning result of neutralizing alternations compared to the non-neutralizing alternations is more likely to be triggered by the phonological neutralization and homophony creation rather than the alternation itself because the same palatalization rule occurred in both the Neutralizing and Non-Neutralizing conditions. Furthermore, participants who learned neutralizing alternations did not perform differently depending on whether the specific stems became homophonous after palatalization. This finding indicates that the fact that some stems became homophonous in the Neutralizing condition caused learners to be less successful at learning the neutralizing alternation overall, but they did not successfully internalize which velar-initial stems resulted in homophony after palatalization.

The findings of the current study are in line with those of Yin and White's study (2018) showing that participants were less successful at learning neutralizing alternations that triggered homophony than at learning non-neutralizing alternations. The similar findings are noteworthy given that there are several differences. In their study, alternations were presented in a deterministic pattern, and learners were exposed to both non-neutralizing and neutralizing alternations. However, in the current study, alternations were probabilistically marked by one of the prefixes in the input, and learners were only exposed to the neutralizing alternation in the Neutralizing condition. Furthermore, learners in the current study were not tested with novel items, but those in Yin & White's study were tested with familiar and novel items. Lastly, learners in Yin and White's study were forced to choose between two options whereas those in the current study were asked to produce a plural form. Despite the differences between the studies, both sets of results suggest that learners were less successful at learning the neutralizing alternation that created homophony compared to the non-neutralizing alternation, which indicates learners' bias against neutralization and homophony creation. Finding the same effect with different methodologies adds to the robustness of this finding.

Additionally, the findings of the current study can be compared to a series of Kapatsinski's studies (2009, 2012, 2013) that showed participants' poor learning results for the alternation that triggered homophony in artificial language learning experiments. A detailed explanation of the studies can be found in section 2.3. Here, I briefly summarize the findings of the studies. Participants learned that velar-final stems were palatalized when followed by the /-i/ plural suffix (/k/, /g/ →[ʧ], [ʤ]/__i). The palatalization was presented in the product- and source-oriented learning tasks. In some conditions, learners were exposed to additional plural forms ending with the /-ʧ/ and /-ʤ/, which triggered the palatalization to be neutralizing and homophony-creating. When there were additional plural forms, learners were more successful at learning the palatalization in the product-oriented task design because additional plural forms supported the product-oriented generalization (i.e., plural forms end with [-ʧi] or [-ʤi]). By contrast, learners were less successful at applying palatalization in source-oriented task design. According to Kapatsinski, this is because the additional plural forms supported a particular source-oriented generalization that prefers adding /-i/ to a stem without stem change (e.g., singular [-k] and a plural [-ki]), which may support the velar-final stems to remain.

The current study did not present participants with additional stems to create homophony. Instead, the stem-initial /t/ and /d/ were correspondingly exchanged with /ʧ/ and /ʤ/ to trigger neutralization and homophony. Hence, the poorer application of the neutralizing alternation compared to the non-neutralizing alternation in the current study cannot be explained by the increased support for the source-oriented generalization because the evidence for the source-oriented generalizations was identical across the Non-Neutralizing and Neutralizing conditions. The exchanged stems could support the product-oriented generalization (i.e., plural forms begin with [niʧ-] or [niʤ-] sequences) in the Neutralizing condition because the exchanged stems increased the number of plural forms starting with [niʧ-] and [niʤ-] sequences in the Neutralizing condition. If the plural forms starting with [niʧ-] and [niʤ-] affected learners' product-oriented generalization in the Neutralizing condition, learners would have palatalized velar-initial stems more frequently in the Neutralizing condition than in the Non-Neutralizing condition. However, I found that learners palatalized velar-initial stems less frequently in the Neutralizing condition than in the Non-Neutralizing condition. This finding indicates that homophony creation affected learning the alternation even with the absence of additional support for a specific source-oriented generalization—the generalization that favours adding an affix without stem change. Accordingly, this finding emphasizes the crucial role of neutralization and homophony on learning a phonological alternation.

2.7.5 Over-generalization of palatalization

In the Neutralizing condition, learners made over-generalization errors by incorrectly palatalizing velar-initials stems after *ba-* across the Prefix Frequency conditions. The largest number of over-generalization errors was found when the alternation-triggering prefix *ni-* was frequent in the input. There are some possible explanations for the over-generalization cases found in the current study.

One possibility is that neutralizing alternations may have triggered learners to be less certain about which prefix triggered the alternation. The experimental results show that English speakers made significantly more over-generalization errors in the Neutralizing condition than in the Non-Neutralizing condition, suggesting that learners were poorer at learning which prefix triggered the alternation in the Neutralizing condition compared to the Non-Neutralizing condition. This is supported by the comments in the debriefing, where a significantly larger number of participants in the Neutralizing condition (61 out of 74, 83%) did *not* mention which prefix triggers the palatalization compared to participants in the Non-Neutralizing condition (36 out of 75, 48%). Thus, learners were likely to make more generalization errors in the Neutralizing condition because they were less certain about the trigger of the neutralizing alternation as well as the neutralizing alternation itself.

Another possible explanation for the over-generalization is product-oriented generalizations. For instance, learners might have acquired two types of generalizations: 1) plural forms should start with [baʧ-] and [baʤ-] sequences or 2) the stem-initial consonant of the plural forms should be /ʧ/ or /ʤ/ (e.g., perhaps as a part of a prosodic template, see Kapatsinski, 2018, ch.7). However, these product-oriented generalizations are not consistent with the full set of results. If the first product-oriented generalization triggered the over-generalization cases, it is expected that the errors would appear most often in the Frequent *ba-* Neutralizing condition, where the plural forms starting with [baʧ-] and [baʤ-] were the most frequently presented. However, learners made more over-generalization errors in the Frequent *ni-* condition where plurals starting with the sequence were presented less often. Furthermore, if the second product-oriented generalization triggered the over-generalization errors of the Neutralizing condition, learners would also have palatalized velar-initial stems after *ni-* as well as after *ba-* in the Neutralizing condition; however, the proportion of times that velar-initial stems were palatalized after *ni-* was significantly lower in the Neutralizing condition than in the Non-Neutralizing condition. In addition, a prominent product-oriented generalization is likely to result in palatalization of all stems regardless of their stem-initial consonants, but participants in the current study rarely over-applied the palatalization to non-velar-initial stems (less than 1.5% of the time across conditions). The design of the current study is also more likely to encourage source-oriented than product-oriented generalization because participants were required to successfully memorize the 10 singular stems (and the matching pictures) before learning the plural forms, which emphasizes the relationship between the singular and plural pairs. Therefore, it is less likely that the over-application of palatalization was triggered by learners' tendency to make product-oriented generalizations.

2.7.6 Potential effects of the scoring system

In the prefix test phase, the scoring system had an unintended coding error. The intended scoring system was that participants would receive the full 10 points when their response matched the alien's (e.g., participant: [niʧimu], alien: [niʧimu]) and a partial 5 points when their response was correct phonologically but did not match the alien's because they picked a different prefix form (e.g., participant: [niʧimu], alien: [bakimu]). In the prefix test phase, participants received the full 10 points for the matching responses, but they unintentionally received a partial 5 points for some error cases. For instance, when the stem of their response matched that of the alien, they received partial points, which could have triggered an over-application (e.g., participant: [baʧimu], alien: [niʧimu]) or under-application (e.g., participant: [nikimu], alien: [bakimu]). However, learners never heard the alien instructor producing an incorrect form, indicating that they always received the correct plural form as their feedback during the prefix test phase. Although the errors could (in principle) have affected participants' learning, in practice, it is unlikely to have had a meaningful impact on the main results. Below, I will outline these potential issues.

First, this scoring system could have affected learning palatalization because learners received partial scores for over- and under-application of the palatalization. This may have led to the over-generalization error cases found in the Neutralizing condition and added noise to the results. However, the scoring system cannot account for different learning results between the Non-Neutralizing and Neutralizing conditions because the same scoring system was implemented in both conditions, yet learners made more over-generalization errors in the Neutralizing condition. Second, the scoring system may have encouraged learners to choose the frequent prefix in their output. If learners successfully learned that palatalization only occurs after *ni-,* the best strategy to get higher scores would be to use the frequent prefix form because they received the full 10 points for matching the plural form and 0 points for unmatching plural forms. If learners did not successfully learn palatalization, they still received more points by choosing the frequent form, which also encouraged using the frequent form in learners' output. Hence, the scoring system may have promoted using the frequent variant, but it would not have systematically preferred one prefix over the other. The experimental results also show that learners did not overly use the frequent prefix form in their output (except for one condition).

Furthermore, these possible effects of the scoring error only applied to the velar-initial stems that were palatalized after *ni-* in the Non-Neutralizing condition and in the Neutralizing condition. For all stems in the No Alternation condition and for the non-velar-initial stems in the Non-Neutralizing and Neutralizing conditions, the intended scoring system was correctly implemented. In addition, a scoring system is effective only when participants pay attention to it. Culbertson et al. (2012) showed that learners' responses were similar when presented with feedback and scores and without them, suggesting that learners could pay less attention to the scores during the experiment. These observations suggest that the potential effect of the error did not have a strong effect on the results.

2.8 Conclusion

In this chapter, I showed that English speakers generally matched the frequency of probabilistic prefix forms in the input to their output. However, they shifted away from the prefix form that triggered a (non-neutralizing) alternation when it was infrequent in the input. They were also less successful at learning the neutralizing alternation than the non-neutralizing alternation. In the next chapter, I study if learners' abundant exposure to neutralization in their native language would affect learning a novel neutralizing alternation.

CHAPTER 3

L1 effect in learning morpho-phonological patterns: Neutralizations in Korean

3.1 Introduction

In the previous chapter, I examined whether learners' biases against phonological alternations and neutralization, together with the frequency of variants in the input, influenced learning a morpho-phonological construction. I found that adult English speakers generally showed probability-matching behaviour, meaning that they replicated the relative frequency of variants presented in the input. However, they shifted away from the variant (*ni-*) that triggered the non-neutralizing alternation when it was infrequent in the input, boosting the use of the more frequent variant (*ba-*) that did not trigger the alternation. Surprisingly, when the alternation became neutralizing and homophony-creating, they no longer shifted away from the prefix form (*ni-*) even when it was infrequent in the input. Instead, they showed frequency-matching behaviour. Learners were also less successful at learning the neutralizing alternation compared to the non-neutralizing alternation, which indicates that triggering homophony can inhibit the learnability of an alternation.

Although adult learners seem to disfavour neutralizing alternations, they are frequently found in languages. For instance, Korean has many neutralizing alternations that can result in homophony (Silverman, 2010; Kaplan, 2011). This chapter investigates whether having extensive prior exposure to neutralizing alternations (e.g., in their L1) facilitates the learning of a novel neutralizing alternation. I replicate the experiment from Chapter 2 with adult native Korean speakers due to the large amount of neutralization in their L1. By comparing the experimental results for Korean speakers to those for English speakers, I test how this L1 effect influences the learnability of a novel neutralizing alternation and the choice of morphological construction.

The specific questions that I test are (a) whether Korean speakers would be more successful at learning a novel alternation that triggers phonological neutralization and homophony than English speakers (who have comparatively less neutralization in their L1), (b) whether Korean speakers have a different tendency to shift away from the alternation-triggering (and neutralization-triggering) variant compared to English speakers, and (c) how the frequency of the alternation-triggering variant in the input influences these effects. To preview the experimental results, Korean speakers were equally good at learning the neutralizing alternation and the non-neutralizing alternation. They also closely matched the probability of prefix forms in the input to their output in all conditions. These results indicate that Korean speakers' substantial exposure to neutralizing alternations in their L1 is likely to facilitate learning the novel neutralizing alternation.

In the remainder of this chapter, I first review previous studies investigating L1 effects during (natural and artificial) language learning. Second, I provide an overview of many neutralizing alternations and derived homophony in Korean to show that Korean speakers are indeed exposed to extensive neutralization in their L1. Finally, I cover the experimental methods, analysis, results, and implications of the current study.

3.1.1 Effect of L1 on natural language acquisition

Learners' previous exposure to languages can often have a substantial effect on learning a new language. Learners often apply phonetic information, phonological rules, and morphosyntactic patterns from previous language experiences to learn a target language (i.e., a second natural language or an artificial language in an experiment), especially in the early stages of learning (Lado, 1957; Gass, 1979; Dulay et al., 1982; Broselow, 1984; Odlin, 1989; Brown, 1998; Sabourin, Stowe, & De Haan, 2006).

This influence of the L1 can obstruct or facilitate the acquisition of a target language (natural language learning: Weinreich, 1953; Brière, 1966; Schachter, 1974; Gass, 1979; Fisiak, 1981; Dulay et al., 1982; Broselow, 1984; Odlin, 1989; artificial language learning: Finn & Hudson Kam, 2008; White et al., 2018; Martin & Culbertson, 2020). Learners' L1 can cause errors in learning a target language (Altenberg & Vago, 1983, p.432). The Contrastive Analysis Hypothesis (CAH) suggests that the errors that learners make while learning a target language can be predicted and explained by the contrast between learners' L1 and the target language (Fries, 1945; Lado, 1957; James, 1980). The CAH views learning as a part of forming habits and indicates that pre-existing linguistic habits (knowledge) obtained from learners' L1 can be *transferred* to or *interfere* with a target language.[16] The terms *transfer* or *interfere*[17] are generally used to indicate learners' tendency to implement and apply existing linguistic knowledge to a target language (Stockwell, Bowen, & Martin, 1965; Wardhaugh, 1970; James, 1977; Gass, 1979). The term *interference* mostly indicates the errors or deviations in a target language that are motivated by habits and knowledge formed from previous language learning experiences (Weinreich, 1953; Dulay et al., 1982; Lott, 1983), and the term tends to focus on the negative effect of the pre-established habits and knowledge in learning a target language. The concept of *transfer* is found when learners apply patterns from previous exposure to languages to learning a target language (Gass, 1979; Odlin, 1989), meaning that transfer effects generally have a broader scope by including both the positive and negative effects of previous learning experiences.

Learners often transfer rules found in their L1 to their L2, which generates errors during L2 learning. For instance, previous studies showed that Hungarian and Polish speakers incorrectly

[16] There are different versions of CAH depending on how much it predicts the errors made in learning a target language and whether the contrasts/similarities between the L1 and a target language facilitate or obstruct learning a target language (see the strong, moderate, and weak CAH in Fries, 1945; Wardhaugh, 1970; Lado, 1957; Broselow, 1984).

[17] The term *transfer* is used differently in literature (see discussions in Murphy, 2003 (pp. 2-3) and Dost and Bohloulzadeh, 2017 (pp. 33-34)).

implement low-level rules from their L1 to their L2 when the L2 has the environment or context where the rules can be applied (Rubach, 1980; Altenberg & Vago, 1983). In Hungarian, the voice feature of an obstruent is assimilated to the *following* obstruent whereas in English, the obstruent becomes [-voice] when *preceded* by a voiceless obstruent. When two assimilation rules compete against one another, native Hungarian speakers who learned English as their L2 often chose to apply the rule from their L1, resulting in errors in English (e.g., *[biyon**t h**iz], for 'beyond his' [bijɑn**d h**ɪz]; Altenberg and Vago, 1983, p. 432). Similarly, Rubach (1980) shows that the palatalizing rule found in Polish can also transfer to English. In Polish, consonants are palatalized[18] when followed by /i/ or /j/, and native Polish speakers who learn English as their L2 often result in the following errors: *[pʲiːp(ə)l] for 'people' [piːp(ə)l], *[lʲiːk] for 'leak' [liːk], or *[tʲiːʧ] for 'teach' [tiːʧ] (Rubach, 1980, p. 367). Notably, the similarity between the voice assimilation rule in Hungarian and the palatalizing rule in Polish is that these rules are low-level rules (also referred to as 'phonetic rules' in Vago, 1980) that simplify pronunciation, reduce articulatory effort, or occur with fewer exceptions in a language (Altenberg & Vago, 1983; Rubach, 1984; Hammarberg, 1990). Additionally, English has an environment where these rules can be applied: English has two concatenating obstruents where Hungarian assimilation can be applied, and a consonant followed by either /i/ or /j/ where Polish palatalization can be applied.

In addition to low-level rules, learners can reflect existing syntactic knowledge from their L1 in learning an L2. In Bhela's study (1999), learners whose L1 varied (Vietnamese, Cambodian, Spanish, and Italian) were asked to write sentences in their L1 and in English (L2) upon seeing sequential pictures. The result showed that for syntactic structures that did not exist in learners' L1, they made errors due to their lack of understanding of the novel syntactic structures in English. Interestingly, when a syntactic structure in English existed in their L1 but when the structure was

---

[18] The palatalization rule in Polish resulted in [Cʲ], which does not change the place or manner of articulation of consonants, except for velar consonants, /k, ɡ, x/, which correspondingly move to post-palatal positions (Rubach, 1980).

used in different forms in their L1, learners found it difficult to learn the structure in English. Furthermore, Türker (2016) showed that the frequency of patterns in the L1 can affect learning a pattern in L2. For metaphorical expressions found in both the L1 and the L2, learners processed the expressions in the L2 more easily when the expression was frequent in their L1 than when it was infrequent.

In this section, I reviewed that learners' established linguistic knowledge, such as low-level rules, a syntactic structure, and the frequency of patterns from their L1, can influence L2 learning. In the following section, I survey the influence of the L1 transfer effect in artificial language learning experiments.

3.1.2 Effect of L1 on artificial language learning

Most prior research has focused on the effects of L1 on natural language acquisition, with relatively few studies using an artificial language learning experiment. Previous studies show that the frequency of patterns, lexical knowledge, and statistical information from learners' L1 can transfer to learning an artificial language (a.o., Onnis & Thiessen, 2013; White et al., 2018; Martin & Culbertson, 2020; Tang & Baer-Henney, 2023). A recent study provided evidence that the predominant morphological construction in the L1 can influence learners' preferences for constructions in an artificial language. Martin and Culbertson (2020) presented sequences of syllables (e.g., ta-ko) to native Kîîtharaka speakers whose L1 predominantly has prefix forms and to native English speakers whose L1 predominantly has suffix forms. Then, both groups of speakers heard two types of novel sequences in a test phase: a prefix-like pattern (where the beginning of a sequence changed) and a suffix-like pattern (where the end of a sequence changed) and were asked which sequence was more similar to what they had already learned. Native Kîîtharaka speakers found that the prefix-like patterns were more similar to what they had learned whereas English speakers found that suffix-like patterns were more similar, indicating that they favoured the prevalent pattern found in their L1 when learning an artificial language. They found a similar pattern when the speakers were tested on a non-linguistic pattern involving shapes.

Similarly, White et al. (2018) suggested that learners' frequent exposure to a morphological structure can affect learning an artificial language. In their study, native speakers of Dutch, English, French, German, and Greek were tested on how often they inferred a local harmony pattern depending on the affix type (either a prefix or a suffix) and the location of stressed vowels (locally or non-locally). In all L1s, suffixed forms were more predominant than prefixed forms. When the input was ambiguous between a local harmony pattern or a nonlocal harmony (vowel dependency) pattern, learners preferred the local harmony pattern. However, their preference for the local pattern was significantly greater when using suffixes than when using prefixes. While this preference can be universally found regardless of learners' L1, the authors indicated that the prevailing suffixed forms in learners' L1 potentially affected their preference for the locality of suffixed forms over prefixed forms.

In addition, the lexical knowledge that is obtained from L1 can affect accepting novel words in an artificial language learning experiment. Tang and Baer-Henney (2023) exposed native German and Mandarin speakers to nonce words whose word-likenesses were manipulated so that some nonce words were more similar to existing words in L1 than others. Learners were more accurate at judging whether or not novel nonce words were a part of the artificial language when the nonce words were less similar to existing words in their L1. Moreover, the study showed a weaker L1 effect for Mandarin speakers than German speakers, which can be partially explained by the structures of Mandarin (monomorphemic words, compounding language, and reduplicated patterns).

Furthermore, Onnis and Thiessen (2013) examined whether learners segment a novel sequence using statistical information used to identify word orders in their L1. In the study, native Korean and English speakers were exposed to two types of continuous sequences: one sequence consisted of nonce words that corresponded to the word order patterns of English, and the other sequence consisted of those that complied with Korean. Korean and English have the opposite word orders for transitive, imperative, and exhortative constructions due to the different locations of a head element of a phrase, suggesting that the languages have different probabilistic information between words. For example, the head appears later in English (e.g., *to school*) and it appears first in Korean (e.g., *school to*). When hearing a pair of novel sequences, learners chose the sequence that complied with the word order patterns in their L1 more frequently than the one that did not. Importantly, the fact that both word order patterns were presented equally often during the training corroborated the finding that learners preferred to segment nonce words in a sequence using statistical information congruent with their L1.

However, learners' previous exposure to L1 does not always affect learning novel patterns in an artificial language. They often do not prefer the pattern in an artificial language that is prevalent in their L1. For example, Culbertson et al. (2012) showed that learners' L1 or L2 experiences did not have a major effect on the choice of noun phrase orders in artificial language learning. Here, I briefly revisit the part of the study examining the potential effect of learners' previous linguistic experiences in choosing the word order in the output (see the detailed summary of Culbertson et al.'s study in section 2.4). During the experiment, native English speakers were exposed to one of the four noun phrase patterns with different noun and modifier orders. One of the orders matched the noun phrase pattern of English. The orders in the patterns were presented with different frequencies in the input (a major and minor order). The results showed that there was no significant difference between the percentage of time that participants chose the major order that matched with their L1 and the one that did not (except for one major order that was less frequently found across languages), suggesting that there was no preference for the order that existed in their L1. Thus, the choice of patterns was more likely to be triggered by universal biases in favour of specific patterns or the language that they were trained on in the experiment rather than by their prior experience from their L1, implying that the L1 had a limited influence in this study. As Culbertson et al.'s (2012) study indicates, there are factors beyond L1 that can impact learning novel linguistic patterns.

3.2 Neutralization in Korean

Korean "is the prototypical 'Linguistics 101' example of neutralization, as neutralization here involves so many contrastive values" (Silverman, 2010, p. 455). Kaplan (2011) also pointed out that studying neutralization and homophony in Korean is convenient because its orthographic system is mostly morphophonemic, which shows syllable structures and presents underlying forms of phonemes. Korean has various neutralizing alternations that merge a number of contrasting values. For instance, the first [m] in a surface form of [ə**m**mu] can potentially be 14 different underlying consonants due to several neutralizations: /n, m, p, p^{h}, lp, lp^{h}, ps, s, s', ʧ, ʧ^{h}, t, t', h/ (Chung, 1980, p. 2). Neutralizations in Korean sometimes result in derived homophony. For example, word-final obstruents are merged to [t], which creates homophony (e.g., /pi**s**/ 'a comb', /pi**ʧ**^{h}/ 'light', or /pi**ʧ**/ 'a debt' – [**pit**], /na**s**/ 'a sickle', /na**ʧ**^{h}/ 'a face' or /na**ʧ**/ 'daytime' – [**nat**]).

Compared to English speakers, Korean speakers are more likely to be frequently exposed to neutralizing alternations that can potentially derive homophony, which makes them suitable participants for an experiment testing whether their extensive exposure to neutralizing alternations from their previous linguistic experience influences learning a novel neutralization. Recall from the experiment reported in Chapter 2 that English speakers were less successful at learning the alternation when it was neutralizing and homophony-creating than when it was not neutralizing. Given Korean speakers have abundant exposure to neutralization in their L1, I test the hypothesis that having excessive neutralizations in their L1 can facilitate Korean speakers' learning of a novel neutralizing alternation.

Table 15. Summary of neutralizing alternations in Korean.

1) Coda neutralization	A coda obstruent in a non-prevocalic position is merged to one of the three unreleased plosives ([t̚], [k̚], or [p̚]). /t, t^h, ʧ, ʧ^h, s, s’, (h)/ → [t̚] / ___$]_\sigma$, C, #, + /k, k^h, k’/ → [k̚] / ___$]_\sigma$, C, #, + /p, p^h/ → [p̚] / ___$]_\sigma$, C, #, +
2) Nasal and liquid neutralization	
2.2) Nasal lateralization	A sequence of a lateral liquid and a coronal nasal consonant becomes a long lateral consonant ([lː]). /l+n/, /n+l/ → [lː]
2.3) Liquid nasalization	A non-coronal obstruent and a lateral liquid in a sequence correspondingly become [+nasal]. /p+l/ → [mn], /k+l/ → [ŋn]
2.4) Nasal assimilation	An obstruent followed by a nasal consonant becomes [+nasal]. $[-\text{sonorant}] \rightarrow [+\text{nasal}]\ /\ \_\_\_\_ \begin{bmatrix}+\text{consonantal}\\ +\text{nasal}\end{bmatrix}$
3) Resyllabification	A coda consonant is resyllabified to the onset of the following syllable starting with a vowel or /h/. VC.(h)V → V.C(h)V
4) Post-obstruent tensing	The laryngeal features of a non-aspirated obstruent become [+tense] when preceded by an obstruent. $\begin{bmatrix}-\text{sonorant}\\ -\text{spread glottis}\end{bmatrix} \rightarrow [+\text{tense}]\ /\ [-\text{sonorant}]\ \_\_\_$
5) [h]-aspiration	A plain stop obstruent is aspirated when adjacent to /h/. $\begin{bmatrix}-\text{sonorant}\\ -\text{continuant}\\ -\text{spread glottis}\\ -\text{tense}\end{bmatrix} \rightarrow [+\text{spread glottis}]\ /\ \_\_\_[\text{h}],\ [\text{h}]\_\_\_$
6) Assibilation	A sequence of a coronal obstruent followed by either /s/ or /s’/ becomes [s’]. $\begin{bmatrix}-\text{sonorant}\\ +\text{coronal}\end{bmatrix} \begin{bmatrix}+\text{strident}\\ -\text{distributed}\\ +\text{continuant}\end{bmatrix} \rightarrow [\text{s'}]$
7) Coda consonant cluster simplification	/lk/ → [k̚], /lm/ → [m], /lp^h/ → [p̚] /lp/ → [l]/[p̚] /ps/ → [p̚], /ks/ → [k̚], /lh/, /ls/, /lt^h/ → [l] /nʧ/, /nh/ → [n]

In the rest of this section, I summarize some of the many examples of neutralization in Korean to illustrate that Korean speakers are indeed frequently exposed to neutralizations in their L1.

Table 15 shows the summary of the nine neutralizing alternations that will be reviewed in this section.

3.2.1 Consonants of Korean

Korean has 19 consonants, and 15 of them are obstruents (see Table 16). The plosive and affricate obstruents have a three-way laryngeal contrast: plain (C), aspiration (C^h), and tense (C’).[23] Fricative obstruents have a two-way contrast (plain and tense). All obstruents in Korean are underlyingly [-voice], and there is no voicing contrast among obstruents in underlying forms. All consonants can be in a syllable-initial position but, in a syllable-final position, only 7 of them, /k, t, p, n, l, m, ŋ/, are realized in a surface form.

Table 16. The consonants in Korean by place and manner of articulation with Korean orthography and IPA.

		bilabial	alveolar	alveo-palatal/ palatal	velar	glottal
Plosive	Plain	ㅂ p	ㄷ t		ㄱ k	
	Aspiration	ㅍ p^h	ㅌ t^h		ㅋ k^h	
	Tense	ㅃ p’	ㄸ t’		ㄲ k’	
Affricative	Plain			ㅈ ʧ[24]		
	Aspiration			ㅊ ʧ^h		
	Tense			ㅉ ʧ’		
Fricative	Plain		ㅅ s			ㅎ h
	Tense		ㅆ s’			
Nasal		ㅁ m	ㄴ n		ㅇ ŋ	
Lateral			ㄹ l			

[23] The terms that refer to the three-way laryngeal contrasts of obstruents in Korean vary in the literature: unaspiration/lax/plain/lenis (C), aspiration (C^h), and reinforce/tense/fortis (C’) (Silverman, 2010; Kim & Jongman, 1996; Kim, 1979; Chung, 1980). In the current thesis, I use the following terms: plain (C), aspiration (C^h), and tense (C’).

[24] The palatal affricatives (ㅈ, ㅊ, ㅉ) in Korean are often transcribed as /c, c^h, c’/, /ʨ, ʨ^h, ʨ’/ or / t͡ʃ, t͡ʃ^h, t͡ʃ’/, and this thesis uses the last transcription.

3.2.2 Coda neutralization

One of the most studied neutralizing alternations in Korean is coda neutralization. When an obstruent in the coda is in a non-prevocalic position—in a syllable-final position (__]σ), at a word and compound boundary (__#/+), or followed by a consonant-initial syllable (__.C)—the obstruent is merged to one of the three unreleased plosives ([t˺], [p˺], or [k˺]) (Martin, 1951; Lee, 1972; Kim-Renaud, 1974; Chung, 1980; Kim & Jongman, 1996; Sohn, 1999; Silverman, 2010; Kaplan, 2011). Coda neutralization affects seven coronal obstruents (/t, t^h, ʧ, $ʧ^h$, s, s' (h)/→[t˺]), two labial obstruents (/p, p^h/→[p˺]), and three dorsal obstruents (/k, k^h, k'/→[k˺]), allowing only three unreleased obstruents ([t˺, p˺, k˺]) to surface in a non-prevocalic position (see Table 17), which can result in a 75% reduction of contrasting obstruents (Silverman, 2010). Silverman showed that 28% of the nouns from the Sejong Project corpus had the unreleased obstruent in a non-prevocalic position (for more examples, see Silverman, 2010, pp. 462-463, Table 6).

3.2.3 Nasal and liquid neutralization

Nasal and liquid consonants are frequently merged in Korean, which sometimes results in derived homophony. In this section, I introduce nasal lateralization, liquid nasalization, and nasal assimilation in turn.

Nasal lateralization

Nasal lateralization is found when a sequence of a lateral and a coronal nasal consonant (i.e., /l+n/, /n+l/, regardless of the order) alternates to a long lateral consonant ([lː]; Kim-Renaud, 1974; Martin, 1992; Davis & Shin, 1999; Silverman, 2010). Nasal lateralization can create derived homophony (see Table 18).

Table 17. Examples of coda neutralization.

<table>
<tr><td colspan="5">Coda neutralization

A coda obstruent in a non-prevocalic position is merged to one of the three unreleased plosives ([t̚], [k̚], or [p̚]).

/t, tʰ, ʧ, ʧʰ, s, s’, (h)/ → [t̚] / ___] σ, C, #, +
/k, kʰ, k’/ → [k̚] / ___] σ, C, #, +
/p, pʰ/ → [p̚] / ___] σ, C, #, +</td></tr>
<tr><td colspan="5">/t, tʰ, ʧ, ʧʰ, s, s’, (h)/ → [t̚]/ ___] σ, C, #, +</td></tr>
<tr><td></td><td>words</td><td>__] σ/#/+</td><td colspan="2">__ .C
(/-to/: affix additive,
/-ta/: declarative marker (DECL))</td></tr>
<tr><td rowspan="3">a)</td><td>빚 /piʧ/ ‘a debt’</td><td rowspan="3">[pit̚]</td><td>/piʧ + to/</td><td rowspan="3">[pit̚.t’o]</td></tr>
<tr><td>빛 /piʧʰ/ ‘light’</td><td>/piʧʰ + to/</td></tr>
<tr><td>빗 /pis/ ‘a comb’</td><td>/pis + to/</td></tr>
<tr><td rowspan="6">b)</td><td>났 /nas’/
‘to happen/grow +pst’</td><td rowspan="6">[nat̚]</td><td>/nas’+ ta/</td><td>[nat̚.t’a]</td></tr>
<tr><td>낫 /nas/ ‘a sickle’</td><td>/nas + to/</td><td rowspan="5">[nat̚.t’o]</td></tr>
<tr><td>낮 /naʧ/ ‘daytime’</td><td>/naʧ + to/</td></tr>
<tr><td>낯 /naʧʰ/ ‘a face’</td><td>/naʧʰ + to/</td></tr>
<tr><td>낟 /nat/ ‘a single grain’</td><td>/nat + to/</td></tr>
<tr><td>낱 /natʰ/
‘a piece/ a unit’</td><td>/natʰ + to/</td></tr>
<tr><td rowspan="2">c)</td><td>갓 /kas/
‘a traditional Korean hat’</td><td rowspan="2">[kat̚]</td><td>/kas + to/</td><td>[kat̚.t’o]</td></tr>
<tr><td>갔 /kas’/ ‘to go+pst’</td><td>/kas’ + ta/</td><td>[kat̚.t’a]</td></tr>
<tr><td>d)</td><td>좋 /ʧoh/ ‘to like’</td><td>[ʧot̚]</td><td>/ʧoh + ta/</td><td>[ʧo.tʰa][25]</td></tr>
<tr><td></td><td colspan="4">/k, kʰ, k’/ → [k̚]/ ___] σ, C, #, +</td></tr>
<tr><td>e)</td><td>국 /kuk/
‘broth’</td><td>[kuk̚]</td><td>/kuk + to/</td><td>[kuk̚.t’o]</td></tr>
</table>

[25] The underlying form /ʧoh+ta/ is realized as[ʧo.tʰa] instead of [ʧot̚.t’a] due to [h]-aspiration, which will be mentioned later in this chapter.

