Morphological Dependencies

by

Guy Tabachnick

A dissertation submitted in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy

Department of Linguistics

New York University

September, 2023

Maria Gouskova

Acknowledgements

Other than the unfortunate circumstance of attending grad school in the middle of a pandemic, I have had great luck in life, in large part thanks to people who have supported me both within and outside of my intellectual life. Here are just a few of them.

Thanks first and most immediately to my advisor, Maria Gouskova. Her incredibly detailed comments addressing points across the spectrum from typos to high-level structural concerns of any piece of writing at any stage of completion are, if overwhelming in the moment, always extremely helpful in sharpening my writing and argumentation. Most of my work that sees the light in the next few years, including, of course, this dissertation, will bear the mark of multiple rounds of Maria's thorough, incisive reading. Her innumerable handouts on all aspects of academic life and many beyond (teaching, writing, knitting, etc.) are a treasure trove to which I will likely return in future years, slowly assimilating the advice within. In particular, I thank Maria for shepherding me through the crucible of my NSF grant application; the funds provided to me by NSF DDRIG BCS-2214315 enabled two of the three studies in this dissertation.

Thanks to the other members of my committee from NYU, Alec Marantz and Gillian Gallagher, and from other departments, Gaja Jarosz and Volya Kapatsinski. Their thoughts in discussions and comments have improved this work greatly and have given me challenges to wrestle with well into the future.

Thanks also to those who have helped me with translating and preparing the Hungarian and Czech studies: Agnes Kolben, Ildikó Szabó, Petr Adámek, and Pavla Šturmová.

Many other members of the NYU linguistics community have had a great impact on my work and life in the past five years as well. My cohort, Alicia Chatten, Kate Mooney, Kimberley Baxter, Anna Alsop, and Jai Peña, who have been wonderful officemates, classmates, and colleagues. Yining Nie and Madeline Gilbert made me feel welcome to the department in my first year and well beyond. Lunchtime and afternoon discussions on morphosyntax with Naomi Lee have always been fascinating, and been accompanied by the pleasant realization that we were working on very similar problems. Hagen Blix and Sandy Abu El Adas were sources of solidarity, most acutely on the picket line, but in many other places as well. Thanks to Selikem Gotah for co-leading the MorphBeer discussion group for the past three years, and to fellow morphology enthusiasts Christine Gu and Stefan Pophristic. I am grateful to the many other friendly faces in the department, including but not limited to José Álvarez Retamales, Nigel Flower, Alicia Parrish, Sarah Phillips, Chiara Repetti-Ludlow, and Ildikó Szabó. Chris Collins and the other members of my second QP committee, Gary Thoms, Stephanie Harves, and Richard Kayne, guided me through the process of learning how to do syntax. Thanks to Juliet Stanton and Lisa Davidson for their support on a number of projects and contribution to an engaging and entertaining evironment on the fifth floor, and to Laurel MacKenzie, who has been a pleasure to work for and with. Finally, thanks to department administrators Teresa Leung and Teresa Colaizzo, who have been nothing but helpful in handling logistical issues, sometimes in a very short timeframe.

Linguists outside of NYU have been gracious in sharing their thoughts on ideas that have ended up in one form or another in this work. Thanks especially to Péter Rebrus, Anne-Michelle Tessier, Laura Kalin, Ruth Kramer, Heidi Harley, Andrea Sims, and Solveiga Armoskaite. Thanks also to my serendipitous colleagues, Thomas Kettig and Marcin Wągiel, for discussions about languages, linguistics, and many other things.

My interest in linguistics began long before coming to NYU. Sheila Krilov, Eliza Kuberska, Claire

Mazzola, and Chris Unruh supported my love of problem-solving and languages. Thanks to Uriel Cohen Priva for advising my undergraduate thesis and giving me useful advice when I went to him with the intention of applying to graduate school. I am especially grateful to Masako Fidler and Carol Rounds for their enthusiasm in teaching me Czech and Hungarian, giving me the necessary background to (eventually) conduct this research.

Thanks to my mostly non-academic friends, especially Ian Eppler and Emma Claire Foley, and other members of my foundational group chats. And thanks to the Crown Heights trad egal community for making a welcoming space where I was, perforce, not working.

I am incredibly lucky that my parents, Yael Mandelstam and Ken Tabachnick, have graced me with boundless love and support, a strong intellectual foundation, and many other qualities necessary on the path that has culminated in this dissertation.

And thanks, finally, to Hannah Feldman, whose impact on this dissertation is indirect but whose impact on my life has been immeasurable.

Abstract

This dissertation investigates morphological dependencies: correlations between two lexically specific patterns, such as selection of inflectional affixes. Previous work has established that such correlations exist in the lexicon of morphologically rich languages (Ackerman et al., 2009; Wurzel, 1989), but has not systematically tested whether speakers productively extend these patterns to novel words. I present a series of corpus and nonce word studies—in Hungarian, Czech, and Russian—testing whether speakers vary their selection of suffixed forms of novel words based on the forms of that word that are presented to them. In all three cases, speakers vary their responses in accordance with the provided stimuli, demonstrating that they have learned and productively apply morphological dependencies from the lexicon.

I present a theoretical account of morphological dependencies that can account for my experimental results, based on the sublexicon model of phonological learning (Allen & Becker, 2015; Becker & Gouskova, 2016; Gouskova et al., 2015). In this model, speakers index lexically specific behavior with diacritic features attached to underlying forms in lexical entries, and learn generalizations over sublexicons defined as words that share a feature. These generalizations are stored as constraints in phonotactic grammars for each sublexicon, enabling speakers to learn phonological and morphological dependencies predicting words that pattern together. This model provides a unified treatment of morphological dependencies and generalizations that are phonological in nature. My studies show a wide range of learned effects, not limited to those that follow an organizational principle like paradigm uniformity. The sublexicon model assumes that speakers can learn arbi-

trary generalizations without restrictions, giving it needed flexibility over more restrictive models which rely on notions of morphophonological naturalness.

Contents

List of Figures

List of Tables

1 Introduction

One striking feature of human capacity for language is *productivity*: speakers are not limited to words and utterances that they have heard before, but are able to generalize learned patterns to create new structures. This dissertation focuses on productivity at the morphological level: how speakers create new forms of words that they have not previously seen. In languages with rich morphology, a word can appear in many different inflected and derived forms, depending on its context, but speakers only see a fraction of possible forms built from any given root (Chan, 2008). Despite this, speakers usually have little trouble producing forms that they have not previously seen. They must have some way of inferring these forms. I investigate how they do this: what patterns are speakers generalizing? The particular cases in this work sit at the intersection of *patterns* that words follow and *exceptions* to those patterns: speakers often cannot say with certainty what a word's form will be unless they have previously seen it, but they can make a good guess.

This dissertation focuses on one specific source of information used by speakers to infer unknown complex forms: knowledge of a word's *known* forms. I refer to these correlations between forms of a word as *morphological dependencies*. Ackerman et al. (2009) call the process of inferring unknown forms from known ones the Paradigm Cell Filling Problem. The goal of this dissertation is to demonstrate and model the psychological reality of these dependencies: speakers learn correlations between morphological patterns from their languages' lexicon and productively apply them to new words.

A typical morphological dependency, as defined in this dissertation, holds between the Czech genitive and locative, which are shown in Table 1.1. This example is studied at length in Chapter 5. Czech nouns inflect for seven cases, three of which are shown below. The suffix *-u* can be used to mark both the genitive and the locative, and both cases have a second allomorph as well: some nouns have *-a* in the genitive, while others have *-ɛ* in the locative.

noun	‘lip’	‘ball’	‘January’	‘forest’
nominative	rɛt-Ø	plɛs-Ø	lɛdɛn-Ø	lɛs-Ø
genitive	rt-u	plɛs-u	lɛdn-a	lɛs-a
locative	rt-u	plɛs-ɛ	lɛdn-u	lɛs-ɛ

Table 1.1: Some Czech genitive and locative case forms

The majority of nouns have *-u* in both cases, like [rɛt] ‘lip’. Nouns that have *-a* in the genitive are also more likely than expected to have *-ɛ* in the locative, like *lɛs* ‘forest’. This is the morphological dependency: a word’s genitive suffix is (somewhat) predictive of its locative suffix, and when speakers are asked to select a locative form for previously unseen words, they are more likely to select locative *-ɛ* for words whose genitive they know to be *-a*.

What counts as a morphological dependency is somewhat dependent on analytical and theoretical considerations. For example, all of the words in Table 1.1 have [ɛ] in the last syllable in the nominative; two ([plɛs] and [lɛs]) retain this vowel in the genitive and locative, while in the other two ([rɛt] and [lɛdɛn]), this vowel is absent in the genitive and locative. In all of these words, the genitive and locative behave the same way: if [ɛ] is present in one suffixed form, it is present in the other. Does this constitute a morphological dependency? Whether a noun undergoes a vowel–zero alternation is certainly lexically specific, in the sense that it is not entirely predictable from a word’s phonology. However, morphophonological analyses of Czech and other Slavic languages (e.g. Gouskova, 2012; Kenstowicz & Rubach, 1987; Lightner, 1965; Rysling, 2016; Yearley, 1995) treat this alternation as a unified property of a noun: if a vowel alternates in one form with an inflectional suffix, it alternates in all such forms. Since this property can be described as a single piece of information that encompasses the genitive and locative, it does not comprise a morphological dependency. The morphological dependencies studied in this dissertation will mostly involve selection of two suffixes, like Czech genitive -*a* and locative -*ɛ*.

1.1 Summary of results

This dissertation includes three case studies—in Hungarian, Czech, and Russian—of morphological dependencies. In each study, I establish by corpus analysis that a morphological dependency exists in the lexicon: two morphological patterns are statistically correlated. I then show that this dependency is active and productive in speakers' grammars with a nonce word study, in which I present suffixed forms of made-up words and ask speakers to select additional suffixed forms. In all three cases, speakers' selection varies with the presented form to match the morphological dependency in the lexicon.

The empirical contribution of this dissertation is thus to establish that speakers learn morphological dependencies from their input, a fact assumed but not directly tested in work on the morphological organization of paradigms like Ackerman et al. (2009), Ackerman and Malouf (2013), and Bonami and Beniamine (2016). The case studies show that speakers have broad abilities to learn these patterns: for example, between inflectional and derivational affixes. These learned patterns can be arbitrary, in the sense that they are not limited to paradigmatic uniformity effects, where forms are linked in the grammar because they share some structure (see Section 2.3.2.1 for discussion). The studies in this dissertation thus show that morphological dependencies are no exception to the extensive body of literature, discussed in Section 2.1.2 and Section 2.1.3, showing that speakers learn and apply gradient patterns from the lexicon (e.g. Albright & Hayes, 2003; Bybee, 1995; Ernestus & Baayen, 2003; Hayes et al., 2009). This fills a gap in the literature: previous work on morphological dependencies (discussed in Section 2.1.5) has focused on the patterns in the lexicon, but as Bonami and Beniamine (2016) note, this is insufficient: to obtain a complete picture of morphological dependencies, we need to explore *whether* and *how* speakers learn these patterns from their input. This work offers a new experimental paradigm for doing just that.

The theoretical contribution of this dissertation is an account of how speakers learn and store morphological dependencies. I extend the sublexicon model of phonological learning (Allen & Becker, 2015; Becker & Gouskova, 2016; Gouskova et al., 2015). In this model, described in Chapter 3, a word's morphological behavior is indexed by symbolic *diacritic features*; speakers learn phonological and morphological generalizations over groups of lexical items that share a feature, which are placed together into a sublexicon defined by each feature and described by a constraint-based sublexical grammar that encodes these generalizations. Speakers then invoke these sublexical grammars in the production of new forms: to determine a novel word's behavior, they must assign it a diacritic feature, and the better the word fits a given feature's sublexical grammar, the more likely it is to be assigned that feature. I discuss the architectural properties of this model at more length in Section 3.1, focusing in particular on one aspect of the model: each feature is defined by *its own* sublexical grammar, which stores patterns and is only invoked in determining the behavior of unfamiliar words. In Section 3.3, I argue that this is a more flexible approach than an alternative in which nonce words are evaluated on a single language-wide phonotactic grammar with constraints indexed to particular lexical items. This model provides a unified treatment of morphological dependencies and generalizations that are phonological in nature: both are encoded in the same module (as constraints in sublexical grammars), and speakers can compound their effects.

In my proposed account of morphological dependencies, I assume that the sublexical grammars sit alongside a generative grammar as described by Distributed Morphology (Halle & Marantz, 1993; Harley, 2014; Harley & Noyer, 1999). In the taxonomy of Hockett (1954), Distributed Morphology is an *item and arrangement* model of morphology: words are comprised of smaller pieces put together in a particular configuration. (In particular, in this theory, words are comprised of phonological *exponents* of syntactic units called *morphemes* which are arranged in the same syntactic structures used to build sentences.) Ackerman and Malouf (2013) argue that in theories like Distributed Morphology, which they call syntagmatic ("emphasiz[ing] the linear combination of constitutive elements"), "surface patterns of both words and networks of words (i.e. relations between surface alternants) are not regarded as proper objects of linguistics analysis, while the abstract elements and operations responsible for constructing these ephemera are" (Ackerman & Malouf, 2013, p. 430). One goal of this dissertation, then, is to demonstrate to skeptics that Distributed Morphology and other piece-based theories of morphology do not preclude taking paradigmatic relations seriously—and, in fact, provide a useful symbolic framework for encoding these relations (namely, lexical diacritic features).

At the same time, the studies in this dissertation are intended to demonstrate to theoretical morphologists the importance of variation and gradient patterns in grammar. As I discuss in Section 2.3.2.2, much of the work on inflectional patterning in Distributed Morphology assumes fixed patterns of lexical items gathered into umbrella inflection classes, ignoring or deemphasizing words that do not fall neatly into these categories. The works in this study show that the notion of an inflection class as a single category that encapsulates a word's entire inflectional behavior is too rough and misses many generalizations; instead, I argue that correlations between individual forms can be more insightful than rigid notions of inflection class. I argue that some of the relations typically hard-coded into grammatical structures are better captured as tight clusters of related inflectional patterns that emerge from learning gradient, variable morphological dependencies. Thus, this work is intended as a synthesis of theoretical and empirical concerns from approaches to morphology that are often pitted against one another.

1.2 Roadmap

Chapter 2 provides an overview of previous theoretical and empirical work on patterns of allomorphy that are conditioned on individual lexical items. Section 2.1 comprises a survey of allomorphy: the factors that can predict how a given word will inflect, and how speakers productively extend these patterns to previously unknown words in nonce word studies. These patterns are generalizations over exceptional and unpredictable lexical items. Accordingly, in Section 2.2, I look at grammatical theories of how such exceptional and unpredictable lexical items are encoded, including lexically conditioned allomorphy. I argue that exceptionality should be grammatically instantiated through diacritic features attached to lexical entries. I focus on work on inflection class diacritics in theories like Distributed Morphology which posit a split lexicon in which phonological, syntactic, and semantic properties are stored separately. I argue, against some previous work, that diacritic features properly belong in phonological underlying forms. Finally, in Section 2.3, I return to phonological and morphological generalizations over lexically conditioned allomorphy, evaluating how different theories learn and encode them (as the sublexicon theory does). I focus on two properties: first, whether the generalizations are *hard-coded* into a generative grammar and thus invoked every time a word is derived, or whether, instead, the generalizations are stored in a separate pattern-matching module that is only invoked for new words that do not have the relevant lexical information stored in the lexicon. I conclude that hard-coding is too rigid to capture many attested phonological and morphological dependencies, and argue for a greater use of pattern-matching modules like sublexical grammars in explaining morphological structure. The second property is *grounding*: can learners only learn patterns that are grounded in some notion of morphphonological naturalness or optimization, or can they also learn arbitrary patterns that are not necessarily grounded? I similarly conclude that requiring learned patterns to be phonologically or morphologically grounded is too limiting to account for the full range of attested generalizations, and a successful model of productivity should not be limited to learning grounded patterns.

In Chapter 3, I present a model for learning morphological dependencies that encodes them in a way that is neither hard-coded nor necessarily grounded. This model is an extension of the sublexicon model of learning *phonological* generalizations (Allen & Becker, 2015; Becker & Gouskova, 2016; Gouskova et al., 2015), which I present in Section 3.1. In the sublexicon model, speakers learn generalizations over sets of words defined by shared behavior, in my implementation a shared diacritic feature. These generalizations are stored in constraint-based grammars that probabilistically assign diacritic features to new words. In Section 3.2, I show that treating these diacritic features like phonological features enables the sublexicon model to also capture *morphological* dependenices without additional changes to the model. Section 3.3 compares the predictions of a sublexicon model with those of a simpler model, in which the same generalizations are stored in a single, language-wide phonotactic grammar. While their predictions are similar, I show that the sublexicon model with multiple grammars is more flexible than the single grammar alternative. Next, in Section 3.4, I present a model using variable, weighted diacritic features to capture lexically and syntactically conditioned variation. In this model, variable features trigger derivational splits and assign scores to candidate derivational paths; one path is probabilistically chosen among the viable candidates. Section 3.4.5 describes how the sublexicon model can assign weights for variable features to new words, extending the earlier sublexicon model which assumed categorical features.

The next three chapters present case studies of morphological dependencies, each of which is structured in a similar way. In Chapter 4, I discuss a dependency in Hungarian between the plural, which shows allomorphy for a small class of nouns called "lowering stems", and the possessive, which has two basic allomorphs. After providing some background information in Section 4.1 and a formal analysis of Hungarian lowering stems and possessive allomorphy in Section 4.2, I present my studies. Section 4.3 shows the results of a corpus study showing that the morphological dependency is active in the lexicon. Section 4.4 presents the results of a nonce word study, showing that Hungarian speakers have learned this dependency and productively apply it: their choice of possessive for nonce words is influenced by the plural form presented for that word. A brief discussion in Section 4.5 concludes the chapter.