	부엌 /pu.əkʰ/ ‘kitchen’	[pu.ək˺]	/pu.əkʰ + to/	[pu.ək˺.t’o]
	닦 /tak’/ ‘to wipe/ to clean’	[tak˺]	/tak’ +ta/	[tak˺.t’a]
	/p, pʰ/→ [p˺]/ ___]σ, C, #, +			
f)	덥 /təp/ ‘hot’	[təp˺]	/təp + ta/	[təp˺.t’a]
	덮 /təpʰ/ ‘to cover’		/təpʰ + ta/	

Table 18. Examples of nasal lateralization.

Nasal lateralization A sequence of a lateral liquid and a coronal nasal consonant becomes a long lateral consonant ([l:]). /l+n/, /n+l/ → [l:]			
a)	찰나 /ʧʰal.na/	‘a moment’	[ʧʰa**l**:a]
b)	찬란 /ʧʰan.lan/	‘glittering, brilliant’	[ʧʰa**l**:an]
c)	분리 /pun.li/	‘separation’	[pu**l**:i]
	불리 /pul.li/	‘a disadvantage’	
d)	날로 /nal.lo/	‘a day + ADV’	[na**l**:o]
	난로 /nan.lo/	‘a heater’	

Liquid nasalization

Liquid nasalization refers to the phenomenon in which a non-coronal obstruent and a lateral liquid in a sequence alternate to corresponding nasal consonants (i.e., /p+l/ → [mn], /k+l/ → [ŋn]), which can trigger derived homophony (see Table 19; Kim-Renaud, 1974; Martin, 1992; Davis & Shin, 1999).

Table 19. Examples of liquid nasalization.

<table>
<tr><td colspan="4">Liquid nasalization

A non-coronal obstruent and a lateral liquid in a sequence correspondingly become [+nasal].

/p+l/ → [mn], /k+l/ → [ŋn]</td></tr>
<tr><td rowspan="4">a)</td><td>답례 /tap.lje/</td><td>‘something in return’</td><td>[tam.nje]</td></tr>
<tr><td>담낭 /tam.naŋ/</td><td>‘a gall (bladder)’</td><td>[tam.naŋ]</td></tr>
<tr><td>압력 /ap.ljək/</td><td>‘pressure’</td><td>[am.njək]</td></tr>
<tr><td>담력 /tam.ljək/</td><td>‘courage, nerve’</td><td>[tam.njək]</td></tr>
<tr><td rowspan="2">b)</td><td>격리 /kjək.li/</td><td>‘isolation’</td><td rowspan="2">[kjəŋ.ni]</td></tr>
<tr><td>경리/kjəŋ.li/</td><td>‘accountancy/ an accountant’</td></tr>
<tr><td rowspan="2">c)</td><td>국립 /kuk.lip/</td><td>‘national’</td><td>[kuŋ.nip]</td></tr>
<tr><td>궁리 /kuŋ.li/</td><td>‘deliberation’</td><td>[kuŋ.ni]</td></tr>
</table>

Table 20. Examples of nasal assimilation.

<table>
<tr><td colspan="4">Nasal assimilation

An obstruent followed by a nasal consonant becomes [+nasal].

$[-\text{sonorant}] \rightarrow [+\text{nasal}] \ / \ \_\_\_\_ \begin{bmatrix}+\text{consonantal}\\ +\text{nasal}\end{bmatrix}$</td></tr>
<tr><td rowspan="2">a)</td><td>낱말 /natʰ.mal/</td><td>‘a word’</td><td>[nan.mal]</td></tr>
<tr><td>난가 /nan.ka/</td><td>‘a disturbed family’</td><td>[nan.ga]</td></tr>
<tr><td rowspan="2">b)</td><td>꽃말
/k’oʧʰ.mal/</td><td>‘symbolic meaning of flower’</td><td>[k’on.mal]</td></tr>
<tr><td>단막 /tan.mak/</td><td>‘one act’</td><td>[tan.mak]</td></tr>
<tr><td rowspan="2">c)</td><td>찻물 /ʧʰas.mul/</td><td>‘tea’</td><td rowspan="2">[ʧʰan.mul]</td></tr>
<tr><td>찬물 /ʧʰan.mul/</td><td>‘cold water’</td></tr>
<tr><td rowspan="4">d)</td><td>작물 /ʧak.mul/</td><td>‘crop’</td><td rowspan="2">[ʧaŋ.mul]</td></tr>
<tr><td>장물 /ʧaŋ.mul/</td><td>‘stolen goods’</td></tr>
<tr><td>학문 /hak.mun/</td><td>‘learning’</td><td rowspan="2">[haŋ.mun]</td></tr>
<tr><td>항문 /haŋ.mun/</td><td>‘anus’</td></tr>
</table>

Nasal assimilation

Nasal assimilation is found when an obstruent is followed by a nasal consonant, and the obstruent merges to [+nasal] (i.e., /tʰ+m/ → [nm], /ʧʰ+m/ → [nm], /s+n/ → [nn], /k+m/ → [ŋm]; Kim-Renaud, 1974; Martin, 1992; Davis & Shin, 1999; for more examples, see Silverman, 2010, pp. 465-469, Table 12). Nasal assimilation was applied to 51% of the nouns in the Sejong Corpus (Hwang, 2008), and it can trigger derived homophony (see Table 20).

3.2.4 Other neutralizations

In this section, I outline some additional neutralizing alternations that do not fall into the categories above: resyllabification, post-obstruent tensing, [h]-aspiration, assibilation, and coda consonant cluster simplification.

Resyllabification

When a coda consonant is followed by a vowel-initial or /h/-initial syllable, the coda consonant is resyllabified into the onset of the following syllable (VC.(h)V →V.C(h)V). The resyllabification can trigger homophony (see Table 21; Kim & Jongman, 1996; Sohn, 1999).

Post-obstruent tensing

In addition, *post-obstruent tensing* (also refers to the tensing rule, partial neutralization, or cluster reinforcement) is frequently found in Korean. When an obstruent-final syllable is followed by a syllable whose onset is a non-aspirated obstruent, the laryngeal features of the onset obstruent merge to [+tense] (Table 22; Kim-Renaud, 1974; Chung, 1980; Sohn, 1999; Martin, 1992; Silverman, 2010).

Table 21. Examples of resyllabification.

<table>
<tr><td colspan="4">Resyllabification

A coda consonant is resyllabified to the onset of the following syllable starting with a vowel or /h/.

VC.(h)V → V.C(h)V</td></tr>
<tr><td rowspan="2">a)</td><td>물이 /mul+i/</td><td>‘water+NOM’</td><td rowspan="2">[mu.ɾi] [26]</td></tr>
<tr><td>무리 /mu.li/</td><td>‘a flock, a pack’</td></tr>
<tr><td rowspan="2">b)</td><td>밤안개
/pam+an.kɛ/</td><td>‘night frog’</td><td>[pa.man.gɛ]</td></tr>
<tr><td>물만두
/mul+man.tu/</td><td>‘a boiled dumpling’</td><td>[mul.man.du]</td></tr>
<tr><td rowspan="2">c)</td><td>곤약 /kon.jak/</td><td>‘devil’s-tongue jelly’</td><td>[ko.njak]</td></tr>
<tr><td>새벽녘
/sɛ.pjək.njəkʰ/</td><td>‘around dawn’</td><td>[sɛ.bjəŋ.njək̚]</td></tr>
<tr><td rowspan="2">d)</td><td>착오 /ʧʰak.o/</td><td>‘a mistake’</td><td rowspan="2">[ʧʰa.go] [27]</td></tr>
<tr><td>차고 /ʧʰa.ko/</td><td>‘a garage’</td></tr>
<tr><td rowspan="2">e)</td><td>집이 /ʧip+i/</td><td>‘a house + NOM’</td><td>[ʧi.bi]</td></tr>
<tr><td>구비/ku.pi/</td><td>‘be handed down orally’</td><td>[ku.bi]</td></tr>
<tr><td rowspan="2">f)</td><td>낮이 /naʧ+i/</td><td>‘daytime+ NOM’</td><td rowspan="2">[na.ʤi]</td></tr>
<tr><td>나지 /na.ʧi/</td><td>‘bare land’</td></tr>
<tr><td rowspan="2">g)</td><td>낯이 /naʧʰ + i/</td><td>‘face/skin+ NOM’</td><td>[na.ʧʰi]</td></tr>
<tr><td>가치 /ka.ʧʰ i/</td><td>‘value, worth’</td><td>[ka.ʧʰi]</td></tr>
<tr><td rowspan="2">h)</td><td>거북이 /kə.puk.i/</td><td>‘a turtle’</td><td>[kə.bu.gi]</td></tr>
<tr><td>계산기 /kje.san.ki/</td><td>‘a calculator’</td><td>[kje.san.gi]</td></tr>
<tr><td rowspan="2">i)</td><td>밖이 /pak’+i/</td><td>‘outside +NOM’</td><td>[pa.k’i]</td></tr>
<tr><td>토끼 /to.k’i/</td><td>‘a rabbit’</td><td>[to.k’i]</td></tr>
<tr><td rowspan="2">j)</td><td>톳이 /tʰos+i/</td><td>‘hijiki +NOM’</td><td rowspan="2">[tʰo.si]</td></tr>
<tr><td>토시 /tʰo.si/</td><td>‘an arm warmer’</td></tr>
</table>

[26] Korean has one phonemic lateral liquid /l/, which has two allophones [ɾ] and [l]. A tap [ɾ] appears in intervocalic and word-initial positions, (e.g., [naɾa] ‘a nation’, [ɾamjʌn] ‘ramen noodle’), and a lateral [l] appears in coda (e.g., [kul] ‘oyster’; Reyes et al., 2017). When /l/ is resyllabified as an onset of the following syllable, it is realized as [ɾ]. When there is a sequence of lateral liquids, it is realised as a long lateral liquid (i.e., [lː]).

[27] A plain obstruent is realized as [+voice] when located intervocalically. The aspirated and tensed obstruents, fricative (/s/ and /s’/) obstruents, and other non-obstruent consonants do not change their voicing feature intervocalically.

Table 22. Examples of post-obstruent tensing.

<table>
<tr><td colspan="4">Post-obstruent Tensing

The laryngeal features of a non-aspirated obstruent become [+tense] when preceded by an obstruent.

$\begin{bmatrix}-\text{sonorant}\\-\text{spread glottis}\end{bmatrix} \rightarrow [+\text{tense}] \,/\, [-\text{sonorant}]$ ___</td></tr>
<tr><td rowspan="2">a)</td><td>접근 /ʧəp.kɯn/</td><td>‘approach’</td><td>[ʧəp˺.k’ɯn]</td></tr>
<tr><td>노끈 /no.k’ɯn/</td><td>‘string, twine’</td><td>[no.k’ɯn]</td></tr>
<tr><td rowspan="2">b)</td><td>돛단배
/toʧ^{h}.tan.pɛ/</td><td>‘a sailing ship’</td><td>[tot˺.t’an.bɛ]</td></tr>
<tr><td>외딴 /we.t’an/</td><td>‘isolated’</td><td>[we.t’an]</td></tr>
<tr><td rowspan="2">c)</td><td>곡식/kok.sik/</td><td>‘crops, grains’</td><td>[kok˺.s’ik]</td></tr>
<tr><td>하나씩/ha.na+s’ik/</td><td>‘one+each (one by one/ each one)’</td><td>[ha.na.s’ik]</td></tr>
<tr><td rowspan="2">d)</td><td>밭두렁
/pat^{h}+tu.ləŋ/</td><td>'a bank around a field’</td><td>[pat˺.t’u.ɾəŋ]</td></tr>
<tr><td>메뚜기/me.t’u.ki/</td><td>‘a grasshopper’</td><td>[me.t’u.gi]</td></tr>
<tr><td rowspan="2">e)</td><td>밧줄/pas.ʧul/</td><td>‘a rope’</td><td>[pat˺.ʧ’ul] [28]</td></tr>
<tr><td>혼쭐 /hon+ʧ’ul/</td><td>‘severe experience’</td><td>[hon.ʧ’ul]</td></tr>
</table>

[h]-aspiration

[h]-aspiration is observed when a plain stop obstruent is adjacent to /h/, and the laryngeal features of the obstruent merge to [+spread glottis] (Table 23; Chung, 1980; Sohn, 1999; Kaplan, 2011).

[28] A coda obstruent followed by an aspirated or a tensed obstruent can sometimes be deleted in surface forms (e.g., /pas.ʧul/ 밧줄, ‘a rope’ can be realized as either [pa.**ʧ’**ul] or [pa**t**˺.**ʧ’**ul). This phenomenon refers to pre-tense/aspirate reduction. It is an optional rule that is often found in casual speech (Sohn, 1999; Kaplan, 2011).

Table 23. Examples of [h]-aspiration.

[h]-aspiration

A plain stop obstruent is aspirated when adjacent to /h/.

$$\begin{bmatrix} -\text{sonorant} \\ -\text{continuant} \\ -\text{spread glottis} \\ -\text{tense} \end{bmatrix} \rightarrow [+\text{spread glottis}] \ / \ \_\_\_\,[\text{h}],\ [\text{h}]\_\_\_$$

a)	곱하기 /kop.ha.ki/	'multiplication'	[ko.**pʰ**a.gi]
	갑판 /kap.pʰan/	'a deck'	[kap˺.**pʰ**an]
b)	묻히다 /mut.hi.ta/	'be buried'	[mu.**ʧʰ**i.da][29]
	비치다 /pi.ʧʰi.ta/	'shine, light up'	[pi.**ʧʰ**i.da]
c)	많다 /manh-ta/	'many-DECL (a lot of, plenty of)'	[man.**tʰ**a]
	맞흥정 /maʧ+hɯŋ.ʧəŋ/	'direct bargain'	[ma.**tʰ**ɯŋ.ʤəŋ]
	맏형 /mat+hjəŋ/	'the eldest brother'	[ma.**tʰ**jəŋ]
	권태 /kwən.tʰɛ/	'get bored with'	[kwən.**tʰ**ɛ]
d)	적히다 /ʧək.hi.ta/	'be written (down)'	[ʧə.**kʰ**i.da]
	가리키다 /ka.li.kʰi.ta/	'point to, indicate (with gesture)'	[ka.ɾi.**k**ʰi.da]

Assibilation

Assibilation occurs when a coronal obstruent is followed by either /s/ or /s'/, and the sequence is realized as [s'] (e.g., /pis-soli/ 'sound of rain' - [pi.**s'**oɾi][30], /təs-sɛ/ 'a resident bird' - [tə.**s'**ɛ], /piʧʰ-sal/ 'light ray', /pis-sal/ 'teeth of a comb' - [pi.**s'**al]; Martin, 1992; Silverman, 2010).

[29] The word /mut-hi-ta/ 묻히다 is realised as [mu.**ʧʰ**i.da] instead of [mu.**tʰ**i.da] due to the palatalization of /tʰ/ to [ʧʰ] before /i/.

[30] Neutralization triggered by assibilation does not always alternate a sequence of a coronal obstruent+/s/ or /s'/ into [s']. Often, the coronal obstruent in coda remains. For instance, /pis.soli/ 'sound of rain' can be realized as either [pi.**s'**oɾi] or [pi**t**˺.**s'**oɾi], and /piʧʰ-sal/ 'light ray' and /pis-sal/ 'teeth of a comb' can be realized as either [pi.**s'**al] or [pi**t**˺.**s'**al].

Coda consonant cluster simplification

Finally, consonant clusters in a coda position are simplified and realized as a single consonant, which can often merge coda obstruents (Sohn, 1999; Kaplan, 2011). In Korean, there are 11 consonant clusters in coda: /lk/, /lm/, /lp/, /lpʰ/, /ps/, /ks/, /lh/, /ls/, /ltʰ/, /nʧ/, and /nh/. The three /l+non-coronal plosive/ clusters (/lk/, /lm/, /lpʰ/) are realized as a corresponding non-coronal plosive ([k˺], [m], [p˺]) in a non-prevocalic position[31], which often creates homophony (e.g., /i**lk**.ta/ 'read-DECL' and /i**k**.ta/ 'to ripe-DECL'– [i**k**˺.t'a], /hɯ**lk**/ 'soil' and / hɯ**k**/ 'black, sobbing sound'– [hɯ**k**˺], /a**lm**/ 'to know' and /a**m**/ 'cancer' – [a**m**][32]). Interestingly, the /lp/ cluster is most often realized as [l] but can also be realized as [p] in some cases (/nəlp.ta/ 'wide and spacious-DECL' – [nə**l**.d'a], /jə.təlp/ 'eight' – [jə.də**l**], /nəlp.ʧuk/ 'long and wide' – [nə**p**˺.ʧ'uk˺], /palp-ta/ 'step on' – [pa**p**˺.t'a]). The rest of the coda clusters are realized as the first consonant: /ps/ – [p˺], /ks/ – [k˺], /lh/, /ls/, /ltʰ/ – [l], /nʧ/, /nh/ – [n].

3.2.5 Homophony created by neutralizations in Korean

The fact that Korean has a long list of neutralizing alternations may lead to the idea that Korean also has many cases of derived homophony. However, while Korean could have more derived homophony than a language such as English, it turns out that Korean has fewer cases of derived homophony than we might expect based on the sheer number of neutralizing alternations alone (Silverman, 2010; Kaplan, 2011). Nonetheless, it is most likely that Korean has more phonological naturalizing alternations than has English.

---

[31] The coda cluster /lk/ is realized as [k] in a word-final or non-prevocalic position (e.g., /ma**lk**+ta/ 'clean-DECL' – [ma**k**.t'a], /il**k**+ta/ 'read-DECL'– [i**k**.t'a]). However, it is realized as [l] when followed by a [k]-initial syllable (e.g., /ma**lk**+ke/ 'clean+to be' – [ma**l**.k'e], /il**k**+kəna/ 'read+or' – [i**l**.k'əna]).

[32] Although /alm/ and /am/ are both realized as [am], the former one is often pronounced with a long vowel [aːm].

For instance, Silverman (2010) compared the amount of derived homophony created by existing neutralizing alternations in Korean to the amount of homophony that could have been derived by a small set of hand-picked hypothetical alternations in Korean. The hand-picked hypothetical neutralizations did not exist in Korean but were similar to existing neutralizations. For instance, an existing nasal lateralization neutralization in Korean alternates the sequence of nasal and liquid consonants to [lː], and the corresponding hypothetical lateral nasalization neutralization alternates the same sequence to [nː]. In the Sejong Corpus, six neutralizing alternations created 42 homophony sets among 35,907 nouns, indicating that existing neutralizing alternations created homophony less than 0.1% of the time. Table 24 shows three existing neutralizing alternations created less homophony than did hypothetical neutralizing alternations in Korean. For example, the hypothetical lateral nasalization neutralizing rule could have created 47 homophony sets in Korean whereas the existing nasal lateralization neutralizing rule created 10 homophony sets. Thus, Silverman concludes that neutralizing alternations have evolved to avoid generating too much homophony that can potentially hinder clear communication.

Kaplan (2011) built on Silverman's study and tested the amount of homophony that could have created by a large set of hypothetical alternations computationally generated by a Monte Carlo simulation. The results replicated Silverman's findings by showing that hypothetical alternations created significantly more homophony sets than did existing neutralizing alternations, echoing the observation that alternations tend not to merge too many lexical contrasts to ensure clear communication.

Table 24. The number of homophony sets created by existing neutralizing rules in Korean and by hypothetical rules (Silverman, 2010, pp.474-476).

Existing Neutralization			Hypothetical Neutralization		
Aplosivization	Stop+nasal assimilation	Nasal lateralization	Word-initial laryngeal neutralization	Nasal + stop assimilation	Lateral nasalization
15 sets	10 sets	10 sets	1912 sets	1072 sets	47 sets

To summarize this section, I have shown that Korean has many neutralizing alternations that merge a number of phonemic contrasts. Silverman's (2010) and Kaplan's (2011) work suggests that homophony avoidance pressures have influenced the development of Korean, resulting in less homophony than would be expected by chance alone. Nevertheless, the observation that Korean has an extensive list of phonological neutralizations means that Korean speakers are more likely to be exposed to and internalize neutralizing alternations in their L1 compared to English speakers, which makes Korean speakers appropriate subjects to examine whether extensive exposure to neutralizing alternations in the L1 facilitates learning a novel neutralizing alternation. To this end, I have replicated the experiment from the previous chapter with Korean speakers to see how the influence from their L1 impacts their learning relative to English speakers.

3.3 Experiment

3.3.1 Experiment Overview

The design and method of the experiment was identical to the experiment described in Chapter 2, except with Korean-speaking participants and without the No Alternation and the 50-50 conditions.[33] During the experiment, Adult Korean speakers were asked to learn singular nouns (10 stems), phonological alternations (non-neutralizing and neutralizing), and two variant prefix forms. Korean speakers first learned 10 stems (CVCV) for singular nouns. Once they successfully acquired the singular words, they learned the plural forms of the singular nouns that they had just learned. The plural forms were made by adding either the *ba-* or *ni-* prefix to the stems (e.g., *ba-*CVCV, *ni-*CVCV). The participants heard both plural forms, but the proportion of times that each plural form was presented during the prefix learning phase differed according to the Prefix Frequency condition. Participants were assigned to one of two Prefix Frequency conditions: 1) a Frequent *ba-* condition: *ba-*CVCV form appeared in two-thirds of trials, and *ni-*CVCV appeared in one-third of trials and 2) a Frequent *ni-* condition: *ni-*CVCV form was presented in two-thirds of trials, and *ba-*CVCV was presented in one-third of trials. Moreover, participants were assigned to one of two Alternation Types: a Non-Neutralizing condition or a Neutralizing condition. In the Non-Neutralizing condition, the *ni-* prefix triggered the palatalization of velar-initial stems (e.g., singular: [kimu], plural: [niʧimu]). In the Neutralizing condition, the palatalization triggered by the *ni-* prefix was neutralizing and homophony-creating (e.g., singular: [kimu], [ʧimu], plural: [niʧimu]). During the prefix test phase, participants were asked to produce the plural form.

---

[33] Korean participants were assigned to **two** Prefix Frequency conditions (a Frequent *ba-* and a Frequent *ni-* condition) and two Alternation Types (a Non-Neutralizing and a Neutralizing condition). In Chapter 2, English participants were assigned to **three** Prefix Frequency conditions (**a 50-50,** a Frequent *ba-*, a Frequent *ni-* condition) and three Alternation Type (**a No Alternation**, a Non-Neutralizing, and a Neutralizing condition). English participants closely replicated the percentage of the time that plural forms were presented in the input to their output in 8 out of 9 conditions. To preview the results, Korean speakers showed frequency matching behaviour in all four conditions. These results suggest that it is highly likely that Korean speakers would also frequency match in the 50-50 and No Alternation conditions. Accordingly, I did not test Korean speakers with the 50-50 and No Alternation conditions assuming that they would show probability matching behaviour in both conditions. During the analysis, I examine the results of two Alternation Types (a Non-Neutralizing and a Neutralizing) and two Prefix Frequencies (a Frequent *ba-* and a Frequent *ni-*) between English and Korean participants.

3.3.2 Method

In this section, I provide only a brief summary of the method because it is a replication of the experiment from Chapter 2 with Korean participants. The full details of the method can be found in Chapter 2 section 2.5.2.

3.3.2.1 Participants

Ninety-nine adult native Korean speakers (64 females, minimum age = 20, maximum age = 59, mean age = 23) completed this experiment. An additional eight participants who failed to reach the 75% accuracy requirement in the stem test phase were excluded from the analysis. Most participants were native South Kyongsang Korean speakers, but some participants spoke other varieties of Korean (Seoul Korean, North Kyongsang Korean, etc.). The experiment was conducted in a research room at Dong-A University, a conference room at Dongeui University, and my private office. Participants were recruited through flyers, advertisements on notice boards, or professors at the universities. Participants received a small amount of monetary compensation.

3.3.2.2 Materials

The materials used in the current experiment were identical to those used in the experiment in Chapter 2. For the stem learning and stem test phases, I used 10 nonce CVCV words for singular noun stems, and each stem was presented with a matching picture of a singular item. For the Non-Neutralizing condition, eight consonants, {p, b, **t, d,** k, ɡ, m, n}, were used to form singular words (e.g., [dapi] 'an egg', [timu] 'a carrot'). In the Neutralizing conditions, /t/ and /d/ were replaced with /ʧ/ and /ʤ/, respectively; thus, {p, b, **ʧ, ʤ,** k, ɡ, m, n} were used to form singular words (e.g., [**ʤ**api] 'an egg' [**ʧ**imu] 'a carrot'**)**. The production of voiceless consonants was aspirated to distinguish them from voiced consonants for English speakers and from plain and tense consonants for Korean speakers. The six consonants in the Non-Neutralizing condition, {p, b, **t, d,** m, n}, and those in the

Neutralizing condition, {p, b, **ʧ, ʤ,** m, n}, were used once in the word-initial position (ˈ**C₁**V₁C₂V₂) of singular noun stems, and the critical two consonants, /k/ and /g/, were used twice in the word-initial position. In the word-medial position (ˈC_1V_1**C_2**V_2), consonants were used a similar number of times in each condition. I used three vowels, {i, a, u}, for stems, and the vowels were used in words a similar number of times (ˈC_1**V_1**C_2**V_2**).

For the prefix learning and prefix test phases, participants learned plural forms of singular noun stems that they had already learned. The plural forms were formed by adding either the *ba-* or *ni-* prefix to the singular words (e.g., *ba*-CVCV or *ni*-CVCV). The *ba-* prefix never triggered alternation in any of the conditions, but the *ni-* prefix triggered palatalization of velar-initial stems (four out of 10 stems; e.g., singular: [kimu] 'a flower'; plural: [bakimu] and [ni**ʧ**imu] 'flowers'). Crucially, in the Neutralizing condition, two stems were exchanged for palatal-initial stems so that the palatalization would be phonologically neutralizing and homophony-creating ([timu] and [dapi] were exchanged for [ʧimu] and [ʤapi]). Hence, in the Neutralizing condition, the palatalizing alternation merged /k, ʧ/ to [ʧ], and /g, ʤ/ to [ʤ], which triggered two velar-initial stems and palatal-initial stems became homophonous in their ni+stem plural forms (e.g., singular stems: [kimu] and [ʧimu]; ni+stem plural form: [ni**ʧ**imu]). The full list of the stimuli used in the current experiment can be found in Appendix A.

3.3.2.3 Procedure

There were four phases in the experiment: a stem learning, a stem test, a prefix learning, and a prefix test phase (revisit section 2.5.2.3 for the details on each phase). Participants were introduced to 10 CVCV singular words during the stem learning phase. They were given 100 self-paced trials where the 10 singular words appeared once in each block in a random order and each block was repeated 10 times. Each singular word was presented with a matching picture of a singular item on the screen. Participants heard the singular word and were asked to verbally repeat what they had heard. In the stem test phase, participants were tested on how well they had learned the singular nouns. Participants were asked to produce the matching word into the microphone upon seeing a picture of a singular item on the screen. After participants provided their answers, an experimenter outside of the recording booth coded the answer. Participants who reached above the 75% accuracy requirement moved to the prefix learning and prefix test phases.

In the prefix learning phase, there were 60 self-paced trials (the 10 plural words were randomly ordered once within each block, repeated six times). A plural form was presented with a picture of plural items. Participants were introduced to both plural forms (*ba*-CVCV and *ni*-CVCV) of singular words that they were already familiar with, meaning they heard both plural forms for each singular word (e.g., singular: *kimu*, plural: *ba-kimu* and *nitʃimu*). The frequency of plural forms differed based on the Prefix Frequency condition (Frequent *ba*- or Frequent *ni*-) that participants were assigned to. In the prefix test phase, participants were required to produce the plural form upon seeing a picture of plural items. Once they produced the answer, an experimenter coded the answer based on the stem, the prefix form, and the alternations of the stem (if required). Participants then heard feedback (the correct plural form) and saw their score change[44] at the bottom of the screen. The frequency of each plural form produced in the feedback in the prefix test phase matched that in the prefix learning phase to which participants were assigned.

3.4 Analysis and results

3.4.1 Analysis plan and predictions

I first tested the hypothesis that Korean speakers would shift away from the alternation-triggering (and neutralization-triggering) prefix form. To test this hypothesis, I examined how often

[44] The full description of the scoring system can be found in section 2.5.2.3.

Korean speakers chose the frequent and infrequent prefix forms in their output during the prefix test phase. Then, I compared the results for Korean speakers to those for English speakers. I predicted that Korean speakers would generally match the relative frequency of prefix forms presented in the input to their output, echoing the general results for English speakers. If they avoided the alternation-triggering (and neutralization-triggering) prefix form, they would shift away from it when it was the infrequent variant in the input.

I further tested the hypothesis that Korean speakers are more successful at learning a novel neutralizing alternation that triggers homophony than are English speakers. I investigated how often Korean speakers applied the neutralizing and non-neutralizing alternations after *ni-* (correct application) and *ba-* (over-generalization errors). Then, I compared the percentage of the time in which Korean and English speakers (from the previous chapter) applied the alternations after *ni-* and *ba-*. I predicted that (unlike English speakers) Korean speakers would be successful at learning both non-neutralizing and neutralizing alternations due to the L1 effect, meaning that they would apply both alternations frequently after *ni-* and less often after *ba* -.

I analyse learners' responses during the prefix test phase. The trials in which participants used an incorrect prefix and stem or failed to answer within the 10-second time limit are excluded from the analysis. All aggregate results are analysed using a mixed effects logistic regression model (Jaeger, 2008) with lme4 packages (Bates et al., 2015) that are implemented in R (R Core Team, 2018). Models generally include the fixed effect of Prefix Frequency (Frequent *ba-* or Frequent *ni-*) and Alternation Type (Non-Neutralizing or Neutralizing), and their interaction. The different maximum random effects structures which are allowed are implemented in each model. I conduct a backward stepwise model comparison using likelihood ratio tests with *anova( )* function in R (Barr et al., 2013) to compare the model to a subset model, removing one effect at a time.

3.4.2 Selection of frequent and infrequent prefix forms

3.4.2.1 Korean speakers' selection of prefix forms

Figure 6 shows how often Korean participants in the current experiment and English participants from Chapter 2 used the frequent prefix form in their output. The experimental results for English speakers are repeated from section 2.6.2. The proportion of trials in which the frequent prefix was presented in the input (66.7%) is indicated as a dashed line to facilitate comparison.

Figure 6. Proportion of trials in which Korean (blue) and English (green; repeated from Figure 3) speakers chose the frequent prefix option in the prefix test phase. Note: Error bars show 95% confidence intervals.

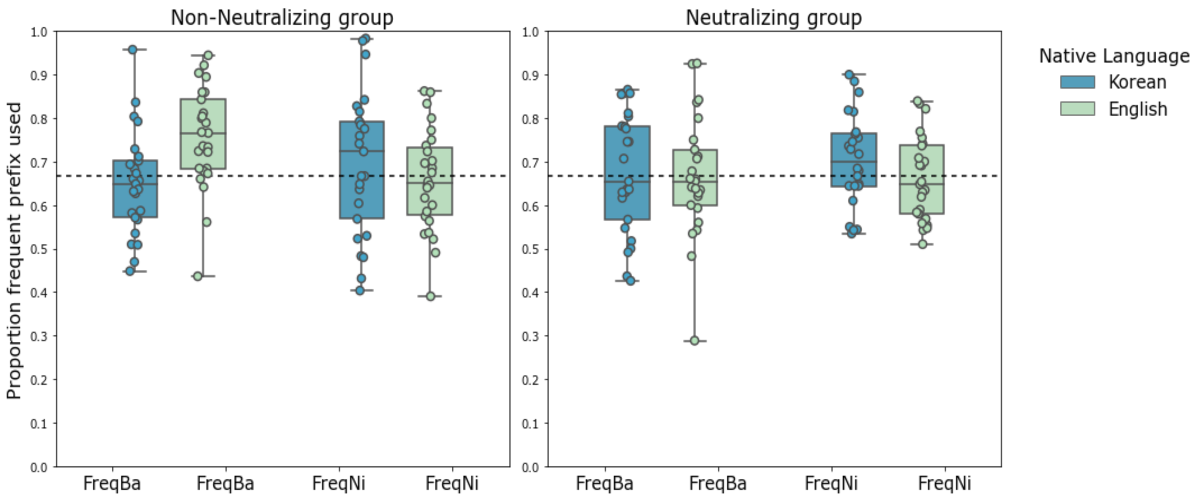

I first tested the proportion of times that Korean speakers chose to use the frequent prefix form in their output during the prefix test phase. The model had the fixed effect of Alternation Type (Non-Neutralizing or Neutralizing) and Prefix Frequency (Frequent *ba-* or Frequent *ni-*) and their interaction. The random slope of Prefix Frequency and Alternation Type by Stem and a random intercept of Subject were also included in the model. The interaction effect of Alternation Type x Prefix Frequency ($\chi^2(1) = 1.29$, p =.26), the simple effect of Prefix Frequency (Frequent *ni-*) ($\chi^2(1) = 0.99$, p =.32), and Alternation Type (Neutralizing) ($\chi^2(1) = 0.10$, p =.75) did not significantly improve the model fit; hence, they were removed from the model. Table 25 presents the summary of the final model's fixed effects. Only the intercept of the model (Frequent *ba-* Non-Neutralizing condition) was significant and positive, indicating that Korean speakers used the frequent variant significantly more often than chance level and responded similarly in all conditions (aggregate results in the Non-Neutralizing condition; Frequent *ba-*: 64.8%, Frequent *ni-*: 69.7%, in the Neutralizing condition; Frequent *ba-*: 67.4%, Frequent *ni-*: 70%). Overall, Korean speakers matched the relative frequency of variants in the input to their output in all conditions.

Table 25. Summary of the fixed effects in the final model predicting selection of the frequent prefix for Korean speakers.

Fixed effect	Estimate	Standard error	Wald *z*	*p*-value
Intercept	**0.85**	**0.08**	**10.29**	**< .001 *****

The reference group (intercept) is the Frequent *ba-* Non-Neutralizing condition, **R code for the final model**: glmer(UsedFrequentPrefix ~ + (1|Subject) + (1+PrefixFreq+AltType|Stem), data=Koreandata, family=binomial). **R code for the full model**: glmer(UsedFrequentPrefix ~ PrefixFrequency*AlternationType + (1|Subject) + (1+PrefixFreq+AltType|Stem), data=Koreandata, family=binomial)

3.4.2.2 Selection of prefix forms between Korean and English speakers

Furthermore, I compared how often Korean speakers used the frequent variant to how often English speakers used it. Here, the comparison focuses on two Alternation Types (Non-Neutralizing and Neutralizing) and two Prefix Frequency conditions (Frequent *ba-* and Frequent *ni-*).[45] The aggregate results were given in Figure 6.

The fixed effects were Alternation Type (Non-Neutralizing or Neutralizing), Prefix Frequency (Frequent *ba-* or Frequent *ni-*), Language (Korean or English), and their interactions. The model had a random intercept of Subject and a random slope of Language by Stem. The following

[45] Recall that Korean speakers were not exposed to a 50-50 and a No Alternation condition because it was assumed that they would closely match the proportion of variants in their output to what they had seen in the input.

interaction effects did not improve the model fit: Alternation Type x Prefix Frequency x Language ($\chi^2(1) = 1.14$, p =.29), Alternation Type x Prefix Frequency ($\chi^2(1) = 0.82$, p =.37), Alternation Type x Language ($\chi^2(1) = 2.42$, p =.12). Additionally, the simple effect of Alternation Type (Neutralizing) ($\chi^2(1) = 1.16$, p =.28) did not significantly improve the model fit. Hence, they were removed from the model. However, the interaction effects of Prefix Frequency x Language significantly improved the model fit, $\chi^2(1) = 6.73$, $p < .01$, and they remained in the model.

Table 26. Summary of the fixed effects in the final model predicting selection of the frequent prefix for Korean and English speakers.

Fixed effect	Estimate	Standard error	Wald *z*	*p*-value
Intercept	**0.70**	**0.09**	**7.45**	**< .001 *******
Frequent *ba-*	**0.27**	**0.13**	**2.15**	**0.03***
Korean	0.22	0.14	1.65	0.1
Frequent *ba-* & Korean	**-0.47**	**0.18**	**-2.61**	**<.01****

The reference group (intercept) is the Frequent *ni-* Non-Neutralizing condition for English speakers, **R code for the final model**: glmer(UsedFrequentPrefix ~ PrefixFrequency*Language + (1|Subject) + (1+Language|Stem), data=data, family=binomial). **R code for the full model**: glmer(UsedFrequentPrefix ~ PrefixFrequency*Language*AlternationType + (1|Subject) + (1+Language|Stem), data=data, family=binomial)

Table 26 presents the summary of the fixed effects from the final model predicting the selection of the frequent prefix form for Korean and English speakers. The intercept of the model (the Frequent *ni-* Non-Neutralizing condition of the English speakers) was significant and positive, indicating that they used *ni-* more often than chance. The significant and positive fixed effect of Frequent *ba-* indicates that English speakers chose the frequent prefix significantly more often in the Frequent *ba-* Non-Neutralizing than in the Frequent *ni-* Non-Neutralizing condition. The interaction effects of Frequent *ba-* x Korean were significant and negative, indicating that Korean speakers (unlike English speakers) did not overly use the frequent prefix more often in the Frequent *ba-*

condition than in the Frequent *ni-* condition (aggregate results in the English Frequent *ba-* Non-Neutralizing: 75.7% *ba-* vs. Korean Frequent *ba-* Non-Neutralizing: 64.8% *ba-*). Thus, this result confirms that whereas English speakers regularized their output in the Frequent *ba-* Non-Neutralizing condition, Korean speakers did not regularize in this (or any other) condition.

In addition, I compared how often Korean speakers chose the *ni-* prefix with velar-initial and non-velar-initial stems to examine whether the prefix choice was a global decision, made regardless of which stem followed. The *ni-* prefix triggered the palatalization of velar-initial stems but did not trigger the palatalization of non-velar-initial stems. Table 27 shows how often Korean speakers chose *ni-* with velar-initial and non-velar-initial stems. The model had a fixed effect of the Stem Type (velar-initial or non-velar-initial stems) and random intercepts for Subject and Stem. The fixed effect of the Stem Type (velar-initial stems) ($\chi^2(1) = 0.56$, $p = .45$) did not significantly improve the model fit, and it was removed from the model. The intercept of the model (non-velar-initial stem) was not significant, suggesting that Korean speakers used *ni-* with non-velar initial stem near chance level and responded comparably with all stem types.

I also compared the usage of *ni-* for velar-initial and non-velar-initial stems between Korean and English participants to see if there were differences between the two language groups. The model had fixed effects of the Stem Type (velar-initial or non-velar-initial stems), Language (Korean or English), and their interaction effect. The random intercepts for Subject and Stem were implemented in the model. The model shows that the interaction effects of Stem Type x Language ($\chi^2(1) = 0.05$, $p = .82$), the simple effects of Language (Korean) ($\chi^2(1) = 2.21$, $p = .14$), and Stem Type (velar-initial stems) ($\chi^2(1) = 0.58$, $p = .45$) did not significantly improve the model fit; thus, they were removed from the model. The intercept (non-velar-initial stem for English speakers) was not significant. These results show that both English and Korean speakers used *ni-* comparatively often with all stem types.

Table 27. Mean usage of *ni-* for velar-initial and non-velar-initial stems, by condition.

Alternation Type	Stem-initial consonant	Prefix Frequency	
		Frequent *ba-*	Frequent *ni-*
Non-Neutralizing	Velar	30%	72%
	Non-velar	39%	68%
Neutralizing	Velar	30%	68%
	Non-velar	35%	71%

3.4.3 Frequency of variants in the input vs. output

In addition, I investigated whether individual participants regularized their output with the frequent prefix form. If there was a tendency to regularize their output with *ba-*, individual participants would have been likely to use *ni-* less often than what they had seen in the input, compensating for the output by overly using *ba-* instead. In other words, they were more likely to regularize their output with *ba-* when it was frequent in the input but less likely to do so with *ni-* when it was frequent in the input. Table 28 below shows the number of Korean participants who regularized their output with either *ba-* or *ni-*.

Following Culbertson et al. (2012), the individual results were analysed using a one-sample sign test. I used a Bonferroni-adjusted alpha level of .0125 (i.e., .05/4) to correct for multiple comparisons. For Korean speakers, there was no significant preference to overuse either prefix form in all conditions, supporting the aggregate results that Korean speakers did not shift away from either prefix form in all conditions. In contrast, English speakers showed a significant preference to overuse *ba-* in the Frequent *ba-* Non-Neutralizing condition, indicating that they shifted away from *ni-* in the condition (see Chapter 2, Table 8).

Table 28. Number of Korean speakers who shifted towards *ba-* and *ni-*.

Condition		Shifted towards *ba-*	Shifted towards *ni-*	*p*-value (sign test)
Non-Neutralizing	Frequent *ba-*	10 / 25 (40%)	15 / 25 (60%)	0.42
	Frequent *ni-*	12/ 25 (48%)	13 / 25 (52%)	1
Neutralizing	Frequent *ba-*	12 / 25 (48%)	13 / 25 (52%)	1
	Frequent *ni-*	10 / 24 (42%)	14 / 24 (58%)	0.54

In summary, the aggregate results showed that, unlike English speakers who shifted away from *ni-* in the Frequent *ba-* Non-Neutralizing condition, Korean speakers closely matched the probability of prefix forms presented in the input to their output in all conditions. This finding is supported by the individual results showing that Korean speakers did not regularize their output with either prefix form.

3.4.4 Application of the palatalizing rule

In this section, I examine how often participants applied the palatalizing rule to velar-initial stems after *ni-* and *ba-*. During the training, stem-initial /k/ and /g/ were palatalized when preceded by *ni-*, meaning that the palatalization of velar-initial stems after *ni-* was a correct case whereas palatalization after *ba-* was an over-generalization error case. I divide the data into two sub-datasets and analyse each dataset separately. The first dataset contains the cases where participants palatalized velar-initial stems after *ni-* (correct cases of palatalization), and the second dataset includes the cases where they palatalized velar-initial stems after *ba-* (incorrect over-generalization error cases of palatalization). I first examine how often Korean speakers palatalized velar-initial stems after *ni-* and *ba-*. I then compare the palatalization of velar-initial stems between Korean and English speakers.

Lastly, I compare if the palatalization of velar-initial stems differed based on whether or not it triggered homophony.

3.4.4.1 Korean speakers' application of the palatalizing rule

Figure 7 shows the percentage of the time in which Korean and English speakers palatalized velar-initial stems after each prefix according to the Alternation Types and Prefix Frequency conditions. I first investigate the cases in which Korean speakers palatalized velar-initial stems after *ni-* and *ba-*. Both models included fixed effects of Alternation Type (Non-Neutralizing or Neutralizing), Prefix Frequency (Frequent *ba-* or Frequent *ni-*), and their interactions. The model that analysed the palatalization after *ni-* had random intercepts for Subject and Stem and the model that analysed the palatalization after *ba-* had a random intercept for Subject as the model failed to converge. Table 29 illustrates the summary of the fixed effects from the final model predicting the palatalization of velar-initial stems after *ni-* for Korean speakers.

A likelihood ratio test showed that Alternation Type x Prefix Frequency interaction effects ($\chi^2(1) = 0.91$, $p = .34$) and Alternation Type simple effect ($\chi^2(1) = 1.61$, $p = .21$) did not significantly improve the model fit; hence, they were removed from the model. In contrast, the fixed effect of Prefix Frequency (Frequent *ba-*) was significant ($\chi^2(1) = 10.18$, $p < .01$) and remained in the model. The intercept of the model was significant and positive, suggesting that Korean speakers palatalized velar-initial stems after *ni-* significantly more often than chance in the Frequent *ni-* Non-Neutralizing condition. The simple effect of Frequent *ba-* was significant and negative and no interaction effects including Frequent *ba-* were significant, indicating that the palatalization of velar-initial stems after *ni-* was significantly less frequent in the Frequent *ba-* condition than in the Frequent *ni-* condition (aggregate results in the Frequent *ni-* condition; Non-Neutralizing: 64.8%, Neutralizing: 64.1% vs. the Frequent *ba-* condition; Non-Neutralizing: 38.7%, Neutralizing: 52.8%).

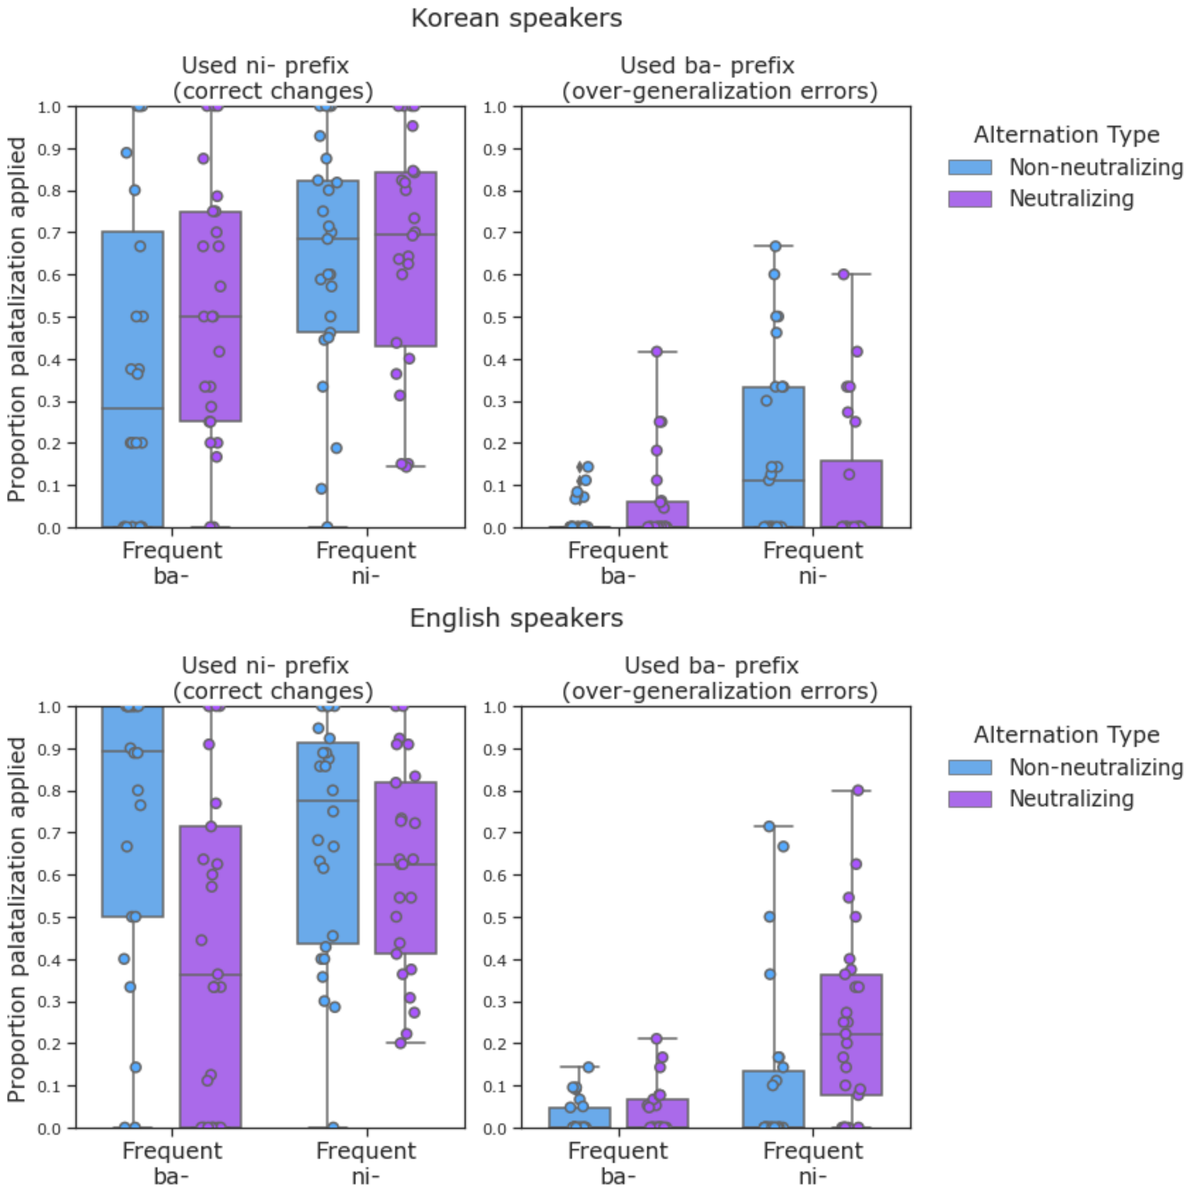

Figure 7. Proportion of times that Korean (top) and English (bottom; repeated from Chapter 2, Figure 4) speakers palatalized velar-initial stems after *ni-* (left) and *ba-* (right) in the Neutralizing condition.

Table 30 shows the summary of the fixed effects for the final model predicting the palatalization after the *ba-* prefix for Korean speakers. The interaction effect of Alternation Type x Prefix Frequency ($\chi^2(1) = 3.24$, p = .07) and the simple effect of Alternation Type ($\chi^2(1) = .08$, p = .78) did not significantly improve the model fit and they were removed from the model. The simple effect of Prefix Frequency significantly improved the model fit ($\chi^2(1) = 12.69$, p < .001), and it remained in the model. The intercept of the model (the Frequent *ni-* Non-Neutralizing condition) was significant and negative, indicating that palatalization after *ba-* was significantly less frequent than chance level in the condition. The significant and negative fixed effect of Frequent *ba-* suggests that palatalization after *ba-* was even less frequent in the Frequent *ba-* condition (aggregate results in the Frequent *ni-* condition: the Non-Neutralizing: 17.5%, the Neutralizing: 11% vs. the Frequent *ba-* condition: the Non-Neutralizing: 2.2%, the Neutralizing: 4.7%).

Table 29. Summary of the fixed effects of the final model predicting application of palatalization rule after *ni-* (correct application) for Korean speakers.

Fixed effect	Estimate	Standard error	Wald z	p-value
Intercept	**0.87**	**0.3**	**2.87**	**< .01****
Frequent *ba-*	**-1.21**	**0.38**	**-3.2**	**< .01****

The reference group (intercept) is the Frequent *ni-* Non-Neutralizing condition of Korean speakers, **R code for the final model**: glmer(AppliedPalatalization ~ PrefixFrequency + (1|Subject) + (1|Stem), data=Koreandata_ni, family=binomial). **R code for the full model**: glmer(UsedFrequentPrefix ~ PrefixFrequency* AlternationType + (1|Subject) + (1|Stem), data=Koreandata_ni, family=binomial)

Table 30. Summary of the fixed effects in the final model predicting application of palatalization rule after *ba-* (over-generalization error) for Korean speakers.

Fixed effect	Estimate	Standard error	Wald z	p-value
Intercept	**-2.54**	**0.39**	**-6.58**	**< .001 *****
Frequent *ba-*	**-1.79**	**0.48**	**-3.71**	**< .001 *****

The reference group (intercept) is the Frequent *ni-* Non-Neutralizing condition of Korean speakers, **R code for the final model**: glmer(AppliedPalatalization ~ PrefixFrequency + (1|Subject), data=Koreandata_ba, family=binomial). **R code for the full model**: glmer(UsedFrequentPrefix ~ PrefixFrequency* AlternationType + (1|Subject), data=Koreandata_ba family=binomial)

Furthermore, I examined how often Korean speakers palatalized non-velar-initial stems (where palatalization should *not* be applied either after *ni-* or *ba-*). There were less than 1% of incorrect cases where Korean speakers palatalized non-velar-initial stems after *ba-* or *ni-* across all conditions, suggesting that they rarely over-applied palatalization to non-velar-initial stems.

To sum up, Korean speakers applied palatalization to velar-initial stems after *ni-* and *ba-* more often in the Frequent *ni-* condition than in the Frequent *ba-* condition. Crucially, there was no difference between how often they palatalized velar-initial stems in the Non-Neutralizing and in the Neutralizing conditions, suggesting that they applied both non-neutralizing and neutralizing alternations to stems relatively often. They rarely over-palatalized non-velar-initial stems.

3.4.4.2 Application of the palatalizing rule between Korean and English speakers

I compare how often Korean and English participants palatalized velar-initial stems. The model included fixed effects for Alternation Type (Non-Neutralizing or Neutralizing), Prefix Frequency (Frequent *ni-* or Frequent *ba-*), Language (Korean or English), and their interactions. The model analysing the palatalization after *ni-* had the random intercepts of Subject and Stem and the random slope of Alternation Type by Stem. The model that analysed the palatalization after *ba-* had the random intercepts of Subject and Stem. The initial models failed to converge; thus, BOBYQA optimization was implanted in the models to increase the maximum number of iterations (maxfun = 2e5; Powell, 2009), which resolved the convergence failure.

Table 31 below shows the summary of the fixed effects from the final model predicting the palatalization after *ni-* for English and Korean speakers. A likelihood ratio test showed that the interaction effects of Alternation Type x Prefix Frequency x Language significantly improved the model fit ($\chi^2(1) = 4.16$, $p = .04$); hence, they remained in the model.

Table 31. Summary of the fixed effects of the final model predicting palatalization after *ni-* (correct application) for Korean and English speakers.

Fixed effect	Estimate	Standard error	Wald *z*	*p*-value
Intercept	**1.23**	**0.39**	**3.17**	**<.01****
Neutralizing	-0.58	0.51	-1.14	.26
Frequent *ba-*	0.31	0.54	0.57	.57
Korean	-0.44	0.48	-0.91	.36
Neutralizing & Frequent *ba-*	-1.4	0.74	-1.88	.06
Neutralizing & Korean	0.72	0.69	1.05	.29
Frequent *ba-* & Korean	**-1.91**	**0.75**	**-2.54**	**.01***
Neutralizing & Frequent *ba-* & Korean	**2.14**	**1.04**	**2.05**	**<.05***

The reference group (intercept) is the Frequent *ni-* Non-Neutralizing condition of English speakers, **R code for the model:** glmer(AppliedPalatalization ~ AlternationType *PrefixFrequency*Language + (1|Subject) + (1+AltType|Stem), data=ni-data_KoreanEnglish, family=binominal, control=glmerControl(optimizer="bobyqa", optCtrl=list(2e5)))

The significant and positive intercept (representing the Frequent *ni-* Non-Neutralizing condition of English speakers) indicates that English speakers palatalized velar-initial stems after *ni-* significantly more often than chance in the Frequent *ni-* Non-Neutralizing condition. The interaction effect of Frequent *ba-* x Korean was negative and significant, indicating that Korean speakers palatalized velar-initial stems after *ni-* significantly less often in the Frequent *ba-* Non-Neutralizing condition than English speakers did in the Frequent *ni-* Non-Neutralizing condition (aggregate results in the English Frequent *ni-* Non-Neutralizing condition: 71.4%, Frequent *ba-* Non-Neutralizing condition: 73.7% vs. the Korean Frequent *ni-* Non-Neutralizing condition: 64.8%, Frequent *ba-* Non-Neutralizing condition: 38.7%). Conversely, the interaction effects of Neutralizing x Frequent *ba-* x Korean were significant and positive, which shows that Korean speakers palatalized the velar-initial stems after *ni-* in the Frequent *ba-* Neutralizing condition similarly often as English speakers did in the Frequent *ni-* Non-Neutralizing condition (aggregate results in the Korean Frequent *ba-*

Neutralizing: 52.8%). No other effects were significant. These results suggest Korean speakers palatalized velar-initial stems after *ni-* significantly less often than English speakers when the *ba-* prefix was frequent and the alternation was non-neutralizing. However, when *ba-* was frequent and the alternation was neutralizing, Korean and English speakers applied the palatalizing rule comparably often to velar-initial stems after *ni-*.

Table 32 shows the summary of the fixed effects of the final model predicting the palatalization after the *ba-* prefix for English and Korean speakers. A likelihood ratio test indicated that only the fixed effect of Prefix Frequency (Frequent *ba-*) significantly improved the model fit ($\chi^2(1) = 40.42$, $p < .001$); hence, all interaction effects and other simple effects were removed from the model.[46] The significant and negative intercept indicates that English speakers palatalized velar-initial stems after *ba-* significantly less often than chance in the Frequent *ni-* Non-Neutralizing condition. Moreover, the simple effect of Frequent *ba-* was significant and negative, and no other interaction effects including Frequent *ba-* were significant, suggesting that both English and Korean speakers palatalized the velar-initial stems significantly less often after *ba-* when *ba-* was infrequent in the input across the Non-Neutralizing and Neutralizing conditions.

To summarize, there was no difference between how often Korean and English speakers palatalized velar-initial stems, except for one case where Korean speakers palatalized velar-initial stems after *ni-* in the Frequent *ba-* Non-Neutralizing condition significantly less often than English speakers did in the Frequent *ni-* Non-Neutralizing condition.

[46] Alternation Type x Prefix Frequency x Language ($\chi^2(1) = 3.84$, $p = .05002$), Alternation Type x Prefix Frequency ($\chi^2(1) = .58$, $p = .45$), Prefix Frequency x Language ($\chi^2(1) = 0.59$, $p = .44$), and Alternation Type x Language ($\chi^2(1) = 3.58$, $p = .058$), Language ($\chi^2(1) = 0.36$, $p = .55$) and Alternation Type ($\chi^2(1) = 3.07$, $p = .08$).

Table 32. Summary of the fixed effects in the final model predicting palatalization after *ba-* (over-generalization error cases) for Korean and English speakers.

Fixed effect	Estimate	Standard error	Wald *z*	*p*-value
Intercept	**-2.16**	**0.22**	**-9.88**	**<.001*** **
Frequent *ba-*	**-1.97**	**0.29**	**-6.67**	**<.001*** **

The reference group (intercept) is the Frequent *ni-* Non-Neutralizing condition of English speakers, **R code for the model:** glmer(AppliedPalatalization ~ PrefixFrequency + (1|Subject) + (1+AltType|Stem), data=ni-data_KoreanEnglish, family=binominal, control=glmerControl(optimizer="bobyqa", optCtrl=list(2e5))) **R code for the full model**: glmer(AppliedPalatalization ~ AlternationType *PrefixFrequency*Language + (1|Subject) + (1+AltType|Stem), data=ni-data_KoreanEnglish, family=binominal, control=glmerControl(optimizer="bobyqa", optCtrl=list(2e5)))

3.4.4.3 Palatalization of Homophony stem types

Finally, I take an even closer look at palatalization to examine if participants' application of palatalization differed based on whether or not it created homophony. To review, in the Neutralizing condition, the *ni-* prefix triggered a neutralizing alternation. Two velar-initial stems and two palatal-initial stems became homophony in their ni+stem plural form (e.g., [kimu], [ʧimu] → [niʧimu]; [gapi], [ʤapi] → [niʤapi]). The *ni-* prefix also triggered palatalization of the other two velar-initial stems and merged distinct phonemes, but these stems did not create homophony ([kuta] → [niʧuta], [gaku] → [niʤaku]). I compared the percentage of times that participants palatalized these two groups of stems, which I correspondingly refer to as Homophony and No Homophony stems (see Figure 8).

I first examine how often Korean speakers palatalized Homophony and No Homophony stems after *ni-*. Models had fixed effects for Homophony (Homophony or No Homophony stems) and Prefix Frequency (Frequent *ba-* or Frequent *ni-*), and their interaction. The model predicting palatalization after *ni-* had random intercepts of Stem and Subject and the model predicting palatalization after *ba-* had a random intercept of Stem. Table 33 presents the summary of the fixed effects from the final model the palatalization of Homophony and No Homophony stems after *ni-* for Korean speakers.

Figure 8. Proportion of times that Korean and English (repeated from Chapter 2, Figure 5) speakers palatalized velar-initial stems that became homophony and that did not after *ni-* (left) and *ba-* (right) in the Neutralizing condition.

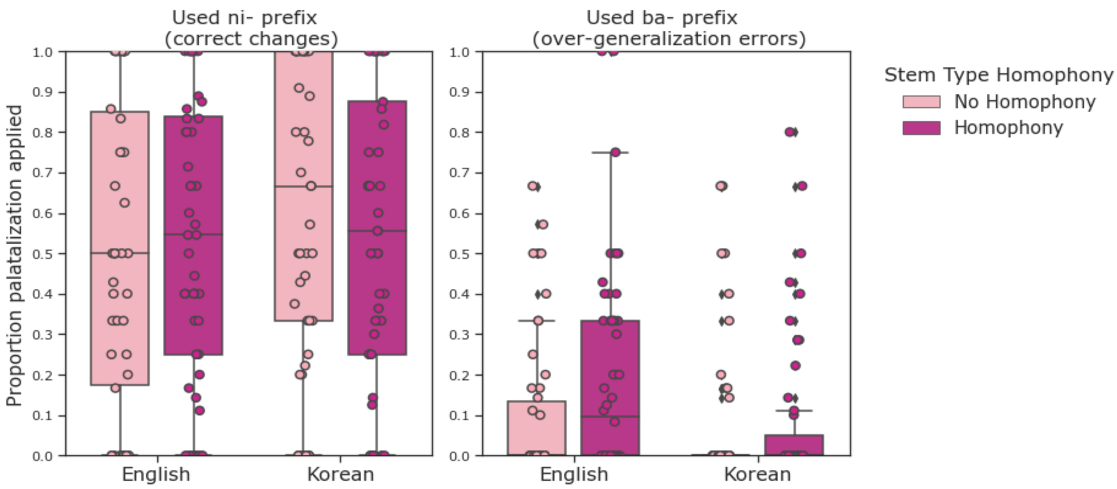

Table 33. Summary of the fixed effects in the final model predicting palatalization of Homophony and No Homophony stems after *ni-* (correct cases) for Korean speakers.

Fixed effect	Estimate	Standard error	Wald *z*	*p*-value
Intercept	0.1	0.17	0.61	.54
Frequent *ni-*	**0.48**	**0.19**	**2.48**	**<.05***

The reference group (intercept) is the Frequent *ba-* No Homophony, **R code for the final model:** glmer(AppliedPalatalization ~ PrefixFrequency + (1|Stem), data=ni_Koreandata, family=binomial), **R code for the full model:** glmer(AppliedPalatalization ~ PrefixFrequency*Homophony + (1|Stem), data=data, family=binomial).

A likelihood ratio test showed that including the interaction effect of Prefix Frequency x Homophony ($\chi^2(1) = 0.08$, p = .78) and the fixed effect of Homophony ($\chi^2(1) = 4e\text{-}04$, p = .98) did not improve the model fit and they were excluded from the model. However, the fixed effect of Prefix Frequency (Frequent *ni-*) was significant, according to both the likelihood ratio test ($\chi^2(1) = 6.15$, $p <.05$) and Wald z test, indicating that Korean speakers palatalized all stems significantly more often after *ni-* when *ni-* was the frequent variant in the input. No other effect was significant. Hence, Korean speakers applied the neutralizing alternation comparably often to stems that triggered homophony and to those that did not.

Table 34. Summary of the fixed effects in the final model predicting palatalization of Homophony and No Homophony stems after *ba-* (over-generalization error cases) for Korean speakers.