Chapter 5 concerns a morphological dependency between minority allomorphs for two case forms in Czech, the genitive and the locative. This case study is marked by the substantial presence of variable lexical items: many nouns appear with both possible genitive/locative forms, in various proportions. Much of this case study is like the previous one. I provide background in Section 5.1 and a formal analysis in Section 5.2, then establish the morphological dependency with a corpus study in Section 5.3 and show that speakers productively apply it with the nonce word study reported in Section 5.4: speakers' choice of locative for nonce words is dependent on the genitive of those words. The widespread variability enables two additional studies further confirming the productivity of the morphological dependency. In Section 5.5, I repeat the behavioral experiment with *real* variable words. This study finds no significant correlation between genitive and locative, confirming my model's prediction that the dependency is a learned variable pattern from the lexicon productively extended to new words, rather than a sort of priming effect active for both nonce words and existing words with stable lexical entries. Similarly, Section 5.6 presents a corpus study looking at variable words in the works of individual authors. Here, too, I find a correlation between the two cases at the individual level. It is plausible, however, that this correlation represents a distribution of individual lexical items faithfully learned from speakers' input rather than one emerging from a grammatical bias between the two cases. If future work can show that these two cases are *not* correlated in the *input* of individual speakers, the results of this author corpus study require an explanation grounded in grammatical biases that crucially requires a multiple grammar model rather than the simpler single grammar alternative. Section 5.7 includes a final discussion and summary of the study.

The third case study, in Chapter 6, looks at dependencies involving a common Russian derivational suffix, the diminutive. One diminutive allomorph is preferred by nouns with particular stress patterns in their inflectional paradigms, as well as nouns that take an irregular plural suffix. I present background information in Section 6.1 and an analysis in Section 6.2, showing that while stress pattern can be encoded phonologically with different configurations of underlying stress, the relationship between the plural and the diminutive is a morphological dependency between two diacritic features. I then establish the morphological dependency in the corpus study in Section 6.3. Unlike previous cases, the evidence for the morphological dependency between plural and diminutive is ambiguous, and speakers are expected to differ in whether they learn and apply it. In the nonce word study, in Section 6.4, I find that speakers productively apply both morphological dependencies: both stress location and suffix for a presented plural form of a nonce word influence the diminutive chosen. I find no evidence that speakers behave differentially—all seem to have learned the plural effect equally. This is evidence for another morphological dependency; unlike the results in previous studies, this correlation cannot be encoded as a paradigm uniformity effect. Thus, the results of this study show that speakers must be able to learn morphological dependencies even when they are not grounded in paradigm uniformity. A discussion in Section 6.5 concludes the chapter.

A brief summary of results and conclusion can be found in Chapter 7.

2 Properties of allomorphy

One important task in learning a language is knowing *arbitrary associations* between words and the patterns of word formation that they follow. For example, English speakers know that the past tense of *blow* [bloʊ] is *blew* [blu], but the past tense of the similar verb *flow* [floʊ] is *flowed* [floʊd], not **flew* *[flu]. Speakers of a language are also able to extend these patterns *productively*: suppose there was a new verb, *prow* [proʊ]. Would people form the past tense as *prew*, like *blow* (and *know* and *grow*), or as *prowed*, like *flow* (and *stow* and *tow*)? As I discuss in this chapter, English speakers often have strong intuitions about what the past tense of such novel verbs should be, and one object of linguistic study is to figure out what is driving these intuitions. That is, what patterns have speakers learned about their language? How have they generalized the arbitrary patterns of known lexical items to unknown ones?

The studies in this dissertation focus on one type of generalization that speakers have learned and productively applied: correlations between arbitrary associations of lexical item and pattern. In English, we can see such a correlation between a verb's past tense and its participle (i.e. *I have Xed*). If speakers learn that the participle of *prow* is *prown* (like *blown*, *known*, and *grown*), they should be more likely to assign a past tense of *prew*; if the perfective is *prowed* (like *flowed*, *stowed*, and *towed*), the past is more likely to be *prowed* as well. This is because most verbs that have -*d* in the past or participle have it in both forms (though there are exceptions, like *show/showed/shown*). Thus, there is a correlation (which I call a *morphological dependency*) between the English past and perfective: both can follow arbitrary patterns and often require speakers to associate individual words with their patterns, but knowing one of the forms can allow a speaker to make a better guess about the other.

In this chapter, I provide the background for my case studies and my theoretical proposal accounting for their results. Each of the sections in this chapter addresses a different aspect of how speakers learn gradient generalizations over arbitrary morphological patterns and apply them productively to new forms. First, in Section 2.1, I provide an overview of the empirical ground: what sorts of information can predict the allomorphs that a word will select? I look at phonology, paradigmatic structure, and syntactic/semantic factors. In all cases, the main point is that selectional pairings of lexical items with their allomorphs, while often somewhat arbitrary, are not totally random. Even when a given word's behavior must be learned individually, it is often more likely to fall into some of the available patterns and less likely to fall into others.

If speakers are drawing generalizations from arbitrary morphological patterns, we must first understand what grammatical constructs, exactly, they are generalizing over. This grammatical encoding of morphological patterns is the subject of Section 2.2, which focuses on theoretical approaches to lexically specific behavior. That is, I discuss how different theories of morphology have implemented lexical marking of the arbitrary associations between lexical items and their patterns. I argue for a *diacritic* approach: lexical items that share a realizational pattern are marked with a symbolic diacritic feature that indexes to a rule of realization encoding the behavior (for example, the insertion of a particular allomorph of an affix). This means that the generalizations described in Section 2.1 can be encoded grammatically as generalizations over lexical items that share a particular diacritic feature.

This is not the only possible way of encoding gradient generalizations, as I discuss in Section 2.3. This section reviews a number of theoretical models that have been proposed to capture phonological and morphological generalizations over arbitrary allomorph selection. I classify these models according to two main criteria. The first is *hard-coding*: are the generalizations directly hard-coded into the generative grammar? In models with hard-coding, the generalization is grammatically active and applies (or, at least, may apply) every time a word is formed; words that do not follow the generalization, even variably, are marked as exceptions. In models with hard-coding, novel words should show the behavior of words without lexical marking to which the baseline grammar applies—if the generalization applies variably, it should do so at a predictable rate for new words. By contrast, in models without hard-coding, the generalizations are stored in a separate pattern-matching module and are not invoked in the course of a normal derivation. Instead, the pattern-matching module only becomes active in determining the behavior of new words when they are encountered. I conclude that models without hard-coding are able to more flexibly capture morphological dependencies. The second criterion is *grounding*. In some theories, generalizations present in the lexicon are only learnable if they are in some way "grounded" or "natural"—that is, if they are phonologically optimizing or expressible in terms of posited universal constraints (such as paradigm uniformity constraints, in the case of morphology). I argue that grounding requirements are too strict: speakers learn at least some generalizations from their input that are not expressible in phonologically or morphologically grounded ways.

The theoretical and empirical background presented in this chapter lay the groundwork for the theoretical model I propose in Chapter 3. The proposed model groups together words that share a diacritic feature and learns gradient and categorical generalizations over such words in a constraint-based grammar that lies in a separate pattern-matching module. These generalizations are not hard-coded and need not be grounded, yielding a more flexible model. When a speaker encounters a new word and wishes to determine its behavior, she evaluates it on the grammars in the pattern-matching module to assign a feature to it; this feature makes it a more completely fleshed out vocabulary item that can then enter derivational processes like other known words. That is, I treat the process of guessing a word's behavior as one of choosing and lexically marking a word's association with one of the possible patterns.

2.1 How to infer unknown forms

2.1.1 Allomorphic relations

Languages often exhibit inflectional and derivational *allomorphy*: the same set of morphosyntactic features can be expressed in multiple forms depending on the stem to which it attaches. The focus of this dissertation is *suppletive* allomorphy, in which the allomorphs cannot plausibly be considered allophonic variants of a single underlying form; instead, suppletive allomorphs have different underlying forms. In this section, I address the question of how a word's pattern of allomorphy can be predicted by other lexical properties like its phonology and its syntactic classification—and, most crucially for this dissertation, how multiple patterns of allomorphy are correlated with one another. That is, how one kind of arbitrary behavior predicts another.

Not all allomorphy is lexically arbitrary, in the sense that the patterning of (at least some) words must be learned individually. In some cases, there are multiple productive allomorphs whose distribution is fully defined by the phonology of the items to which they attach (Paster, 2006; Wolf, 2008)—this is discussed in more detail in Section 2.3.1.1. Wolf (2008) lists one such example from Moroccan Arabic (Harrell, 1962): the 3SG masculine possessive marker is *-h* after vowels and *-u* after consonants, as shown in (1):

(1) *Phonologically conditioned possessive allomorphy in Moroccan Arabic (Harrell, 1962; Wolf, 2008)*

a. *V-final stem* xtˤa-h 'his error'

b. *C-final stem* ktab-u 'his book'

These are separate allomorphs, with distinct underlying forms, because there is no plausible path to relate them as allophones of the same underlying form. In these cases, allomorphy is *systematic* and *predictable*: a word's phonology fully determines the selected allomorph.

The cases of allomorphy discussed in this dissertation are not fully systematic or predictable. That is, the allomorph selected by a given word cannot always be inferred from its phonology or other properties like its gender. In at least some cases, the realization of individual items must be memorized. In Hungarian, the possessive suffix has two allomorphs, *-ɒ* and *-jɒ* (abstracting away from the language's vowel harmony; see Chapter 4 for more details). For example, the possessive of [paːr] 'pair' is [paːr-jɒ], while the possessive of [kaːr] 'damage' is [kaːr-ɒ]. Both endings are quite common: according to a web corpus study by Rácz and Rebrus (2012), *-ɒ* is used in 63% of types and 81% of tokens. The choice of allomorph is partially predictable from phonology: nouns that end in vowels, like [kɒpu] 'gate' categorically take *-jɒ* ([kɒpu-jɒ]), while nouns ending in palatals and sibilants categorically take [-ɒ] (e.g. [aːɟ-ɒ] 'her bed' and [haːz-ɒ] 'her house'). However, the possessive form of other nouns must be learned individually, and even very similar nouns (like [paːr] and [kaːr]) can take different possessives. I call this pattern *lexical* or *lexically conditioned* variation.

Patterns like the Hungarian possessive require speakers to associate allomorphs with individual words. However, this does not mean that the forms are distributed randomly. There are often gradient tendencies in the lexicon—for example, I show in Chapter 4 that in Hungarian, nouns that end in non-sibilant coronals have possessive *-jɒ* more often than nouns ending in other consonants (see also Rácz & Rebrus, 2012). Speakers are clearly capable of learning the correct associations of individual words and their selectional properties, because lexical variation can be learned with relative stability. Do they also learn the gradient tendencies over these associations? Answering this question requires more targeted studies. I discuss these in the next section.

2.1.2 Allomorphy and productivity

In order to probe the patterns that speakers have learned, we must put them in a situation where they cannot rely on their existing associations with individual words. This is done most effectively using a *nonce word study* or *wug test*, innovated by Berko (1958). In the original wug test, Berko asked children (aged 4–5) and adults to form the plural of nonce nouns like *wug* [wʌg]; both children and adults successfully applied the dominant regular pattern: [wʌg-z].

When tested on lexically variable patterns (e.g. Albright & Hayes, 2003; Becker et al., 2011; Ernestus & Baayen, 2003; Gouskova et al., 2015), and artificial language studies (e.g. Hudson Kam & Newport, 2005), adults usually follow what Hayes et al. (2009) call the Law of Frequency Matching:

(2) *Law of Frequency Matching (Hayes et al., 2009, p. 826)*
Speakers of languages with variable lexical patterns respond stochastically when tested on such patterns. Their responses aggregately match the lexical frequencies.

In particular, speakers match the frequencies of *types* in the lexicon—that is, they arrive at their distributions by counting lexical items without regard to (or, at least, independently from) their frequency (Albright & Hayes, 2003; Bybee, 1995, 2001; Hayes & Wilson, 2008; Hayes et al., 2009; Pierrehumbert, 2001). In English, for example, the regular plurals -z, -s, and -əz are used on the vast majority of words, even if many of them are very infrequent (like *sackbut*). Thus new words tend to take one of these predominant plurals. On the other hand, the plural pattern shown in *child*/*children* is not extended to new words: even though *child* is a very frequent word, it is the only word that follows this plural formation pattern.

The "law" does not hold universally: Hayes et al. (2009) note that adult speakers do not always blindly apply all patterns from the lexicon, but seem to be better equipped to learn some patterns than others (see Becker et al., 2011; Pertsova, 2004). In addition, some studies suggest that children may not match frequencies the way adults do (see Yang (2016), Schuler et al. (2021), and Jarosz (2022) for discussion). If this discrepancy proves robust, it requires explanation; however, the present work is concerned with studying and explaining the behavior of adult speakers, and the behavior of children is beyond its scope. Thus, I adopt the Law of Frequency Matching as a starting hypothesis for how adult speakers extend patterns of lexical variation to novel words in wug tests: they should assign allomorphs to these new words stochastically, roughly matching the distribution of the lexicon. I assume that this applies both to the lexicon as a whole and to well-defined subsets of it: speakers should also match the frequencies for consonant-final words, monosyllabic words, etc.

2.1.3 Phonological predictors of allomorphy

In the cases of arbitrary lexical patterns discussed so far, like the Hungarian possessive, the active gradient generalizations predicting a word's morphological patterning have been *phonological* in nature. Most prior wug tests have addressed phonological dependencies: how is speakers' inflection of novel words influenced by the words' phonological form? In this section, I summarize research in the phonological domain. Next, I move onto other sources of information predicting inflected forms, which have been less thoroughly tested experimentally but are no less present in the lexicon, such as syntactic and semantic factors. Among these is the subject of this dissertation's study: morphological dependencies, or the correlations between two arbitrary morphological patterns.

One of the most frequently studied cases of lexically specific allomorphy is the English past tense: most English verbs form the past tense by suffixing *-d*, *-d*, or *-əd* (which are distributed according to the phonology of the verb to which they attach), but some classes of verbs instead mark the past tense with an alternation in the root vowel: thus, we have present–past pairs like *run* [ɾʌn] and *ran* [ræn], *hang* [hæŋ] and *hung* [hʌŋ], etc. These two patterns, in particular, are sometimes extended to new verbs in nonce word experiments alongside the regular pattern with the suffix *-d* (Albright & Hayes, 2003; Bybee & Moder, 1983; Prasada & Pinker, 1993). That is, given a nonce verb like *pling* [plɪŋ], English speakers may form its past tense as any of [plɪŋd], [plæŋ], or [plʌŋ] (*plinged*, *plang*, *plung*).

The English past tense, however, is quite complicated, showing a wide array of possible realizations that can apply in overlapping circumstances. While this example is by no means unique—one similar and well-studied example is the German plural (Marcus et al., 1995)—simpler cases can yield a clearer picture of cases of lexical variation. For example, Ernestus and Baayen (2003) and Becker et al. (2011) look at very similar voicing alternations in Dutch and Turkish, respectively. In both languages, obstruents alternated between voiced in word-medial position and voiceless in word-final position. Words that have unvoiced final consonants in bare, unsuffixed forms can behave in two different ways when given a vowel-initial suffix: some have a consistent unvoiced stem-final consonant, while other stems show a voicing alternation. A Turkish example is shown in (3): both of the words have unvoiced stops in unsuffixed form; when taking a vowel-initial suffix like the possessive, [anatʃʰ] 'female cub' is unchanged, while [amatʃʰ] 'target' instead voices its final consonant.

(3) *Variable voicing alternation in Turkish bare stems and possessives (cf. Becker et al., 2011, p. 85)*

   a. anatʃʰ   anatʃʰ-i   'female cub'
   b. amatʃʰ   amadʒ-i   'target'

While speakers must learn individually which words alternate, there are patterns: in Dutch, for example, fricatives [s f χ] alternate more frequently than stops [p t]; in a wug test, Dutch speakers matched the frequencies of alternations in the lexicon for nonce words ending in different consonants.

The Turkish and Dutch studies, like many others (e.g. Becker & Gouskova, 2016; Bybee & Pardo, 1981; Linzen et al., 2013), look at lexical variation in a morphophonological alternation—that is, stems can have several alternants related by a phonological opposition (voiced vs. voiceless, vowel vs. zero, monophthong vs. diphthong, etc.), and the task of the learner is to figure out which lexical items undergo the alternation and which do not. Gouskova et al. (2015) look at lexical variation in another domain, which plays a role in the English past tense and German plural: allomorph selection. Russian masculine nouns have three productive diminutive suffixes: *-ók*, *-ʲik*, and *-tʃʲik*. Russian speakers produce diminutives for nonce words in accordance with the phonological characteristics of words that take the respective suffixes: for example, nouns ending in clusters take *-ʲik* rather than *-ók* and *-tʃʲik*, and speakers likewise preferred to assign *-ʲik* to nonce words ending in clusters (I study this case in Chapter 6).

The theory of lexical specification developed in Chapter 3 handles lexically variable morphophonological alternations using the same mechanism as lexically variable allomorph selection: diacritic features. Frequency matching effects have been found in both types of lexical variability, and it is convenient to be able to treat them as variants of the same general phenomenon.

This concludes a summary of the types of phonological effects found in typical wug tests. I now move on to other potential sources of predictability whose productivity has not been studied as systematically.

2.1.4 Syntactic and semantic predictors of allomorphy

A speaker's knowledge of a lexical item goes beyond its phonology and its association with lexically specific allomorphy. It also includes syntactic features associated with the word—for example, the gender of a noun—and its meaning. Although linguists have frequently observed correlations between a word's syntactic properties and its inflectional realization, active grammatical connections between them have rarely been tested empirically. I discuss syntactic and semantic factors influencing inflectional patterns in this section before moving on in Section 2.1.5 to yet another type of predictability, which is the main object of my study: correlations *between* related forms built from the same root.