Fixed effect	Estimate	Standard error	Wald *z*	*p*-value
Intercept	**-3.08**	**0.31**	**-9.86**	**<.001*** **
Frequent *ni-*	**0.95**	**0.35**	**2.73**	**<.01** **

The reference group (intercept) is the Frequent *ba-* No Homophony, **R code for the final model:** glmer(AppliedPalatalization ~ PrefixFrequency + (1|Stem), data=ba_Koreandata, family=binomial), **R code for the full model:** glmer(AppliedPalatalization ~ PrefixFrequency*Homophony+ (1|Stem), data= ba_Koreandata, family=binomial).

Table 34 shows the summary of the fixed effects from the final model predicting the palatalization of Homophony and No Homophony stems after *ba-* for Korean speakers. A likelihood ratio test shows that the interaction effects of Homophony x Prefix Frequency ($\chi^2(1) = 1.2$, $p = .27$), and the fixed effect of Homophony ($\chi^2(1) = 2.3$, $p = .13$) did not significantly improve the model fit and they were removed from the model. The fixed effect of Prefix Frequency ($\chi^2(1) = 7.16$, $p <.01$) remained in the model as it significantly improved the model's fit. The significant and negative intercept shows that Korean speakers palatalized No Homophony stems after *ba-* significantly less often than chance in the Frequent *ba-* condition. The fixed effect of the Frequent *ni-* was significant and positive, indicating that the percentage of times that they palatalized stems after *ba-* was significantly more frequent in the Frequent *ni*- condition than in the Frequent *ba*- condition. Hence, these results show that whether or not the palatalization created homophony did not significantly affect how often Korean speakers applied the palatalization to velar-initial stems.

I then compared how often Korean and English speakers palatalized two groups of stems. Models had the fixed effects of Homophony (Homophony or No Homophony stems), Language (English or Korean) and their interaction. The model that predicted palatalization after *ni*- had the random intercepts for Subject and Stem, and the model that predicted palatalization after *ba*- had a random intercept for Subject.

Table 35. Summary of the fixed effects in the final model predicting palatalization of Homophony and No Homophony stems after *ni*- (correct cases) for Korean and English speakers.

Fixed effect	**Estimate**	**Standard error**	**Wald *z***	***p*-value**
Intercept	0.34	0.22	1.53	.13

The reference group (intercept) is the No Homophony English speakers, **R code for the final model:** glmer(AppliedPalatalization ~ + (1|Subject) + (1|Stem), data=ni-KoreanEnglishdata, family=binomial), **R code for the full model:** glmer(AppliedPalatalization ~ Homophony*Language+ (1|Subject) + (1|Stem), data= ni-KoreanEnglishdata, family=binomial).

Table 35 presents the summary of the fixed effects from the final model predicting palatalization of Homophony and No Homophony stems after *ni*- for Korean and English speakers. The interaction effects of Homophony x Language ($\chi^2(1) = 0.15$, p = .28), the fixed effect of Homophony ($\chi^2(1) = 0.25$, p = .61), and Language ($\chi^2(1) = 1.09$, p = .3) did not affect the model's fit; hence, they were not included in the model (see Table 35). The intercept did not reach significance. These results indicate that both English and Korean speakers palatalized stems that triggered homophony and those that did not similarly often after *ni*-.

Table 36 shows the summary of the fixed effects from the final model predicting palatalization of Homophony and No Homophony stems after *ba-* for Korean and English speakers. The interaction effect of Homophony x Language ($\chi^2(1) = 0.69$, $p = .41$) did not significantly improve the model fit and was removed from the model. However, fixed effects of Homophony ($\chi^2(1) = 12.36$, $p < .001$) and Language ($\chi^2(1) = 4.77$, $p = .03$) significantly improved the model fit and remained in the model. The intercept reached significance. The fixed effect of Homophony (Homophony stem) was significant and positive and no interaction effects involving Homophony stem were significant, suggesting that English and Korean speakers palatalized stems that created homophony more frequently than stems that did not after the *ba-* prefix. The significant and negative fixed effect of Korean indicates that they generally palatalized stems after *ba-* less often than English speakers did. These results suggest that overall English and Korean speakers made more over-generalization errors when the palatalization triggered homophony than when it did not trigger homophony. The results also indicate that Korean speakers generally made fewer over-generalization errors than English speakers.

Table 36. Summary of the fixed effects in the final model predicting palatalization of Homophony and No Homophony stems after *ba-* (over-generalization error cases) for Korean and English speakers.

Fixed effect	Estimate	Standard error	Wald *z*	*p*-value
Intercept	**-3.14**	**0.37**	**-8.53**	**<.001****
Homophony stem	**0.86**	**0.25**	**3.44**	**<.001****
Korean	**-1.02**	**0.48**	**-2.13**	**.03***

The reference group (intercept) is the No Homophony English Speakers, **R code for the final model:** glmer(AppliedPalatalization ~ Homophony+Language + (1|Subject), data=ba_KoreanEnglishdata, family=binomial), **R code for the full model:** glmer(AppliedPalatalization ~ Homophony*Language + (1|Subject), data=ba_KoreanEnglishdata, family=binomial).

3.4.5 Debriefing analysis

After terminating the experiment, an experimenter debriefed participants by asking them to describe their thought processes during the experiment and their response strategies during the test phases. An experimenter recorded the comments, and I later coded them into five categories based on the level of detail that participants were able to express about the phonological rules that they had learned. The responses were categorized using an ordinal scale from 1 (the least detailed) to 5 (the most detailed). Table 37 shows the number of responses from Korean and English participants in each category, along with the criteria used to code the responses. The number of responses from English speakers was replicated from Chapter 2, Table 14.

I analysed the distribution of the comments along the scale to examine how well participants learned and internalized the phonological alternation presented during the training. The analysis was focused on comparing the number of responses in each category between Non-Neutralizing and Neutralizing conditions for Korean speakers. In the Non-Neutralizing condition, Korean speakers provided more explicit information about the alternation (median = 3) than did those in the Neutralizing condition (median = 2). Specifically, in the Non-Neutralizing condition, 16 out of 49 (32.7%) participants provided detailed explanations of the alternation by mentioning some specific information about the target or the trigger of the alternation (level 4 and above on the scale) while 11 out of 47 (23.4%) participants did so in the Neutralizing condition. However, a Mann-Whitney U test did not find a significant difference between the scale of comments between the Non-Neutralizing and the Neutralizing conditions ($U = 1206.5$, $p = 0.67$), suggesting that Korean speakers in both conditions explained the alternation with a similar level of detail. Thus, the debriefing analysis supports the conclusion that Korean speakers learned the neutralizing alternation that can create homophony as well as the non-neutralizing alternation. Unlike Korean speakers, the debriefing results for English speakers showed that the scale of comments in the Neutralizing condition was significantly lower than that of the Non-Neutralizing condition, implying that they were poorer at learning the neutralizing rule that creates homophony than the non-neutralizing rule (see section 2.6.5 for the full report on debriefing analysis for English participants).

Table 37. Breakdown of Korean and English speakers' comments during debriefing along with coding criteria.

Scale value	Coding criteria		Non-Neutralizing condition			Neutralizing condition		
			Fre-quent *ba-*	Fre-quent *ni-*	**total**	Fre-quent *ba-*	Fre-quent *ni-*	**total**
5.	Explicitly mentioned at least one of the phonological changes (k→ʧ or g→ʤ).	Korean	6	5	**11**	8	3	**11**
		English	4	2	**6**	0	4	**4**
4.	Mentioned that the *ba-* prefix does not trigger an alternation but the *ni-* prefix triggers an alternation without describing the explicit phonological changes.	Korean	2	3	**5**	0	0	**0**
		English	10	11	**21**	4	1	**5**
3.	Mentioned that noun stems are changed, but nothing about the specific sounds or triggers involved	Korean	4	5	**9**	2	4	**6**
		English	1	4	**5**	6	9	**15**

2.	Mentioned two different options for plural forms but did not specifically mention that a change occurred	Korean	5	10	**15**	12	14	**26**
		English	7	6	**13**	10	7	**17**
1.	Did not mention any phonological changes or two options for plural forms.	Korean	7	2	**9**	2	2	**4**
		English	3	3	**6**	5	3	**8**

3.5 Discussion

In this chapter, I tested adult native Korean speakers to examine whether the more frequent exposure to neutralizations in their L1 would facilitate learning a novel neutralizing alternation. During the experiment, Korean speakers were required to learn a probabilistic prefix pattern in which one of the prefix forms triggered alternations in an artificial language learning paradigm. The results indicate that Korean speakers replicated the relative frequency of prefix forms presented in the input to their responses in all conditions; they did not shift away from the prefix form that triggered alternations in any of the conditions. Moreover, unlike the English speakers in Chapter 2, Korean speakers applied the neutralizing alternation to velar-initial stems as often as the non-neutralizing alternation. Both English and Korean speakers applied the neutralizing alternation similarly often to stems that created homophony and those that did not after *ni-*. Korean speakers generally made fewer over-generalization errors compared to English speakers but both speakers made more over-generalization errors when stems triggered homophony than when they did not. The following sections provide implications of these findings.

3.5.1 L1 effect on learning a novel neutralizing alternation

I found that Korean speakers were equally successful at learning the neutralizing alternation and the non-neutralizing alternation whereas English speakers were poorer at learning the neutralizing alternation than the non-neutralizing alternation. These different learning results cannot be caused by the design of the experiment as the experiment was identical for both groups. Furthermore, the progressive palatalization presented in the experiment (k→ʧ/i__, ɡ→ʤ/i__) is not found in either English or Korean. Thus, the nature of the alternation being learned also cannot explain the difference between the two groups of speakers. Instead, it appears that Korean speakers find it easier to learn a novel neutralizing alternation relative to English speakers; presumably, this is due to the extensive exposure to neutralization that they encounter in Korean.

The idea that previous exposure to the L1 affects learning novel patterns is closely linked to the findings of Martin and Culbertson's (2020) study (revisit section 3.1.2 for a detailed explanation) which shows that the frequently appearing patterns in the L1 can affect learners' judgement on novel patterns. In their study, Kîîtharaka speakers who are extensively exposed to prefix forms than suffix forms found the sequences of nonce syllables that resemble prefix forms more similar to each other than the ones similar to suffix forms (e.g., nonce syllable: ta, ko – a prefix-like form: **mo,** ta, ko vs. a suffix-like form: ta, ko, **mo**). Martin and Culbertson further found that this effect can apply to learning non-linguistic patterns. When the learners were exposed to sequences of shapes (e.g., a star, an arrow, a square, etc.), they also found that the sequences of shapes that are like prefix forms (e.g., a star, a square – **an arrow,** a star, a square) more similar to each other than the sequences that resemble suffix forms (e.g., a star , a square –a star, a square, **an arrow**). However, English speakers who are more frequently exposed to suffix forms showed the opposite results. These findings imply that when people are previously more frequently exposed to a certain pattern and system, their cognitive system may favour the familiar pattern and system and facilitate learning them.

The experimental results also show the effect of both L1 and learners' general tendencies on language learning. Learners' frequency-matching behaviour is a universal tendency that is domain-general (not specific to a language). This general tendency to frequency match has been found in the experimental results for English and Korean speakers who mostly matched the relative frequency of prefix forms in the input to their output. Moreover, the tendency to disfavour homophony can be universal and not limited to a specific language. It is likely that children initially disfavour homophonous patterns and would avoid creating homophony. Children prefer a one-to-one relationship where a label exclusively refers to one item (Markman & Wachtel, 1988); however, in homophony, a label refers to multiple items. The strength of the bias against homophony can be changed depending on learners' exposure to their L1. As children grow up, they may encounter much evidence of neutralization and homophony in their L1, which can decline the strength of the homophony avoidance bias. In this chapter, Korean-speaking adults learned the novel neutralizing alternation as successfully as the non-neutralizing alternation, which is likely to be affected by their frequent exposure to neutralizations and homophony in Korean. In Chapter 2, however, English-speaking adults were poor at learning the neutralizing alternation, indicating that, without abundant exposure to neutralizations and homophony in their L1, the strength of their bias against homophony is likely to remain strong. The different learning results of the neutralizing alternation between Korean and English speakers suggest the L1 effect on learning a biased pattern. In the next chapter, I will model the difference between English and Korean speakers in terms of the reduction in this homophony avoidance tendency due to the L1 effect.

Additionally, the finding that Korean speakers applied the neutralizing alternation as often as the non-neutralizing one could also have implications for newly emerging words in Korean. The number of existing neutralizing alternations is likely to remain in Korean. The existing neutralizing alternations are also likely to be productively applied to new words. For example, in a newly emerged word, /gas.sɛŋ/ 갓생 'diligent and hard-working life' – [ga**t**˺.**s'**ɛŋ], Korean speakers apply the coda neutralization that merges two underlyingly distinct phonemes, /s/ and /t/ to [t˺] in the coda of the first syllable. They also apply post-obstruent tensing rule that merges contrasting phonemes /s,s'/ to [s'] in the onset of the second syllable. Similarly, they also apply the post-obstruent tensing rule that triggers phonological neutralization in the onset of the second (/ʧ, ʧ'/ →[ʧ']) and the third syllable (/s, s'/ →[s']) of the new word /sik.ʧip.sa/ 식집사 'plant mom' – [sik˺.**ʧ'**ip.**s'**a]. The nasal assimilation merges underlyingly distinct phonemes /k/ and /ŋ/ to [ŋ] in the coda of the first syllable in the new word /mək.no.mɛ/ 먹노매 'catch the critical moment' – [mʌ**ŋ**.no.mɛ]. These observations suggest that Korean speakers are likely to continue applying existing neutralizing alternations to words that emerge in Korean, which implies that the number of existing neutralizing alternations is likely to remain in Korean.

3.5.2 Choice of morphological constructions

The aggregate and individual results show that Korean speakers matched the relative frequency of prefix forms presented in the input to their output in all Alternation Type and Prefix Frequency conditions; they did not shift away from the infrequent variant that triggered alternations in their output. In Chapter 2, I also found that English speakers generally matched the relative frequency of prefix forms in the input to their output except for one condition. When the infrequent prefix form triggered the non-neutralizing alternation, English speakers shifted away from the infrequent prefix form and regularized their output with the frequent prefix form. In this chapter, I have tested Korean speakers in 4 conditions, and they showed frequency-matching behaviour in all of the conditions. This indicates participants in 12 out of 13 conditions (4 conditions for Korean speakers and 9 conditions for English speakers) matched the input frequency to their output. This finding implies the possibility that the regularization found in the Frequent *ba-* Non-Neutralizing condition for English speakers could have been a sampling error.

The probability-matching behaviour found in Korean speakers is consistent with previous studies showing adult learners' tendency to replicate the relative probability of variants presented in the input to their output (probabilistic artificial languages: Hudson Kam & Newport, 2005, 2009; Austin, 2010; Schuler, 2017; probabilistic native language patterns: Ernestus & Baayen, 2003; Hayes et al., 2009; probabilistic non-linguistic patterns: Gardner, 1957; Weir, 1972).

Furthermore, the findings that adult Korean speakers did not avoid using the construction that triggered alternation imply the different learning results between children and adults. Do (2018) showed that child Korean speakers used alternative morphological structures or added morphemes to avoid alternating base forms that are different from the correct construction used by adult Korean speakers (revisit section 2.2 for a detailed description of the study). Unlike Korean children, the adult speakers in Do's study and the current study did not avoid using the morphological construction that triggered a phonological alternation.

One possible explanation for the diverging result between child and adult speakers is that adult speakers have more experience with alternations in languages than child speakers, which could lead them to expect more alternations in languages (see Reyes et al., 2017 for the relationship between exposure time and learnability). Another factor that could cause different results is the different designs of the experiments. Do's study tested Korean children's knowledge of their L1, and the current study examined Korean adult's acquisition of an artificial language. Adult learners can be more conscious of the fact that they are learning a novel language and focus more on correctly matching the input. For future work, it would be worth replicating the current experiment with Korean children.

3.6 Conclusion

The first main conclusion of this study is that Korean speakers learned the non-neutralizing alternation as well as the neutralizing alternation. In the previous chapter, I showed that English speakers who are less frequently exposed to alternations that are neutralizing and homophony-creating in their L1 were poorer at learning the neutralizing alternation than the non-neutralizing alternation. These different learning results of neutralizing alternations imply that excessive exposure to neutralizations in Korean could facilitate learning a novel neutralizing alternation presented in the experiment. This indicates that the L1 effect can affect learners to be less biased against homophony-creating patterns. However, I found that Korean speakers did not avoid using the prefix that triggered phonological alternation in their responses. A potential limitation of this study is that it is possible that Korean and English speakers have different phonological grammars coming into the experiment and the study cannot account for the possible differences in their phonological grammars.[47]

[47] I would like to thank Adam Albright for discussing this particular point with me during the poster session and conversations at CreteLing (2022).

CHAPTER 4

Homophony avoidance in learning: A discounted input approach

4.1 Introduction

In the experiments described in Chapters 2 and 3, I examined whether learners' choice of probabilistic prefix form in their output is affected by the relative frequency of prefix forms in the input and by the fact that one of the prefix forms triggered a phonological alternation. In Chapter 2, the experiment results showed that English speakers generally replicated the relative frequency of prefix forms in the input to their output; however, they shifted away from the infrequent variant when it triggered a non-neutralizing alternation. Furthermore, English speakers were less successful at learning the neutralizing alternation that can create homophony than the non-neutralizing alternation in all Prefix Frequency conditions. In Chapter 3, Korean speakers closely matched the frequency of the prefix forms presented in the input in all conditions. However, Korean speakers in the aggregate applied the neutralizing alternation as often as the non-neutralizing alternation, which indicates that a weaker neutralization and homophony avoidance effect is found among Korean speakers than English speakers.

In this chapter, I present a probabilistic model that accounts for the different learning results for neutralizing and non-neutralizing alternations. I demonstrate that a Discount model, which implements the avoidance effect as a discounting of inputs that trigger homophony, successfully captures the different learning results of neutralizing alternations between English and Korean speakers and provides a straightforward account for the avoidance effect. Furthermore, I also show that the model successfully predicts Yin and White's (2018) experiment results showing the different learnability between neutralizing alternations that did and did not trigger homophony. Finally, I consider an alternative model that implements the avoidance effect as a constraint in learners' grammars (instead of discounting the input data that trigger homophony), but I argue that this *Neutralization model does not accurately account for the different learning results between Korean and English speakers.

In the rest of this chapter, I first review previous literature that has looked at homophony avoidance effects within phonological systems (particularly anti-homophony blocking effects), and I discuss the ways in which these phenomena are distinct from the effects being studied in the current experiment. Then, I introduce the architecture of the Discount model, including an overview of a Maximum Entropy (MaxEnt) learning model, and report the model's predictions. After, I introduce the alternative *Neutralization model and its predictions. I also train the two models with an additional dataset from Yin and White's study (2018) and compare the models' predictions. Finally, I discuss the findings and limitations of the proposed framework.

4.2 Anti-homophony blocking

In this section, I review how homophony avoidance was examined with different constraints in learners' grammar. Previous studies suggest that anti-homophony blocking which occurs when an expected phonological alternation does not occur if it would create homophony is motivated by learners' innate grammar to avoid homophony creation. (Crosswhite, 1999; Crosswhite & Jun, 2001; Kenstowicz, 2002; Kawahara, 2003; Ito & Mester, 2004; Ichimura, 2006). The majority of these studies suggest that anti-homophony blocking only occurs among words within the same paradigm (but cf. Ichimura, 2006). Words and phrases are in an *intra-paradigmatic* relationship (often referred to as being within the same paradigm) when they share the same root. As they share the same root, they are morphologically related and are the same lexeme. By contrast, words and phrases are in a *trans-paradigmatic* relationship when they do not share a root; hence, they are *not* morphologically related and are distinct lexemes. For instance, in English, the words {*go, going, gone*, *went}* are part of the same paradigm as they share the root *go* but the words {*go, drink, make*} are not within the same paradigm. In the following sections, I first introduce studies that proposed different constraints exclusively blocking homophony among words in the same paradigm (Crosswhite, 1999; Kawahara, 2003; Ito & Mester, 2004). Then, I review Ichimura's study (2006) which suggests that anti-homophony blocking can also occur among words in a trans-paradigmatic relationship.

4.2.1 Anti-homophony blocking within a paradigm

Crosswhite (1999) introduced a ANTI-IDENT constraint to explain the blocking of vowel reduction that can create homophony in Trigrad Bulgarian. In Trigrad Bulgarian, the mid, back vowels /-o/ or /-ɔ/ are reduced to [a] when unstressed (e.g., [ˈ**r**o**gave**] 'horns' vs. [r**a**gaˈvete] 'the horns'). However, the reduction of the word-final vowel in /ˈzɔrn**o**/ 'a grain' is blocked (*[zɔrn**a**]) presumably because it would be homophonous with the plural form [ˈzɔrn**a**] 'grains'. The ANTI-IDENT constrain is defined as follows: "for two forms, S_1 and S_2, there must be some segment α which is a member of S_1 such that α is not identical to its correspondent in S_2" (Crosswhite, 1999, p.8). The ANTI-IDENT constraint blocks reduction of the word-final vowel in /ˈzɔrn**o**/ 'a grain' because the change would result in homophonous pattern between words within the same paradigm. However, ANTI-IDENT does not block the vowel reduction when it creates homophony among words in trans-paradigmatic relations. For example, in [ˈblag**o**] 'benefit' – [ˈblag**a**] 'benefits' and [ˈblag**o**] 'sweet' (predicative) – [ˈblag**a**] 'sweet'(attributive), homophony is blocked between words within a paradigm, but it is not blocked between those across paradigms.

Similarly, Kawahara (2003) proposed a *MERGE constraint,[51] arguing that it can explain the anti-homophony blocking in exceptional fusion cases of Zoque (Zoque-Mixe, Southern Mexico). In Zoque, root faithfulness generally outranks affix faithfulness. For example, the stem-final /ŋ/ maintains its place (i.e., /<u>ma**ŋ**</u>+tɯʔ+u/ 'he intended to go' – [ma**ŋ**dɯʔu]) but the affix-final /m/ undergoes nasal place assimilation (i.e., /<u>tih</u>+ta**m**+tɯʔ+u/ 'we/you were about to arrive' – [tihta**n**dɯʔu]). However, when an affix /N-/ precedes a /w/-initial stem, this sequence can be fused. In this fusion case, the orality of stem-initial /w/ is not preserved and undergoes nasalization (e.g., /N+<u>wenu</u>/ 'my breaking (1st person progressive)' – [**w̃**enu] instead of *[wenu]). If the orality of the stem was preserved (i.e., [wenu]), it would be homophonous with another word in the same paradigm (e.g., /wen + u/ 'it broke (3rd person perfective)' – [wenu]). This fusion case can be explained by the *MERGE constraint that is described as follows: "underlyingly distinct forms within a paradigm must receive different phonological exponence" (Kawahara, 2003, p.13). The *MERGE constraint blocks the preservation of orality in the stem and prevents the merge of two distinct forms within a paradigm (see Kawahara, 2003, Tableau (31)).

While the constraints proposed by Crosswhite (1999) and Kawahara (2003) account for the blocking of a phonological rule that can create homophony, Ito and Mester (2004) suggested a PARADIGM CONTRAST (PARCONTRAST) constraint that explains changes to a different allomorph when a sequence of affixes would create homophony within a paradigm. In Japanese, '*ra*-dropping' can block homophony creation. The potential suffix /-rare/ is realized as [-re] by dropping the first syllable *ra*-, which prohibits the creation of homophony within the same paradigm (e.g., /tabe-rare/ 'eat-POTENTIAL' – [tabe-**re**] and /tabe-rare/ 'eat-PASSIVE' – [tabe-**rare**]; see Ito & Mester, 2004,

[51] Kawahara's *MERGE constraint is a modified version of the *MERGE constraint initially suggested by Padgett (2003a, 2003b). The *MERGE constraint suggested by Padgett (2003a, 2003b) is stemmed from the Dispersion Theory (Flemming, 2002). The constraints in Dispersion Theory prefer perceptually more noticeable contrasts over less noticeable ones between sounds.

Tableau 28, p.14). The PARCONTRAST constraint prohibits a pair of words within a paradigm from having identical forms, which accounts for the change to a different allomorph ([-re]).

4.2.2 Anti-homophony blocking across paradigms

Unlike the studies mentioned in the previous section, Ichimura (2006) proposed a Minimal Pair Analysis (MPA) and a CONTRAST constraint to account for homophony blocking among words across paradigms. Typically, in Japanese, a vowel preceding the negative suffix (/**V**-nai/) and a vowel in the allomorph of the negative suffix (/-**a**nai/) are dropped when they are preceded by /r/ (i.e., …r**V**-nai and …r-**a**nai → …r-nai). Then, /r/ is assimilated to [n], which results in the [nn] sequence (e.g., /ku**re**-nai/ 'give (me)+NEG' → [ku**n**nai], /waka**r**- **a**nai/ 'understand+NEG'→ [waka**n**nai]). However, when the nasal assimilation triggers homophony between lexically distinct words, the alternation in the sequence of rV+nai is blocked (e.g., /waka**r**- **a**nai/ 'understand+NEG'→ [waka**n**nai], and /waka**re**-nai/ 'get separated+NEG'→ [waka**re**nai] instead of *[waka**nn**ai]).

Ichimura (2006) showed that MPA and CONTRAST can account for the anti-homophony blocking between contracted forms across paradigms in Japanese. MPA is a monitoring mechanism that considers a pair or triplet of words as a set of inputs and evaluates a set of output, which is used to analyse the interaction of phonological alternations in multiple words in intra- or trans-paradigmatic relationships (e.g., see Ichimura, 2006, Tableau 50, p.102). The correspondence among outputs in a set is evaluated by CONTRAST which is defined as the following in (3):

(3) CONTRAST (Ichimura, 2006, (107), p.97)

> *Contrastiveness*[52] in underlying forms between words with the same major lexical category must be maintained in surface forms.

---

[52] The precise definitions of the contrastiveness and major lexical category can be found in Ichimura (2006; (107), (108), p. 97).

The CONTRAST constraint prevents the pair, /waka**r**-**a**nai/ and /waka**re**-nai/, from having an identical output [waka**nn**ai] and blocks the alternation and homophony creation between the words in the pair (see Ichimura, 2006, Tableau 51, p.101). Thus, Ichimura (2006) shows that words across paradigms can be explained by implementing constraints in learners' grammars.

4.2.3 Relation of anti-homophony blocking to the current study

Although the studies mentioned in this section so far explore homophony avoidance effects in phonology, they are focused on an issue that is somewhat distinct from the focus of the studies in this dissertation. The studies in this section mostly explain synchronic homophony-blocking cases where the expected application of an existing rule is blocked in specific lexical items; they are not focused on how homophony avoidance impacts the learnability of a new phonological alternation. By contrast, in this dissertation, I focus on the question of whether the learning of a new phonological rule is impaired when that rule creates homophony. While it is possible that synchronic anti-homophony blocking effects should be limited to pairs of words within a paradigm (e.g., as proposed by Crosswhite, 1999; Kawahara, 2003; Ito & Mester, 2004; but c.f. Ichimura, 2006), the experimental results from Chapter 2 suggest that homophony avoidance effect can inhibit the learnability of a rule even when it creates homophony among words that are not part of a single paradigm.

Hence, the models that I propose in this chapter focus on explaining the different learnability between a neutralizing alternation that triggers homophony and a non-neutralizing alternation. They also account for the different learnability of the neutralizing alternation between Korean and English speakers. I consider two models: a Discount model and a *Neutralization model. The Discount model explains the decreased learnability of novel homophony-creating alternations by reducing the effect of the input data that exhibit derived homophony. In this way, the model explains the homophony avoidance effect on phonological rule learning across paradigms with a general mechanism that is external to learners' phonological grammars. In an alternative model, the *Neutralization model, I include a *Neutralization constraint within learners' grammars which penalizes cases in which a surface form undergoes a phonological change that makes it identical to a different surface form. The *Neutralization constraint is inspired by the earlier studies that included a similar type of constraint to account for synchronic anti-homophony blocking effects (particularly Ichimura's study, 2006, since it included a constraint penalizing homophony creation across paradigms).

4.3 Architecture of the Discount model

This section presents the Discount model, a probabilistic approach to implement the avoidance of neutralization and homophony in a learning process. The Discount model applies the avoidance effect by reducing the influence of the observations that trigger homophony in the training data instead of creating a constraint penalizing neutralization in the grammar. I first describe the MaxEnt model in general as well as the specific implementation of the Discount model used in the study. I then describe the structure of the model and compare the model's predictions to the experimental results for Korean and English speakers.

4.3.1 Maximum Entropy models

I used the framework of the MaxEnt learning model. The MaxEnt learning model is a statistical model that has been used in various fields. The model uses entropy to find the probability of the observed data. Entropy is an information theory-based measurement that calculates the uncertainty (often referred to as the unpredictability or randomness) of a set of events given the amount of information to be successfully delivered (Shannon, 1948; Cover & Thomas, 1991). The

MaxEnt learning model is a weighted constraint-based model that returns the probability distribution over candidates. Starting with Goldwater and Johnson (2003), the MaxEnt learning model has ample precedent in the phonological literature (Jäger, 2004; Wilson, 2006; Hayes & Wilson, 2008; White, 2017; Hughto et al., 2019; Baird, 2021; Hayes, 2022). For example, the MaxEnt learning model has been used to account for free variation (Goldwater & Johnson, 2003), phonotactic learning (Hayes & Wilson, 2008), saltatory alternations (White, 2017), and frequency-based patterns (Baird, 2021). The MaxEnt model learns a set of weights for constraints to match the input data, and it predicts the probability of each output candidate for the given input using the learned weights. Instead of generating one winning output, the MaxEnt model provides a probability distribution of all possible outputs. The goal of the model is to find the weight for each constraint that generates the highest probability of observed data.

4.3.1.1 Assigning probabilities to outputs

Rather than categorically choosing an output, MaxEnt models assign a probability which is proportional to its constraint violation profile of harmony to a candidate. Namely, they produce a probability distribution over output candidates for given inputs. The probability distribution is calculated based on the weights of constraints that output candidates violate. The likelihood of an output candidate (o) for a given input (i) is defined in (4), where

w_j is the weight of the *j*th constraint,
$C_j(o,i)$ is the number of violations the *j*th constraint assigns the candidate
$\sum_{j=1}^{n}$ is the summation over all constraints (C_1, C_2, …, C_n),
$o \in O(i)$ is an output candidate (o) in a set of possible outputs for a given input (O(i))

$$(4) \quad \Pr(C) = \frac{e^{\left(-\sum_{j=1}^{n} w_j C_j(o,i)\right)}}{\sum_{o \in O(i)} e^{\left(-\sum_{j=1}^{n} w_j C_j(o,i)\right)}}$$

To calculate the *harmony*[53] ($\sum_{j=1}^{n} w_j C_j\ (o,i)$), the weight of the jth constraint (w_j) is multiplied by the number of times that an output candidate violates the jth constraint ($C_j(o,i)$), and this value is summed over all constraints (C_1, C_2, …, C_n).[54] Then, e (the base of the natural logarithm) is raised to the negative harmony ($e^{-harmony}$). Finally, the $e^{-harmony}$ of each candidate is divided by the sum of all output candidates given the input ($\sum_{o \in O(i)} e^{-\ harmony}$). An output candidate that violates a constraint with a higher weight will have a lower probability compared to an output candidate that violates a constraint with a lower weight, all else being equal.

4.3.1.2 Learning the constraint weights

The task of learning the set of weights for a MaxEnt model is defined as learning the weights that maximize the likelihood of the data (Pr(D)) given the model. To avoid underflow errors associated with multiplying probabilities, the log likelihood is used instead.[55] The log likelihood of the data (Pr(D)) is defined as the sum of the log likelihood of the candidates (see (5)). As one pair of input and output can appear multiple times in the observed data, a pair of input and output that appears more often in the observed data has a greater effect on the model than a pair that appears less frequently in the observed data.

---

[53] The term harmony matches with the Harmonic value used in Harmonic Grammar (Smolensky, 1986; Smolensky & Legendre, 2006). Harmony is also closely related to the *score* (h(x)) used in Hayes & Wilson (2008, p. 383).

[54] Unlike classical Optimality Theory (OT; Prince & Smolensky, 1993/2004) in which the candidate that does not violate the higher-ranked constraints wins, the candidate that violates the constraint with the highest weight still can be the winner in the MaxEnt model. As in Harmonic Grammar (Legendre et al., 1990), the MaxEnt model sums the violations of weighted constraints, which allows multiple constraints of lower weights to "gang up" and demote a candidate that does not violate the constraint with higher weight (Hayes & Wilson, 2008). A similar effect can also be found in Linear OT (Keller, 2000) and stochastic OT (see Hayes & Londe, 2006, p. 81).

[55] The probability of observed data (Pr(D)) is calculated by the product of the probability of the output candidate. However, multiplying probabilities results in a very small number; hence, the log probability is used instead of the product of probabilities.

(5) $$\log \Pr(\mathrm{D}) = \sum_{c \in C} \log\ \Pr(C)$$

4.3.1.3 Restricting the weights using a prior

A *prior* can be introduced to apply a bias during learning and to avoid overfitting the training data. Following Goldwater and Johnson (2003), I used the Gaussian distribution prior to restrict the weights of constraints. The Gaussian distribution prior has two elements: the preferred weight (the mean) of each constraint (μ) and the standard deviation (σ^2). The formula in (6) shows the calculation of the Gaussian prior. The assigned weight (w_j) is subtracted by the preferred weight (μ_j), and the difference is squared. Then the value is divided by $2\sigma^2$. The model aims to find a set of weights that maximizes the probability of the observed data given the prior. The goal of learning is to maximize the objective function in (7) by balancing a high likelihood for the data with a penalty on how far the weights of constraints can deviate from their preferred values; thus, the result of (6) is subtracted from the log probability of the observed data (the function in (5)).

(6) $$\sum_{j=1}^{n} \frac{\left(w_j - \mu_j\right)^2}{2\sigma^2}$$

(7) $$\text{Objective} = \left[\sum_{c \in C} \log\ \Pr\,(C)\right] - \left[\sum_{j=1}^{n} \frac{\left(w_j - \mu_j\right)^2}{2\sigma^2}\right]$$

A constraint is penalized when the learned weight (w_j) differs from the preferred weight (μ), suggesting that the greater the difference between the learned weight and its preferred weight is, the greater penalty is imposed on the constraint. Additionally, the standard deviation (σ^2) indicates how strictly the weight of a constraint should follow its preferred weight. A low value of σ^2 imposes a higher penalty for constraints that deviate from their preferred weight, forcing constraints to be closer to their preferred weight. A low value of σ^2 would also slow the process of learning because a constraint needs more observations from the input to deviate from the preferred weight. This type of prior helps prevent overfitting because it encourages weights to be distributed across the constraints. Without exceedingly large number of observations in the input, constraints are unlikely to have relatively high weights.

In the implementation used in the current study, the preferred weight (μ) of the PU constraint, the μ of the *Neutralization constraint (only in the *Neutralization model), and the standard deviation (σ^2) were treated as free parameters. These parameters were set to the values that best fit the experimental results. To set the parameters for each model, I used a grid search method to find the parameter values that maximized the log likelihood of the results given the model. Although σ^2 is fit as a free parameter, the model was restricted to having the same σ^2 for each constraint. For the grid search, I used a range of μ values from 0 to 15 at 1 interval and a range of σ^2 values from 0.1 to 9.1 at 0.2 intervals.

Weights were learned using a gradient descent algorithm as implemented in pyMEG developed by James White. Instead of trying all possible combinations of weights, the model starts each constraint at its preferred weight for the first iteration and conducts an iterated calculation until it finds the set of weights for each constraint that generates the maximum likelihood of the observed data. Based on the current stage of iteration, the model finds the direction to ascend in the next stage of iteration (see Della Pietra, Della Pietra, & Lafferty, 1997 and Hayes & Wilson, 2008). The algorithm ends once the ascending gradient is approximately 0, indicating that it has reached the peak of the likelihood surface and thus the optimal set of weights.

4.3.2 Structure of the Discount model

In the Discount model, I account for the avoidance of neutralization and homophony by discounting the inputs that create homophony in the training data. In the model, it is assumed that the selection of a plural prefix (either *ba-* or *ni-*) is already completed, and the model only predicts the choice of stem types for the given input of prefix+stem. The structure of the models is similar to the second stage of the two-stage model proposed by Kapatsinski (2009, Chapter 5, p. 180). Kapatsinski proposed the one-stage and two-stage models to account for learning affixed forms (see section 2.7.3). In the one-stage model, the selection of an affix and a base form occurs at the same stage, suggesting that whether or not the base needs to be modified is likely to affect the choice of an affix. In the two-stage model, the choice of an affix occurs first, and then the selection of a base form (e.g., with or without palatalization) occurs next. As the choices of an affix and a base form are in two separate stages, whether the base needs to be alternated or not cannot affect the choice of an affix.

The two-stage model is better aligned with the experimental results which suggest that prefix selection was completely independent from which stem was present. Specifically, in Chapters 2 and 3, the percentage of the time that the *ni-* prefix (that triggered the alternation) was chosen was similar across non-velar-initial stems and velar-initial stems in all Prefix Frequencies and Alternation Types, even though only the velar-initial stems alternated. Hence, assuming that the plural prefix is already selected, the Discount model focuses on predicting the choice of the stem (whether the stem was alternated or not) given the input. For instance, the model predicts the choice between ba+faithful stem (where the stem is not alternated) or ba+palatalized stem (where the stem is alternated) for the given input of ba+stem.

4.3.3 Constraints and discount implementation

The Discount model takes a constraint-based approach in the spirit of Optimality Theory (OT; Prince & Smolensky, 1993/ 2004); it has two OT-style constraints: Paradigm Uniformity (PU) and *Ni-[+DORSAL]. In this section, I introduce these two constraints, and I describe how observations that trigger homophony are discounted in the Discount model.

*4.3.3.1 Paradigm Uniformity and *Ni-[+DORSAL]*

Paradigm Uniformity (PU) was implemented in the Discount model to evaluate whether the shared root in the singular and plural forms has a matching surface form. The PU constraint used in the models follows the notion of paradigm uniformity (Steriade, 1994, 2000; Hayes, 1997). To be explicit, I describe PU as follows:

(8) Paradigm Uniformity (PU)[56]

Assign a violation for every morpheme if its surface realization differs from its surface realization in a morphologically related form.

In the current models, the output of prefix+palatalized stem for the given input prefix+stem violates the PU constraint because the stem is alternated and does not have an identical surface form across singular and plural forms. For example, PU is violated when /ni-kimu/ surfaces as [niʧimu] and /ba-kimu/ surfaces as [baʧimu] but not when those underlying forms surface as [nikimu] and [bakimu].

Furthermore, the model has one markedness constraint that penalizes velar-initial stems surfacing after the prefix *ni-*, which is defined below (9). In the Non-Neutralizing and Neutralizing

[56] Correspondence theory establishes relationships between related forms and requires that they are as similar as possible. The PU constrain can be formalized as an output-output correspondence (OO-faithfulness) constraint (Burzio, 1994, 1998; McCarthy & Prince, 1995; Benua, 1995) but I left the constraint general in this thesis. Also, using an input-output faithfulness constraint would work identically in the model. However, there is a reason to believe that learners place an input-output faithfulness constraint below a markedness constraint in the initial stage of learning (see discussion Gnanadesikan, 1995; Smolensky, 1996; Prince & Tesar, 1999/ 2004; Boersma & Levelt, 2000; Curtin & Zuraw, 2002; Hayes, 2004; White, 2017).

conditions, the *ni-* prefix triggered palatalization of velar-initial stems (e.g., [kimu]→[niʧimu], [gapi]→[niʤapi]). The pressure to palatalize is formalized as the constraint *Ni-[+DORSAL].

(9) *Ni-[+DORSAL] [57]

Assign a violation when the *ni-* prefix is followed by a velar-initial stem.

To illustrate how the two constraints interact, Tableaux 10 and 11 show the analysis using OT tableaux. When *Ni-[+DORSAL] is ranked above PU (Tableau 10), candidate (10b) wins because the markedness constraint enforces the palatalization of velar-initial stems. When PU outranks *Ni-[+DORSAL] (Tableau 11), candidate (11a) wins because PU prohibits the palatalization of the stems. Recall that because the actual model is probabilistic in nature, the candidates will each have a probability of being an output rather than one being the strict winner as in OT.

(10)

/kimu/ - /ni+kimu/	*Ni-[+DORSAL]	PU
(a) [kimu] - [nikimu]	*!	
☞ (b) [kimu] - [niʧimu]		*

(11)

/kimu/ - /ni+kimu/	PU	*Ni-[+DORSAL]
☞ (a) [kimu] - [nikimu]		*
(b) [kimu] - [niʧimu]	*!	

[57] It is written as *ni-[+DORSAL] for ease of reading but it can be considered a general markedness constraint such as *iK. Changing *ni-[+DORSAL] to a general markedness constraint would not change the model's predictions for the purposes of this study.

4.3.3.2 Discounting observed data

In the Discount model, I account for the avoidance of neutralization and homophony by reducing the influence of input data that present derived homophony. The motivation for this approach stems from learners' bias against a many-to-one relationship. In section 2.3, I reviewed previous literature showing that learners are biased against a many-to-one relationship, which plays a general role in multiple areas of learning (e.g., morphological, semantic, and lexical learning). The derived homophony presented in the experiments in Chapters 2 and 3 is a many-to-one relationship because two singular forms are mapped to an identical plural from (e.g., /ʧimu/ 'a carrot' and /kimu/ 'a flower' - [niʧimu]).

In the experiments, learners needed to set up a paradigm between a singular form and their plural prefix forms (e.g., a singular form: [kimu], 'a flower', and plural forms: [bakimu]/[niʧimu], 'flowers'), meaning that they were required to understand these two forms were related and in the same paradigm. Setting up this relationship between morphologically related forms is an important step for learners to discover and acquire the relevant alternation, where required (e.g., /k/ → [ʧ] in [niʧimu]). The experimental results generally showed that learners were able to successfully set up the paradigmatic relationship between singular and plural pairs and apply the alternation. However, in Chapter 2, I showed that English speakers were less successful at learning the neutralizing alternation that triggered homophony than the non-neutralizing alternation.

I propose that learners were slower to set up the paradigmatic relationship between the singular and plural pair that required an alternation (e.g., [kimu] and [niʧimu]) in the cases where it created derived homophony due to interference from another pair that already had a matching form (e.g., [ʧimu] and [niʧimu]). In the presence of the homophonous plural form (e.g., [niʧimu]), learners would prefer to map the plural form [niʧimu] to [ʧimu] and pay less attention to (or discard) building a paradigmatic relation between [niʧimu] and [kimu] (see Figure 9). It makes the pair to be weaker evidence for the phonological alternation than it would in the absence of the homophonous plural form.[58] The difficulty setting up these critical paradigmatic relationships between pairs that present the alternation resulted in poorer learning of the alternation overall in the Neutralizing condition.

In the Discount model, I implement this tendency by discounting the singular and plural pairs that exhibit derived homophony. The precise amount of the discounting is fit based on the experimental results. In the current study, the palatalizing alternation triggered neutralization and created two homophonous outputs only in the Neutralizing condition: [niʧimu] (derived from /ni-ʧimu/ and /ni-kimu/) and [niʤapi] (derived from /ni-ʤapi/ and /ni-gapi/). Hence, the output of [niʧimu] for the given input of /ni-kimu/ and the output of [niʤapi] for the given input of /ni-gapi/ are discounted in the Neutralizing condition.

Figure 9. The strong correspondence between *niʧimu* and *ʧimu* (solid line) and the weak correspondence between *niʧimu* and *kimu* (dashed line).

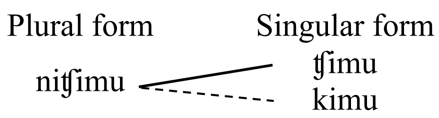

The rate of discounting the input that creates homophony is implemented as a parameter, which categorically ranged from a 0% to a 100% discount (0%, 25%, 50%, 75%, and 100%). The rate of discounting must be flexible to account for the different learning results between Korean and

[58] In theory, setting a paradigm between singular and plural forms can be affected by the order of plural forms presented during the experiment. However, in the experiments introduced in Chapters 2 and 3, the order of presentation is less likely to affect building the paradigm. This is because, in the stem learning phase, learners first learned all singular stems including velar-initial and palatal-initial stems (e.g., [kimu], [ʧimu], [gapi], and [ʤapi]). Once they successfully learned the singular stems, they were exposed to plural forms, suggesting that they had known singular forms and already had a preference to build a paradigm between particular singular and plural forms (e.g., [ʧimu] and [niʧimu]).

English speakers. The best fitting discount for a given experiment was found using a grid search and reported with the log likelihood (see Table 40). By testing many levels of discount, I can examine to what extent learners discount the input when it triggers homophony and address the hypothesis that Korean speakers differ systematically from English speakers (i.e., Korean speakers discount homophonous patterns less than English speakers). The number of homophonous plural forms in the training data before and after discounting is found in Table 38. For instance, in the 50-50 Neutralizing condition, homophonous plural forms were presented six times during the experiment. Discounting the plural forms 75% of the time means that the plural forms were presented only 1.5 times (6 x 0.25) in the training data. Other inputs that did not trigger homophony were *not* discounted in the Discount model.

Table 38. The number of observations that create homophony before discounting and the number of observations after discounting in the training data.

Prefix Frequency condition	Before discounting	After discounting				
		100%	75%	50%	25%	0%
Frequent *ba-*	4	0	1	2	3	4
50-50	6	0	1.5	3	4.5	6
Frequent *ni-*	8	0	2	4	6	8

4.4 Testing the Discount model

In this section, I present how the parameters of the Discount model fit the experimental results and compare the model's prediction to the experimental results. I briefly review the findings of the experimental results reported in Chapters 2 and 3 that are predicted by the model. In Chapter 2, the experimental results showed that English speakers were less successful at learning the neutralizing alternation than the non-neutralizing alternation. They palatalized velar-initial stems less often after *ni-* (the correct cases) and more often after *ba-* (incorrect over-generalization cases) compared to the non-neutralizing alternation. In Chapter 3, unlike English speakers, Korean speakers were equally successful at applying the neutralizing alternation and the non-neutralizing alternation.

4.4.1 Overview of testing the Discount model

The model aims to mimic the process that participants experienced during experiments. At the beginning of learning, the model considers two possible outputs as potential candidates for a given input because participants do not know which of the two outputs is the correct plural form for the given input. Recall that the model predicts the stem for a given input of the prefix+stem. In all Alternation Types, there are two possible outputs: prefix+faithful stem (that is not alternated) and prefix+palatalized stem (where the stem is palatalized). However, the actual outputs that participants hear during the experiment—and therefore the plural forms that are presented in the data to train the model—differ by the Alternation Types and stem types. The possible outputs and actual outputs presented in the training data according to the Alternation Types and stem types are given in Table 39 below. In the No Alternation condition, the actual outputs presented in the training data are either ba+stem or ni+stem for all inputs because there are no alternations of stems. In the Non-Neutralizing and Neutralization conditions, the actual outputs differ by stem type. For ba+non-velar-initial stem and ni+non-velar-initial stem, the outputs presented in the training data are correspondingly ba+stem and ni+stem. In the crucial cases where the input is either ba+velar-initial stem or ni+velar-initial stem, the actual outputs in the training data are correspondingly ba+stem and ni+palatalized stem. During training, the model's task is to observe the winning outputs for given inputs and to learn the set of weights that best account for these data. At the test phase, the model then predicts the probability of the winner for the given input (e.g., predict the probability between [bakimu] and [baʧimu] for given input /ba+kimu/).

Table 39. The possible outputs and actual outputs (shaded cells) for given inputs in the training data by Alternation Type.

Input	No Alternation condition	Non-Neutralizing and Neutralizing conditions
ba+non-velar-initial stem	ba+stem	ba+stem
	ba+palatalized stem	ba+palatalized stem
ni+non-velar-initial stem	ni+stem	ni+stem
	ni+palatalized stem	ni+palatalized stem
ba+velar-initial stem	ba+stem	ba+stem
	ba+palatalized stem	ba+palatalized stem
ni+velar-initial stem	ni+stem	ni+stem
	ni+palatalized stem	ni+palatalized stem

4.4.2 The effect of different parameters

Parameters were fit to the experimental results for English and Korean speakers. In the Discount model, the values of three parameters were fit to match each dataset: the preferred weight (μ) of PU, the value of standard deviation (σ^2) (constrained to be identical for all constraints), and different levels of discount. Assigning a μ above 0 for a constraint indicates that learners have an a priori bias in favour of the constraint. I propose learners have a prior bias in favour of PU (with the precise strength of that bias being left as a free parameter). Hence, the μ for PU started with the possibility of having a positive value. The *Ni-[+DORSAL] constraint is left at a μ of 0, indicating that learners do not have a prior bias in favour of the constraint and the weight of the constraint will only be increased after seeing the evidence of it being active. The parameters were set according to log likelihood based on a grid search. For the grid search, I used a range of μ values from 0 to15 at 1 interval, a range of σ^2 values from 0.1 to 9.1 at 0.2 intervals, and different levels of discount (discounted by 0%, 25%, 50%, 75%, and 100%). The values that best fit to match each dataset (English and Korean speakers) were chosen for the respective model. The log likelihood indicates how closely the models matched the experimental results. A log likelihood that is closer to 0 indicates a better model fit for the experimental results. The detailed results of the grid search are reported in Appendix B. Table 40 shows the best model at each discount level, as well as its fit parameters and log likelihood. Shaded cells indicate the model that best fits the English and Korean datasets.

Table 40. The log likelihood according to different sets of parameter values for the English and Korean datasets.

Experimental results	Log likelihood	Discounting	μ of PU	σ^2
English speakers	-1604.3145	100%	7	3.3
	-1600.1038	**75%**	**7**	**2.7**
	-1601.9202	50%	7	2.5
	-1606.9479	25%	7	2.3
	-1613.5829	0%	7	2.1
Korean speakers	-1090.1587	100%	7	2.7
	-1072.7389	75%	7	2.1
	-1061.6084	50%	7	1.9
	-1054.8724	25%	8	1.9
	-1051.4754	**0%**	**8**	**1.7**

The Discount model most closely matched the English dataset when the homophony-creating inputs were discounted by 75%. In contrast, the model best fit the Korean dataset when the inputs were discounted by 0% (i.e., when they were <u>not</u> discounted). Importantly, these findings reflect the different neutralization and homophony avoidance effects between English and Korean speakers: Korean speakers did not discount the input that triggers homophony whereas English speakers discounted the input by 75%.

4.4.3 Predictions of the Discount model

In this section, I examine how well the Discount model predicted the experimental results. I compare the models' prediction to the results for English speakers and Korean speakers in turn.

4.4.3.1 Comparison to English dataset

Table 41 shows the initial preferred weights (μ) and final weights of constraints before and after training. In the No Alternation condition, *Ni-[+DORSAL] did not change from its initial weight of 0 because there was evidence in the input that the velar-initial stems occur after *ni-*. The final weight of PU also did not change much from its initial weight (7→7.07) because there was no evidence that stems alternated. In the Non-Neutralizing and Neutralizing conditions, the initial weight of *Ni-[+DORSAL] increased to motivate the palatalization of velar-initial stems after *ni-*. The weight of PU decreased in the Non-Neutralizing and Neutralizing conditions because the input data presented evidence of non-unified paradigms.

Notably, the differences between the initial and final constraint weights were smaller in the Neutralizing condition compared to the Non-Neutralizing condition. The weights of *Ni-[+DORSAL] increased less across the three prefix frequency conditions in the Neutralizing condition than in the Non-Neutralizing condition (Table 41; the Non-Neutralizing condition, *Ni-[+DORSAL] μ: 0→ Frequent *ba-*: 4.35, 50-50: 4.68, Frequent *ni-*: 4.87 vs. the Neutralizing condition, *Ni-[+DORSAL] μ: 0→ Frequent *ba-*: 3.88, 50-50: 4.27, Frequent *ni-*: 4.5). Also, the weights of PU decreased less across the three prefix frequency conditions in the Neutralizing condition than in the Non-Neutralizing condition (Table 41; the Non-Neutralizing, PU μ: 7→ Frequent *ba-*: 3.96, 50-50: 3.78, Frequent *ni-*: 3.64 vs. the Neutralizing, PU μ: 7→ Frequent *ba-*: 4.18, 50-50: 3.95, Frequent *ni-*:3.80). The smaller changes in the Neutralizing condition compared to the Non-Neutralizing condition were motivated by discounting the observations that created homophony in the Neutralizing condition. In the training data, the homophonous ni+palatalized stem was discounted by

75% of the time in the Neutralizing condition (see Table 40), meaning that there was less (effective) evidence in the training data to support palatalization. Furthermore, in the Non-Neutralizing and Neutralizing conditions, the final weights of *Ni-[+DORSAL] and PU varied somewhat across the three Prefix Frequency conditions because the amount of palatalization varied. In the test phase, the Discount model used these final weights to calculate the predicted probability of each output candidate. An example calculation is given in Table 42 below.

Table 41. The initial preferred weights (μ) and the final weights of constraints in the Discount model (fit to the English dataset).

No Alternation condition				
Constraint	μ (preferred weight)	Final weight		
		Frequent *ba-* condition	50-50 condition	Frequent *ni-* condition
*Ni-[+DORSAL]	0	0	0	0
PU	7	7.07	7.07	7.07
Non-Neutralizing condition				
Constraint	μ (preferred weight)	Final weight		
		Frequent *ba-* condition	50-50 condition	Frequent *ni-* condition
*Ni-[+DORSAL]	0	4.35	4.68	4.87
PU	7	3.96	3.78	3.64
Neutralizing condition				
Constraint	μ (preferred weight)	Final weight		
		Frequent *ba-* condition	50-50 condition	Frequent *ni-* condition
*Ni-[+DORSAL]	0	3.88	4.27	4.5
PU	7	4.18	3.95	3.80

Table 42. The example of calculating predicted probabilities in the Frequent *ba-* Neutralizing condition.

/kimu/ /ni+kimu/	*Ni- [+DORSAL] 3.88	PU 4.18	Penalty score	$e^{(\text{-penalty})}$	Predicted outcome (predicted probability)
[kimu] [nikimu]	1		3.88	.02065	.57
[kimu] [niʧimu]		1	4.18	.0153	.43
				= 0.03595	

To examine the performance of the model, the predicted probabilities of outputs are compared to the aggregate experimental results of English speakers reported in Chapter 2. Based on the experimental results, the model was expected to capture two main points. First, English speakers learned that velar-initial stems were palatalized only after *ni-* in the Non-Neutralizing and Neutralizing conditions.[59] Second, English speakers were less successful at learning palatalization of velar-initial stems when it was neutralizing and homophony-creating than when it was not neutralizing. The complete comparisons between the model's predictions and the experimental results of English speakers in all Alternation Types and Prefix Frequency conditions are reported in Appendix C. In Table 43 below, I show the critical cases which are the models' prediction of the output (prefix+stem or prefix+palatalized stem) for the given input of prefix+velar-initial stems in the Non-Neutralizing and Neutralizing conditions.

[59] This also indicates that they learned that non-velar-initial stems were not palatalized either after *ba-* or *ni-*, velar-initial stems were palatalized after *ni-* but not after *ba-*, and no stems were palatalized in the No Alternation condition.

Table 43. The predicted probabilities of the Discount model and the aggregate experimental results for English speakers. Note: The numbers indicate the percentage of times (%) that the output was chosen in the model and in the experimental results (see Appendix C for the complete predictions of the model).

Non-Neutralizing condition							
		Frequent *ba-*		50-50 condition		Frequent *ni-*	
Input	Output	Model prediction	Experimental result	Model prediction	Experimental result	Model prediction	Experimental result
ba+velar-initial stem	ba+stem	98	98	98	96	97	89
	ba+ palatalized stem	2	2	2	4	3	11
ni+velar-initial stem	ni+stem	40	26	29	26	23	29
	ni+ palatalized stem	60	74	71	74	77	71

Neutralizing condition							
		Frequent *ba-*		50-50 condition		Frequent *ni-*	
input	output	Model prediction	Experimental result	Model prediction	Experimental result	Model prediction	Experimental result
ba+velar-initial stem	ba+stem	98	96	98	92	98	76
	ba+ palatalized stem	2	4	2	8	2	25
ni+velar-initial stem	ni+stem	57	55	42	41	33	38
	ni+ palatalized stem	43	45	58	59	67	62

In the Non-Neutralizing condition, the model correctly predicts that palatalization of velar-initial stems only occurs after *ni-*. The model predicted the ni+palatalized stem more frequently than the ni+stem for the given input of ni+velar-initial stem and the ni+palatalized stem more frequently than the ba+palatalized stem (see shaded cells).[60] This follows from the change in weights shown in Table 41 above; the increase in the weight of *Ni-[+DORSAL] (along with the decline in the weight of PU) resulted in more palatalization after *ni-*.

Importantly, the model accurately captured the fact that the English speakers were poorer at learning the alternation when it was neutralizing (Neutralizing condition) than when it was non-neutralizing (Non-Neutralizing condition). The percentages of times that the model chose the ni+palatalized stem for the given ni+velar-initial stem in the Neutralizing condition (shaded cells in Table 43) were always lower than those in the Non-Neutralizing condition, which matches the aggregate experimental results for English speakers. This is a direct result of the 75% discount applied to input forms that caused derived homophony, reflecting the difficulty that English speakers appear to have in terms of learning from homophonous input forms. On the other hand, the model did not successfully predict that English speakers made more over-generalization errors in the Neutralizing condition than in the Non-Neutralizing condition (see the Discussion section for a detailed explanation).

4.4.3.2 Comparison to Korean dataset

Table 44 shows the initial preferred weights (μ) of constraints before training and the final weights after training.61 The most important finding is that the final weights were identical across the Non-Neutralizing and Neutralizing conditions. This is because the Discount model that most

---

[60] In the Frequent *ba-* Non-Neutralizing condition, the Discount model slightly under-predicted the ni+palatalized stem for the given input of ni+velar-initial stem (model's prediction: 60% vs. experimental result: 74%). This is because there were fewer examples of ni+palatalized stems in the input in the Frequent *ba-* condition; hence, the model was slower to learn that velar-initial stems were palatalized after *ni-* in the Frequent *ba-* condition compared to other Prefix Frequency conditions. The under-prediction cases were also found in the Discount model fit to match the Korean dataset and the *Neutralization models fit to match the English and Korean datasets. However, the crucial point is that these under-prediction cases did not have a meaningful influence on the models' performance: the Discount model successfully captured the different experimental results between English and Korean speakers whereas the *Neutralization model did not.

[61] Recall that Korean participants were not exposed to the 50-50 and No Alternation conditions, and the model was likewise not trained on these conditions.

accurately matched the Korean dataset did not discount the observations that created homophony in the training data (see section 4.4.2). The weight of *Ni-[+DORSAL] increased (0→the Frequent *ba*-: 4.06, the Frequent *ni*-: 4.66) to motivate the palatalization of velar-initial stems after *ni*-, and the weight of PU decreased (8→ the Frequent *ba*-: 4.45, the Frequent *ni*-: 4.01) as the training data showed that non-unified paradigms were acceptable in both conditions.

Table 44. The initial preferred weights (μ) and the final weights of constraints in the Discount model (fit to the Korean dataset).

Non-Neutralizing condition			
Constraint	μ (preferred weight)	Final weight	
		Frequent *ba*- condition	Frequent *ni*- condition
*Ni-[+DORSAL]	0	4.06	4.66
PU	8	4.45	4.01
Neutralizing condition			
Constraint	μ (preferred weight)	Final weight	
		Frequent *ba*- condition	Frequent *ni*- condition
*Ni-[+DORSAL]	0	4.06	4.66
PU	8	4.45	4.01

The predictions of the model are compared to the experimental results for Korean speakers in Table 45 below. A full report of the model's predictions along with the aggregate results for Korean speakers is found in Appendix C.

Table 45. The predicted probabilities of the Discount model and the aggregate experimental results for Korean speakers. Note: The numbers indicate the percentage of times (%) that the output was chosen in the model and in the experimental results.

Non-Neutralizing condition					
		Frequent *ba-*		Frequent *ni-*	
Input	Output	Model prediction	Experimental result	Model prediction	Experimental result
ba+velar-initial stem	ba+stem	99	98	98	83
	ba+ palatalized stem	1	2	2	17
ni+velar-initial stem	ni+stem	60	61	34	35
	ni+ palatalized stem	40	39	66	65

Neutralizing condition					
		Frequent *ba-*		Frequent *ni-*	
Input	Output	Model prediction	Experimental result	Model prediction	Experimental result
ba+velar-initial stem	ba+stem	99	95	98	89
	ba+ palatalized stem	1	5	2	11
ni+velar-initial stem	ni+stem	60	47	34	36
	ni+ palatalized stem	40	53	66	64

In Table 45, I show the crucial cases of the model's predictions and the aggregate experimental results for Korean speakers. In both the Non-Neutralizing and Neutralizing conditions, the model successfully predicted palatalization after *ni-*, which matched the aggregate experimental results for Korean speakers. Furthermore, the model successfully captured the fact that Korean speakers learned the neutralizing and non-neutralizing alternation comparably well. The model chose the ni+palatalized stem equally often in the Non-Neutralizing and Neutralizing conditions (see shaded cells in both conditions), which is motivated by the constraint weights being equal in the two conditions (see Table 44). The model did not successfully predict that Korean speakers made over-generalization errors, particularly in the Frequent *ni-* condition.

In this section, I showed that the model accurately predicted learners' acquisition of the alternations and captured the different learning results between English and Korean speakers. Nonetheless, the model was unable to predict the over-generalization of palatalization after *ba-*. I also showed that the predictions of the Discount model provide a straightforward and effective explanation for the avoidance of neutralization and homophony creation.

4.5 Architecture of the *Neutralization model

In this section, I consider an alternative model that accounts for the neutralization and homophony avoidance effect using a constraint in the grammar (instead of discounting the input that results in homophony as we saw in the Discount model). In the *Neutralization model, I implemented the *Neutralization constraint which penalizes neutralizing a phonemic contrast. To preview the result, the *Neutralization model failed to capture the different learning results between English and Korean speakers. I briefly describe the structure of the *Neutralization model and the *Neutralization constraint. Then, I compare the model's predictions to the experimental results for English and Korean speakers.

4.5.1 Structure of the *Neutralization model

The structure of the *Neutralization model is identical to that of the Discount model except for the additional *Neutralization constraint (and no discount for homophony in the training data). The *Neutralization model also uses the framework of the MaxEnt learning model (see section 4.3.1) and focuses on predicting the alternation of stem-initial consonants rather than the choice of the prefix form (see section 4.3.2).

Inspired by the previous studies that account for neutralization and homophony avoidance with constraints in learners' grammars (Crosswhite, 1999; Kawahara, 2003; Ito & Mester, 2004; Ichimura, 2006), the *Neutralization model implements the avoidance effect as a constraint, *Neutralization. The *Neutralization constraint is defined below (12). Note that *Neutralization penalizes the neutralization of phonemic contrasts (neutralization at the phonetic level) rather than exclusively penalizing the neutralization of lexical contrasts (neutralization at the lexical level).

(12) *Neutralization

> Assign a violation to a segment in a surface form that has a faithfulness violation and is also identical to a segment in a different surface form.

In the Non-Neutralizing and Neutralizing conditions, velar-initial stems were palatalized after the *ni-* prefix. In the Non-Neutralizing condition, no stems underlyingly began with palatals, and palatalization of velar-initial stems did not merge any phonemic contrasts, satisfying *Neutralization. In the Neutralizing condition, however, both palatal- and velar-initial stems surfaced with initial palatals when they took the *ni-* prefix, violating *Neutralization. Tableaux 13 and 14 illustrate *Neutralization using OT tableaux. *Neutralization penalizes [ʧ] in [kimu] – [niʧmu] because it shows a faithfulness violation ([k] was palatalized to [ʧ] after *ni-*) and is identical to [ʧ] in another form ([ʧmu] – [niʧmu]) in the Neutralizing condition. *Neutralization always occurs alongside a faithfulness violation (i.e., PU) because both constraints prohibit a surface form that has a faithfulness violation. In classical OT, *Neutralization would be redundant, but the MaxEnt model is a type of Harmonic Grammar with the 'ganging up' property, suggesting that *Neutralization can impact the probability of the outputs. *Neutralization does play a role in the Non-Neutralizing condition because there are *no* palatal-initial stems (e.g., [ʧmu] – [niʧmu]) that create identical segments in surface forms (tableau 14). Accordingly, palatalization of velar-initial stems is penalized by *Neutralization only in the Neutralizing condition.