In languages that partition the set of nouns into multiple agreement classes (genders) and inflection classes (groups of nouns that show the same realizations in inflected forms), the two types of classes are often closely aligned. That is, nouns with the same grammatical gender (agreement class) often share a set of inflectional endings, and nouns with the same inflectional pattern also usually share an agreement class, meaning that they trigger the same agreement on adjectives and past-tense verbs and are coreferenced by the same gendered set of pronouns (whether grounded in natural gender or not). For example, Russian has three genders: masculine, feminine, and neuter. Like other gender systems, Russian gender is built around a common semantic core (male animates usually take masculine agreement, etc.) with inanimates going into all three genders. Russian also has four main inflection classes, shown (in phonemic transcription) in Table 2.1.

class	I	II	III	IV
example	‘law’	‘school’	‘bone’	‘wine’
nominative	zakon	ʂkol-a	kost^{j}	vin-o
accusative	zakon	ʂkol-u	kost^{j}	vin-o
dative	zakon-u	ʂkol-e	kost^{j}-i	vin-u
genitive	zakon-a	ʂkol-i	kost^{j}-i	vin-a
instrumental	zakon-om	ʂkol-oj	kost^{j}-ju	vin-om
locative	zakon-e	ʂkol-e	kost^{j}-i	vin-e

Table 2.1: Inflection classes with singular case forms for Russian nouns (from Corbett, 1982)

Class I includes masculine inanimate nouns like [zakon] ‘law’, masculine animate nouns like [sv^{j}okor] ‘father-in-law’, and animate nouns that can take either masculine or feminine agreement like [doktor] ‘doctor’. Most masculine nouns fall into class I. Class II includes feminine animate and inanimate nouns ([ʐenʃ^{j}ːina] ‘woman’, [ʂkola] ‘school) and a sizable group of masculine animate nouns, many of which are very common (like [muʃ^{j}ːina] ‘man’ and [d^{j}ad^{j}a] ‘uncle’). Most feminine nouns are in class II, though a substantial minority belong to class III, which is almost entirely feminine. Finally, class IV is the primary class for neuters, although it also includes some derived forms of masculine nouns like [xlʲebuʂko], a diminutive of [xlʲeb] ‘bread’. Thus, in Russian, there is not a perfect alignment between gender and inflection class, but inferences can be made in both directions. (See Corbett (1982) and Gouskova and Bobaljik (2022) for additional details about the intersection of gender and inflection class in Russian.)

Rodina and Westergaard (2012) show that even very young children (aged 2.5–4) have learned correlations between a noun’s inflection class and its gender agreement properties. They tested how children produce gender agreement on adjectives and past tense verbs for several groups of nouns with atypical gender agreement, including masculine class II nouns like [dʲadʲa] ‘uncle’, which take masculine agreement, and class I diminutives of female names like [lʲenok], from [lʲena] ‘Lena’, which can take either masculine or feminine agreement. The children usually correctly assigned masculine agreement to nouns like [dʲadʲa], but sometimes incorrectly used feminine agreement in a minority of trials (53/696, or 7.6%). Of the 25 young children in the experiment, 11 made at least one error, and all correctly produced masculine agreement some of the time. The presence of these mistakes suggests that children associate class II nouns with feminine agreement, even if they are able to sort out most of the exceptions. On the other hand, the young children overwhelmingly produced *masculine* agreement for class I diminutives of female names like [lʲenok]: 197 of 223 trials (85.7%) had masculine agreement, and 15 of the 25 children exclusively used masculine agreement endings for these nouns. For this class of nouns, children use inflection class as a cue to determine gender agreement overriding the semantic cue that the referents of these nouns are female. The importance of inflection class as a cue to gender is not necessarily a given. Corbett (1991) predicts that Russian children should use semantic gender agreement when possible for animates—that is, [dʲadʲa] ‘uncle’ should always get masculine agreement, while [lʲenok] should always get feminine agreement. Thus, Rodina and Westergaard (2012) show that Russian children use the correlation between inflection class and gender as an organizing property of their morphosyntax, even at times when such cues are ambiguous and less ambiguous clues are available.

The connection between gender and inflection is less clearly fundamental to other languages' morphology. For example, one of the factors influencing a German noun's choice of plural is its gender, which is marked on articles and cannot be readily inferred from the noun itself. As I discuss in more detail in Section 2.2.2.4, feminine nouns tend to take the plural suffix *-(ə)n*, and nouns that take *-e* or no suffix are more commonly masculine or neuter (Elsen, 2002; McCurdy et al., 2020). Nonce word studies of German testing whether speakers correlate gender and plural realization have had mixed results. For example, Marcus et al. (1995) asked German speakers to rate potential plurals of nonce words; they reported no effect of gender on how participants scored a given plural. However, Zaretsky and Lange (2016) had speakers *produce* plural forms of nonce words (the same stimuli used by Marcus et al. (1995)), and they found that feminine nouns were treated differently from masculine and neuter ones. Thus, there is some experimental evidence that speakers learn correlations present in the lexicon between gender and lexically specific allomorphy.

I am not aware of any nonce word studies testing the effects of lexical semantics on inflection class in particular—although Gagliardi and Lidz (2014) test the effect of semantics on assignment of nonce words to *agreement* classes in Tsez, a Nakh-Dagestani language.[1] This is presumably due (at least in part) to the fact that it is difficult to reliably assign semantics to words that do not exist (though not impossible: Gagliardi and Lidz use pictures to identify a nonce word as referring to an animal, for example). However, some semantic effects do exist in the lexicon. These can be very specific. Czech nominal paradigms are roughly divided into two broad classes: "hard-stem" and "soft-stem", mostly depending on a noun's final consonant. As described in more detail in Chapter 5, masculine inanimate hard-stem nouns show lexically conditioned allomorphy in the genitive (which can be *-u* or *-a*) and locative (which can be *-u* or *-ɛ*). My corpus includes 33 words that (fully or nearly) categorically take *-a* in the genitive and *-u* in the locative. Of these,

[1] One study with a similar goal is Magomedova (2017), discussed in Section 6.1.3, who tested whether certain Russian diminutives are preferred in pejorative or affectionate contexts. This study, however, was intended as a test of the influence of the context in which a word is used, not of its inherent meaning.

21 mark units of time,[2] including 8 of the 9 hard-stem month names (like [uːnor] 'February'), nouns derived from time adverbials like 'today' ([dnɛʃ-ɛk], cf. adverb [dnɛs]) and 'earlier period' ([dr̝iːv-jɛjʃ-ɛk], from the adverb [dr̝iːv] 'earlier'), and times of the day ([vɛtʃɛr] 'evening'). Whether speakers have actually learned and productively apply the association between time terms (whether monomorphemic or derived with *-ɛk*) and this inflection pattern is a separate question.

Similar—if perhaps less specific—effects of meaning have been noted in the assignment of gender, which (as previously discussed) is often strongly correlated with inflectional behavior. Corbett (1991, c. 2–3) gives examples of agreement class systems that are wholly or partially determined by semantics. For example, Tamil nouns are divided into three semantically determined genders containing male rational creatures (humans, gods, anthropomorphized animals, etc.), female rational creatures, and non-rational creatures; English nouns follow a similar division. Many languages, like Russian, take natural gender as the core of their agreement class system, but most inanimate nouns are given masculine or feminine gender as well. Other languages show more complicated patterns: for example, Archi, a Nakh-Dagestani language, divides non-rationals into two genders, III and IV: gender III tends to contain domestic animals and birds, larger wild animals and birds, insects, mythical beings, musical instruments, cereals, trees, water phenomena, and astronomical and meteorological phenomena; gender IV tends to contain young animals and birds, smaller wild animals and birds, tools and cutting instruments, cloth and clothing, metals, liquids, and abstract concepts.

Many of the semantic generalizations about agreement and inflection described above are somewhat subjective analyses by linguists. A more objective, quantitative approach to meaning uses semantic vectors, which operationalize a word's meaning as its distribution. Semantic vectors are calculated by taking the contexts in which a word appears; two words will have similar semantic

[2]Ten of the remaining 12 words are generally treated as animate, which introduces a confound, since *animate* nouns of this class always have *-a* in the genitive. Thus, there are really only two non-time nouns that show this inflection pattern: [jɛtʃmɛn] 'barley' (which is usually declined as a soft-stem noun anyway) and [vʃɛhomiːr], an antiquated word for 'universe'.

vectors if they cooccur with similar words in their contexts. Semantic vectors correspond surprisingly well with our intuitive sense of meaning, and are widely used in natural language processing. Williams et al. (2020) find that the semantic vectors of Czech and German nouns are moderately informative of their inflectional paradigm (in Czech, where nouns are inflected for case and number) and plural form (in German, where nouns are not inflected for case), even once gender is taken into account. Moreover, meaning is not equally predictive for all classes: some are better defined semantically than others.

On the other hand, Guzmán Naranjo (2020) finds that a neural network trained on semantic vectors of Russian nouns performs only marginally above chance in predicting their inflectional paradigms. This discrepancy may be due to a difference in the languages involved (though this is unlikely, since Czech and Russian are both Slavic languages and share a substantial portion of their vocabulary) or to differences in the design of the studies: Williams et al. (2020) calculate a measure of mutual information between semantic vectors and inflection class given gender, while Guzmán Naranjo (2020) trains a neural network on semantic vectors and has it predict individual paradigm cells. The relatively poor performance of the Russian neural network may also be attributable to the finer resolution of its input: whereas Guzmán Naranjo's Russian data set groups nouns into 108 distinct inflectional classes (many of which differ in a small number of paradigm cells), Williams et al. (2020) have only 13 Czech classes and 16 German classes, since they at times combine several Czech variants (some of which could be treated as phonological variants of a single set of underlying forms) into a single class.

While further research is needed to find the reason for these studies' differing results, the tentative conclusion is that semantic vectors can be somewhat predictive of inflectional patterns in the lexicon, at least in some languages. However, it would be very difficult to use semantic vectors to study the effect of semantics on speakers' inflection of nonce words in a wug test: semantic vectors are based on a word's existing distribution, and nonce words, by definition, have not previously appeared in the language, so there is no easy way to calculate semantic vectors for nonce words.

2.1.5 Implicative structure in inflectional paradigms

So far in this section, I have surveyed external factors influencing the distribution of morphological patterns: phonology, syntax, and semantics. I now turn to the *internal* structure of morphological paradigms: correlations between morphologically complex forms built from the same stem. In general, morphologically related words do not stand completely independent of one another. Rather, paradigms have an *implicative structure* (Wurzel, 1989, p. 114): inflected forms are partially or wholly predictive of one another. In other words, knowing some members of an inflectional paradigm allows speakers to (with greater or lesser reliability) infer unknown forms. The extent to which inflectional systems license these inferences is studied as the Paradigm Cell Filling Problem (e.g. Ackerman & Malouf, 2013; Ackerman et al., 2009; Bonami & Beniamine, 2016; Parker & Sims, 2020). Ackerman and Malouf (2013) posit that inflectional systems, no matter how large their set of exponents, are relatively learnable because these exponents are distributed in such a way that speakers can infer unknown inflected forms of a given word from known ones (the low conditional entropy conjecture). The literature on the Paradigm Cell Filling Problem has focused on the information contained within inflectional paradigms themselves, without testing whether and how speakers actually use this information. Thus, Bonami and Beniamine (2016) conclude: "It should be stressed that this paper only established that speakers are exposed to relevant information and that this information is helpful; the next step, of course, is to establish experimentally that speakers do indeed rely on joint prediction when addressing predicting the form of unknown words. [...] Until such studies [are] available, though, there is no reason to doubt that speakers make use of what information is available to them." The studies in this dissertation are intended to show experimentally that speakers do use known inflected forms in the inflection of novel inflected forms.

The discussion in this section also serves as a preview to the remainder of the chapter, in that I discuss not just examples themselves, but also generative analyses of them. Both Müller (2004) (for Russian, discussed in Section 2.1.5.1) and Halle and Marantz (2008) (for Polish, discussed in Section 2.1.5.2) encode arbitrary morphological patterns through *diacritic features* associated with individual lexical items. I discuss this theoretical tool at length in Section 2.2.2.3 and adopt it throughout the dissertation. The two analyses also differ in how they encode correlations between morphological patterns, the subject of Section 2.3.2. Using one of the criteria I develop in that section, I classify the generalizations in Müller (2004) as *hard-coded* and those in Halle and Marantz (2008) as not hard-coded. I compare the two in that section and argue that the type of learned generalization proposed by Halle and Marantz (2008) is necessary to account for morphological dependencies, and I adopt a very similar proposal in my model in Chapter 3.

2.1.5.1 Syncretism

Ackerman and Malouf (2013) and their interlocutors are not working in a generative framework, and often criticize piece-based theories like Distributed Morphology for not treating morphological paradigms as discrete units of analysis. In Distributed Morphology and other syntactically grounded theories of morphology, words are composed of smaller units that comprise morphosyntactic feature sets: a word is a whole made up of pieces (morphemes), whereas in paradigmatic approaches, a word is a part in a large whole (a paradigm). The piece-based perspective is only well-equipped to deal with one very particular type of paradigmatic relation: identity. That is, theories of morphology in which words are composed of smaller syntactic pieces—such as Distributed Morphology and Nanosyntax (see Baunaz & Lander, 2018)—largely address paradigmatic relations in the form of *syncretism*, namely when two paradigm cells (that is, morphosyntactic feature sets) share a realization. Accordingly, much of the work in these theories has the goal of accounting for why two affixes for case, number, gender, etc. look the same in certain circumstances. Let us look at some typical examples of syncretism using the basic Russian inflection classes in Table 2.1, repeated here:

class	I	II	III	IV
example	'law'	'school'	'bone'	'wine'
nominative	zakon	ʂkol-a	kost^{j}	vin-o
accusative	zakon	ʂkol-u	kost^{j}	vin-o
dative	zakon-u	ʂkol-e	kost^{j}-i	vin-u
genitive	zakon-a	ʂkol-i	kost^{j}-i	vin-a
instrumental	zakon-om	ʂkol-oj	kost^{j}-ju	vin-om
locative	zakon-e	ʂkol-e	kost^{j}-i	vin-e

Table 2.2: Inflection classes with singular case forms for Russian nouns (from Corbett, 1982)

Müller (2004) notes two kinds of syncretism in these paradigms: first, in *intra-paradigmatic* syncretism, a noun takes the same ending in multiple cases. Thus, [ʂkola] and other class II nouns have *-e* in both the genitive and the locative, while class III nouns like [kost^{j}] have *-i* in the dative, genitive, and locative. In *trans-paradigmatic* syncretism, different noun classes share a case ending. Some examples include the dative, genitive, instrumental, and locative in class I and IV (the last of which, *-e*, is also shared by class II) and genitive *-i* in class II and III (which is realized as [ɨ] for words like [ʂkola] 'school' through a regular phonological process).

Syntactic theories of morphology have two natural ways to account for syncretism. The first is shared features: the syncretic forms share some set of morphosyntactic features, even if they differ in others; the differences between them are simply not expressed. Relatedly, syncretic forms can be defaults that apply to all cases where there is no more specific affix available. For example, Müller (2004) decomposes cases and inflection classes into combinations of features that can be underspecified in cases of syncretism: the six cases break down into [±subj(ect)], [±gov(erned)], and [±obl(ique)], and the four inflection classes decompose into [$\pm\alpha$] and [$\pm\beta$]. The feature composition of each case and inflection class is shown in Table 2.3.

case	[±subj]	[±gov]	[±obl]
nominative	+	−	−
accusative	−	+	−
dative	−	+	+
genitive	+	+	+
instrumental	+	−	+
locative	−	−	+

class	[±α]	[±β]
I	+	−
II	−	+
III	−	−
IV	+	+

Table 2.3: Composition of case and inflection class features in Müller (2004)

In (4), we see how these featural decompositions can be used to account for syncretism through underspecified rules of realization. First, the instrumentals of class II and III are non-syncretic: they share their realization with no other paradigm cells. Accordingly, their rules of realization, in (4a) and (4b), respectively, are fully specified. A less specified rule, in (4c), captures the intra-paradigmatic syncretism in class II: the dative and locative are both -*e*, and both share the feature specification of [−subj, +obl]. These features are also shared by the instrumental, but rule (4c) does not apply to the instrumental because it is overriden by the more specific rule in (4a) according to the Subset Principle (Halle, 1997), which states that the most specific rule of realization applies in any given circumstance. Finally, the rule in (4d) is the most general and captures both the intra-paradigmatic syncretism in class III (-*i* is shared between the dative, genitive, and locative) and the trans-paradigmatic syncretism of the genitive in class II and III. This is because the rule is very general and applies to all cases except for the nominative and accusative in class II and III—except those which are overridden by the other, more specific rules.

(4) *Rules of realization for syncretic Russian suffixes in class II and III (cf. Müller, 2004, p. 204)*

a. [$-$subj, $-$gov, $+$obl] $\leftrightarrow$ oj / [$-\alpha$, $+\beta$] ___

b. [$-$subj, $-$gov, $+$obl] $\leftrightarrow$ ju / [$-\alpha$, $-\beta$] ___

c. [$-$subj, $+$obl] $\leftrightarrow$ e / [$-\alpha$, $+\beta$] ___

d. [$+$obl] $\leftrightarrow$ i / [$-\alpha$] ___

The analysis in (4), as is typical of Distributed Morphology and related theories, mix morphosyntactic features like [$-$subj] with diacritic features like [$-\alpha$]. Part of the learning process is to induce the proper rules and features that can combine to account for the surface patterns. In Section 2.3.2.2, I propose a different way of encoding the lexically specific inflectional properties indexed by [$\pm\alpha$] and [$\pm\beta$] that does not require speakers to learn these unified inflection class diacritics.

2.1.5.2 Beyond syncretism

Although shared featural content can provide elegant analyses for syncretism, this tool has nothing to say about implicational relationships between related forms that are not identical. Müller (2004), as is typical for such analyses, assumes that the nouns pattern into four inflection classes and encodes these in inflection class features [$\pm\alpha$, $\pm\beta$]. These features provide an analytical explanation for why nouns with *-oj* in the instrumental also have *-a* in the nominative: both of these suffixes are inserted for nouns that have the class II featural specification [$-\alpha$, $+\beta$]. These umbrella diacritics, however, are very blunt instruments and ignore many of the intricacies of actual morphological systems. In Russian, for example, nouns can show many inflectional patterns beyond these four broad classes (see Parker & Sims, 2020). Relations between these subpatterns cannot be captured through shared featural identity and umbrella inflection class features; instead, as we will see shortly, they require relations *between* more narrowly targeted diacritic features. These correlations have no syntactic content whatsoever, only morphological content. They are truly paradigmatic, and thus fall outside the natural domain of morphological theories based on syntactic pieces, as argued by Ackerman and Malouf (2013).