(13)

Also exists: [ʧimu] – [ni+ʧimu] /kimu/ – /ni+kimu/	*Ni-[+DORSAL]	PU	*Neutralization
(a) [kimu] – [nikimu]	*		
☞ (b) [kimu] – [niʧimu]		*	*

(14)

/kimu/ – /ni+kimu/	*Ni-[+DORSAL]	PU
(a) [kimu] – [nikimu]	*	
☞ (b) [kimu] – [niʧimu]		*

4.6 Testing the *Neutralization model

4.6.1 Overall fit and effect of different parameters

In the *Neutralization model, the values of three parameters were set using a grid search: the initial preferred weights (μ) for PU and *Neutralization, as well as the value of standard deviation (σ^2) for the constraints (constrained to be the same for all constraints as in the previous model). For the grid search, I used a range of μ values from 0 to 15 at 1 interval and a range of σ^2 values from 0.1 to 9.1 at 0.2 intervals. The values that best fit the results for each dataset (English and Korean speakers) were chosen for the respective model. The log likelihood indicates how closely the models matched the experimental results. Table 46 shows the log likelihood of the models that best fit the English and Korean datasets. A range of initial preferred weights of *Neutralization resulted in the best-performing model (see the footnote). [62]

Table 46. The log likelihood of the best *Neutralization model for the English and Korean datasets.

Experimental results	Log likelihood	μ of PU	μ of *Neutralization	σ^2
English speakers	-1613.58	7	0	2.1
Korean speakers	-1053.74	8	0	1.7

The log likelihood of the *Neutralization model that best fit the English and Korean data is lower than the log likelihood of the Discount model (*Neutralization model English: -1613.58, Korean: -1053.74, vs. Discount model English: -1600.10, Korean: -1051.48). Although the complexity and the assumptions of the models are different, the lower log likelihood score of the *Neutralization model implies that the model did not fit the experimental results as well as the Discount model overall.

4.6.2 Predictions of the *Neutralization model

In the following sections, I examine how well the model's predictions capture the experimental results for English and Korean speakers.

[62] A range of preferred initial weights for *Neutralization resulted in the same log likelihood of the best-fitting models for the English and Korean datasets. Here, I reported the smallest of these weights (prior weight of 0) in tables.

4.6.2.1 Comparison to English dataset

Table 47 shows the initial preferred weight (μ) and final weights of the constraints after training the model with experimental input. In the No Alternation condition, the final weights of constraints closely matched their preferred weight because there was no evidence of palatalization in the training input. In both the Non-Neutralizing and Neutralizing conditions, *Ni-[+DORSAL] gained weight and PU lost weight, motivating palatalization after *ni-*. In the Non-Neutralizing condition, the final weight of *Neutralization matched the preferred weight because there was no evidence of the neutralizing alternation.

Notably, in the Neutralizing condition (where it was relevant), the final weight of *Neutralization remained at 0. This is because *Neutralization and PU partially overlap in violations and the MaxEnt model spreads the weight out amongst the constraints that favour observed forms. Hence, as soon as *Neutralization picks up any weight, the weight of PU decreases. Reducing the final weight of PU would diminish the overall performance of the model when measured with log likelihood because PU penalizes a large number of filler items in the input data. In the training data used for the *Neutralization model, there were more inputs of the filler items (prefix+non-velar-initial stem: 36 items) than those of the velar-initial stems (prefix+velar-initial stem: 24 items). Also, in the experimental results, there were more observations of the filler cases than of the prefix+velar-initial stem. *Neutralization only penalized the prefix+palatalized stems for given inputs of the prefix+velar-initial stems but PU penalized both filler cases and prefix+palatalized stems for given inputs of the prefix+velar-initial stems, indicating that even the slight change of the weight of PU had a great effect on the fit of the model (see Appendix D for a detailed explanation).

As an alternative, I used the squared correlation coefficient (r^2) value to examine the fit of the *Neutralization model because, unlike log likelihood, r^2 is not affected by the number of items in the training data. However, it was complicated to select the parameter values of the best-performing model for the English dataset because there was a wide range of initial preferred weights that resulted in the same r^2 (see Appendix E for a detailed explanation). This highlights the complexity of the interactions between constraints in the *Neutralization model, which makes the *Neutralization model less straightforward than the Discount model.

Table 47. The initial preferred weights (μ) and the final weights of constraints in the *Neutralization model (fit to the English dataset).

No Alternation condition				
Constraint	μ (preferred weight)	Final weight		
		Frequent *ba*-condition	50-50 condition	Frequent *ni*-condition
*Ni-[+DORSAL]	0	0	0	0
PU	7	7.05	7.05	7.05
*Neutralization	0	0	0	0
Non-Neutralizing condition				
Constraint	μ (preferred weight)	Final weight		
		Frequent *ba*-condition	50-50 condition	Frequent *ni*-condition
*Ni-[+DORSAL]	0	4.04	4.39	4.59
PU	7	3.97	3.76	3.62
*Neutralization	0[63]	0	0	0
Neutralizing condition				
Constraint	μ (preferred weight)	Final weight		
		Frequent *ba*-condition	50-50 condition	Frequent *ni*-condition
*Ni-[+DORSAL]	0	4.04	4.39	4.59
PU	7	3.97	3.76	3.62
*Neutralization	0	0	0	0

[63] In the Non-Neutralizing condition. The final weight of *Neutralization was identical to its initial preferred weight because no outputs violated *Neutralization in the condition. For instance, if the initial preferred weight of *Neutralization was 2, the final weight was also 2.

The predictions of the model using these final weights on the data are given in Table 48 below. The complete predictions of the *Neutralization model are reported in Appendix C. The model correctly captured the fact that English speakers successfully learned that *ni-* triggered palatalization. In the Non-Neutralizing and Neutralizing conditions, palatalization is predicted more often after *ni-* (shaded cells) than after *ba-*. Also, in both conditions, palatalization is predicted more frequently than the (faithful) stem after *ni-*. However, the model failed to predict that English speakers were poorer at learning the neutralizing alternation than the non-neutralizing alternation. The model predicted identical rates of palatalization after *ni-* in the Non-Neutralizing and Neutralizing conditions. The identical prediction of both conditions reflects that the final weights were identical across the two conditions. The model further failed to capture the significantly higher over-generalization error cases in the Neutralizing condition compared to the Non-Neutralizing condition. Recall that the Discount model also failed to predict the over-generalization error cases. Overall, the *Neutralization model was unsuccessful at capturing the main experimental results for English speakers while also being less straightforward than the (more successful) Discount model.

Table 48. The predicted probabilities of the *Neutralization model and the aggregate experimental results for English speakers.

Non-Neutralizing condition							
		Frequent *ba-*		50-50 condition		Frequent *ni-*	
Input	Output	Model prediction	Experimental result	Model prediction	Experimental result	Model prediction	Experimental result
ba+velar-initial stem	ba+stem	98	98	98	96	97	89
	ba+palatalized stem	2	2	2	4	2	11
ni+velar-initial stem	ni+stem	48	26	35	26	27	29
	ni+palatalized stem	52	74	65	74	73	71

Neutralizing condition							
		Frequent *ba-*		50-50 condition		Frequent *ni-*	
Input	Output	Model prediction	Experimental result	Model prediction	Experimental result	Model prediction	Experimental result
ba+velar-initial stem	ba+stem	98	96	98	92	97	76
	ba+ palatalized stem	2	4	2	8	3	25
ni+velar-initial stem	ni+stem	48	55	35	41	27	38
	ni+ palatalized stem	52	45	65	59	73	62

4.6.2.2 Comparison to Korean dataset

Table 49 shows the initial preferred weights and final weights of the constraints. The final weights of *Ni-[+DORSAL] and PU were identical in the Non-Neutralizing and Neutralizing conditions. In both conditions, the final weight of *Ni-[+DORSAL] increased and that of PU decreased to motivate palatalization of velar-initial stems after *ni-*. In the Non-Neutralizing condition, the final weight of *Neutralization matched its initial preferred weight (see the footnote 66). This is because there were no observations of neutralization in the training data. However, the final weight of *Neutralization constraint also remained 0 in the Neutralizing condition where the neutralizing alternations were observed in the training data (see section 4.6.2.1).

Table 49. The initial preferred weights (μ) and the final weights of constraints the *Neutralization model (fit to the Korean dataset).

Non-Neutralizing condition			
Constraint	μ (preferred weight)	Final weight	
		Frequent *ba-* condition	Frequent *ni-* condition
*Ni-[+DORSAL]	0	4.06	4.66
PU	8	4.45	4.01
*Neutralization	0	0	0
Neutralizing condition			
Constraint	μ (preferred weight)	Final weight	
		Frequent *ba-* condition	Frequent *ni-* condition
*Ni-[+DORSAL]	0	4.06	4.66
PU	8	4.45	4.01
*Neutralization	0	0	0

Table 50. The predicted probabilities of the *Neutralization model and the aggregate experimental results for Korean speakers.

Non-Neutralizing condition					
		Frequent *ba-*		Frequent *ni-*	
Input	Output	Model prediction	Experimental result	Model prediction	Experimental result
ba+velar-initial stem	ba+stem	99	98	98	83
	ba+palatalized stem	1	2	2	17
ni+velar-initial stem	ni+stem	60	61	34	35
	ni+palatalized stem	40	39	66	65

Neutralizing condition					
		Frequent *ba-*		Frequent *ni-*	
Input	Output	Model prediction	Experimental result	Model prediction	Experimental result
ba+velar-initial stem	ba+stem	99	98	98	83
	ba+palatalized stem	1	2	2	17
ni+velar-initial stem	ni+stem	60	47	34	36
	ni+palatalized stem	40	53	66	64

The prediction of the model using the final constraint weights is reported in Table 50. The model's predictions were compared to the experimental results for Korean speakers. In both the Non-Neutralizing and Neutralizing conditions, the model accurately captured the fact that learners were successful at palatalizing the velar-initial stems after *ni-*. The percentage of the time that the ni+palatalized stem (shaded cells) was chosen was more frequent than the ba+palatalized stem. In both conditions, palatalization is predicted more frequently than the (faithful) stem after *ni-*. Furthermore, the model successfully captured the fact that Korean speakers were comparably successful at applying the neutralizing alternation and the non-neutralizing alternation because the model's predictions for underlyingly velar-initial stems after *ni-* were identical in the Non-Neutralizing and the Neutralizing conditions.

4.7 Application of the models to Yin and White (2018)

In this section, I examine whether the Discount and *Neutralization models can account for another set of data that also investigates learners' avoidance of neutralization and homophony. Yin and White (2018) showed that learners were less successful at learning neutralizing alternations that created homophony compared to non-neutralizing alternations; hence, I train the Discount and

*Neutralization models with Yin and White's experimental input and compare the models' predictions to the experimental results. I demonstrate that the Discount model again predicts Yin and White's experimental results more accurately than the *Neutralization model. In the following sections, I briefly review Yin and White's study and report the performance of both models.

4.7.1 Review of Yin and White's (2018) study

Yin and White examined the learnability of neutralizing and non-neutralizing alternations using an artificial grammar learning paradigm. During the learning phase, native English speakers were exposed to singular nonce-words and their plural forms, which were marked by adding a suffix /-i/ to singular words. The /-i/ suffix triggered palatalization of alveolar-final stems (/-t, -d, -s, -z/), which surfaced as [-ʧ, -ʤ, -ʃ, -ʒ] before /-i/. These alternations could be either non-neutralizing or neutralizing depending on the language that learners were exposed to.[64] Learners were exposed to both non-neutralizing and neutralizing alternations in all conditions; however, the percentage of time that neutralizing alterations created homophony varied. Depending on conditions, the neutralizing alternation triggered homophony either 0%, 50%, or 100% of the time. During the test phase, participants heard two plural forms and were asked to choose the correct plural form for a given singular noun. The test items had familiar and novel words. The percentage of time that English speakers chose the correct plural form for novel test items in the test phase is presented in Table 51.

[64] The experiment had two languages (Language A and B) to exclude the potential innate preference for certain palatalization rules. In each language, the neutralizing and non-neutralizing alternations were counterbalanced. For instance, in Language A, the palatalization of stem-final /t/ and /d/ to [ʧ] and [ʤ] before *-i* was neutralizing whereas the palatalization of stem-final /s/ and /z/ to [ʃ] and [ʒ] before *-i* was non-neutralizing (and *vice versa* in Language B).

Table 51. The percentage of the time that participants chose the correct plural form for novel test items in the test phase according to the types of alternation and amount of homophony (modified from Yin & White, 2018, Table 7).

	0% Homophony	50% Homophony	100% Homophony
Neutralizing alternation	66.9%	57.9%	47.5%
Non-neutralizing alternation	64.4%	69.2%	69.4%

The results show that learners were less successful at learning the neutralizing alternation than the non-neutralizing alternation when the neutralizing alternation created homophony. The percentage of the time that learners successfully chose the correct plural in the 0% homophony condition was significantly higher than in the 100% homophony condition for the neutralizing alternation. In the 0% homophony condition, learners chose the correct plural for the neutralizing alternation as successfully as the non-neutralizing alternation.

4.7.2 Constraints of the models

The Discount and *Neutralization models were trained with the same data used in Yin and White's experiments. In the training phase, the models considered two potential outputs, palatalized stem+i (where the stem was palatalized) or (faithful) stem+i (where the stem was not alternated). In the observed training data, the non-alveolar-final stem surfaced faithfully, and the alveolar-final stem palatalized.

The models trained with Yin and White's experimental input have two markedness constraints that can trigger alternations: *{t/d}i and *{s/z}i. During the experiment, learners were exposed to both alternations (i.e., /-t, -d/ → [-ʧi, -ʤi] and /-s, -z/ → [-ʃi, -ʒi]) but which alternations were neutralizing or non-neutralizing was counterbalanced according to the language that they were exposed to. For instance, in Language A, *{s/z}i could trigger the neutralizing alternation and *{t/d}i could trigger the non-neutralizing alternation (and *vice versa* in Language B). I refer to these constraints as the non-neutralizing markedness and neutralizing markedness constraints, which would stand in for *{t/d}i or *{s/z}i in whichever counterbalancing languages that learners were assigned to. Using these names for constraints, I abstract over the fact that one constraint triggered the non-neutralizing alternation and one triggered the neutralizing alternation. The non-neutralizing markedness constraint penalizes the output where the non-neutralizing alternation is *not* applied. For instance, stems ending with /s/ or /z/ should be palatalized when followed by the /-i/ plural suffix, triggering a non-neutralizing alternation (e.g., [-s] → [-ʃi]) in the cases where there was no /ʃ, ʒ/ as separate phonemes. When the non-neutralizing alternation is not applied before /-i/ (e.g., [-s] →[-si]), it violates the non-neutralizing markedness constraint. Similarly, the neutralizing markedness constraint penalizes the output where the neutralizing alternation is *not* applied. Both non-neutralizing and neutralizing markedness constraints received an initial preferred weight (μ) of 0. PU was also implemented in the model and received a preferred weight.

The Discount and *Neutralization models were trained with the experimental input of Yin and White's study and had the same set of parameters as in section 4.3.1.3. Moreover, the values of the parameters in the models trained with Yin and White's study followed those in the models fit to match the experimental results for English speakers by default (revisit sections 4.4.2 and 4.6.1 for selecting the parameters).

4.7.3 Testing the Discount model trained with Yin and White's study

Table 52 shows the initial preferred weight (μ) and final weights of constraints. In all conditions, the non-neutralizing markedness and the neutralizing markedness constraints picked up weight, motivating the palatalization. The weight of PU decreased after training because the input contained non-unified paradigms. Within the three homophony conditions, the weight of PU decreased less in the 100% homophony condition compared to the 0% homophony condition because the 100% homophony condition had fewer cases of alternation in the input due to the discounting of observations that trigger homophony.

Table 52. The initial preferred weights (μ) and the final weights of constraints in the Discount model (fit to the dataset of Yin and White).

Constraint	μ (preferred weight)	Final weight		
		0% homophony	50% homophony	100% homophony
Non-neutralizing markedness	0	3.47	3.55	3.74
Neutralizing markedness	0	3.47	3.03	2.02
PU	7.3	2.72	2.83	3.1

Table 53 below illustrates the prediction of the model using its final weights. The model accurately captured that learners were less successful at learning the neutralizing alternations than the non-neutralizing alternations when it triggered homophony. Importantly, the probability of applying the neutralizing palatalization was the lowest in the 100% homophony condition and the highest in the 0% homophony condition, which matches the aggregate experimental results.[65]

[65] Note that the free parameters were set to the values that best fit the experimental results in this thesis rather than to the values that best fit the Yin & White (2018) data. This could explain the slight mismatch between the models' prediction and the experimental results from Yin & White's paper (2018) in some conditions. However, the important point in this analysis is the decline in palatalization as homophony increases, which matches the Yin & White's results.

Table 53. The predicted probabilities of the Discount model and the aggregate experimental results for Yin and White's study.

		0% Homophony		50% homophony		100% homophony	
Input	Output	Model prediction	Experimental result	Model prediction	Experimental result	Model prediction	Experimental result
alveolar-final stem+i (non-neutralizing)	palatalized stem+i	68	64.4	67	69.2	65	69.4
alveolar-final stem+i (neutralizing)	palatalized stem+i	68	66.9	55	57.9	25	47.5

4.7.4 Testing the *Neutralization model trained with Yin and White's study

Table 54 shows the initial preferred weight (μ) and final weights of constraints in the *Neutralization model. The final weights of all constraints were identical across all homophony conditions. In all conditions, the final weight of PU reduced, and that of the non-neutralizing markedness and the neutralizing markedness increased, which encourages the application of alternation to alveolar-final stems before [-i]. The final weight of the *Neutralization constraint remained 0.

Table 55 shows the model's prediction of palatalization and the aggregate experimental results of Yin and White's study. The most noticeable finding is that the model's predictions of outputs were identical across given inputs in all homophony conditions, due to its identical learning weights across conditions. This finding shows that the *Neutralization model failed to account for the homophony avoidance effect (i.e., lower performance on the neutralizing alternations in the 100% homophony condition but not in the 0% homophony condition).

Table 54. The initial preferred weights (μ) and the final weights of constraints in the *Neutralization model (fit to the dataset of Yin and White).

Constraint	μ (preferred weight)	Final weight		
		0% homophony	50% homophony	100% homophony
Non-neutralizing markedness	0	3.2	3.2	3.2
Neutralizing markedness	0	3.2	3.2	3.2
PU	7	2.71	2.71	2.71
*Neutralization	0	0	0	0

Table 55. The predicted probabilities of the *Neutralization model and the aggregate experimental results for Yin and White's study.

		0% Homophony		50% homophony		100% homophony	
Input	Output	Model prediction	Experimental result	Model prediction	Experimental result	Model prediction	Experimental result
alveolar-final stem+i (non-neutralizing)	palatalized stem+i	62	64.4	62	69.2	62	69.4
alveolar-final stem+i (neutralizing)	palatalized stem+i	62	66.9	62	57.9	62	47.5

4.8 Discussion

In this chapter, I argued that the Discount model accounted for learners' avoidance of neutralization and homophony in a simple and straightforward way. Based on learners' general preference for a one-to-one relationship, the inputs that created homophony were discounted in the data used to train the Discount model. The model also accurately captured the different learnability of the neutralizing alternation between English speakers and Korean speakers. The model that best fit the English dataset showed that English speakers discounted the inputs that triggered homophony by 75%. At this level of discount, the model successfully matches the experimental results showing that English speakers were less successful at learning the neutralizing alternation than the non-neutralizing alternation. In contrast, the best-performing model for the Korean dataset showed that the observations were *not* discounted, reflecting the experimental result that Korean speakers were equally good at learning the neutralizing alternation and the non-neutralizing alternation. Moreover, when the Discount model was trained with another dataset from Yin and White's study, it successfully accounted for the fact that English speakers' performance on neutralizing alternations decreased as the amount of homophony triggered by those alternations increased.

Additionally, I considered an alternative model that accounts for the same avoidance effect using a constraint against neutralization. I showed that the *Neutralization model did not accurately capture the experimental results. The model that best fit the English and Korean datasets had a wide range of parameters and it failed to capture the different learning results of the neutralizing alternation between Korean and English speakers. When trained on data from Yin and White's study, the *Neutralization model also did not derive the difference in learnability of neutralization in any of the conditions.

The highly accurate predictions of the Discount model show the strength of this straightforward approach to the issue of homophony avoidance in learning. Instead of implementing a constraint that penalizes neutralization, the Discount model modified the number of observations that triggered homophony to reflect the fact that learners give less attention to the input that trigger homophony due to their preference for one-to-one patterns. The model not only accurately captured the difference in learning neutralizing and non-neutralizing alternations but also the difference between English and Korean speakers.

4.8.1 The L1 effect on learners' preference for one-to-one relationships

The different discount levels in the best models for English and Korean speakers (75% and 0%, respectively) indicate that prior exposure to neutralizing alternations can reduce learners' bias against homophony-creating neutralizations that present many-to-one relationships. In Chapter 3, I showed that Korean speakers are more likely to be frequently exposed to various neutralizing alternations that can trigger homophony in their L1 compared to English speakers. Neutralizations present a many-to-one relationship because multiple phonemes or lexical items have one surface form, which means that Korean speakers are likely to have more exposure to many-to-one relationships in their L1. The finding that a 0% discount best matched the Korean dataset supports the conclusion that Korean speakers are more accepting of novel many-to-one relationships during learning, presumably due to their extensive exposure to neutralization in their L1.

However, the model that best fit the English dataset discounted the homophony patterns by 75%. This result suggests that, without significant prior exposure to neutralization, learners are more strongly affected by the bias against a many-to-one relationship. Accordingly, they may need more pieces of evidence and a longer time to learn a many-to-one relationship than a one-to-one relationship. This finding is also in line with the idea that children are likely to ignore observations showing a many-to-one relationship, particularly in the early stage of learning (Pinker, 1996), which implies that more time and exposure to many-to-one relationships is needed for them to successfully learn such relationships compared to a one-to-one relationship.

4.8.2 Over-generalization error cases

While the Discount model accurately predicted the experimental results compared to the *Neutralization model, both models failed to capture the fact that English speakers made more over-generalization errors (palatalization after *ba-*) when the alternation was neutralizing than when it was non-neutralizing. However, the model chose palatalization after *ba-* equally often in both conditions. This finding is not particularly surprising given that the model was not implemented with any constraint that could account for this over-generalization (i.e., nothing in the model motivated palatalization after *ba-*). One possible solution for this issue could be implementing an additional product-oriented constraint that penalizes the output that does not palatalize. The constraint can motivate the palatalization of velar-initial stem after *ba-* and can potentially help the model to predict the over-generalization cases in the experimental results. This issue will not be addressed further in this thesis and will be left for future study.

4.8.3 Learning the discounts

The Discount models that best fit English and Korean datasets had different levels of discount for input data causing derived homophony (English: 75% discount vs. Korean 0% discount). This observation raises the question of how this difference in discount level comes to be. The discount in the Discount model is a type of prior that is based on learners' prior assumption about the data. I propose that children would initially have a strong prior bias against derived homophony, which would be formalized in the Discount model as a moderately high discount for homophonous patterns between 75% and 100% (but less than 100%, else learning would not be possible as all relevant data would be ignored). Over time, exposure to neutralizing patterns in the L1 would gradually reduce this prior bias (i.e., lowering the amount of the discount). By adulthood, a speaker exposed to a language with abundant neutralization (e.g., Korean) would have a very weak prior bias whereas a speaker exposed to a language with less neutralization (e.g., English) would still have a stronger one.

Both English and Korean children are likely to start with a high level of discount for homophony (near 100% discount). In childhood, Korean children's discount for homophony is still high and they are biased against homophony and neutralization. As Korean speakers grow up, they are exposed to a lot of neutralization in their L1, which begins to reduce their discount level. By adulthood, Korean speakers' discount level reaches 0%, as with the participants in the experiments. English speakers are exposed to some neutralization in their L1 but less than in Korean, which slightly lowers their discount level but maintains it higher into adulthood, as with the English participants in the experiments.

Under this hypothesis, all children start with a universal bias against derived homophony (i.e., a high discount level). This strong universal bias can give subtle pressure on languages over generations and gradually shift languages to avoid homophonous patterns. This explains how even in a language like Korean, there is diachronic pressure to avoid homophony (Silverman, 2010; Kaplan, 2011) driven by children's prior bias. However, with much exposure to neutralization and remaining homophony in the L1, the bias reduces by the time Korean speakers are adults.

CHAPTER 5

General Conclusion

5.1 Summary of the dissertation

This dissertation examined how various factors affect the acquisition of a probabilistic morpho-phonological pattern. I focused particularly on the effects of the following factors: learners' avoidance of phonological alternations, neutralization, and homophony, as well as their previous exposure to their native language (L1). I also tested how these factors interact with the relative frequency of variants in the input. The individual effect of each factor has been frequently studied; however, whether these factors interact and have a cumulative effect on language learning has been understudied (except for recent studies done by Kapatsinski and colleagues; Kapatsinski, 2009, 2010, 2012, 2013, 2018; Stave et al., 2013; Smolek & Kapatsinski, 2018; Smolek, 2019). In this thesis, I examined the interaction of these factors by controlling individual factors and testing multiple factors through variations on the same experiment.

In Chapter 2, the experimental results for adult native English speakers showed that they were biased against neutralization that generates derived homophony. I showed that English speakers were less successful at learning the alternation when it was neutralizing and homophony-creating than when it was non-neutralizing. This finding shows that their biases against neutralization and homophony can obstruct learning a novel neutralizing pattern. Furthermore, learners generally matched the frequency of prefix forms presented in the input to their output, showing probability-matching behaviour, except for one condition. Learners shifted away from the prefix that triggered the non-neutralizing alternation when it was already infrequent in the input. This interaction suggests that paradigm uniformity can cause learners to prefer morphological constructions that do not trigger a phonological alternation, but mainly when the alternation-triggering constructions are already infrequent. This could play a role in shaping the likelihood that novel morphological (and perhaps lexical) variants enter a language over time.

In Chapter 3, the same experiment was replicated with adult native Korean speakers. I showed that Korean speakers learned the neutralizing and non-neutralizing alternations equally well. This finding indicates that Korean speakers were more accepting of a novel neutralizing alternation compared to English speakers, which I argue is due to their extensive prior exposure to neutralizations in their L1. In terms of prefix selection, Korean speakers closely matched the probability of prefix forms presented in the input to their output in all conditions, repeating the finding that adults generally show frequent-matching behaviour.

In Chapter 4, I compared two approaches to account for the avoidance of neutralization and homophony. In the Discount model, the avoidance effect was implemented as a discounting of the input data that triggered derived homophony. The model represented a simple and straightforward way to explain learners' avoidance of neutralization and homophony. The model also successfully predicted the different experimental results between English and Korean speakers. The approach taken in the Discount model suggests a way to model a cognitive preference without creating a specific constraint in learners' grammars by placing the preference in the learning process rather than in the grammars. Additionally, in the *Neutralization model, the avoidance effect was implemented as a *Neutralization constraint within the grammars that penalized phonological neutralization. The *Neutralization mode failed to predict the experimental results.

The general findings suggest that various factors have influence on the learnability of phonological alternations and the choice of morphological constructions. Overall, the findings of this dissertation contribute to the understanding of how various factors can motivate biased language learning and show the importance of considering multiple factors when examining language learning. The following section offers suggestions for future studies.

5.2 Future research plan

The findings of this dissertation raise a further follow-up question: would Korean-speaking children behave differently in the experiment than Korean-speaking adults, and if so, how? Several factors might be relevant here. First, Korean-speaking adults in the current study did not show reduced learning of the novel neutralizing alternation, presumably due to their abundant prior exposure to neutralizations in their L1; however, Korean children have much less exposure to their L1. This briefer amount of exposure might be not enough to overcome their natural bias against neutralization that creates derived homophony. Secondly, Korean children are still within the critical period for language acquisition and therefore may rely on somewhat different learning mechanisms than adult learners (e.g., see discussion in Schuler, 2017). Lastly, children have lower working memory and less well-developed cognitive capacity compared to adults, which may impact their ability to track patterns presented in the input (Ullman et al., 1997; Ullman, 2001b; Schuler, 2017).

We know that from research that children often behave differently than adults when it comes to language learning tasks. For instance, Do (2018) reported on a particularly relevant example involving Korean-speaking children and adults. When asked to produce specific constructions that required applying phonological alternation, Korean children instead used alternative morphological structures as a way to avoid phonological alternations, a strategy that adults never used. In the current study, participants were required to learn a novel alternation while also having the option to choose between two morphological constructions, one that triggered the alternation and one that did not. Due to the factors above, Korean children might behave differently than adults. For instance, they may be more likely to choose the construction that avoids the alternation or more likely to regularize in general (Ross, 2001; Senghas & Coppola, 2001; Singleton & Newport, 2004; Hudson Kam & Newport, 2005, 2009). Additionally, they can be poorer at learning a novel neutralizing alternation (i.e., more like English-speaking adults) compared to Korean-speaking adults due to their baseline assumptions about mutual exclusivity (Markman & Wachtel, 1988) and less frequent exposure to neutralizing alternations in Korean. An advantage of the Discount modelling approach developed in this thesis is that it provides a straightforward way to measure how much the children are affected by homophony avoidance; specifically, the level of discount in the best fitting model for children can be compared to that of the English-speaking adults (75% discount) and Korean-speaking adults (0% discount).

Another question to be answered in a future study is what prompts adult learners to regularize their output. Multiple elements can affect adult learners' acquisition of probabilistic morpho-phonological patterns, such as the probability of the frequency variant in the input. In the current study, there were two variants presented in the input, and the frequent variant was presented 66.7% of the time. The learners in the current study generally matched the relative frequency of variants in the input to their output. In Culbertson et al.'s study (2012), the input also had two variants, and the frequent variant was presented 70% of the time. Learners in their study generally regularized their output with the frequent variant unless they were biased against the frequent variant, implying that a slightly higher probability of the frequent variant (70% vs. 66.7%) in the input can potentially induce regularisation (Culbertson p.c.). Additionally, the number and consistency of variants in the input can also influence learning a probabilistic pattern. For example, in Hudson Kam and Newport's study (2009), there was one frequent variant (presented 60% of the time) and a number of infrequent variants (two to 16 variants), and learners regularized their output with the frequent variant when the various infrequent variants were presented inconsistently (probabilistically appeared across lexical items). Finally, other factors such as experimental design (forced choice or production), syntax (prefix or suffix), or the task in the experiment (learning sentences, word orders, morphological structures, or allomorphs) might affect whether adults regularize in a given experiment. Although many factors seem to influence whether adult learners regularize their outputs, it is less clear whether each factor has a particular threshold to initiate the regularization, which individual factor has the strongest effect, and whether a certain combination of factors is required. These questions can be answered in future studies by modifying individual elements and by collectively testing them within the same experiment.

Reference

Albright, A. (2005). The morphological basis of paradigm leveling. In Downing, L., Hall, T., A., & Raffelsiefen, R. (Eds.), *Paradigms in Phonological Theory*. Oxford University Press.

Altenberg, E. P., & Vago, R. M. (1983). Theoretical implications of an error analysis of second language phonology production. *Language learning*, *33*(4), 427–447. https://doi.org/10.1111/j.1467-1770.1983.tb00943.x

Aronoff, M. (1976). Word formation in generative grammar. *Linguistic Inquiry Monographs Cambridge, Mass*, (1), 1–134.

Austin, A. C. (2010). *When children learn more than what they are taught: Regularization in child and adult learners*. University of Rochester.

Baayen, R. H., Piepenbrock, R., & Gulikers, L. (1995). The CELEX lexical database (release 2). *Distributed by the linguistic data consortium*, University of Pennsylvania.

Baird, M. (2021). Deriving frequency effects from biases in learning. *Proceedings of the Linguistic Society of America*, *6*(1), 514–525. https://doi.org/10.3765/plsa.v6i1.4986

Barr, D. J., Levy, R., Scheepers, C., & Tily, H. J. (2013). Random effects structure for confirmatory hypothesis testing: Keep it maximal. *Journal of memory and language*, *68*(3), 255–278. https://doi.org/10.1016/j.jml.2012.11.001

Bateman, N. (2007). *A crosslinguistic investigation of palatalization*. University of California, San Diego.