This does not mean that paradigmatic relations are *incompatible* with theories like Distributed Morphology. Indeed, in the following example, Halle and Marantz (2008) capture correlations between paradigm cells in the context of a Distributed Morphology analysis. However, there is an important difference between the constructs used by Müller (2004) to account for syncretism and those used by Halle and Marantz (2008) to account for non-syncretic paradigmatic relations. In the typology used in Section 2.3, the syncretic relations of Müller (2004) are *hard-coded*—in that they are built into the rules of realization in (4) that are active within the course of a derivation—and morphosyntactically *grounded*, meaning that they emerge from natural properties of the morphosyntactic feature sets used in the analysis. By contrast, Halle and Marantz (2008) express the paradigmatic relations as generalizations over sets of lexical items whose behavior speakers have learned individually. They are not hard-coded into rules of realization, and are thus not active in the course of derivations. They are also not morphosyntactically grounded: they express relations in the lexicon that are strong but that do not emerge from any natural properties of the features in question. The approach of Halle and Marantz (2008) is, in spirit, very similar to the one that I adopt. Grammars in Distributed Morphology can capture paradigmatic relations, but only *outside* of the components that are usually taken to comprise the core of the theory, namely rules of realization and other rules mediating between syntax and phonology. Instead, paradigmatic relations are stored alongside other generalizations over arbitrary patterns in a pattern-learning module described in Chapter 3.

Halle and Marantz (2008) analyze the seven Polish masculine noun paradigms in Table 2.4 (see also Cameron-Faulkner & Carstairs-McCarthy, 2000). All seven of these paradigms are cognate to class I in Russian (as well as the Czech inflectional suffixes described in Chapter 1 and Chapter 5).[3] There are three main differences between these Polish paradigms and Russian class I as

[3]In Halle and Marantz (2008), Polish is presented orthographically and class 2 is split into 2a and 2b; otherwise, the data are the same. I also add glosses.

presented in Table 2.2. First, Polish has an additional vocative case, which Russian has lost. Second, the dative has an additional variant, *-ɔvʲi*. Third, *-u*, which only appears in the dative in the Russian paradigms in Table 2.2, also appears for some nouns in every case except the nominative and instrumental. (In fact, the Russian paradigms presented by Corbett (1982), as shown in Table 2.2, paper over similar variation: some Russian class I nouns also have *-u* in the genitive and/or locative.) This presentation also omits the accusative, which, as in Russian, is identical with the nominative for inanimate nouns like [dvur] 'court' and with the genitive for animate nouns like [pɔlak] 'Pole'.

class	1	2	3	4	5	6	7
example	'professor'	'Pole'	'dog'	'gentleman'	'merchant'	'court'	'country'
nominative	prɔfɛsɔr-Ø	pɔlak-Ø	pʲɛs-Ø	pan-Ø	kupʲɛts-Ø	dvur-Ø	kraj-Ø
genitive	prɔfɛsɔr-a	pɔlak-a	ps-a	pan-a	kupts-a	dvɔr-u	kraj-u
dative	prɔfɛsɔr-ɔvʲi	pɔlak-ɔvʲi	ps-u	pan-u	kupts-ɔvʲi	dvɔr-ɔvʲi	kraj-ɔvʲi
instrumental	prɔfɛsɔr-ɛm	pɔlakʲ-ɛm	ps-ɛm	pan-ɛm	kupts-ɛm	dvɔr-ɛm	kraj-ɛm
locative	prɔfɛsɔʐ-ɛ	pɔlak-u	pɕ-ɛ	pan-u	kupts-u	dvɔʐ-ɛ	kraj-u
vocative	prɔfɛsɔʐ-ɛ	pɔlak-u	pɕ-ɛ	paɲ-ɛ	kuptʂ-ɛ	dvɔʐ-ɛ	kraj-u

Table 2.4: Polish masculine noun inflection classes (Cameron-Faulkner & Carstairs-McCarthy, 2000; Halle & Marantz, 2008)

Ignoring the stem alternations[4], one could analyze these data by saying that Polish masculine nouns fall into one of seven classes labelled 1–7, and that each class is associated with a full set of case endings. However, this would not be very enlightening: unlike in the (oversimplified) Russian case, where the trans-paradigmatic syncretisms (that is, overlaps between the classes) are quite rare, the seven Polish classes each share only two exponents for each of the genitive, dative,

[4]The vowel–zero alternations in words like [pʲɛs] 'dog' are limited to particular nouns (see e.g. Rubach, 1984), and also appear in Czech (see Chapter 1), Russian (see Section 2.2.1 and Section 2.2.2.4), and other Slavic languages. The stem-final consonant alternations triggered by *-ɛ* and *-ɛm* are fully regular.

locative, and vocative: -*u* can represent all four cases; -*ɛ*, the locative and vocative; -*a*, the genitive, and -*ɔv*^{j}*i*, the dative.

The first part of the analysis in Halle and Marantz (2008) comprises the rules of realization, known in Distributed Morphology as vocabulary items, in (5) (here I only include rules relevant to the analysis). Most of the rules directly spell out one of the cases. The exception is (5e), which operates as a default: if none of the other rules apply, spell out the case ending as -*u*.

(5) *Vocabulary items for Polish masculine case forms (Halle & Marantz, 2008, p. 68)*

 a. GEN ↔ a

 b. DAT ↔ ɔv^{j}i

 c. LOC ↔ ɛ

 d. VOC ↔ ɛ

 e. ↔ u

If the rules in (5) were the entire analysis, we would have a problem: no noun would ever take -*u*, because the relevant case features are indexed in more specific rules. Thus, the default -*u* is accessed by means of rules of *impoverishment*, an operation that removes a syntactic feature in a particular context. The rules of impoverishment in (6) delete case features in the context of corresponding lexical diacritic features. That is, the genitive case feature is deleted in the context of a [−Gen] feature, and so on.

(6) *Rules of impoverishment for Polish masculine case forms (Halle & Marantz, 2008, p. 69)*

 a. GEN → Ø / [−Gen] ___

 b. DAT → Ø / [−Dat] ___

 c. LOC → Ø / [−Loc] ___

 d. VOC → Ø / [−Voc] ___

If the genitive feature is deleted by the rule in (6a), then the case suffix will be spelled out as default

-u. This means that nouns with *-u* for a given case form have the corresponding feature triggering impoverishment in their lexical entry. (I discuss lexical diacritic features, a core component of my model of morphological dependencies, at greater length in Section 2.2.2.3.) Class 1 nouns like [prɔfɛsɔr] 'professor', with no *-u* anywhere in the paradigm, are totally unmarked, while class 7 nouns like [kraj] 'country' have [−Gen], [−Loc], and [−Voc].

With this analysis in place, we can proceed to the part of the analysis that is of interest here: accounting for the relations of cooccurrence between the realizations of different classes, which are implemented, in this case, as cooccurrence between shared features on lexical entries. With two exponents for each of four cases, there are sixteen possible classes of masculine nouns. However, only seven of these actually occur. This means that certain combinations of case endings (that is, features) cluster together, while other logically possible combinations do not occur. (Ackerman and Malouf (2013) make a similar point about much more complicated inflectional systems.) For example, no noun has *-u* for both the genitive and the dative. For Halle and Marantz (2008), this is expressed as the generalization that no noun has both the [−Gen] feature and the [−Dat] feature (*[−Gen, −Dat]). Thus, if a speaker knows that a given masculine noun takes *-u* in the genitive (class 6 and 7 in Table 2.4), she can infer that it also takes *-ɔv^{j}i* in the dative, and vice versa: dative *-u* (class 3 and 4) implies genitive *-a.*

The analysis of Halle and Marantz (2008) emphasizes two important points for this dissertation. First of all, it is (often) more productive to conceptualize a word's inflectional behavior as a network of related forms than as one of a set of mutually exclusive inflection classes (Ackerman et al., 2009). I return to this point in Section 2.3.2.2. Even more foundationally, known forms of a word, individually or together, can predict unknown forms. This dissertation follows in their footsteps, showing that these generalizations can in fact be facilitated by a syntactic, piece-based theory of morphology: lexically arbitrary morphological behavior is indexed by diacritic features on lexical items that play active roles in the derivation, and paradigmatic relations are encoded as cooccurrence relations between these diacritic features. Thus, before we can encode morphological dependencies, we need a theory of how grammars encode lexical exceptionality—for example, with diacritic features like the Polish [−Gen]. I discuss diacritic features and alternative ways of grammatically associating lexical items with their arbitrary morphological behavior in Section 2.2.

Before moving onto the theoretical exploration that will lay the groundwork for a theory of morphological dependencies, I make one note about the work described in this section. Halle and Marantz (2008), like Ackerman and Malouf (2013) and colleagues discussed above, study the structure of the lexicon rather than what speakers actually do. A key goal of this dissertation is to fill the gap in the literature by demonstrating that speakers do, in fact, productively apply morphological dependencies. The published nonce word study most similar to those in this dissertation is Bybee and Pardo (1981), who test the productivity of various stem alternations in Spanish. Like the studies in this dissertation, Bybee and Pardo provided participants with two forms of a nonce word and elicited a third. However, their study did not really test the productivity of morphological dependencies as such; rather, they are mostly concerned with whether speakers extend lexically conditioned stem alternations presented in one form to other forms that are expected to show the *same* alternation.

In reading the next theoretical section, we should keep in mind that the ultimate goal is to account for patterns of behavior demonstrated experimentally; the theoretical concerns described are useful to us insofar as they are accurate and facilitate the learning of morphological dependencies demonstrated in the case studies in later chapters.

2.2 Lexical variation in the grammar

In Section 2.1, I showed how lexically variable patterns of allomorphy observe gradient generalizations: a word's behavior is often neither fully predictable nor fully random, subject to different kinds of generalizations that hold often but not universally. In this section, I take a step back and discuss theories of how speakers know which individual words follow which patterns. I argue in favor of theories that encode lexical exceptionality in lexical diacritic features stored in underlying forms of roots and affixes, as used by Halle and Marantz (2008) in Polish and Müller (2004) in Russian in the analyses described in Section 2.1.5. These can then serve as the basis for encoding morphological dependencies as coccurrences between diacritic features; in Section 2.3, I survey ways of grammatically encoding the generalizations discussed in Section 2.1 above, adopting an approach very similar to that of Halle and Marantz (2008) described in Section 2.1.5.2.

2.2.1 Morphophonological exceptionality

Lexical items can show exceptionality in many morphophonological domains. My main focus is on morphological patterning: specifically, my case studies mostly involve lexically conditioned *allomorph selection*. The inflection classes in Russian and Polish surveyed in Section 2.1 differ in the suffixes that they select to represent combinations of case and number. Similarly, (some) "irregular" plurals like English *oxen* differ in their choice of plural suffix. There are other ways in which lexical items can be exceptional, however. Some of these arguably have a morphological component: for example, in Chapter 6 I show the effects of lexically determined stress patterns in Russian. Others are more squarely located in the phonology. I discuss these types of exceptionality below.

One possibility is that different corners of the lexicon show more or less permissive phonotactics. For example, Japanese divides its vocabulary into four recognizable strata that define a hierarchy of phonotactic restrictions: native (Yamato) words are the most restricted, followed by long-established loans from Chinese (Sino-Japanese), more recent assimilated foreign words, and unassimilated foreign words (Ito & Mester, 2008; Itô & Mester, 1999). Each layer, in turn, allows sounds and sequences that previous layers do not, as shown in Table 2.5.

stratum	voiced obstruent geminates?	non-geminate [p]?	voiceless obstruents after nasals?
Yamato	no	no	no
Sino-Japanese	no	no	yes
assimilated foreign	no	yes	yes
unassimilated foreign	yes	yes	yes

Table 2.5: The phonotactics of Japanese lexical strata (from Itô & Mester, 1999, p. 69)

These lexical classes are active in the grammar in other ways as well: for example, Sino-Japanese roots tend to form compounds with other Sino-Japanese roots (Itô & Mester, 1999, p. 64). The authors propose a nested lexicon, wherein Yamato words form the lexical core and successively more permissive layers are placed at more peripheral strata. These layers are implemented as faithfulness constraints with different rankings relative to a fixed set of markedness constraints (like *p): Yamato words have the weakest faithfulness, while unassimilated foreign words have the strongest.

Phonotactic restrictions like those in Japanese are detectable not just by the shape of morphemes themselves, but by alternations in derived environments. As shown in Table 2.5, Yamato words differ from others in that obstruents after nasals must be voiced. Accordingly, Yamato morphemes beginning with voiceless obstruents, like the gerund *-te*, voice after nasals: /sin-te/ 'die (gerund)' surfaces as [ɕinde]. By contrast, the Sino-Japanese compound /sin-tai/ 'body' surfaces with faithful voicing, as [ɕintai].

The alternation triggered by phonotactic differences in Japanese leads us to another category of lexical exceptionality: lexical items can exceptionally undergo or be exempt from phonological processes in certain environments. In Japanese, this susceptibility to phonological alternations can be recast as the consequence of phonotactic restrictiveness, as Itô and Mester (1999) do, but other lexically specific phonological alternations cannot necessarily be reduced to phonotactics. One such example is the "fleeting vowels" found in many Slavic languages (though similar patterns are found in unrelated languages as well): some nouns show vowel–zero alternations in the final syllable, with a vowel appearing when the word is unaffixed. The examples in Table 2.6, from Russian, show that alternating words contrast in both directions: on the one hand, with words that have vowels throughout their paradigm, and on the other, with words that have clusters throughout their paradigm.

word	*pattern*	nominative	genitive
'wind'	fleeting vowel	v^jét^jir	v^jétr-ə
'motor boat'	vowel throughout	kát^jir	kát^jir-ə
'meter'	no vowel throughout	m^jétr	m^jétr-ə

Table 2.6: Russian alternating and non-alternating nouns (from Gouskova, 2012, p. 82)

The fleeting vowel alternation is arguably phonologically optimizing (see Section 2.3.1)—for example, Gouskova (2012) attributes it as satisfying a general preference in Russian against mid vowels. However, even if the alternation is driven by a phonological pressure, this pressure is not absolute: phonologically similar words may pattern differently. Recognizing this fact only raises more questions for the analysis: is the alternation deletion or epenthesis? Is alternation a default process (from which non-alternators must be marked as exempt) or exceptional (such that alternators must be marked)? The answers may differ from one Slavic language to another: Gouskova (2012) argues that the Russian alternation is exceptional deletion and that alternators are marked, while Rysling (2016), using a similar theoretical framework, argues that Polish alternation is default epenthesis and that words ending in non-alternating clusters are exceptionally exempt.

2.2.2 Grammatical representations of exceptionality

The common link behind the phenomena described in this section is that lexical items show morphophonological behavior that is not fully predictable from their phonological form—these are the patterns over which speakers extract generalizations, as described in Section 2.1. Thus, in order for the grammar to properly derive the correct output forms for every lexical item, lexical items that behave differently must somehow differ in their representation. This difference can be grammatically encoded in several ways. One possibility is that all surface forms are stored in full, and speakers only create forms productively when they have no stored version. Alternately, regular forms are derived while exceptional forms are stored in full. A third possibility is that *all* forms are derived and that forms that differ in their behavior have some abstract difference in their lexical representation. This can take the form of either a symbolic diacritic marker that indexes a lexical item's behavior, or some structural difference that affects the derivation but does not appear on the surface. I discuss each of these possibilities in turn.

2.2.2.1 Connectionist and analogical models

The first cut to be made between models is whether they allow for derivations that include formal symbolic rules. Although I focus mostly on models that do involve symbolic rules, I first briefly discuss connectionist, or parallel distributed processing, models like Rumelhart and McClelland (1986), which store neither rules nor words. Instead, mappings between bases and derived forms are encapsulated in a network of nodes that combine to represent sequences of phonemes. For example, in their model, words are represented as groups of triplets of phonological features called Wickelfeatures, which include a feature value for each phoneme, its predecessor, and its successor. For example, the verb [ɾʌn] includes a Wickelfeature encoding that [ʌ] is high (their feature set only includes a two-way height distinction), its predecessor [ɾ] is a continuant, and its successor [n] is also a continuant. This is just one of the many Wickelfeatures in this verb's representation, alongside others that represent various phonological properties of the [r ʌ n] triplet of phonemes and Wickelfeatures representing the word's other triplets of phonemes: [# r ʌ] and [ʌ n #], where # represents a word boundary. The mapping between base verbs and their past-tense forms, then, is a series of weighted connections between input and output. In the case of [rʌn∼ræn], all Wickelfeatures are mapped faithfully except those marking the height of the vowel, which is high in the input and low and in the output.

In their model, learning is a task of adjusting the weights of connections between the input nodes and output nodes in response to input–output pairs. If the weights are properly calibrated, they can capture both regular and irregular forms: regular forms follow a few general patterns that can be productively extended to new words (for example, suffixing *-d*, which in this model involves adding new Wickelfeatures centered around [d] and changing Wickelfeatures that end in word boundaries to those that end in the features for [d]), while irregular forms follow highly specialized pathways specific to the individual lemmas in question. However, both familiar and unfamiliar words are derived anew each time: neither regular nor irregular forms are stored, and the model will occasionally output the wrong past tense even for very familiar verbs. Any generalization of past tense formation patterns emerges from the feature weights; none are stored as rules.

In network models like Bybee (1995) and other analogical models, both base and derived forms are stored in full, and generalizations of varying productivity emerge from patterns of similarity (often called *schemata*) between forms that behave in the same way. For example, in Bybee (1995), English bare present-tense verbs are linked to their past-tense form by virtue of a shared base—this is one part of the network. The links go further: forms that share morphosyntactic features are also linked together. Thus, for example, the English present-tense verbs *cling*, *sling*, *sting*, and *stick* are connected, as are their past tenses *clung*, *slung*, *stung*, and *stuck*; each base verb is also linked to its past tense. Thus, the verbs form a dense network of associations in several dimensions that is reinforced by the forms' phonological similarity.

In models like that of Bybee (1995), when a speaker wishes to produce a known inflected form, she retrieves it from her lexicon. If she needs to generate a new form, she does so by analogy to previously existing forms. For example, the network described above features verbs ending in [ɪŋ] or [ɪk] whose past tenses have [ʌ]. Thus, when faced with a new verb like [splɪŋ], speakers compare it to existing schemata for similar verbs and form the past tense by analogy—given the network of verbs described above, this may be [splʌŋ]. Alternatively, speakers may place it into the dominant schema of verbs that add *-d* to form the past tense and produce [splɪŋd]. Thus, all productive processes are treated in the same way, whether "regular" or "irregular": through analogy to existing forms. In the analogical modeling approach of Skousen (1989), the connections are less explicit than in the network model but new forms are created in a similar way: to determine the behavior of a new word, a speaker compares it with existing exemplars in its phonological neighborhood. In both types of models, any productive defaults emerge from the network of connections or clusters of similar forms; there are no default rules encoded in the system.

Connectionist and analogical models allow for what Albright and Hayes (2003) call *variegated similarity*: analogies can be formed on the basis of any kind of similarity, not just those involving segments at the locus of allomorphy (for the English past tense, the right edge or the root vowel). Thus, Bybee and Moder (1983) argue that verbs that mark their past tense by changing the stem vowel to [ʌ] are best described by a prototype: ending in a velar nasal, starting with an [s]-initial cluster, and having the vowel [ɪ]. The closer a nonce verb is to this prototype, the more likely speakers are to form the past tense with [ʌ]: Bybee and Moder (1983) found that verbs like [splɪŋ], which perfectly fit the prototype, were more likely to be placed into this class (with a past tense like [splʌŋ]) than verbs like [plɪŋ], which satisfy some but not all of the characteristics of the prototype.

Bybee (1995) argues specifically for a multidimensional network of related forms because these can capture generalizations that are both *source-oriented* and *product-oriented*. A source-oriented schema refers to properties of the base form: for example, verbs ending in [ɪŋ] form a morphological class. On the other hand, a product-oriented schema refers to the inflected forms: in this case, the class would be defined by past tenses ending in [ʌŋ]. Here, the source- and product-oriented schemas describe a similar set of verbs, but this need not be the case. For example, Bybee (1995) suggests that regular verbs affixed with *-əd* form a product-oriented pattern: the majority of English verbs end with a lax vowel followed by an alveolar stop; this product-oriented generalization, in turn, also applies to verbs that have identical present and past forms like *put*, *set*, and *spread*. Gouskova et al. (2015) show that speakers can apply both source-oriented and product-oriented generalizations to novel forms in a given inflectional pattern (in their case, Russian diminutive allomorphy; see Chapter 6).

Connectionist and network-based analogical models are successful at capturing the sorts of prototype effects that speakers show in nonce word studies like Bybee and Moder (1983). However, they sometimes overgenerate, making errors that human speakers do vanishingly rarely, if at all. For example, the analogical model tested by Albright and Hayes (2003), developed from the Generalized Context Model (Nakisa et al., 2001; Nosofsky, 1990), formed the past tense of *render* and *whisper* with *-əd* and *-t*, respectively: *[ɹɛndəɹəd] and *[wɪspəɹt] instead of [ɹɛndəɹd] and [wɪspəɹd]. These allomorphs of the regular past tense suffix would rarely if ever be applied by adult speakers to these verbs, who have learned that *-əd* attaches only to verbs ending in alveolar stops and *-t*, to voiceless obstruents.

2.2.2.2 Dual route models

Connectionist models of full storage, then, are not fully successful when faced with patterns that show true default behavior. One alternative is to limit full storage to a smaller set of forms that are unpredictable from general, language-wide processes. These are *dual route* models (e.g. Bermúdez-Otero, 2012; Clahsen, 1999; Pinker & Prince, 1988), so named because derivations can go through one of two paths: forms are either derived by rule from component parts or plucked fully formed from the stored lexicon. Prasada and Pinker (1993) argue for a dual route approach to the English past tense: verbs that take the regular past tense variants (*-d*, *-t*, and *-əd*) do not have stored past tense forms, and their past tenses are derived online by rule every time. Other verbs, such as those that express the past tense through a vowel change, have their past tense forms stored in full and retrieved as a unit. The dual route method thus creates a sharp dividing line between *regular* words (derived by general morphophonological rules) and *irregular* ones. As in the connectionist models, stored (irregular) forms are linked in an associative network that licenses analogy to similar forms; if a novel word sufficiently resembles an irregular form in this network, an irregular plural could be produced by analogy.

Dual route models make several predictions for productivity. The first is that analogical similarity effects should be found for irregulars but not regulars. That is, in the case of the English past tense, novel words should be more likely to get an irregular past tense if they resemble other past tense words, but *-d* should be equally likely for everything else. Albright and Hayes (2003) dispute this, arguing that "islands of reliability", where a past-tense formation process is more likely due to a density of phonologically similar forms that undergo it, exist for both regulars and irregulars. In particular, all English verbs ending in voiceless fricatives take the regular past suffix *-d* (like *miss* and *laugh*), and speakers were more likely to assign *-d* to nonce words ending in voiceless fricatives as well.

A similar prediction is that novel verbs that do not sound like any English verbs at all, like *ploamph*, should form regular plurals. This is because these nonce words have no words in their phonological neighborhood, let alone irregular ones that could serve as the foundation for an analogically constructed form. Indeed, Prasada and Pinker (1993) found that speakers generate and highly rate regular forms like *ploamphed* for such verbs. This prediction, too, has been challenged. Kapatsinski (2005) shows that these two properties of a default regular—resistance to analogical effects and attachment to "weird" stems—are dissociated in the case of theme vowels of Russian verbs: participants in a nonce word study preferred to use one theme vowel, *-i*, for nonce verb roots ending in coronals that did not resemble other verb roots, and this theme vowel was also preferred by stimuli similar to existing stimuli that took *-i*. On the other hand, another theme vowel *-a*, did not show phonological neighborhood effects: participants used the theme vowel *-a* for coronal-final nonce roots without regard to their similarity to existing roots. These results are problematic for a dual route model assuming a single productive default.

In some dual route models (e.g. Bermúdez-Otero, 2012; Hay, 2003; Schreuder & Baayen, 1995), the two derivational paths compete: the speaker begins to derive the desired form through a regular pathway while also searching the lexicon for a complete form with the same morphosyntactic properties. Whichever path produces a form first is chosen. Assuming that lexical access is inversely correlated with frequency (that is higher-frequency lexical items are accessed more quickly), the first path takes the time of lookup proportional to the frequency of the base form plus the time it takes to perform derivational processes on the base form, while the time of the second path is proportional to the frequency of the derived form, with no further derivational processes needed. This predicts that there can be variation (within and between lexical items) between regular and irregular forms, and that irregular forms should be more likely when the derived form is frequent relative to the base (making the lookup time of the derived form relatively short compared to the lookup time of the base form). These are known as "parallel race" models.

Bermúdez-Otero (2012) gives an example, from Kraska-Szlenk (2007), where token frequency mediates lexicalized variation between regular and irregular processes: the stress pattern of English deverbal nouns suffixed with *-ation*. The example itself is quite complicated, but the basic expected pattern is that these nominalizations inherit the stresses of the verbs from which they are derived, as in (7a): while the suffix takes the main stress, the primary stress on the verb remains secondary in the derived form. However, some nominalizations show variation, as in (7b), or consistently remove the main stress of the verb before *-ation*, as in (7c).

(7) *Expected, unexpected, and variable stress patterns in English nominalizations (Bermúdez-*

Otero, 2012, p. 37)

a. *expected stress* cond[ɛ́]mn cònd[ɛ̀]mnátion

b. *variable stress* cond[ɛ́]nse cònd[ɛ̀∼ə]nsátion

c. *unexpected stress* cons[ɜ́]rve còns[ə]rvátion

The relevant factor in this case is the relative frequency of the verb and the nominalization: *condemn* is considerably more frequent than *condemnation* (7.09 vs. 2.57 words per million in the Corpus of Contemporary American English), so the latter will be derived by regular process. On the other hand, *conservation* is much more frequent than *conserve* (9.11 vs. 1.65 words per million), so the former will be stored as a unit and consistently win the lookup race. Finally, *condense* and *condensation* have very similar frequencies (.28 vs. .22 words per million), so the two routes will each win some of the time, leading to variation.

Bermúdez-Otero (2012) argues that traditional dual route models, which offer a choice between regular derivation and whole form storage, do not predict these sorts of variable frequency effects: either a form is accessed or derived, but not both. Thus, parallel race models offer more flexibility and empirical coverage in their derivational processes. However, in a reanalysis of word frequency effects, Lignos and Gorman (2014) argue that base and derived frequency effects are not as clear-cut as reported in earlier work, and as Marantz (2023) discusses, *all* words seem to show effects of both base and derived frequency. Thus, the distinction between lookup of base form with derivation and lookup of derived form is not as categorical as assumed by Bermúdez-Otero (2012), and the effects of "derived frequency" can also be captured as the transition probability between base and affix, as shown by (Solomyak & Marantz, 2010). Thus, the frequency effects in (7) do not unambiguously point to a parallel race model; they are also likely compatible with some models without full storage like those I discuss later in this section and later adopt.

One question for dual route models with productive defaults, like Pinker and Prince (1988) and Bermúdez-Otero (2012), is how these productive rules are formed. How do children learn to form a productive rule for English plurals ending in *-z* but not for irregular plurals like *oxen*? One approach is the Tolerance Principle (Yang, 2016), which proposes a precise limit in the number of exceptions that can be tolerated for a productive rule. In particular, the Tolerance Principle argues that learners will turn a pattern whose context is satisfied by N verbs into a productive rule when there are no more than $\frac{N}{\ln N}$ exceptions—that is, words to which the rule could conceivably apply but fails to.

Returning to our example of the English past tense, Yang (2016) found 1022 past-tense verbs in child-directed speech. For the regular suffix *-d* to be learned as productive under the Tolerance Principle, there would need to be at most $\frac{1022}{\ln 1022} \approx 147$ exceptions (assuming that *-t* and *-əd* are allophones rather than separate allomorphs with distinct underlying forms). In fact, he found 127, so a universal default (context-free) rule inserting *-d* is predicted to be used productively. This seems to be the case: for example, children often have a period of overapplying the productive rule to irregular verbs, like *bringed*. What about rules for irregular verbs that apply in a narrower set of circumstances? One such potential rule is the pattern of verbs like *sing* and *ring*, where [ɪ] changes to [æ] before [ŋ] (past tense *sang*, *rang*). Of the eight verbs in the corpus that end in [ɪŋ], and thus satisfy the structural description, only three (including *spring*) follow this past tense pattern, while the other five (*bring*, *fling*, *sting*, *swing*, and *wing*) do not. This, like the other irregular past tense patterns, does not satisfy the Tolerance Principle, so children are predicted to memorize verbs following this pattern individually and not apply it productively. Indeed, children overapply irregular verb patterns extremely rarely, and only in a very limited set of circumstances.

The Tolerance Principle thus provides an account of how rules can come to be productive in dual route models. However, Schuler et al. (2021) raise two issues for classical dual route models that limit their efficacy in explaining nonce word results. The first is that the distribution of allomorphs may be balanced enough that *no* productive rule emerges, because no one general rule passes the Tolerance Principle threshold. In this case, children would memorize the inflectional behavior of every word individually, and we would not expect to see hallmarks of productivity like overapplication in acquisition. Dąbrowska (2001) looks at acquisition of the Polish genitive: as shown in Table 2.4 in Section 2.1.5.2 above, most masculine nouns take either *-a* or *-u* in the genitive. Both are quite frequent, with *-a* appearing in 70–80% of both types and tokens. However, children do not treat *-a* as a productive default—in particular, overapplication rates of *-a* are much lower than they are for English children learning the past tense, and also for *the same* children learning the genitive plural in Polish, which does show a default suffix, *-uv*. The Tolerance Principle predicts the possibility of a system without a default, as seems to be the case for the Polish genitive singular. Dual route models predicated on the presence of a default in every scenario do not predict the existence of cases like the Polish genitive singular.[5]

The second issue acknowledged by Schuler et al. (2021) is that, while children seem to obey the Tolerance Principle, adults do not; instead, in nonce word tests (e.g. Becker et al., 2011; Hayes et al., 2009) and artificial language studies (e.g. Hudson Kam & Newport, 2005), adult speakers tend to extend gradient patterns at roughly the same rates as exist in the lexicon (the Law of Frequency Matching described in Section 2.1.2). Schuler et al. (2021, p. 29) propose that children and adults may behave differently because “children are in a maturational state that is optimized for forming simple, clean, productive generalizations, while adults are optimized for something else”. Whatever the reason for this discrepancy, the upshot is that dual route models make the wrong predictions for the experimental task in this dissertation: nonce word studies on adults. I discuss better-suited models of the task in Section 2.3; these fit best with grammatical models in which all forms are derived. I describe such models in the next section.

[5] Schuler et al. (2021) refer to the lack of default in the Polish genitive singular as a gap. This may mean that they predict that speakers are unable to perform a nonce word task at all—which is unlikely, given that Polish is perfectly capable of assimilating new words without paradigm gaps in the genitive singular.

2.2.2.3 Exceptionality without full storage: diacritics

The third alternative is that all complex forms, both "regular" and "irregular", are derived by the grammar. Of course, since two words may differ unpredictably in the morphological patterns they follow, they must differ somehow in the derivational pathways they take. There are two main ways to create differential outcomes from similar surface forms. The first, which I discuss in this section and adopt in my model in Chapter 3, is to limit the application of inflectional patterns or phonological processes to a subset of the lexicon marked with a diacritic feature. We have already seen some uses of diacritic features in Section 2.1.5: Müller (2004) marks inflectional classes in Russian with the diacritic features [$\pm\alpha$] and [$\pm\beta$], while Halle and Marantz (2008) use diacritic features [$-$Gen], [$-$Dat], [$-$Loc], and [$-$Voc] to mark Polish nouns that take *-u* in the genitive, dative, locative, and vocative, respectively. The second approach, which I discuss in Section 2.2.2.4, is to equip certian underlying forms with abstract phonological structure that affects the derivational process but does not itself surface. As we saw in Section 2.1.5, the use of diacritic features can facilitate the storage of morphological dependencies, which can be treated as correlations between features. For this reason, I adopt a diacritic feature analysis and clarify the precise theoretical nature of such features, concluding that they are properties of phonological underlying forms of exponents rather than syntactic features.

One early generative use of diacritic features is found in the analysis of Russian phonology by Lightner (1965). In certain environments, alveolar and velar consonants alternate with stridents. However, alveolar stops can undergo two different alternations. This is shown in (8): [t] alternates with either [tʃʲ] or [ʃʲː], depending on the verb.

(8) *Russian verbs with alveolar–strident alternations in the infinitive and first person singular (Lightner, 1965, p. 90)*

a. otvetit^j otvetʃ^{j}u 'answer (inf./1sg.)'

b. vozvratit^j vozvraʃ^{j}ːu 'return (inf./1sg.)'

Lightner (1965) attributes the [t∼tʃ^j] alternation, and others like it, to a pair of rules: the first turns alveolars into palatalized velars before /j/ (which is later deleted and often highly abstract in any case), then the second turns the velars to stridents before front vowels and glides. Thus, /t k/ both end up as [tʃ^j] before underlying /j/. He does not provide an analysis for the alternation in (8b). His main point, however, is that the alternations are one manifestation of a split in the Russian vocabulary that he marks with a feature [±R]: [+R] words like [otvetit^j] 'answer' are subject to one set of rules, while [−R] words like [vozvratit^j] 'return' undergo another. This difference roughly corresponds to native inherited East Slavic words ([+R]) and borrowings from a liturgical language, Church Slavonic ([−R]), but Lightner argues that the split is not just an inherited relic of a diachronic distinction: these markings are active in the synchronic grammar and are diagnosed by a number of phenomena including the strident alternations in (8).

These features work as follows: "all morphemes are (somehow) associated either with the marker {+R} or with the marker {−R}; each segment of a morpheme is specified [+R] or [−R] by application of a general rule which associates the morpheme marker with individual segments of the morpheme" (Lightner, 1965, p. 90). That is, diacritic features are properties of *morphemes* in the lexicon, but their implementation in phonology occurs on the level of *segments*, which inherit these features from the morphemes in which they occur. Thus, the rules responsible for bringing /t/ to [tʃ^j] would be marked as applying to segments with the [+R] feature in addition to their phonological specification (in Lightner's feature set, [+obstr, −grave]; in modern terms, [+obstr, coronal]).

Although these diacritic features exist in the phonology, they are formally different from phonological distinctive features that define segments. Chomsky and Halle (1968, p. 374) call diacritic features "nonphonological" and specify that only nonphonological features are distributed to individual segments of marked morphemes. They relate features like Lightner's [±R], which divide the lexicon for use in rules that apply idiosnycratically, to inflection class features that "account for the phonetic realization of the gender, number, and case features" (Chomsky & Halle, 1968, p. 373). Thus, Russian nouns are specified not just for [±R], but also for their inflection class as shown above in Table 2.2.

Many modern phonological theories, in particular Optimality Theory (McCarthy & Prince, 1993b; Prince & Smolensky, 2004), dispense with rule-based grammars in favor of derivations in which potential surface forms are evaluated by grammars comprising ranked constraints, and the candidate with the least egregious violations of these constraints is the winner. How do diacritic features account for lexical exceptionality in generative constraint-based theories? The most common approach is *lexically indexed constraints* (Flack, 2007; Pater, 2006, 2010). In this approach, constraints come in a language-general form and a version that is indexed to a diacritic feature and thus only applies to morphemes marked with that feature.