Bates, D., Maechler, M., Bolker, B., & Walker, S., (2015). Fitting Linear Mixed-Effects Models Using lme4. *Journal of Statistical Software*, *67*(1), 1–48. doi:10.18637/jss.v067.i01.

Benua, L. (1995). Identity effects in morphological truncation. In Beckman, J., Urbanczyk, S., & Walsh, L. (Eds.), *University of Massachusetts Occasional Papers in Linguistics 18: Papers in Optimality Theory* (pp.77–136), Graduate Linguistic Student Association, Amherst, Massachusetts.

Bhela, B. (1999). Native language interference in learning a second language: Exploratory case studies of native language interference with target language usage. *International Education Journal*, *1*(1), 22–44.

Birdsong, D. (1999). Introduction: Whys and why nots of the critical period hypothesis for second language acquisition. *Second language acquisition and the critical period hypothesis*, 1–22.

Blevins, J., & Wedel, A. (2009). Inhibited sound change: An evolutionary approach to lexical competition. *Diachronica, 26*(2), 143–183. doi: https://doi.org/10.1075/dia.26.2.01ble

Boersma, P., & Levelt, C. (2000). Gradual constraint-ranking learning algorithm predicts acquisition order. In *Proceedings of Child Language Research Forum* 30 (pp. 229–237).

Breiss, C. (2021). *Lexical Conservatism in phonology: theory, experiments, and computational modeling*. University of California, Los Angeles.

Bresnan, J., & Ford, M. (2010). Predicting syntax: Processing dative constructions in American and Australian varieties of English. *Language*, *86*(1), 168–213. doi:10.1353/lan.0.0189

Brière, E. J. (1966). An Investigation of Phonological Interference. *Language*, *42*(4), 768–796. https://doi.org/10.2307/411832

Broselow, E. (1984). An investigation of transfer in second language phonology. *22*(4), 253–270. https://doi.org/10.1515/iral.1984.22.4.253

Brown, C. A. (1998). The role of the L1 grammar in the L2 acquisition of segmental structure. *Second Language Research*, *14*(2), 136–193. https://doi.org/10.1191/026765898669508401

Burzio, L. (1994). *Principles of English stress* (No. 72). Cambridge University Press.

Burzio, L. (1998). Multiple correspondence. *Lingua*, *104*(1-2), 79–109. https://doi.org/10.1016/S0024-3841(97)00025-9

Bybee, J. (2000b). The phonology of the lexicon: Evidence from lexical diffusion. In Michael, B., & Suzanne, K. (Eds.), *Usage-based models of language* (pp. 65–85), Stanford, CSLI Publications.

Bybee, J. (2001). Language Use as Part of Linguistic Theory. In *Phonology and Language Use* (Cambridge Studies in Linguistics, pp. 1–18). Cambridge: Cambridge University Press. doi:10.1017/CBO9780511612886.00

Bybee, J. & Brewer, M. A. (1980). Explanation in morphophonemics: changes in Provençal and Spanish preterite forms. *Lingua*, *52*(3-4), 201–242. https://doi.org/10.1016/0024-3841(80)90035-2

Chung, K. (1980). *Korean neutralization* [Doctoral dissertation, University of Texas at Austin].

Cover, T. M., & Thomas, J. A. (1991). Entropy, relative entropy and mutual information. *Elements of information theory*, *2*(1), 12–13.

Cowan, N. (1997). *The development of memory in childhood.* Cowan, N. (Ed.), Hove, East Sussex: Psychology Press.

Crosswhite, K. (1999). Intra-paradigmatic homophony avoidance in two dialects of Slavic. *UCLA working papers in linguistics*, *1*(3), 48–67.

Crosswhite, K., & Jun, A. (2001). *Vowel reduction in optimality theory*. Psychology Press.

Culbertson, J. (2023) Artificial language learning. In Sprouse, J. (Ed.), *Oxford Handbook of Experimental Syntax*. Oxford University Press, New York, NY.

Culbertson, J., Smolensky, P., & Legendre, G. (2012). Learning biases predict a word order universal. *Cognition*, *122*(3), 306–329. https://doi.org/10.1016/j.cognition.2011.10.017

Curtin, S., & Zuraw, K. (2002). Explaining constraint demotion in a developing system. In *Proceedings of the Boston University conference on language development* (26, pp. 118-129).

Davis, S., & Shin, S. H. (1999). The syllable contact constraint in Korean: An optimality-theoretic analysis. *Journal of East Asian Linguistics*, *8*(4), 285–312. https://doi.org/10.1023/A:1008335414055

Della Pietra, S., Della Pietra, V., & Lafferty, J. (1997). Inducing features of random fields. *IEEE transactions on pattern analysis and machine intelligence*, *19*(4), 380–393. doi: 10.1109/34.588021.

Diamond, A., & Doar, B. (1989). The performance of human infants on a measure of frontal cortex function, the delayed response task. *Developmental Psychobiology: The Journal of the International Society for Developmental Psychobiology*, *22*(3), 271–294. https://doi.org/10.1002/dev.420220307

Dinnsen, D. A. (1985). A re-examination of phonological neutralization1. *Journal of linguistics*, *21*(2), 265–279. https://www.jstor.org/stable/4175789

Do, Y. (2018). Paradigm uniformity bias in the learning of Korean verbal inflections. *Phonology*, *35*(4), 547–575. doi:10.1017/S0952675718000209

Dost, I. N., & Bohloulzadeh, G. (2017). A review of contrastive analysis hypothesis with a phonological and syntactical view: a cross-linguistic study. *The Buckingham Journal of language and linguistics*, *10*, 32–41. https://doi.org/10.5750/bjll.v10i0.1482

Dulay, H., Burt, M. K., & Krashen, S. D. (1982). *Language two*. Oxford University Press. New York

Elman, J. L. (1993). Learning and development in neural networks: The importance of starting small. *Cognition*, *48*(1), 71–99. https://doi.org/10.1016/0010-0277(93)90058-4

Ernestus, M., & Baayen, R. H. (2003). Predicting the unpredictable: Interpreting neutralized segments in Dutch. *Language*, *79*(1), 5–38. doi:10.1353/lan.2003.0076.

Finn, A. S., & Hudson Kam, C. L. (2008). The curse of knowledge: first language knowledge impairs adult learners' use of novel statistics for word segmentation. *Cognition*, *108*(2), 477–499. https://doi.org/10.1016/j.cognition.2008.04.002

Fischer, J. L. (1958). Social Influences on the Choice of a Linguistic Variant, *Word*, *14*(1), 47–56, doi: 10.1080/00437956.1958.11659655

Fisiak, J. (1981). *Contrastive Linguistics and the Language Teacher*. Fisiak, J. (Ed.), New York: Pergamon.

Flemming, E. S. (2002). *Auditory Representations in Phonology* (1st ed.). Routledge. *https://doi.org/10.4324/9781315054803*

Fredriksson, A., Schröder, N., Eriksson, P., Izquierdo, I., & Archer, T. (2000). Maze learning and motor activity deficits in adult mice induced by iron exposure during a critical postnatal period. *Developmental Brain Research*, *119*(1), 65–74. https://doi.org/10.1016/S0165-3806(99)00160-1

Fries, C. C. (1945). *Teaching and learning English as a foreign language*.

Gardner, R. A. (1957). Probability-Learning with Two and Three Choices. *The American Journal of Psychology*, *70*(2), 174–185, https://doi.org/10.2307/1419319

Gass, S. (1979). Language transfer and universal grammatical relations. *Language Learning*, *29(2)*, 327–44. https://doi.org/10.1111/j.1467-1770.1979.tb01073.x

Gathercole, S. (1998). The Development of Memory. *The Journal of Child Psychology and Psychiatry and Allied Disciplines, 39*(1), 3–27. doi:10.1017/S0021963097001753

Gnanadesikan, A. (1995). Markedness and faithfulness constraints in child phonology. Ms., University of Massachusetts, Amherst

Goldwater, S., & Johnson, M. (2003). Learning OT constraint rankings using a maximum entropy model. In *Proceedings of the Workshop on Variation within Optimality Theory* (pp. 111–120). http://homepages.inf.ed.ac.uk/sgwater/papers/OTvar03.pdf

Greenberg, J. H. (1963). Some universals of grammar with particular reference to the order of meaningful elements. *Universals of language*, *2*, 73–113.

Hammarberg, B. (1990). Conditions on transfer in second language phonology acquisition. In *The 1990 Amsterdam Symposium on the Acquisition of Second-Language Speech*. University of Amsterdam.

Hawkins, J. A. (1983). *Word order universals*. New York: Academic Press.

Hayes, B. (1997). Anticorrespondence in Yidiɲ. Ms., UCLA.

Hayes, B. (2004). Phonological acquisition in Optimality Theory: the early stages. In Kager, R., Pater, J., & Zonneveld. W. (Eds.), *Constraints in phonological acquisition* (pp. 158–203). Cambridge University Press.

Hayes, B. (2022). Deriving the wug-shaped curve: A criterion for assessing formal theories of linguistic variation. *Annual Review of Linguistics*, *8*, 473–494.

Hayes, B., & Londe, Z. (2006). Stochastic phonological knowledge: The case of Hungarian vowel harmony. *Phonology, 23*(1), 59–104. doi:10.1017/S0952675706000765

Hayes, B., & Wilson, C. (2008). A maximum entropy model of phonotactics and phonotactic learning. *Linguistic inquiry*, *39*(3), 379–440. doi:https://doi.org/10.1162/ling.2008.39.3.379

Hayes, B., Siptár, P., Zuraw, K., & Londe, Z. (2009). Natural and Unnatural Constraints in Hungarian Vowel Harmony. *Language*, *85*(4), 822–863. http://www.jstor.org/stable/40492955

Hock, H. H. (1991). *Principles of historical linguistics* (2nd ed.). The Hague: Mouton de Gruyter. https://doi.org/10.1515/9783110219135

Hockett, C. F. (1967). The quantification of functional load. *Word*, *23*(1-3), 320–339. doi: 10.1080/00437956.1967.11435484

Hooper, J. B. (1976). *An introduction to natural generative phonology*. New York: Academic Press.

Houston, A. C. (1985). *Continuity And Change In English Morphology: The Variable (ing)* (Order No. 8515390). [Doctoral dissertation, University of Pennsylvania, Philadelphia]. ProQuest Dissertations & Theses Global; ProQuest One Literature. (303373625). https://www.proquest.com/dissertations-theses/continuity-change-english-morphology-variable-ing/docview/303373625/se-2

Hudson Kam, C. L., & Newport, E. L. (2005). Regularizing unpredictable variation: The roles of adult and child learners in language formation and change. *Language learning and development*, *1*(2), 151–195. https://doi.org/10.1080/15475441.2005.9684215

Hudson Kam, C. L., & Newport, E. L. (2009). Getting It Right by Getting It Wrong: When Learners Change Languages. *Cognitive Psychology, 59*(1), 30–66. https://doi.org/10.1016/j.cogpsych.2009.01.001

Hughto, C., Lamont, A., Prickett, B., & Jarosz, G. (2019). Learning exceptionality and variation with lexically scaled MaxEnt. In *Proceedings of the Society for Computation in Linguistics (SCiL)* (2019, pp. 91–101).

Hurford, J. (2003). The interaction between numerals and nouns. In Plank, F. (Ed.), *Noun phrase structure in the languages of Europe* (pp. 561–620). New York: Mouton de Gruyter.

Hwang, J. (2008). Neutralizing nasalization in Korean. Ms., San José State University.

Ichimura, L. K. (2006). *Anti-homophony blocking and its productivity in transparadigmatic relations* [Doctoral dissertation, University of Massachusetts Boston].

Ito, J., & Mester, A. (2004). Morphological contrast and merger: ranuki in Japanese. *Journal of Japanese Linguistics*, *20*(1), 1–18. https://doi.org/10.1515/jjl-2004-0103

Jaeger, T. F. (2008). Categorical data analysis: Away from ANOVAs (transformation or not) and towards logit mixed models. *Journal of memory and language*, *59*(4), 434–446. https://doi.org/10.1016/j.jml.2007.11.007

Jäger, G. (2004). *Maximum Entropy models and Stochastic Optimality Theory*. Rutgers Optimality Archive 625.

Jakobson, R. (1931). Prinzipien der historischen Phonologie. Travaux Du Cercle *Linguistique de Prague*, *4*, 246–267.

James, C. (1980). *Contrastive analysis*. Longman.

James, J. (1977). Language transfer reconsidered. *Interlanguage studies bulletin*, *2*(3), 7–21.

Janacsek, K., Fiser, J., & Nemeth, D. (2012). The best time to acquire new skills: Age-related differences in implicit sequence learning across the human lifespan. *Developmental science*, *15*(4), 496–505. https://doi.org/10.1111/j.1467-7687.2012.01150.x

Johnson, J. S., & Newport, E. L. (1989). Critical period effects in second language learning: The influence of maturational state on the acquisition of English as a second language. *Cognitive psychology*, *21*(1), 60–99. https://doi.org/10.1016/0010-0285(89)90003-0

Kapatsinski, V. (2009). *The architecture of grammar in artificial grammar learning: Formal biases in the acquisition of morphophonology and the nature of the learning task* [Doctoral dissertation, Indiana University].

Kapatsinski, V. (2010). Velar palatalization in Russian and artificial grammar: Constraints on models of morphophonology. *Laboratory Phonology*, *1*(2), 361–393. https://doi.org/10.1515/labphon.2010.019

Kapatsinski, V. (2012). What statistics do learners track? Rules, constraints and schemas in (artificial) grammar learning. In Gries, S., & Divjak, D. (Eds.)*, 1. Frequency effects in language learning and processing* (pp. 53–82). Berlin, Boston: De Gruyter Mouton. https://doi.org/10.1515/9783110274059.53

Kapatsinski, V. (2013). Conspiring to mean: experimental and computational evidence for a usage-based harmonic approach to morphophonology. *Language*, *89*(1), 110–148. http://www.jstor.org/stable/23357723

Kapatsinski, V. (2018). *Changing minds changing tools: From learning theory to language acquisition to language change*. MIT Press.

Kaplan, A. (2011). How much homophony is normal? *Journal of Linguistics, 47*(3), 631–671. doi:10.1017/S0022226711000053

Kawahara, S. (2003). Root-controlled fusion in Zoque: Root-faith and neutralization avoidance. Ms., University of Massachusetts, Amherst. [Rutgers Optimality Archive 599-0403].

Kazazis, K. (1969). Possible evidence for (near-) underlying forms in the speech of a child. In *Proceedings of the Chicago Linguistic Society* (*5*, pp.382–388). Chicago: CLS.

Keating, P., & Lahiri, A. (1993). Fronted velars, palatalized velars, and palatals. *Phonetica*, *50* (2), 73–101. https://doi.org/10.1159/000261928

Keller, F. (2000). *Gradience in grammar: Experimental and computational aspects of degrees of grammaticality* [Doctoral dissertation, University of Edinburgh].

Kenstowicz, M. (2002). Paradigmatic uniformity and contrast. *MIT Working Papers in Linguistics 42*, 141–163.

Kerkhoff, A. O. (2007). *Acquisition of morpho-phonology: The Dutch voicing alternation* [Doctoral dissertation, LOT, Utrecht University].

Kim-Renaud, Y. K. (1974). *Korean consonantal phonology* [Doctoral dissertation, University of Hawaii].

Kim, C. W. (1979). Neutralization in Korean revisited. In *Studies in the Linguistic Sciences* (*9*(1), pp. 147–156).

Kim, H., & Jongman, A. (1996). Acoustic and perceptual evidence for complete neutralization of manner of articulation in Korean. *Journal of Phonetics*, *24*(3), 295–312. https://doi.org/10.1006/jpho.1996.0016

King, R. D. (1967). Functional Load and Sound Change. *Language*, *43*(4), 831–852. https://doi.org/10.2307/411969

Krug, M. (1998). String Frequency: A Cognitive Motivating Factor in Coalescence, Language Processing, and Linguistic Change. *Journal of English Linguistics*, *26*(4), 286–320. https://doi.org/10.1177/007542429802600402

Lado, R. (1957). *Linguistics across cultures: Applied linguistics for language teachers*. University of Michigan Press.

Lee, K. M. (1972). Kwukosakesul. In *An introduction to the history of the Korean language*. Seoul: Tower Press.

Legendre, G., Miyata, Y., & Smolensky, P. (1990). Harmonic Grammar–A formal multi-level connectionist theory of linguistic well-formedness: An application. In *12th Annual Conf. CSS Pod* (pp. 884–891). Psychology Press.

Lincoln, P. C. (1976a). *Describing Banoni, an Austronesian Language of Southwest Bougainville.* [Doctoral dissertation, Honolulu: University of Hawaii].

Linzen, T., Kasyanenko, S., & Gouskova, M. (2013). Lexical and phonological variation in Russian prepositions. *Phonology, 30*(3), 453–515. doi:10.1017/S0952675713000225

Lott, D. (1983). Analysing and counteracting interference errors. *ELT Journal*, *37*(3), 256–261. https://doi.org/10.1093/elt/37.3.256

Markman, E. M., & Wachtel, G. F. (1988). Children's use of mutual exclusivity to constrain the meanings of words. *Cognitive psychology*, *20*(2), 121–157. https://doi.org/10.1016/0010-0285(88)90017-5

Martin, A., & Culbertson, J. (2020). Revisiting the suffixing preference: Native-language affixation patterns influence perception of sequences. *Psychological Science*, *31*(9), 1107–1116. https://doi.org/10.1177/09567976209311

Martin, S. E. (1951). Korean Phonemics. *Language*, *27*(4), 519–533. https://doi.org/10.2307/410039

Martin, S. E. (1992). *A reference grammar of Korean: A complete guide to the grammar and history of the Korean language*. Tuttle Publishing.

Martinet, A. (1952). Function, Structure, and Sound Change, *WORD*, *8*(1), 1–32, doi:10.1080/00437956.1952.11659416

Mathesius, V. (1931). Zum Problem der Belastungs- und kombinationsfähigkeit der Phoneme. *Travaux Du Cercle Linguistique de Prague*, *4*, 148–152.

McCarthy, J. J. (1998). Morpheme structure constraints and paradigm occultation. In Gruber, M. C., Higgins, D., Olson, K., & Wysocki, T. (Eds.), *CLS 32, Part 2: The Panels* (pp. 123–150). Chicago, IL: Chicago Linguistic Society.

McCarthy, J. J., & Prince, A. (1995). Faithfulness and reduplicative identity. *Linguistics Department Faculty Publication Series*, 10.

McMullin, K., & Hansson, G.Ó. (2019). Inductive learning of locality relations in segmental phonology. *Laboratory Phonology, 10*(1), 14. https://doi.org/10.5334/ labphon.150

Moreton, E. (2008). Analytic bias and phonological typology. *Phonology, 25*(1), 83–127. doi:10.1017/S0952675708001413

Moulton, W. G. (1962). The vowels of Dutch: phonetic and distributional classes. *Lingua*, *11*, 294–312. https://doi.org/10.1016/0024-3841(62)90038-4

Murphy, S. (2003). Second language transfer during third language acquisition. *Studies in Applied Linguistics and TESOL*, *3*(2). https://doi.org/10.7916/salt.v3i2.1625

Newport, E. (1990). Maturational constraints on language learning. *Cognitive Science, 14*(1), 11–28. https://doi.org/10.1016/0364-0213(90)90024-Q

Odlin, T. (1989). Phonetics, phonology, and writing systems. In *Language Transfer: Cross-Linguistic Influence in Language Learning* (pp. 112–128), Cambridge Applied Linguistics, Cambridge: Cambridge University Press. doi:10.1017/CBO9781139524537.009

Onnis, L., & Thiessen, E. (2013). Language experience changes subsequent learning. *Cognition*, *126*(2), 268–284. https://doi.org/10.1016/j.cognition.2012.10.008

Padgett, J. (2003a). Contrast and post-velar fronting in Russian. *Natural Language & Linguistic Theory*, *21*(1), 39-87. https://doi.org/10.1023/A:1021879906505

Padgett, J. (2003b). The emergence of contrastive palatalization in Russian. In Holt. D. E. (Ed.), *Optimality Theory and Language Change* (pp. 307–335). Dordrecht: Kluwer Academic Press.

Phillips, B. S. (2001). Lexical diffusion, lexical frequency and lexical analysis. In Bybee, J., & Hopper, P. (Eds.), *Frequency and the emergence of linguistic structure* (pp. 123–136). Amsterdam, John Benjamins.

Pinker, S. (1996). *Language learnability and language development: with new commentary by the author*. Harvard University Press.

Powell, M. J. (2009). *The BOBYQA algorithm for bound constrained optimization without derivatives.* Cambridge NA Report NA2009/06, University of Cambridge, Cambridge, 26.

Prince, A., & Smolensky, P. (1993). *Optimality theory: constraint interaction in generative grammar.* Cambridge: Blackwell. [Published in 2004. First circulated in 1993].

Prince, A., & Tesar, B. (1999). Learning phonotactic distributions. *Constraints in phonological acquisition (2004,* pp. 245–291).

R Core Team (2018) R: A Language and Environment for Statistical Computing. *R Foundation for Statistical Computing, Vienna*. https://www.R-project.org

Reyes, A. M., Arechabaleta-Regulez, B., & Montrul, S. (2017). The acquisition of rhotics by child L2 and L3 learners. *Journal of Second Language Pronunciation*, *3*(2), 242–266. https://doi.org/10.1075/jslp.3.2.04rey

Rijkhoff, J. (1998). Word order in the languages of Europe: The simple noun phrase. In Siewierska, A. (Ed.), *Constituent order in the languages of Europe* (pp. 321–382). New York: Mouton de Gruyter.

Roberts, J. R. (1994). *Acquisition of variable rules:(-t, d) deletion and (ing) production in preschool children* [Doctoral dissertation, University of Pennsylvania].

Roelofs, A. (1996). Serial order in planning the production of successive morphemes of a word. *Journal of Memory and Language*, *35*(6), 854–876. https://doi.org/10.1006/jmla.1996.0044

Roettger, T. B., Winter, B., Grawunder, S., Kirby, J., & Grice, M. (2014). Assessing incomplete neutralization of final devoicing in German. *Journal of Phonetics*, *43*, 11–25. https://doi.org/10.1016/j.wocn.2014.01.002

Ross, D. S. (2001). *Disentangling the nature–nurture interaction in the language acquisition process: Evidence from deaf children of hearing parents exposed to non-native input.* University of Rochester.

Ross, D. S., & Newport, E. L. (1996). The development of language from non-native linguistic input. In *Proceedings of the 20th annual Boston University Conference on language development* (*2*, pp. 634–645). Boston: Cascadilla Press.

Rubach, J. (1980). Rule ordering in phonological interference. In Fisiak, J. (Ed.), *Contrastive linguistics: prospects and problems* (pp. 365–377). https://doi.org/10.1515/9783110824025

Rubach, J. (1984). Rule typology and phonological interference. In *Theoretical issues in contrastive phonology: Studies in descriptive linguistics* (pp. 37–50).

Sabourin, L., Stowe, L. A., & De Haan, G. J. (2006). Transfer effects in learning a second language grammatical gender system. *Second Language Research*, *22*(1), 1–29. https://doi.org/10.1191/0267658306sr259oa

Schachter, J. (1974). An error in error analysis 1. *Language learning*, *24*(2), 205–214. https://doi.org/10.1111/j.1467-1770.1974.tb00502.x

Schlaug, G. (2001). The brain of musicians. *Annals of the New York Academy of Sciences*, *930*(1), 281–299. https://doi.org/10.1111/j.1749-6632.2001.tb05739.x

Schuler, K. D. (2017). *The acquisition of productive rules in child and adult language learners*. Georgetown University.

Senghas, A., Coppola, M., Newport, E. L., & Supalla, T. (1997). Argument structure in Nicaraguan Sign Language: The emergence of grammatical devices. In *Proceedings of the 21st annual Boston university Conference on Language Development* (*21*(2), pp. 550–561), Boston: Cascadilla Press.

Senghas, A., & Coppola, M. (2001). Children Creating Language: How Nicaraguan Sign Language Acquired a Spatial Grammar. *Psychological Science*, *12*(4), 323–328. https://doi.org/10.1111/1467-9280.00359

Shannon, C. E. (1948). A mathematical theory of communication. *The Bell system technical journal*, *27*(3), 379–423. 10.1002/j.1538-7305.1948.tb01338.x

Silverman, D. (2010). Neutralization and anti-homophony in Korean. *Journal of Linguistics*, *46*(2), 453–482. doi:10.1017/S0022226709990247

Singleton, J. L., & Newport, E. L. (2004). When learners surpass their models: The acquisition of American Sign Language from impoverished input. *Cognitive Psychology*, *49*, 370–407. https://doi.org/10.1016/j.cogpsych.2004.05.001

Slowiaczek, L. M., & Dinnsen, D. A. (1985). On the neutralizing status of Polish word-final devoicing. *Journal of phonetics*, *13*(3), 325–341. https://doi.org/10.1016/S0095-4470(19)30763-6

Smolek, A. (2019). *Teaching papa to cha-cha: How change magnitude, temporal contiguity, and task affect alternation learning*. University of Oregon. https://scholarsbank.uoregon.edu/xmlui/handle/1794/25273

Smolek, A., & Kapatsinski, V. (2018). What happens to large changes? Saltation produces well-liked outputs that are hard to generate. *Laboratory Phonology, 9*(1), 10. https://doi.org/10.5334/labphon.93

Smolensky, P. (1986). *Information processing in dynamical systems: Foundations of harmony theory*. Colorado University at Boulder Dept of Computer Science.

Smolensky, P. (1996). The initial state and 'richness of the base' in Optimality Theory. Ms., Baltimore: Johns Hopkins University, Rutgers Optimality Archive, 293.

Smolensky, P., & Legendre, G. (2006). *The harmonic mind: From neural computation to optimality-theoretic grammar (Cognitive architecture).* MIT Press.

Sohn, H. M. (1999). *The Korean language*. Cambridge University Press.

Starling, S. J. (2013). *Learning and behavior in an uncertain world: Probability learning in children and adults*. University of Rochester.

Stave, M., Smolek, A., & Kapatsinski, V. (2013). Inductive bias against stem changes as perseveration: Experimental evidence for an articulatory approach to output-output faithfulness. *Proceedings of the Annual Meeting of the Cognitive Science Society, 35*, 3454–3459.

Steriade, D. (1994). Positional neutralization and the expression of contrast. *Ms., UCLA*.

Steriade, D. (2000). Paradigm uniformity and the phonetics-phonology boundary. *Papers in Laboratory Phonology V: Acquisition and the Lexicon*, 313–334.

Stockwell, R. P., Bowen, J. D., & Martin, J. W. (1965). *The grammatical structures of English and Spanish.* 4. University of Chicago Press.

Tang, K. & Baer-Henney, D., (2023). Modelling L1 and the artificial language during artificial language learning, *Laboratory Phonology, 14*(1). doi: https://doi.org/10.16995/labphon.6460.

Tessier, A. M. (2012). Testing for OO-Faithfulness in the acquisition of consonant clusters. *Language acquisition*, *19*(2), 144–173. https://doi.org/10.1080/10489223.2012.660552

Thompson-Schill, S. L., Ramscar, M., & Chrysikou, E. G. (2009). Cognition Without Control: When a Little Frontal Lobe Goes a Long Way. *Current Directions in Psychological Science*, *18*(5), 259–263. https://doi.org/10.1111/j.1467-8721.2009.01648.x

Tiersma, P. M. (1982). Local and General Markedness. *Language*, *58*(4), 832–849. https://doi.org/10.2307/413959

Trubetzkoy, N. S. (1939). Thoughts on the Indo-European problem. *Acta linguistica*, *1*(1), 81–89. https://doi.org/10.1080/03740463.1939.10410851

Türker, E. (2016). The role of L1 conceptual and linguistic knowledge and frequency in the acquisition of L2 metaphorical expressions. *Second Language Research*, *32*(1), 25–48. https://doi.org/10.1177/0267658315593336

Ullman, M. T. (2001b). A neurocognitive perspective on language: The declarative/procedural model. *Nature reviews neuroscience*, *2*(10), 717–726. https://doi.org/10.1038/35094573

Ullman, M. T., Corkin, S., Coppola, M., Hickok, G., Growdon, J. H., Koroshetz, W. J., & Pinker, S. (1997). A neural dissociation within language: Evidence that the mental dictionary is part of declarative memory, and that grammatical rules are processed by the procedural system. *Journal of cognitive neuroscience*, *9*(2), 266–276. doi: 10.1162/jocn.1997.9.2.266

Vago, R. M. (1980). *The sound pattern of Hungarian*. Georgetown University, School of Language.

van de Vijver, R., & Baer-Henney, D. (2014). Developing biases. *Frontiers in psychology*, *5*, 634.

Walton, K., Lieberman, D., Llinas, A., Begin, M., & Llinas, R. (1992). Identification of a critical period for motor development in neonatal rats. *Neuroscience*, *51*(4), 763–767. https://doi.org/10.1016/0306-4522(92)90517-6

Wardhaugh, R. (1970). The Contrastive Analysis Hypothesis. *TESOL Quarterly*, *4*(2), 123–130. https://doi.org/10.2307/3586182

Wedel, A., Kaplan, A., & Jackson, S. (2013). High functional load inhibits phonological contrast loss: A corpus study. *Cognition*, *128*(2), 179-186. https://doi.org/10.1016/j.cognition.2013.03.002

Weinart, F. E., & Schneider, W. (Eds.) (1995). *Memory performance and competencies: Issues in growth and development.* Mahwah, NJ: Lawrence Erlbaum Associates Inc.

Weinreich, U. (1953). *Languages in contact. Findings and problems*. The Hague.

Weir, M. W. (1964). Developmental changes in problem-solving strategies. *Psychological Review, 71*(6), 473–490. https://doi.org/10.1037/h0041785

Weir, M. W. (1972). Probability Performance: Reinforcement Procedure and Number of Alternatives. *The American Journal of Psychology*, *85*(2), 261–270. https://doi.org/10.2307/1420666

Wexler, K., & Culicover, P. W. (1980). *Formal principles of language acquisition*. MIT Press. Cambridge, Mass.

White, J. (2013). *Bias in phonological learning: Evidence from saltation* [Doctoral dissertation, UCLA].

White, J. (2014). Evidence for a learning bias against saltatory phonological alternations, *Cognition, 130*(1), 96–115, https://doi.org/10.1016/j.cognition.2013.09.008.

White, J. (2017). Accounting for the learnability of saltation in phonological theory: a maximum entropy model with a P-map bias. *Language*, *93*(1), 1–36. http://www.jstor.org/stable/26630343

White, J., Kager, R., Linzen, T., Markopoulos, G., Martin, A., Nevins, A., Peperkamp, S., Polgárdi, K., Topintzi, N., & van De Vijver, R. (2018). Preference for locality is affected by the prefix/suffix asymmetry: Evidence from artificial language learning. *NELS 48: Proceedings of the Forty-Eighth Annual Meeting of the North East Linguistic Society* (3, pp. 207–220). Amherst, MA: GLSA

Wilson, C. (2006). Learning phonology with substantive bias: An experimental and computational study of velar palatalization. *Cognitive Science, 30,* 945–982. https://doi.org/10.1207/s15516709cog0000_89

Wolansky, M. J., Cabrera, R. J., Ibarra, G. R., Mongiat, L., & Azcurra, J. M. (1999). Exogenous-NGF alters a critical motor period in rat striatum. *NeuroReport*, *10*(13), 2705–2709.

Woods, H. B. (1979). *A socio-dialectology survey of the English spoken in Ottawa: A study of sociological and stylistic variation in Canadian English* [Doctoral dissertation, University of British Columbia].