In Section 2.2.1, I discussed the analysis of Russian fleeting vowels in Gouskova (2012). The data, shown in Table 2.6, is repeated in Table 2.7 below: some nouns like [vʲétʲir] 'wind' show vowel–zero alternations in the last syllable when suffixed, while others, like [kátʲir] 'motor boat', do not. There are also nouns with similar CC clusters throughout the paradigm, like [mʲétr] 'meter'.

word	*pattern*	nominative	genitive
'wind'	fleeting vowel	vʲétʲir	vʲétr-ə
'motor boat'	vowel throughout	kátʲir	kátʲir-ə
'meter'	no vowel throughout	mʲétr	mʲétr-ə

Table 2.7: Russian alternating and non-alternating nouns (from Gouskova, 2012, p. 82)

Gouskova (2012) assumes that alternating words have an underlying mid vowel and are marked with a diacritic feature *L*, whereas non-alternating words like [kátʲir] are unmarked: /vʲetʲer$_L$/ vs.

/katʲer/. In unmarked (non-alternating) words, unstressed /e/ reduces by raising to [i]. This is driven by a dispreference for mid vowels, *MID, that outranks the constraint against vowels changing their quality, IDENT. Thus, in (9), the faithful candidate [kátʲerə] incurs a fatal violation of *MID and loses to the candidate with a reduced vowel, [kátʲirə]. Vowel deletion is not considered as a repair for the *MID violation because the constraint against deletion, MAX-V, outranks IDENT.

(9) *Non-indexed constraints force reduction of /e/ (cf. Gouskova, 2012, p. 98)*

/katʲer/-a	*MID$_L$	IDENT$_L$	MAX-V	*MID	IDENT
a. kátʲerə				*!	
☞ b. kátʲirə					*
c. katrə			*!		

In alternating words marked with the feature *L*, the constraint ranking must be reversed: the constraint against vowel deletion, MAX-V, must be outranked by *MID, to penalize the mid vowel, and *IDENT, to remove raising as a viable option. This is achieved by having higher-ranked copies of these two constraints indexed to *L*: if a mid surface vowel corresponds with an underlying segment marked *L*, it violates both the general *MID constraint and its higher-ranked indexed copy, *MID$_L$. An example of this is shown in (10): the faithful candidate [vʲétʲerə] fatally violates the indexed markedness constraint *MID$_L$, while the version with reduction, [vʲétʲirə], violates IDENT$_L$. The best candidate, then, is [vʲétrə], which only violates the (now) lower-ranked constraint against deletion, MAX-V.

(10) *Indexed constraints force exceptional deletion of /e/ (cf. Gouskova, 2012, p. 98)*

/vʲetʲer$_L$/-a	*MID$_L$	IDENT$_L$	MAX-V	*MID	IDENT
a. vʲétʲerə	*!			*	
b. vʲétʲirə		*!			*
☞ c. vʲétrə			*		

This example shows how constraints indexed to diacritic features can handle phonologically based lexical exceptionality in constraint-based theories, just as rules referencing diacritic features can in rule-based theories. Though the derivations are very different, the lexical entries are the same: exceptional lexical items are marked with diacritic features.

The Russian diacritics of Gouskova (2012) and Lightner (1965) both involve lexical exceptionality in the morphophonological domain: lexically triggered phonological alternations handled in the phonological module of the grammar. Accordingly, these features must, by definition, be accessible to phonology. What about inflection class features like the Russian $[\pm\alpha]$ and $[\pm\beta]$ from Müller (2004), discussed in Section 2.1.5.1? These must be visible at the point where phonological material is inserted—which, in many generative theories, strictly precedes phonology. On the other hand, inflection classes "generally have no syntactic function" according to Chomsky and Halle (1968, p. 373); it is descriptively useful to define inflection classes by their lack of syntactic effect (as opposed to agreement classes like gender, which must be syntactic).[6] In the remainder of this section, I explore the question of where, exactly, diacritic features indexing allomorphy are stored. I conclude that, in theories like Distrtibuted Morphology that store phonological and syntactic lexical information separately, these diacritics are phonological, not syntactic. This argument is fairly long, so those who wish to skip it can proceed to Section 2.2.2.4. However, I linger on it for two reasons. First, this conclusion is the opposite of that frequently assumed in the literature—though Gouskova and Bobaljik (2022) reach the same conclusion that I do. Second, my model of morphological dependencies in Chapter 3 assumes that morphological diacritic features are present in phonological underlying forms. The following argument serves to justify this assumption.

In some generative approaches, the question of storage location is moot, since they make no modular division of the lexicon. For Lieber (1980, p. 66), lexical entries comprise phonological and semantic representations, lexical category and subcategorization frames (for building up word structure), insertion frames (for placing words into sentences), and diacritics like inflection

[6] See Arsenijević (2021) for a dissenting view.

class markers and [±Latinate], a diacritic for English that serves a similar purpose to [±R] in Russian of marking segments of the vocabulary as being subject to certain phonological processes. Complex words are built up through binary branching word syntax, and features of all types percolate up from their heads. Thus, the English adjectivizing suffix *-able* is marked [+Latinate], so words formed from this suffix, like *breakable*, inherit both the adjective category and their [+Latinate] marking from the suffix. Similarly, in languages with more extensive grammatical gender like German, suffixes can be marked for a gender that percolates up to derived forms. Thus, German diminutives formed with *-çən* are always neuter, overwriting any gender features of the base (Lieber, 1980, p. 85). For Lieber, then, diacritic features and syntactic properties are active (through percolation) with a word syntax that constructs complex words out of morphemes, and lexical entries include phonological, syntactic, and semantic information alongside one another.

These omnibus lexical entries are explicitly rejected by like Distributed Morphology (Halle & Marantz, 1993; Harley, 2014; Harley & Noyer, 1999), in which phonological material is inserted for abstract syntactic objects after (at least some) syntactic operations have already taken place (known as "late insertion" models). In Distributed Morphology, lexical information is stored in several places. First, the Lexicon contains syntactic terminals that are marked for syntactic features. Thus, the German diminutive would be, roughly speaking, a terminal that selects nouns and has neuter gender. These syntactic terminals are then spelled out as phonological material through vocabulary items, pairings between terminals (or feature sets) and underlying forms.

In this dissertation, I assume Distributed Morphology and its split view of lexical storage. Thus, we must resolve the issue of whether diacritic features are phonological (in underlying forms) or syntactic (in feature bundles of syntactic terminals). Features like [±R] and [±Latinate], discussed above, are unambiguously phonological, as they are required to delimit the scope of phonological processes and evaluations. Thus, they should be inserted as part of underlying forms, not as part of syntactic feature bundles. The location of inflection class features, at least those that handle selection of affix allomorphs, is less clear. Consider the difference between Russian feminine nouns in class II and III, discussed in Section 2.1.5. Class II nouns like [tʲotʲa] have *-oj* in the instrumental (which reduces to [əj] when unstressed, [tʲotʲəj]), while class III nouns like [pletʲ] have *-ju* ([pletʲju]). In Distributed Morphology, the difference is located in vocabulary items: the two instrumental suffixes are exponents of the same syntactic terminal in different contexts, as shown in (4) and repeated below.[7] Recall that, in the analysis of Müller (2004), class II and III share a $[-\alpha]$ feature and differ in their $[\pm\beta]$ feature: class II roots like [tʲotʲ] are marked $[+\beta]$ while class III roots like [pletʲ] are given $[-\beta]$. Similarly, instrumental case decomposes into the features [−subj, −gov, +obl]. Thus, the rules in (4) spell out the instrumental case suffix for class II and III nouns, respectively.

(4) *Rules of realization for the Russian instrumental in class II and III (cf. Müller, 2004, p. 204)*

a. [−subj, −gov, +obl] ↔ oj / $[-\alpha, +\beta]$ ___

b. [−subj, −gov, +obl] ↔ ju / $[-\alpha, -\beta]$ ___

The rules in (4) are not syntactic, but they make reference to syntactic features that appear in their context; likewise, they can reference previously spelled out phonological material in their context as well—for example, a vocabulary insertion rule could apply when the adjacent phonological context ends in a consonant. Should inflection class diacritics be grouped with the former or the latter? Much previous work in Distributed Morphology (e.g. Embick & Halle, 2005; Privizentseva, 2022) has assumed that inflection class features are syntactic and can undergo syntactic operations.[8] I

[7]An alternative that is largely notationally equivalent to assigning diacritic features to lexical items is placing lexical items in lists that appear in the contexts of rules like those in (4). Thus, (4a) would contain in its context a list of all class II nouns, and so on. This is the approach taken by Embick and Marantz (2008).

[8]Müller (2004) proposes that inflection class features are uninterpretable in syntax, in analogy with features that drive syntactic movement in Minimalist theories of syntax. In his account, inflection causes these features to be valued and thus made inert. This explains why these features (by definition) are never active in the syntax—he argues, contrary to the assumptions of Distributed Morphology, that morphology precedes syntax, and inflection features have already been removed from the derivation (through inflection) at the point that syntactic operations like agreement take place. Müller (2004, p. 222) suggests that this may be an argument against late insertion: "At the point where a late insertion approach needs an inflection class feature, the feature has long been deleted." However, this argument

believe that this approach is inferior for three reasons, which I will now lay out. I reach the same conclusion as Gouskova and Bobaljik (2022), who also argue that diacritic features in Distributed Morphology are properties of exponents, not syntactic terminals.

The first reason to keep allomorph selection features out of the syntax is that, as previously defined, they are never active in the syntax, while they do have an effect on phonological representations. This alone is a conceptual reason to place them in the phonology. However, there is an additional theoretical argument: inflection class features are language-specific, so if syntactic features are drawn from a universal set, as argued by Chomsky (2004), Cinque (2013), Caha (2021), and others, these cannot include inflection class features.

A second argument against the location of inflection features in the syntax involves the assumption of cyclic spellout first proposed within Distributed Morphology by Bobaljik (2000). Under this assumption, spellout begins with the root and proceeds outward from there. Thus, allomorphy of an inner morpheme should not be conditioned on the *phonological* properties of a more peripheral morpheme, because the outer morpheme has not yet been spelled out at the point of spellout of the inner morpheme. A further (though not necessary) assumption made by Bobaljik (2000) is that spellout *replaces* syntactic features: once a syntactic terminal is spelled out, its syntactic properties are no longer accessible in the context of spellout of future morphemes. This predicts that allomorphy of an outer morpheme should not be conditioned on the *syntactic* properties of morphemes that are closer to the root. (This prediction has been argued to be false by Carstairs-McCarthy (2001) and Gribanova and Harizanov (2017), among others.) Inflection class features are properties of inner morphemes that condition allomorphy on affixes further out. Under a strict assumption of root-outward spellout, inner morphemes have already been spelled out when allomorphs of outer affixes are selected, so they have to be present in the *phonological* context of spellout for these outer morphemes. If inflectional features were syntactic, they would have to be replaced in spell-

also simply *assumes* that inflection features must be syntactic features like gender; he does not consider the possibility that they are phonological, as I argue.

out of roots and other inner morphemes before they have a chance to condition spellout of outer morphemes.

The third argument for placing inflection features in the phonology comes from root suppletion (see Harley, 2014), where a root can take two contextually dependent suppletive allomorphs. If affix selection features are placed in the syntax, then they should be properties of *syntactic roots*, and suppletive allomorphs should always inflect the same way because they spell out the same root. However, if inflectional features are placed in the phonology, they are properties of exponents, not syntactic roots. In this case, suppletive allomorphs are predicted to be able to differ in their inflection without a corresponding change in meaning. I now present two examples from Czech showing that suppletive allomorphs can differ in their inflectional paradigms.

The first case involves root suppletion for verbal aspect. Aspect is an organizing component of the Czech verbal system: verbs can be perfective or imperfective, and come in aspectual pairs. These pairs can be related through a number of derivational pathways, examples of which are shown in Table 2.8. The verbs in this table are presented in the masculine singular past tense, which allows for a less complicated presentation of the pairs than the usual citation form, the infinitive. Each verb includes the past-tense marker *-l* and a theme vowel, which is *-a* for all three of these verbs. The most common verb pattern is exemplified by [psal] 'write': the imperfective form has the basic verb stem, and a perfective is formed using one of a number of lexical prefixes, in this case *na-*. Other prefix–verb combinations instead do double duty, changing both the aspect (to perfective) and the meaning. For example, the prefix *pod(ɛ)-* attaches to the verb [psal] to yield [podɛpsal], which is the perfective of the verb meaning 'sign'. For a small number of verbs, the perfective form is the basic root and the imperfective is derived from it. This is the case with [dal] 'give', which is perfective; to form the imperfective, one inserts the imperfective suffix *-aːv* after the root (but before the theme vowel). Finally, there are pairs where the perfective and imperfective are built off different roots. One such verb, 'take', will be our focus: when not attached to any lexical prefixes like *pod(ɛ)-*, the root is [br] in the imperfective and [vz] in the perfective.

meaning	*imperfective*	*perfective*	*relation*
'write'	ps-a-l	na-ps-a-l	imperfective basic, perfective prefixed
'give'	d-aːv-a-l	d-a-l	perfective basic, imperfective suffixed
'take'	br-a-l	vz-a-l	root suppletion

Table 2.8: Patterns of Czech aspectual pairs

If [br] and [vz] are truly suppletive allomorphs of the same syntactic root, the choice of allomorph should not have any effect on the semantics. That is, we would expect contextually dependent idiosyncratic meanings available for one form to also be available for the other, so long as the meaning is semantically compatible with both aspects. This is indeed the case: for example, with the reflexive clitic [sɛ], this verb takes on the meaning of 'get married' in both aspects: [bral sɛ] (imperfective) and [vzal sɛ] (perfective).

With this background on aspectual root suppletion, I now show that suppletive allomorphs can vary in their inflectional patterns. In particular, we will look at the Czech passive (presented here, again for reasons of practicality, with the masculine singular adjectival suffix *-iː*). The imperfective root of 'take', [br], takes the usual passive suffix *-n* (which is followed by the adjectival suffix); the perfective root, [vz], instead has the passive *-t*, which appears with a much smaller number of verbs. This choice is not due to the aspectual difference alone: when [br] is given a lexical prefix like [vɪ] (which changes both its meaning and its aspect), the resulting perfective form takes passive *-n*; similarly, there are underived imperfective verbs like [mlɛl] 'grind' (comprised of a root [ml] and a theme vowel *-ɛ*) that combine with passive *-t*.

meaning	*aspect*	*past*	*passive*
‘take’	imperfective	br-a-l	br-a-n-iː
‘take’	perfective	vz-a-l	vz-a-t-iː
‘choose’	perfective	vɪ-br-a-l	vɪ-br-a-n-iː
‘grind’	imperfective	ml-ɛ-l	ml-ɛ-t-iː

Table 2.9: Allomorphy in the Czech passive

The passive allomorphy in Table 2.9 shows that the choice of passive suffix is a selectional property of the verb root to which it attaches (or, at least, to the combination of verb root and theme vowel, which is also lexically specific), marked in this theory by a diacritic feature. Since suppletive allomorphs can select different passives, inflection features must be associated with individual suppletive allomorphs, not abstract syntactic roots.

A similar pattern can be found in a suppletive noun, which takes inflectional suffixes directly without the complication of theme vowels present on verbs. The noun [rok] ‘year’ rarely appears in oblique cases in the plural; instead, it is replaced with forms borrowed from [lɛːto] ‘summer’. In particular, the very common genitive plural is almost always [lɛt]. The noun [rok] is masculine, while [lɛːto] is neuter, and the suppletive genitive plural shows the null case ending typical of neuter and feminine nouns and a very small number of masculine nouns. By contrast, most masculine nouns have *-uː* in the genitive plural, so the expected non-suppletive form would be [rokuː]. (Interestingly, the cases that show suppletion are precisely those that show no gender distinction in agreement in the plural in Standard Czech.)

How do we know that [rok] and [lɛt] are truly suppletive allomorphs? One criterion is *complementarity*: the two root allomorphs divide up the space of derived adjectives. Adjectives built off of the base noun typically use the [rok] allomorph with a [k∼tʃ] consonant alternation ([kaʒdorotʃɲiː] ‘annual’, cf. [kaʒdiː rok] ‘every year), while those built off of genitive numeral phrases use [lɛt] ([tr̝̊iːlɛtiː] ‘three-year-old’, cf. [tr̝̊iː lɛt] ‘three.GEN years.GEN).[9] The suppletive forms can also be used in constructions with idiosyncratic meaning. One example is the construction [uɦɛrskiː rok], literally ‘Hungarian year’, which indicates a long time period, usually in the idiom ‘once in a Hungarian year’, meaning ‘once in a blue moon’. For semantic reasons, this construction very rarely occurs in the genitive plural, but when it does, we can find it with the suppletive form [lɛt], as in these examples from the Czech National Corpus (Křen et al., 2022): ‘if something substantial happened once in ten Hungarian [lɛt]’ (in Czech, numbers five or greater take genitive plural complements), ‘the Ebens make a record at intervals of proverbial Hungarian [lɛt]’. This is to be expected if [rok] and [lɛt] are suppletive allomorphs of the same underlying root.

A more common construction is [svjɛtɛlniː rok] ‘light year’, which Czech and English both calqued from German—cf. Czech [svjɛtlo] ‘light’. In the genitive plural, this construction is attested with both the suppletive ([svjɛtɛlniːx lɛt]) and non-suppletive ([svjɛtɛlniːx rokuː]) allomorphs, though the former is about eight times more common in the Czech National Corpus (Křen et al., 2022). I provide no account here of how the non-suppletive root is sometimes available in this construction; however, it provides the crucial contrast for our purposes, which is that [rok] always takes *-uː*, while [lɛt] always takes a null ending. We never find genitive plural forms where the root allomorphs take different endings ([lɛtuː] or [rok]). Thus, for speakers who allow the non-suppletive root allomorph [rok] in some circumstances in the genitive plural, it always takes the genitive plural suffix *-uː*; when these same speakers use the suppletive allomorph in other circumstances, the form is always the unsuffixed [lɛt].[10] This shows that, when the suppletive allomorphs of the root for ‘year’ overlap slightly in their usage, each allomorph is associated with its own genitive plural suffix. This is expected if inflectional features are phonological, associated with individual exponents, but

[9] Adjectives like [tr̝̊iːrotʃɲiː] ‘three-year’ are attested but extremely rare—much rarer, in fact, than non-suppletive oblique plural forms of [rok] like genitive [rokuː]. Thus, occasional uses of such adjectives can be grouped together with other rare examples of the regular root allomorph [rok] in cases where the specalized allomorph [lɛt] is expected.

[10] To provide one example from the Czech National Corpus, Sylva Daníčková’s book *Skrytá poselství vědy* (*The Hidden Missions of Science*) contains 137 tokens of genitive plural [lɛt], of which 5 are in the construction ‘light years’, and 6 tokens of genitive plural [rokuː], of which 3 are used in ‘light years’, and no examples of genitive plural [letuː] or [rok].

is unexpected under the assumption that inflectional features are properties of syntactic objects which may have multiple exponents.

Given these arguments, I assume that inflection class features are properties of exponents active in the phonological grammar, just like more clearly phonological diacritic features like the [±R] of Russian in Lightner (1965). This makes them distinct from syntactic features like gender—although, as discussed in Section 2.1.4, inflection class is often highly correlated with such features. This is theoretically desirable for another reason as well: the line between morphological processes and morphologically conditioned phonological processes is unclear and highly contested, and a treatment of diacritic features that groups these two together means that this debate has no major consequences for the treatment of lexical exceptionality. Phonological and morpho(phono)logical exceptionality are marked using the same theoretical construct, diacritic features on vocabulary entries (phonological underlying forms).

2.2.2.4 Exceptionality without full storage: abstract structure

One criticism of diacritic features (see e.g. Bermúdez-Otero, 2013) is that they are phonologically arbitrary: grouping words together through diacritics is often explicitly a recognition that there is no phonologically natural way to define that class of words. (In Section 2.3, I argue in favor of arbitrariness in phonology.) In some cases, linguists instead aim to account for exceptionality through a set of phonological rules that apply generally, paired with abstract or defective underlying forms that are subject to independently necessary rules. I call this the *abstract structure* approach.

The use of diacritics and abstract structure are not necessarily mutually exclusive, and indeed, many works use both for different purposes. Accordingly, our early example of abstract structure comes from Lightner (1965), whose use of diacritics like [±R] has already been discussed. The phenomenon in question is the Russian vowel–zero alternation discussed in Section 2.2.1 and Section 2.2.2.3. As shown in Table 2.7 above, Russian has nouns like [v^jét^jir] ‘wind’ that show vowel–zero alternations when suffixed, as in the genitive, [vʲétr-ə]. These nouns contrast with verbs with stable vowels, like [kátʲir] 'motor boat' (genitive [kátʲir-ə]).

In the analysis of Gouskova (2012) described in Section 2.2.2.3, both words have mid vowels in their underlying form, but the alternating word is marked with a diacritic feature *L* marking it as an alternator: /katʲer/, /vʲetʲer$_L$/. For unmarked words, a relatively low-ranked constraint *MID forces the unstressed vowel to reduce to high [i], outranking a constraint *IDENT which penalizes the change in vowel quality. However, marked words trigger violations of higher-marked *MID$_L$ and *IDENT$_L$, which together force a different repair: deletion of the unstressed mid vowel. Only one segment alternates, but it is the entire noun that is marked with a diacritic feature.

A more common approach to alternating Slavic vowels instead marks *individual segments* differently: alternating /e/ somehow contrasts underlying with non-alternating /e/. Historically, these alternating vowels, known as *yers*, which I will represent as [ĭ ŭ], were originally short high vowels (front and back, respectively) that deleted in a historical process known as Havlík's Law (Gouskova, 2012; Kiparsky, 1979): going from right to left in the word, every other yer deleted, starting with the first. Thus, [vʲétʲir] was historically [vetĭrŭ], and genitive [vʲétrə] was [vetĭra]. In the nominative singular, the rightmost yer is the word-final one, which deletes, and the stem-internal one remains; in the genitive, the only yer is the one inside the stem, so it deletes.

Lightner's analysis recapitulates this diachronic process in the synchronic grammar: alternating vowels are underlyingly lax high vowels. However, they cannot surface as such: first, rule (11a) lowers alternating yers (that is, yers preceding another yer separated by a consonant) to mid vowels. Then, rule (11b) deletes all remaining yers that were not previously lowered. (Afterwards, a separate vowel reduction process, not shown here, re-raises unstressed [e] to the surface [i].)

(11) *Rules for Russian vowel–zero alternations (Lightner, 1965, p. 28)*[11]

[11] Lightner (1965) formats the rules somewhat differently than I do; in particular, he uses the feature [±diffuse] to distinguish high vowels from mid vowels where I use [±high].

a. $[-\text{tense}, +\text{high}] \rightarrow [-\text{high}]$ / ___C_1 $[+\text{tense}, +\text{high}]$

b. $[-\text{tense}, +\text{high}] \rightarrow \emptyset$

The feature [±tense] is "an abstract feature which does not appear in phonetic representations" (Lightner, 1965, p. 26); in Lightner's system, it is used not just for alternating vowels but also to distinguish /o/ ([−tense]) from /a/ ([+tense]).

The lowering analysis of yers offers at least three claims to phonological "naturalness" over a diacritic analysis: first, it recapitulates a historical process, so some Slavic speakers at some point plausibly had rules like those in (11) active in their grammar. Second, phonological exceptionality is handled through phonological means, using features that are contentful and plausibly part of a universal set. Finally, the [±tense] feature is used elsewhere in the phonological system: the opposition between alternating and non-alternating vowels is not narrowly targeted for this specific purpose, but is integrated into the Russian vowel space more generally. With that being said, there are yer-like vowel alternations in Russian that require a different analysis that does not invoke Havlík's law anyway (see e.g. Linzen et al., 2013).

The analysis of alternating vowels as vowels that cannot surface because of their features, or alternately, insufficient structure (for example, lacking a mora or timing slot) has been applied to many Slavic languages (e.g. Kenstowicz & Rubach, 1987; Yearley, 1995). However, it does leave some aspects of the phenomenon unexplained: for example, why is it that only mid vowels alternate? Gouskova (2012) motivates this with a general constraint against mid vowels, *MID. In this sense, the abstract segment analysis misses something that a marked morpheme analysis can capture.

The marked morpheme analysis of Gouskova (2012) is conceptually quite different from the lowering analysis of yers, both that of Lightner (1965) and later ones that tie alternations to defective structure. However, the line between abstract underlying forms and diacritic features is not always so clear. Kiparsky (1982) criticizes "purely diacritic use of phonological features" in works of early generative phonology like Chomsky and Halle (1968). He provides the example of antiharmonic stems in Hungarian, shown in (12): most words with front unrounded vowels [ɛ eː i iː] take suffixes with front vowels, as is the case with [keːʃ] 'knife' in (12a). However, some such words instead show back harmony, taking the back equivalents of these suffixes. One such word is [heːj] 'rind', shown in (12b).

(12) *Harmonic and antiharmonic stems in Hungarian (cf. Kiparsky, 1982, p. 128)*

   a. *harmonic stem* keːʃ-ɛm 'my knife'

   b. *anti-harmonic stem* heːj-ɒm 'my rind'

One way to mark this distinction would be to claim that there are really two [ɛ] phonemes that are identical on the surface but trigger different harmony patterns.[12] Thus, for example, [keːʃ] 'knife' would be underlyingly /keːʃ/, while [heːj] could have an underlying non-front vowel: /həːj/. This /ə/ would trigger back suffixes before fronting to neutralize with /e/ on the surface. While this analysis accounts for the antiharmonic stems, Kiparsky (1982) points out that the proposed antiharmonic vowels never appear on the surface, and serve no purpose other than to diacritically mark certain words as behaving exceptionally. In this, the analysis is functionally equivalent to one marking words like [heːj] with a diacritic feature: /heːj$_{[+\text{antiharmonic}]}$/.

Of course, many abstract phonological analyses are not simply making diacritic use of otherwise unused features. Modern analyses can get very complex, breaking down segments into smaller parts and combining them in particular ways to capture generalizations that would otherwise have to be overtly stipulated. One recent example is the analysis of German plurals in Trommer (2021). The system is quite complicated: there are several suffixes (including -ə, -ɐ, -n, and no suffix), some of which can cooccur with umlaut (fronting of a stem vowel). Examples of these combinations are found in Table 2.10.

[12] More recent work has shown that vowels in harmonic and antiharmonic stems actually do show minute phonetic differences (Benus & Gafos, 2007; Szeredi, 2016).

suffix	*no umlaut*		*umlaut*	
Ø	tsoːbl̩ ∼ tsoːbl̩	‘sable(s)’	foːgl̩ ∼ føːgl̩	‘bird(s)’
-ə	tʀoːn ∼ tʀoːnə	‘throne(s)’	zoːn ∼ zøːnə	‘son(s)’
-*n*	zaːt ∼ zaːtn̩	‘seed(s)’		
-ɐ			ʀɔs ∼ ʀœsɐ	‘horse(s)’

Table 2.10: Plural realizations in German (from Trommer, 2021)

Although a noun’s plural is generally not predictable from its phonological form and gender alone, there are some strong generalizations:

(13) *Central generalizations on plural umlaut and suffix allomorphy (Trommer, 2021, p. 614)*

Generalization I: Umlaut and plural -*n* are in complementary distribution

Generalization II: Feminine nouns strongly prefer plural -*n* in contexts where non-feminine nouns do not

Generalization III: Noun roots ending in [ə] always take plural -*n* (and consequently never show umlaut in the plural)

Generalization IV: Nouns with plural -ɐ *always* umlaut in the plural

The goal of Trommer (2021) is to derive these generalizations from three plural suffixes with fixed underlying forms: one attached to all plurals, and additional plural suffixes for feminine and neuter nouns. The feminine suffix is a floating nasal feature (represented as ⓝ)—this drives Generalization II, since the nasal feature often surfaces as [n]. The general plural suffix, on the other hand, comprises an underspecified root note (that is, a segment timing slot, represented as •) and a coronal vocalic place feature (represented as ⓒ), which characterizes front vowels like [e/ɛ] and [ø/œ].

Explaining all of the details of the analysis of Trommer (2021) would be too complicated, so to give a sense of the argumentation, I sketch out the explanation of Generalization I using the feminine nouns [naːt] 'seam' and [zaːt] 'seed', whose plurals are [nɛːtə] (umlaut, no *-n*) and [zaːtn̩] (*-n*, no umlaut). Umlauting nouns differ from non-umlauting vowels in that the former are underspecified for place (roughly, backness). Thus, the [a] in [zaːt] is underlyingly a fully specified /a/, with a characteristic pharyngeal place feature and height features [−high, +low], while the [a] in [naːt] has height features but no place feature (represented as /A/). Thus, the underlying forms of the two words' plurals are as shown in (14):

(14) *Underlying representations and surface forms of German feminine plurals (Trommer, 2021, p. 606)*

a. seed-PL-F.PL /zaːt-ⓒ•-ⓝ/ [zaːtn̩$_{c,n}$]

b. seam-PL-F.PL /nAːt-ⓒ•-ⓝ/ [nɛː$_{c}$təⓝ]

If the root vowel has no place feature, as in /nAːt/, the floating coronal feature from the plural attaches to the root vowel (indicated by the subscript c in (14a)) and fronts it, causing umlaut. In this case, the nasal cannot attach to the leftover timing slot (due to a constraint penalizing coronal consonants) and is left floating—that is, unpronounced. Thus, the timing slot from the plural surfaces as the default vowel, [ə]. By contrast, in the plural of [zaːt], the root vowel already has a place feature, so the floating coronal feature cannot attach to it. Floating coronal features are penalized even more heavily than coronal consonants, so the floating coronal and nasal features together dock onto the timing slot (indicated by the subscript c and n in (14b)) to make syllabic [n̩].[13]

Trommer (2021) succeeds in developing an analysis that derives the generalizations in (13), as well as other aspects of German plural inflection. However, this analysis requires a large amount of theoretical machinery: a complex feature geometry, abstract and defective underlying representations, narrowly tailored constraints, an additional layer of structure between underlying form and surface realization wherein phonological material may surface phonologically but not phonetically, and

[13]It is not clear from the relevant tableau (Trommer, 2021, p. 617) why a candidate like [zaːtɛ$_{c}$ⓝ], in which the coronal docks to the timing slot but the nasal does not, is suboptimal.

the accompanying duplication of markedness constraints such that they can refer to either phonological structure (realized or unrealized) or phonetic structure (always realized). The reliance on abstract structure means that the generalizations are not phonologically optimizing in the sense of reducing some universal measure of markedness. Such a system would also presumably be very difficult for children to learn (Gouskova, 2012; Pater, 2006), and it is not clear why they should arrive at this particular combination of abstract structures rather than some other.

By contrast, in the approach developed in this dissertation, various plural realizations would be marked on noun roots through diacritic features, and the generalizations in (13) are learned from the surface forms as phonologically arbitrary but transparent correlations. The analysis in Trommer (2021) and the sublexicon model described in Chapter 3 represent two extremes on the spectrum of encoding phonological and morphological generalizations: in Trommer (2021), all generalizations are *hard-coded* into the derivational grammar and *grounded* in phonological markedness as the output of properly ranked markedness constraints. In the sublexicon model, generalizations are *not hard-coded* into the grammar that derives forms online, but rather stored in a separate module dedicated to analogical pattern matching; these patterns are not necessarily *grounded* in principles of markedness or naturalness. In Section 2.3, I focus on these two criteria—hard-coding and grouding—in an overview of approaches to encoding phonological and morphological generalizations.

2.3 Hard-coding and grounding

In Chapter 3, I present a model that encodes lexically specific patterns with diacritic features, and captures generalizations over the distribution of these features with a constraint-based grammar characterizing each feature. In this model, generalizations are neither hard-coded into the grammar nor necessarily phonologically grounded. In this section, I discuss theories that assume one or both of these criteria and show that they are too restrictive. In Section 2.3.1, I discuss the criteria of hard-coding and arbitrariness with respect to phonological generalizations. In Section 2.3.2, I extend the use of these criteria to morphological dependencies. This approach represents an alternative to standard hard-coded inflection classes, in which inflection classes are instead emergent.

2.3.1 Phonological generalizations

Previous work on gradient generalizations over morphological patterning has focused on *phonolgical* generalizations, so I begin my discussion of theoretical accounts of these generalizations here. I start with cases where the allomorph that a given stem takes is determined *entirely* by its phonology. While these are not necessarily phonologically optimizing, there is a strong theoretical tendency to have such generalizations be hard-coded. That is, the derivational grammar chooses an affix allomorph by reading directly off the phonology of its host, rather than storing categorical generalizations over a diacritic feature. I then move on to variable generalizations that include some degree of lexical conditioning. While notions of grounding and hard-coding must be relaxed slightly, some theories still use them. After discussing one such theory, from Becker (2009), at length, I argue that even these relaxed conceptions are too restrictive, leading to my choice of a model in which generalizations are neither hard-coded nor necessarily phonologically grounded.

2.3.1.1 Categorical allomorph selection

The strongest form of lexical tendencies in allomorphy are fully categorical ones where the choice of allomorph selected by a stem is dependent entirely on its phonology. This is known as *phonologically conditioned suppletive allomorphy* (PCSA). Paster (2015) distinguishes two approaches to PCSA. In the first approach, P$\gg$M (e.g. Kager, 1996; McCarthy & Prince, 1993a, 1993b), allomorph selection is placed into a constraint-based Optimality Theoretic phonological grammar alongside purely phonological considerations. In this model, phonological constraints (P) favoring one allomorph can outrank ($\gg$) morphological constraints (M) expressing a default preference for another allomorph. Paster (2015) presents the example of the genitive in Djabugay, which is *-n* after vowels and *-ŋun* after consonants (Patz, 1991, p. 269):

(15) *Phonologically conditioned genitive allomorphy in Djabugay (Patz, 1991, p. 269)*

	V-final stems			*C-final stems*	
a.	guludu-n	'dove (gen.)'	d.	girgir-ŋun	'bush canary (gen.)'
b.	guraː-n	'dog (gen.)'	e.	gaɲal-ŋun	'goanna (gen.)'
c.	ɟama-n	'snake (gen.)'	f.	bibuj-ŋun	'child (gen.)'

Kager (1996) points out that this distribution is phonologically optimizing: if *-n* is attached to consonant-final stems like [gaɲal], the result ends in a cluster (*[gaɲaln]), and Djabugay categorically bans complex syllable codas. Accordingly, in his analysis, a phonological constraint against complex codas, *COMPLEX, outranks a morphological constraint preferring *-n* as the default selection, GENITIVE=*-n*. He suggests that this morphological constraint is a special case of a universal constraint preferring the minimal (phonologically shortest) exponent for a given morphological category.

(16) *Phonology outranks morphology in Djabugay (cf. Kager, 1996, p. 2)*

/gaɲal-{n, ŋun}/	*COMPLEX	GENITIVE=*-n*
☞ a. ga.ɲal.ŋun		*
b. ga.ɲaln	*!	

In Optimal Interleaving (Wolf, 2008), a P≫M model, the morphological constraints are more morphosyntactically grounded: candidate allomorphs are associated with different morphosyntactic features, and morphological constraints enforce feature matching between morphosyntactic structures and their exponents. For example, Spanish has two singular definite articles: [el], which is usually masculine, and [la], which is feminine. A small class of feminine nouns beginning in stressed [á] also takes [el], e.g. [el árma] 'the weapon', not the expected *[la árma]. Wolf (2008) connects this pattern to the way that vowel–vowel sequences are handled in Spanish: in general, two identical vowels coalesce into one, but this is not allowed when the second vowel is stressed.

Thus, /la arma/ cannot surface as *[lárma], and the faithful *[la árma] has a marked hiatus. The solution is to avoid hiatus by inserting the consonant-final definite article, [el], at the cost of a gender mismatch.