Yin, S. H., & White, J. (2018). Neutralization and homophony avoidance in phonological learning. *Cognition*, *179*, 89–101. https://doi.org/10.1016/j.cognition.2018.05.023

Appendix

Appendix A: Stimuli list presented during experiment

stem	*ba*-form	*ni*-form		
		No Alternation condition	Non-Neutralizing condition	Neutralizing condition
puti	ba-puti	ni-puti	ni-puti	ni-puti
bagu	ba-bagu	ni-bagu	ni-bagu	ni-bagu
timu/ **ʧ**imu	ba-timu/ ba-**ʧ**imu	ni-timu	ni-timu	ni-**ʧimu**
dapi/ **ʤ**api	ba-dapi ba-**ʤ**api	ni-dapi	ni-dapi	ni-**ʤapi**
mipa	ba-mipa	ni-mipa	ni-mipa	ni-mipa
niba	ba-niba	ni-niba	ni-niba	ni-niba
kimu	ba-**k**imu	ni-**k**imu	ni-**ʧ**imu	ni-**ʧimu**
gapi	ba-**g**api	ni-**g**api	ni-**ʤ**api	ni-**ʤapi**
kuta	ba-**k**uta	ni-**k**uta	ni-**ʧ**uta	ni-**ʧ**uta
gaku	ba-**g**aku	ni-**g**aku	ni-**ʤ**aku	ni-**ʤ**aku

Note: In the Neutralizing condition, singular stems, /ʧimu/ and /ʤapi/, were correspondingly presented instead of /timu/ and /dapi/ (shaded cells) during the experiment. In the No Alternation and Non-Neutralizing conditions, singular stems /ʧimu/ and /ʤapi/ were not presented during the experiment.

Appendix B: Grid search for parameters

Figure 10. Grid search showing the log likelihood of the Discount model according to different combinations of parameter values (fit to the English dataset). Note: The level of discounting; 100% (top), 75% (middle), and 50% (bottom).

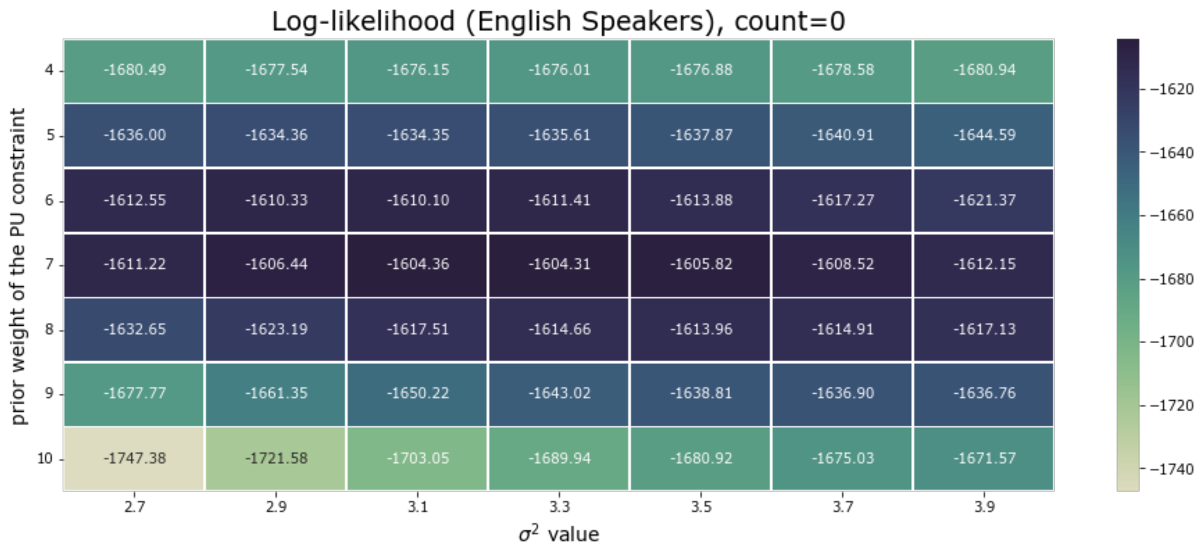

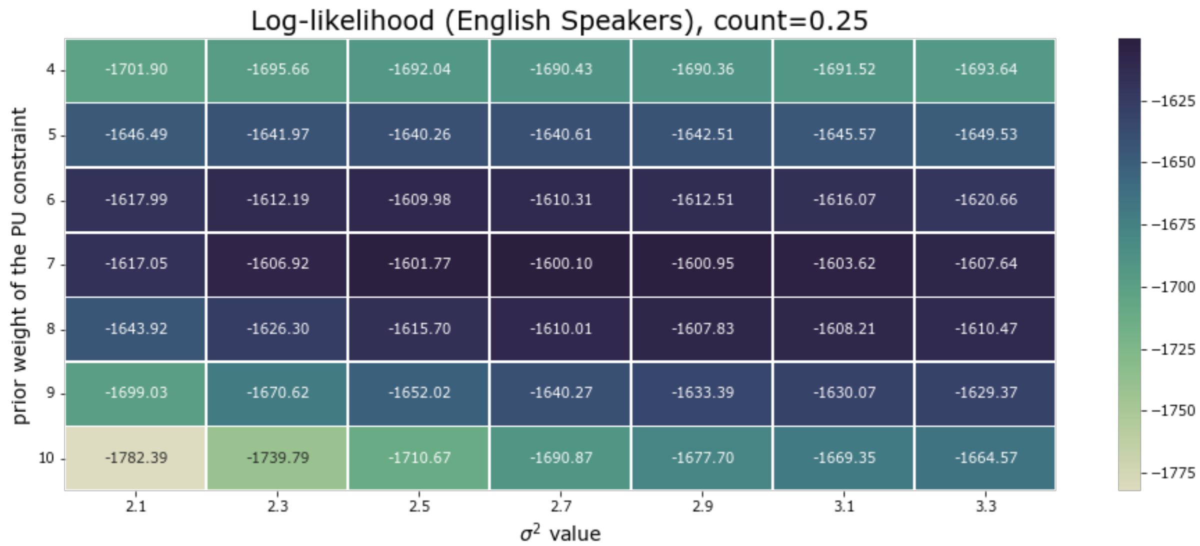

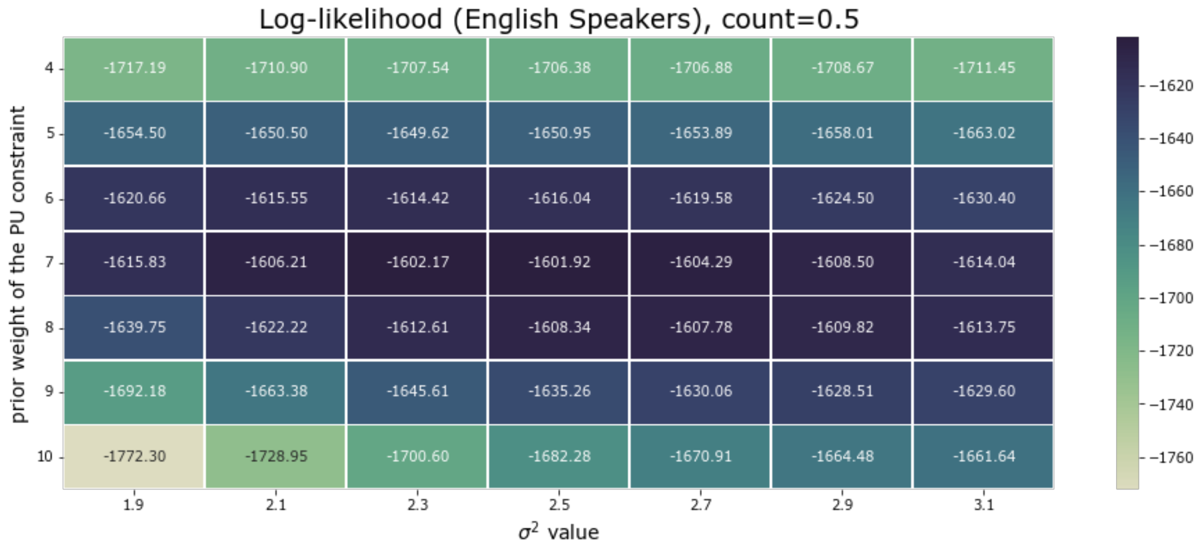

Figure 11. Grid search showing the log likelihood of the Discount model according to different combinations of parameter values (fit to the Korean dataset). Note: The level of discounting; 25% (top), and 0% (bottom).

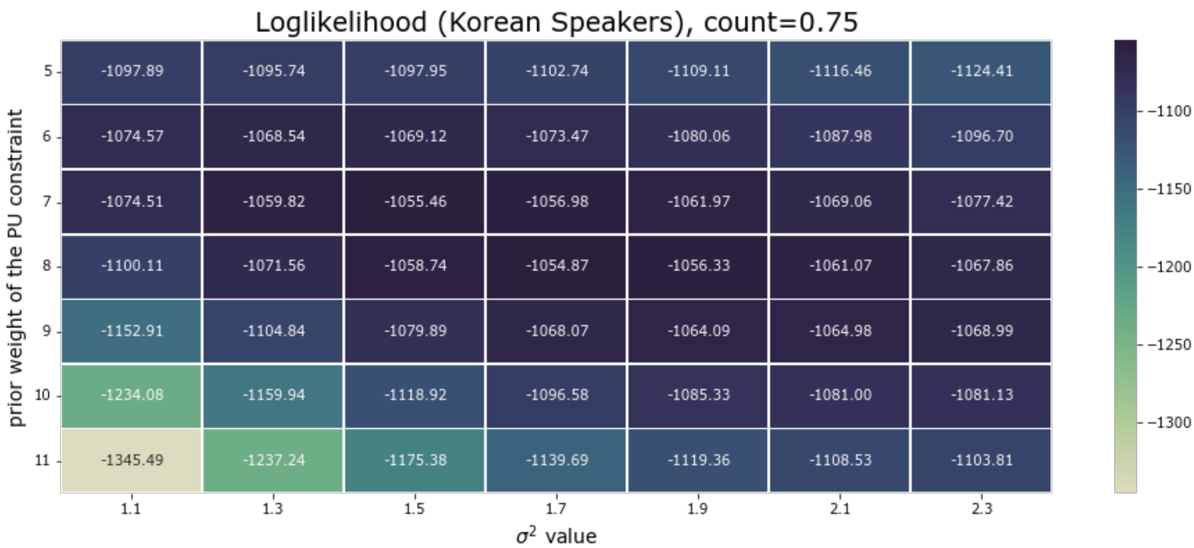

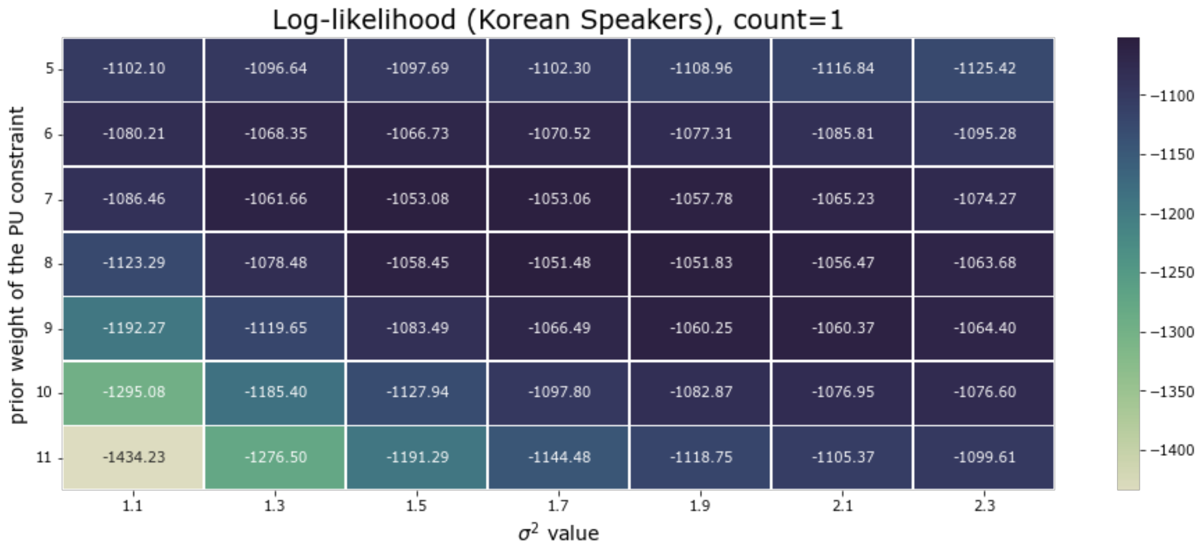

Appendix C: The complete comparison between models' predictions (log likelihood) to the experimental results

Table 56. The complete predicted probabilities of the Discount model and the aggregate experimental results for English speakers. Note: The numbers indicate the percentage of the time (%) that the output was chosen in the model and in the experimental results.

No Alternation condition							
		Frequent *ba-*		50-50 condition		Frequent *ni-*	
Input	Output	Model prediction	Experimental result	Model prediction	Experimental result	Model prediction	Experimental result
ba+non-velar-initial stem	ba+stem	100	100	100	100	100	100
	ba+ palatalized stem	0	0	0	0	0	0

ni+non-velar-initial stem	ni+stem	1000	100	100	100	100	100
	ni+ palatalized stem	0	0	0	0	0	0
ba+velar -initial stem	ba+stem	100	100	100	100	100	100
	ba+ palatalized stem	0	0	0	0	0	0
ni+velar-initial stem	ni+stem	100	100	100	100	100	100
	ni+ palatalized stem	0	0	0	0	0	0
Non-Neutralizing condition							
		Frequent *ba-*		50-50 condition		Frequent *ni-*	
Input	Output	Model predic-tion	Experi-mental result	Model predic-tion	Experi-mental result	Model predic-tion	Experi-mental result
ba+non-velar-initial stem	ba+stem	98	100	98	100	97	100
	ba+ palatalized stem	2	0	2	0	3	0
ni+non-velar-initial stem	ni+stem	98	100	98	100	97	100
	ni+ palatalized stem	2	0	2	0	3	0
ba+velar -initial stem	ba+stem	98	98	98	96	97	89
	ba+ palatalized stem	2	2	2	4	3	11
ni+velar-initial stem	ni+stem	40	26	29	26	23	29

	ni+ palatalized stem	60	74	71	74	77	71
Neutralizing condition							
		Frequent *ba-*		50-50 condition		Frequent *ni-*	
Input	Output	Model predic-tion	Experi-mental result	Model predic-tion	Experi-mental result	Model predic-tion	Experi-mental result
ba+non-velar-initial stem	ba+stem	98	100	98	100	98	99
	ba+ palatalized stem	2	0	2	0	2	1
ni+non-velar-initial stem	ni+stem	98	100	98	99	98	100
	ni+ palatalized stem	2	0	2	1	2	0
ba+velar-initial stem	ba+stem	98	96	98	92	98	76
	ba+ palatalized stem	2	4	2	8	2	25
ni+velar-initial stem	ni+stem	57	55	42	41	33	38
	ni+ palatalized stem	43	45	58	59	67	62

Table 57. The complete predicted probabilities of the Discount model and the aggregate experimental results for Korean speakers.

Non-Neutralizing condition					
		Frequent *ba-*		Frequent *ni-*	
Input	Output	Model prediction	Experimental result	Model prediction	Experimental result
ba+non-velar-initial stem	ba+stem	99	100	98	100
	ba+palatalized stem	1	0	2	0
ni+non-velar-initial stem	ni+stem	99	100	98	100
	ni+palatalized stem	1	0	2	0
ba+velar-initial stem	ba+stem	99	98	98	83
	ba+palatalized stem	1	2	2	17
ni+velar-initial stem	ni+stem	60	61	34	35
	ni+palatalized stem	40	39	66	65
Neutralizing condition					
		Frequent *ba-*		Frequent *ni*	
Input	Output	Model prediction	Experimental result	Model prediction	Experimental result
ba+non-velar-initial stem	ba+stem	99	100	98	100
	ba+palatalized stem	1	0	2	0
ni+non-velar-initial stem	ni+stem	99	100	98	100
	ni+palatalized stem	1	0	2	0
ba+velar-initial stem	ba+stem	99	95	98	89
	ba+palatalized stem	1	5	2	11
ni+velar-initial stem	ni+stem	60	47	34	36
	ni+palatalized stem	40	53	66	64

Table 58. The complete predicted probabilities of the *Neutralization model and the aggregate experimental results for English speakers.

No Alternation condition							
		Frequent *ba-*		50-50 condition		Frequent *ni-*	
Input	Output	Model prediction	Experimental result	Model prediction	Experimental result	Model prediction	Experimental result
ba+non-velar-initial stem	ba+stem	100	100	100	100	100	100
	ba+ palatalized stem	0	0	0	0	0	0
ni+non-velar-initial stem	ni+stem	100	100	100	100	100	100
	ni+ palatalized stem	0	0	0	0	0	0
ba+velar-initial stem	ba+stem	100	100	100	100	100	100
	ba+ palatalized stem	0	0	0	0	0	0
ni+velar-initial stem	ni+stem	100	100	100	100	100	100
	ni+ palatalized stem	0	0	0	0	0	0
Non-Neutralizing condition							
		Frequent *ba-*		50-50 condition		Frequent *ni-*	
Input	Output	Model prediction	Experimental result	Model prediction	Experimental result	Model prediction	Experimental result
ba+non-velar-initial stem	ba+stem	98	100	98	100	97	100
	ba+ palatalized stem	2	0	2	0	3	0
ni+non-velar-	ni+stem	98	100	98	100	97	100

initial stem	ni+ palatalized stem	2	0	2	0	3	0
ba+velar -initial stem	ba+stem	98	98	98	96	97	89
	ba+ palatalized stem	2	2	2	4	3	11
ni+velar-initial stem	ni+stem	48	26	35	26	27	29
	ni+ palatalized stem	52	74	65	74	73	71
Neutralizing condition							
		Frequent *ba-*		50-50 condition		Frequent *ni-*	
Input	Output	Model prediction	Experimental result	Model prediction	Experimental result	Model prediction	Experimental result
ba+non-velar-initial stem	ba+stem	98	100	98	100	97	99
	ba+ palatalized stem	2	0	2	0	3	1
ni+non-velar-initial stem	ni+stem	98	100	98	99	97	100
	ni+ palatalized stem	2	0	2	1	3	0
ba+velar -initial stem	ba+stem	98	96	98	92	97	76
	ba+ palatalized stem	2	4	2	8	3	25
ni+velar-initial stem	ni+stem	48	55	35	41	27	38
	ni+ palatalized stem	52	45	65	59	73	62

Table 59. The complete predicted probabilities of the *Neutralization model and the aggregate experimental results for Korean speakers.

Non-Neutralizing condition					
		Frequent *ba-*		Frequent *ni-*	
Input	Output	Model prediction	Experimental result	Model prediction	Experimental result
ba+non-velar-initial stem	ba+stem	99	100	98	100
	ba+palatalized stem	1	0	2	0
ni+non-velar-initial stem	ni+stem	99	100	98	100
	ni+palatalized stem	1	0	2	0
ba+velar-initial stem	ba+stem	99	98	98	83
	ba+palatalized stem	1	2	2	17
ni+velar-initial stem	ni+stem	60	61	34	35
	ni+palatalized stem	40	39	66	65
Neutralizing condition					
		Frequent *ba-*		Frequent *ni*	
Input	Output	Model prediction	Experimental result	Model prediction	Experimental result
ba+non-velar-initial stem	ba+stem	99	100	98	100
	ba+palatalized stem	1	0	2	0
ni+non-velar-initial stem	ni+stem	99	100	98	100
	ni+palatalized stem	1	0	2	0
ba+velar-initial stem	ba+stem	99	95	98	89
	ba+palatalized stem	1	5	2	11
ni+velar-initial stem	ni+stem	60	47	34	36
	ni+palatalized stem	40	53	66	64

Appendix D: Final weight of *Neutralization

The *Neutralization model performed the best when the initial preferred weight (μ) for the *Neutralization constraint was between 0 and 3. This range of μ for the *Neutralization constraint resulted in the final weight of 0 for the constraint. Once the final weight of the *Neutralization constraint started to gain any weight, the MaxEnt model compensated for the weight of *Neutralization by decreasing the weight of PU. This is because *Neutralization and PU partially overlapped and penalized the same outputs. In Tableau 15 (repeated from Tableau 13 in the main text), the candidate [nitʃimu] for the given input /nikimu/ was penalized by both PU and *Neutralization constraints in the Neutralizing condition. Hence, the weights are spread between the two constraints. Tableau 16 (repeated from Tableau 14 in the main text) shows that [nitʃimu] for the given input /nikimu/ was only penalized by PU in the Non-Neutralizing condition.

(15)

Also exists: [tʃimu] – [ni+tʃimu] /kimu/ - /ni+kimu/	*Ni- [+DORSAL]	PU	*Neutralization
(a) [kimu] - [nikimu]	*		
☞ (b) [kimu] - [nitʃimu]		*	*

(16)

/kimu/ – /ni+kimu/	*Ni- [+DORSAL]	PU
(a) [kimu] – [nikimu]	*	
☞ (b) [kimu] – [nitʃimu]		*

However, even a slight reduction of the weight for PU greatly affected the model's performance because PU penalized a large number of outputs. Specifically, there were more inputs of the prefix+non-velar-initial stem (filler casese:36 items) than those of the prefix+velar-initial stem (critical cases: 24 items) in the training data. Similarly, in the experimental results, there were more observations of the filler cases than of the prefix+velar-initial stem. While *Neutralization only penalized the outputs of the prefix+palatalized stems for given inputs of the prefix+velar-initial stems, PU penalized the outputs of the prefix+palatalized stem for given inputs of the prefix+velar-initial stem and the prefix+non-velar-initial stems (filler cases), suggesting that a change in weight of PU had a great effect on the overall fit of the model.

Furthermore, the calculation of log likelihood is sensitive to the number of observations in the experimental results, meaning that outputs found more often in the experimental results have a greater effect on log likelihood than the outputs found less often. Hence, the final weight of *Neutralization was likely to remain 0 to maintain the final weight of PU high in the best performing *Neutralization model. This shows the complexity driven by the interactions between constraints in the *Neutralization model, which makes the model less straightforward to account for the experimental results.

Appendix E: Find models' fit with r^2

I consider a different way to test how well parameters fit experimental results using the squared correlation coefficient (r^2), which represents how much variants the model accounts for. The fact that r^2 is not sensitive to the number of observations in the experimental results allows us to test the models' fit without being too strongly affected by the filler cases. A higher r^2 indicates a better fit of the model. Table 60 shows the r^2 of the Discount models according to different values of parameters. The model fit the English dataset the best when the observations that created homophony were discounted 100% of the time, which indicates that the observation was *never* counted in the training data. In contrast, the model best fit the Korean dataset when the observations were discounted 0% of the time in the training data. Although the level of discounting homophony for English dataset was slightly different between the models fit with r^2 (100% discount) and those fit with log likelihood (75% discount), both models show that English speakers heavily discounted the homophony-creating inputs whereas Korean speakers did not.

Table 61 shows the prediction of the Discount model for English speakers using its final weights. The model accurately predicted that English speakers were less successful at applying the neutralizing alternation compared to the non-neutralizing alternation. Table 62 illustrates the prediction of the Discount model for Korean speakers using its final weights. The model correctly predicted that Korean speakers applied both non-neutralizing and neutralizing alternations relatively often in their output.

Table 60. The r^2 according to different sets of parameter values for the English and Korean datasets.

Experimental results	r^2	discounting	μ of PU	σ^2
English speakers	0.98	100%	3	2.5
	0.979	75%	3	1.9
	0.977	50%	3	1.7
	0.974	25%	3	1.5
	0.971	0%	4	1.5
Korean speakers	0.952	100%	4	2.5
	0.964	75%	3	1.7
	0.973	50%	3	1.5
	0.978	25%	3	1.3
	0.981	0%	3	1.1

Table 61. The complete predicted probabilities of the Discount model and the aggregate experimental results for English speakers (fit with r^2). Note: The numbers indicate the percentage of the time (%) that the output was chosen in the model and in the experimental results.

No Alternation condition							
		Frequent *ba-*		50-50 condition		Frequent *ni-*	
Input	Output	Model prediction	Experimental result	Model prediction	Experimental result	Model prediction	Experimental result
ba+non-velar-initial stem	ba+stem	98	100	98	100	98	100
	ba+ palatalized stem	2	0	2	0	2	0
ni+non-velar-initial stem	ni+stem	98	100	98	100	98	100
	ni+ palatalized stem	2	0	2	0	2	0
ba+velar-initial stem	ba+stem	98	100	98	100	98	100
	ba+ palatalized stem	2	0	2	0	2	0
ni+velar-initial stem	ni+stem	98	100	98	100	98	100
	ni+ palatalized stem	2	0	2	0	2	0
Non-Neutralizing condition							
		Frequent *ba-*		50-50 condition		Frequent *ni-*	
Input	Output	Model prediction	Experimental result	Model prediction	Experimental result	Model prediction	Experimental result
ba+non-velar-initial stem	ba+stem	95	100	94	100	93	100
	ba+ palatalized stem	5	0	6	0	7	0

ni+non-velar-initial stem	ni+stem	95	100	94	100	93	100
	ni+ palatalized stem	5	0	6	0	7	0
ba+velar -initial stem	ba+stem	95	98	94	96	93	89
	ba+ palatalized stem	5	2	6	4	7	11
ni+velar-initial stem	ni+stem	35	26	26	26	20	29
	ni+ palatalized stem	65	74	74	74	80	71
Neutralizing condition							
		Frequent *ba-*		50-50 condition		Frequent *ni-*	
Input	Output	Model predic-tion	Experi-mental result	Model predic-tion	Experi-mental result	Model predic-tion	Experi-mental result
ba+non-velar-initial stem	ba+stem	96	100	95	100	94	99
	ba+ palatalized stem	4	0	5	0	6	1
ni+non-velar-initial stem	ni+stem	96	100	95	99	94	100
	ni+ palatalized stem	4	0	5	1	6	0
ba+velar -initial stem	ba+stem	96	96	95	92	94	76
	ba+ palatalized stem	4	4	5	8	6	25
ni+velar-initial stem	ni+stem	56	55	56	43	34	38
	ni+ palatalized stem	44	45	44	57	66	62

Table 62. The complete predicted probabilities of the Discount model and the aggregate experimental results for Korean speakers (fit with r2).

Non-Neutralizing condition					
		Frequent *ba-*		Frequent *ni-*	
Input	Output	Model prediction	Experimental result	Model prediction	Experimental result
ba+non-velar-initial stem	ba+stem	93	100	91	100
	ba+palatalized stem	7	0	9	0
ni+non-velar-initial stem	ni+stem	93	100	91	100
	ni+palatalized stem	7	0	9	0
ba+velar-initial stem	ba+stem	93	98	91	83
	ba+palatalized stem	7	2	9	17
ni+velar-initial stem	ni+stem	55	61	34	35
	ni+palatalized stem	45	39	66	65
Neutralizing condition					
		Frequent *ba-*		Frequent *ni*	
Input	Output	Model prediction	Experimental result	Model prediction	Experimental result
ba+non-velar-initial stem	ba+stem	93	100	91	100
	ba+palatalized stem	7	0	9	0
ni+non-velar-initial stem	ni+stem	93	100	91	100
	ni+palatalized stem	7	0	9	0
ba+velar-initial stem	ba+stem	93	95	91	89
	ba+palatalized stem	7	5	9	11
ni+velar-initial stem	ni+stem	55	47	34	36
	ni+palatalized stem	45	53	66	64

Table 63 shows the r^2 of the *Neutralization models that best fit the English and Korean dataset. The best fit of the *Neutralization model for the English dataset had extremely high μ for PU (14) and *Neutralization (19) as well as the value of σ^2 (5). The best-performing model for the Korean speakers had a relatively small μ of 2 for PU and 0 for *Neutralization. Furthermore, an extremely large range of μ of constraints resulted in the best-performing models. The fit of the *Neutralization model for the English dataset improved as the values of μ and the standard deviation (σ^2) increased. I tested the model's fit for the experimental results with moderately high and extremely high values of parameters and found that the r^2 slightly increased with the extremely high values of parameters (see Table 64), suggesting that, in theory, the performance of the model can continuously improve if the values of parameters increase. It is possible that the ratio among values of parameters has a greater effect on the model's fit rather than the definite value of each parameter due to the way MaxEnt learns weights. This emphasizes the complexity of selecting the *Neutralization model that best fits experimental results.

Table 63. The r^2 of the best *Neutralization model for the English and Korean datasets.

Experimental results	r^2	μ of PU	μ of *Neutralization	σ^2
English speakers	0.975	14	19	5
Korean speakers	0.981	2	0	1

Table 64. The r^2 of the *Neutralization model according to different parameter values (fit to the English datasets).

(μ of*ni-[+DORSAL], PU, *Neutralization), σ^2 value	r^2
(0, 14, 19), 5	0.9751
(0, 271, 500), 105	0.9761

The prediction of the *Neutralization model accurately captured that English speakers were more successful at learning the non-neutralizing alternation compared to the neutralizing alternation (Table 65). The model also correctly predicted that Korean speakers successfully learned both neutralizing and non-neutralizing alternations (Table 66).

Table 65. The complete predicted probabilities of the *Neutralization model and the aggregate experimental results for English speakers (fit with r^2).

No Alternation condition							
		Frequent *ba-*		50-50 condition		Frequent *ni-*	
Input	Output	Model prediction	Experimental result	Model prediction	Experimental result	Model prediction	Experimental result
ba+non-velar-initial stem	ba+stem	100	100	100	100	100	100
	ba+ palatalized stem	0	0	0	0	0	0
ni+non-velar-initial stem	ni+stem	100	100	100	100	100	100
	ni+ palatalized stem	0	0	0	0	0	0
ba+velar-initial stem	ba+stem	100	100	100	100	100	100
	ba+ palatalized stem	0	0	0	0	0	0
ni+velar-initial stem	ni+stem	100	100	100	100	100	100
	ni+ palatalized stem	0	0	0	0	0	0

Non-Neutralizing condition							
		Frequent *ba-*		50-50 condition		Frequent *ni-*	
Input	Output	Model prediction	Experimental result	Model prediction	Experimental result	Model prediction	Experimental result
ba+non-velar-initial stem	ba+stem	100	100	100	100	100	100
	ba+ palatalized stem	0	0	0	0	0	0
ni+non-velar-initial stem	ni+stem	100	100	100	100	100	100
	ni+ palatalized stem	0	0	0	0	0	0
ba+velar-initial stem	ba+stem	100	98	100	96	100	89
	ba+ palatalized stem	0	2	0	4	0	11
ni+velar-initial stem	ni+stem	37	26	25	26	20	29
	ni+ palatalized stem	63	74	75	74	80	71

Neutralizing condition							
		Frequent *ba-*		50-50 condition		Frequent *ni-*	
Input	Output	Model prediction	Experimental result	Model prediction	Experimental result	Model prediction	Experimental result
ba+non-velar-initial stem	ba+stem	98	100	98	100	98	99
	ba+ palatalized stem	2	0	2	0	2	1
ni+non-velar-	ni+stem	98	100	98	99	98	100

initial stem	ni+ palatalized stem	2	0	2	1	2	0
ba+velar -initial stem	ba+stem	100	96	100	92	100	76
	ba+ palatalized stem	0	4	0	8	0	25
ni+velar-initial stem	ni+stem	57	55	39	41	30	38
	ni+ palatalized stem	43	45	61	59	70	62

Table 66. The complete predicted probabilities of the *Neutralization model and the aggregate experimental results for Korean speakers (fit with r^2).

Non-Neutralizing condition					
		Frequent *ba-*		Frequent *ni-*	
Input	Output	Model prediction	Experimental result	Model prediction	Experimental result
ba+non-velar-initial stem	ba+stem	91	100	88	100
	ba+palatalized stem	9	0	12	0
ni+non-velar-initial stem	ni+stem	91	100	88	100
	ni+palatalized stem	9	0	12	0
ba+velar-initial stem	ba+stem	91	98	88	83
	ba+palatalized stem	9	2	12	17
ni+velar-initial stem	ni+stem	53	61	33	35
	ni+palatalized stem	47	39	67	65
Neutralizing condition					
		Frequent *ba-*		Frequent *ni*	
Input	Output	Model prediction	Experimental result	Model prediction	Experimental result
ba+non-velar-initial stem	ba+stem	91	100	88	100

	ba+palatalized stem	1	0	12	0
ni+non-velar-initial stem	ni+stem	99	100	88	100
	ni+palatalized stem	1	0	12	0
ba+velar-initial stem	ba+stem	99	95	88	89
	ba+palatalized stem	1	5	12	11
ni+velar-initial stem	ni+stem	53	61	33	35
	ni+palatalized stem	47	39	67	65