Wolf (2008) formalizes this analysis as follows: hiatus avoidance is driven by a high-ranking constraint *HIATUS. This is dominated by UNIFORMITY/$\acute{\sigma}$, which penalizes segments of stressed syllables that coalesce with other segments. This constraint prevents the usual hiatus resolution of the two segments coalescing into one. Both of these constraints outrank morphological feature matching constraints. Following Becker (2009), I call this feature ϕ-MATCH: the ϕ-features associated with the suffix (person, number, and gender) must match the morphosyntactic features active in the derivation.[14] This constraint is violated when the feminine definite article is spelled out as [el], but [el] is still the best option because it avoids hiatus, which is penalized by a higher-ranked constraint. The tableau in (17) places this analysis into a standard Optimality Theory framework, ignoring many of the architectural features of Optimal Interleaving. Here the numerical subscripts index segments for correspondence purposes, while morphosyntactic features are also subscripted: the definite article complex is associated with a [fem] feature, and each candidate allomorph has its own gender feature as well.

(17) *Phonology outranks morphosyntactic feature matching in Spanish (cf. Wolf, 2008, p. 119)*

/$\{$el$_{[masc]}$, la$_{1\,[fem]}\}_{[fem]}$ a$_2$rma$_{[fem]}$/	UNIFORMITY/$\acute{\sigma}$	*HIATUS	ϕ-MATCH
☞ a. elá$_2$rma			*
b. la$_1$á$_2$rma		*!	
c. lá$_{1,2}$rma	*!		

In the P≫M approach, categorical phonological generalizations are hard-coded into the grammar: the choice of *-n* and *-ŋun* in Djabugay, for example, is derived for every word from a language-

[14] Wolf (2008) himself formalizes morphological spell-out in a different way from what I am assuming here; for him, feature matching is governed by faithfulness constraints like DEP-M-masc and MAX-M-fem, which penalize the addition of a misculine feature and the deletion of a feminine feature, respectively. Both of these are violated by the winning candidate in (17).

wide constraint ranking. They is also grounded: the choice of allomorph follows language-wide markedness conditions, and is thus phonologically (and morphologically) optimizing.[15]

The other model described by Paster (2015) is subcategorization: different allomorphs have the same morphosyntax but different subcategorization frames. In the theory of Distributed Morphology assumed in this dissertation, rules of realization (known as vocabulary items) spell out morphosyntactic features as phonological material. For example, the Djabugay genitive described in (15) would have the following rules of realization:

(18) *Rules of realization for the Djabugay genitive*

a. GEN ↔ n / V___

b. GEN ↔ ŋun / C___

The vocabulary items in (18) have the same empirical coverage as the constraints (16). However, they differ in their explanatory power: the P≫M analysis derives the allomorph selection as a natural consequence of the language's phonotactics (specifically, its ban on complex clusters). In the subcategorization approach, the alignment between allomorph choice and phonotactics is incidental: nothing rules out an opposite-Djabugay in which vowel-final nouns select for *-ŋun* and consonant-final nouns, for *-n* (which would then require a later phonological repair like vowel epenthesis to avoid a complex coda). This account is also redundant: the selectional properties of the Djabugay genitive allomorphs recapitulate its general phonology, so the same distribution is encoded in the grammar twice (in phonological constraints and in morphological vocabulary items). The subcategorization approach does not imply that the distribution of allomorphs is a *coincidence*—it is perfectly compatible with diachronic or functional explanations for why things are the way they are. However, diachronic and functional pressures are not encoded in the syn-

[15] Wolf (2008) acknowledges the existence of patterns that appear arbitrary on the surface, such as ergative allomorphy in Dyirbal (see his Section 2.3). However, his theory requires him to ground such arbitrary preferences in morphosyntactic feature matching (even if such features are not overtly distinguished on the surface) and plausibly universal phonological constraints—thus, within his proposed analyses, allomorph choice is optimizing, even if it does not obviously appear to be so on the surface.

chronic grammar. Thus, the subcategorization approach is phonologically *ungrounded*. However, in this approach, generalizations are still *hard-coded* into the generative grammar: in the course of a derivation, the morphophonology directly reads off a noun's final segment to determine its genitive.

The subcategorization model is less restrictive in that it does not require phonologically conditioned suppletive allomorphy to be phonologically optimizing. This seems to be necessary: indeed, in some cases, allomorphs are selected in a way that creates a *more marked* structure—Paster (2015) calls these "perverse" patterns. One example is the suffixal determiner in Hatian Creole, which has two allomorphs, *-a* and *-la*. Here, *-a* attaches to words ending in vowels—creating marked hiatus, which is then often repaired with glide insertion—while *-la* goes with consonant-final words, creating a (lightly) marked cluster (Bonet et al., 2007; Hall Jr., 1953):

(19) *Anti-optimizing phonologically conditioned allomorphy of the Haitian Creole determiner (Bonet et al., 2007, p. 908)*

	C-final stems			*lax V-final stems*	
a.	liv-la	'the book'	e.	papa-a	'the father'
b.	ʃat-la	'the cat'		*tense V-final stems w/glide insertion*	
c.	malad-la	'the sick one'	f.	lapli-ja	'the rain'
d.	bagaj-la	'the thing'	g.	tu-wa	'the hole'

Because the subcategorization approach allows for arbitrary selectional properties, this pattern is no harder to capture than Djabugay, using the rules of realization in (20).

(20) *Rules of realization for the Haitian Creole determiner*

a. DEF $\leftrightarrow$ a / V___

b. DEF $\leftrightarrow$ la / C___

Such a pattern is problematic for P≫M, because there is no language-wide markedness constraint that *prefers* to create hiatus (a universally marked structure) or insert a glide where one could avoid doing so. Indeed, surveys of phonologically conditioned suppletive allomorphy like Paster (2015) and Kalin (2022) (who focuses specifically on infixes) conclude that it is, in general, not necessarily optimizing, and that approaches like the subcategorization model, which place allomorph selection strictly before phonology, make the correct predictions crosslinguistically.

2.3.1.2 Variable allomorph selection

Some of the cases described in this dissertation include categorical phonological conditions on suppletive allomorphy: for example, the Hungarian possessive has two allomorphs, *-ɒ* and *-jɒ*; nouns ending in palatals and sibilants always choose *-ɒ*, while nouns ending in vowels always take *-ɒ* (see Chapter 4). Most of the generalizations discussed in this dissertation, however, are *gradient*—for example, in Hungarian, words that end in non-sibilant alveolar consonants prefer *-jɒ*, but many take *-ɒ* as well: for example, the possessive of [kaːr] 'damage' is [kaːr-ɒ]. In these cases, phonological subcategorization is not enough: for at least some words, the choice of possessive must be memorized and stored individually—I assume that this is done using a lexical diacritic feature. Thus, [kaːr] would be marked with a feature [+j].

Becker (2009) discusses a similar case in Hebrew. Like the Russian analysis of Gouskova (2012), which I described in Section 2.2.2.3, he uses copies of constraints indexed to nouns according to their behavior. Becker's innovation is that nouns can be indexed to *different* constraints that offer multiple pathways for a word to behave in a non-default way. These pathways are, in turn, dependent on phonological properties of the root. Accordingly, if some copies of constraints are invoked more often than others, speakers will extend this preference to new nouns that have the phonological properties associated with those constraints. This is a looser conception of phonological grounding than that discussed in the previous section: grammatically active gradient patterns of allomorphy do not necessarily need to be *optimizing*, but they must be expressible in terms of phonological properties of a stem connected to a phonologically motivated constraint. While this notion of phonolgical grounding allows for many phonological generalizations to be learned, I will show that it is still too restrictive: speakers learn gradient generalizations over lexically variable patterns that Becker (2009) argues they should not. In Chapter 3, I opt for a less restrictive model in which speakers are able to learn any generalization over lexically specific allomorphic patterns, not just those that are grounded in one of a particular set of constraints.

Let us now go through the Hebrew case and its analysis in Becker (2009) in more detail. Nouns can take one of two plural suffixes, *-im* and *-ot*. Most masculine nouns take *-im* (which is also the plural suffix used for masculine adjectival agreement), while most feminine nouns take *-ot* (also used for feminine agreeing plural adjectives). However, there are exceptions in both directions: some masculine nouns pluralize with *-ot*, and some feminine nouns, with *-im*. Examples are shown below. In (21), the stem [jelad] 'child' takes the plural suffix associated with its natural gender, as do agreeing adjectives and verbs:

(21) *Hebrew nouns with "matching gender" plurals*

a. jelad-ím ktan-ím ʃar-ím
boy-PL little-M.PL sing.PRES-M.PL
'Little boys are singing.'

b. jelad-ót ktan-ót ʃar-ót
girl-PL little-F.PL sing.PRES-F.PL
'Little boys are singing.' (Becker, 2009, p. 76)

Nouns with the unexpected plural suffix are shown in (22). In these cases, the noun's gender is indicated by the plural suffixes on the agreeing adjective and noun: [xalon] 'window' is masculine despite taking plural *-ot*, while [nemal] 'ant' is feminine despite taking plural *-im*.

(22) *Hebrew nouns with "mismatched gender" plurals*

a. xalon-ót gdol-ím niftax-ím
window-PL big-M.PL open.PRES-M.PL
'Big windows are opening.'

b. nemal-ím ktan-ót nixnas-ót
ant-PL little-F.PL enter.PRES-F.PL
'Little ants are coming in.' (Becker, 2009, p. 77)

The group of masculine nouns taking plural *-ot* is lexically determined, but follows a strong phonological tendency: 63% of such nouns have the vowel [o] in their last syllable. This tendency in the lexicon is also active in the grammar of Hebrew speakers: in nonce word studies, speakers assign *-ot* more often to masculine nouns with [o] in the stem than those with other vowels.

The analysis of Becker (2009) is a mix of the P≫M and subcategorization approaches—some amount of subcategorization is necessary, given that allomorphs in lexically variable patterns must be associated with individual lexical items. The generalizations pattern with P≫M along the two criteria at issue: patterns are hard-coded into the grammar and are not phonologically arbitrary. However, the relation to phonology is somewhat different than in the P≫M cases discussed in Section 2.3.1.1: rather than necessarily being phonologically *optimizing*, the model presented in Becker (2009) has a weaker requirement of phonological *grounding*. I explain the difference in the following discussion of his analysis.

The active morphological constraint in Becker's analysis of Hebrew is ϕ-MATCH, also used in the analysis of Spanish in (17): *-im* and *-ot* have masculine and feminine features, respectively, and a violation is assessed if the gender feature in the suffix does not match the underlying gender feature on the noun. For masculine nouns like [xalon] that take *-ot*, this morphological constraint must be dominated by a phonological constraint preferring *-ot*. Becker (2009) notes that it is specifically *native* nouns with [o] in the stems that prefer *-ot*. One salient difference between native words and loans is that the former, but not the latter, shift stress onto the plural suffix—thus, the plural of the loan words [blog] is [blóg-im], not *[blog-ím] or *[blog-ót]. Accordingly, the relevant phonological constraint should specifically disprefer *-im* after unstressed [o]. This constraint is LOCAL(o), which requires that [o] must be licensed, either by bearing stress or by being autosegmentally associated with a stressed [o] in an adjacent syllable, such as is provided by stressed plural *-ot*. This constraint penalizes plural forms with unstressed [o] in the stem and no [o] in the plural—that is, native nouns with stem [o] pluralized with *-im*. If LOCAL(o) $\gg$ ϕ-MATCH, plural *-ot* will be chosen.

Many native masculine nouns with [o] do take *-im*: for example, the plural of [alon] 'oak' is [alon-im]. In such words, the morphological constraint, ϕ-MATCH, outranks LOCAL(o). The way to achieve this is to clone the constraint LOCAL(o) (Pater, 2006), so that the grammar has two copies, one ranked above ϕ-MATCH and one below it. For Becker (2009), each of these versions is lexically associated with the words that observe it. Thus, using the version of lexical indexation described in this dissertation (see Section 2.2.2.3), we can say that regular words like [alon] are marked with a feature *I*, while words like [xalon] that take *-ot* are marked with a feature *O*; this pairs with a constraint ranking LOCAL(o)$_O$ $\gg$ ϕ-MATCH $\gg$ LOCAL(o)$_I$.

We can see the effects of this ranking in the tableaux below. First, regular nouns like /alon$_I$/ follow the derivation in (23): since this noun is marked *I*, the violation of LOCAL(o) associated with the winning candidate, [alon-ím], is tallied by the lower-ranked version of the constraint, and this violation is not enough to overcome the morphological preference.

(23) *Regular Hebrew masculine nouns: morphosyntactic feature matching outranks phonology (cf. Becker, 2009, p. 93)*

/alon$_{I[\text{masc}]}$-{im$_{[\text{masc}]}$, ot$_{[\text{fem}]}$}/	LOCAL(o)$_O$	ϕ-MATCH	LOCAL(o)$_I$
☞ a. aloním			*
b. alonót		*!	

By contrast, in (24) we see the derivation of a masculine noun that takes *-ot*, like /xalon$_O$/. This word is marked with *O*, so the violation of LOCAL(o) in [xalon-ím] is attributed to the higher-ranked version of the constraint, and this forces the selection of gender-mismatched *-ot*.

(24) *Irregular Hebrew masculine nouns: specific phonology outranks morphosyntactic feature matching (cf. Becker, 2009, p. 94)*

/xalon$_{O\text{[masc]}}$-{im$_{\text{[masc]}}$, ot$_{\text{[fem]}}$}/	LOCAL(o)$_O$	ϕ-MATCH	LOCAL(o)$_I$
a. xaloním	∗!		
☞ b. xalonót		∗	

In the data of Bolozky and Becker (2006), 377 masculine nouns with [o] in the last syllable take *-im* (72.1%), while 146 take *-ot* (27.9%). Since both of the clones of LOCAL(o) are lexically marked and neither is default, speakers must assign new items to one of the constraints—in our terms, each new stress-shifting masculine noun with [o] in the final syllable must get assigned one of the features *I* and *O*. Since 72.1% of existing nouns have *I*, a new noun will be placed into the *I* class, and thus take *-im*, 72.1% of the time.

Of course, not every masculine noun that takes *-ot* has an [o] in the stem: for example, the plural of [ʃem] 'name' is [ʃemot]. Becker (2009) offers two possible solutions. The first is that there is *no* phonological constraint that prefers [ʃemot] over *[ʃemim], and since ϕ-MATCH prefers *[ʃemim], there is no way for the phonology to generate [ʃemot]. In this case, the preference for *-ot* must be exceptionally listed and exempt from normal morphophonological processes (see also Tessier, 2011). This listing is entirely unproductive, and will not have any effect on the treatment of new words. That is, in this approach, no new masculine word without [o] will ever be assigned *-ot*, no matter how many exceptional forms are listed in the lexicon, because they are all placed outside of normal phonological processes. (This is a very different approach to lexical listing than the one I describe in Chapter 3, in which lexical listing serves as the basis for drawing generalizations rather than being exempt from them.)

The second possibility, which Becker (2009) adopts, is that there is a phonological constraint that universally prefers *-ot* to *-im* by penalizing a feature of the latter suffix itself. This constraint is $\acute{\sigma}$/HIGH, which penalizes stressed high vowels. Kenstowicz (1997) and de Lacy (2004) propose similar constraints to capture the cross-linguistic hierarchy that mid non-central vowels like [e o] bear stress more readily than high vowels like [i u]; thus, *-ím* is universally more marked than *-ót*. As before, this constraint must be cloned: for most masculine nouns, the feature match outweighs the markedness of stressed [i]. We can assume that regular masculine nouns like [alon] 'oak', which pluralize with *-im*, are marked with a feature *M* associated with low-ranking $\acute{\sigma}$/HIGH$_M$, while nouns like [ʃem] 'name' that take *-ot* have a feature *T* associated with higher-ranking $\acute{\sigma}$/HIGH$_T$. The derivation for such words is shown in (25) (here LOCAL and HIGH are abbreviated as LOC and HI):

(25) *Irregular Hebrew masculine nouns: general phonology outranks morphosyntactic feature matching (cf. Becker, 2009, p. 98)*

/ʃem$_{T\,[\text{masc}]}$-{im$_{[\text{masc}]}$, ot$_{[\text{fem}]}$}/	LOC(o)$_O$	$\acute{\sigma}$/HI$_T$	ϕ-MATCH	LOC(o)$_I$	$\acute{\sigma}$/HI$_M$
a. ʃemím		*!			
☞ b. ʃemót			*		

Becker (2009) discusses one necessary task for learners in this model: ensuring that words are associated with the most specific constraints. That is, the *T* feature could apply to all masculine nouns that take *-ot*, while *O* can only apply to masculine nouns with [o] in the last syllable. Words like [xalon] 'window', which take plural *-ot* and have [o], must thus be associated with LOCAL(o)$_O$, not $\acute{\sigma}$/HIGH$_T$. Otherwise, the phonological generalization will not be learned, since the specific phonology of the stem is not relevant for $\acute{\sigma}$/HIGH. Becker (2009, pp. 104–108) describes the learning process; I do not go into it in detail here.

Let us summarize how the cloned constraints analysis of Becker (2009) captures phonological generalizations over lexically specific allomorphy like the Hebrew plural. First, we identify the morphologically preferred form (the *default*) and the morphologically dispreferred form. In the case of Hebrew, the default is *-im*, whose gender feature matches that of masculine nouns; *-ot* is dispreferred for masculine nouns because it is inherently feminine. A morphological constraint,

ϕ-MATCH, enforces this preference. Next, there must be some phonological constraint that prefers the non-default form. If no such constraint exists, then the idiosyncratic choice is lexically listed and is not productive: no new nouns will follow this pattern. In Hebrew, we have two such phonological constraints that prefer *-ot* over *-im*: LOCAL(o), which penalizes unstressed [o] without an adjacent stressed [o], and $\acute{\sigma}$/HIGH, which penalizes stressed high vowels. The latter constraint applies to all masculine nouns (since it penalizes the *-im* suffix itself), while the former only applies to nouns with [o] in their final syllable. These constraints are cloned, such that one copy is ranked above the morphological constraint preferring *-im*, ϕ-MATCH, while the other is ranked below it. Words that match the conditions for the more specific constraint (here, LOCAL(o)) are associated with one copy of that constraint, while other words are associated with one copy of the general constraint (here, $\acute{\sigma}$/HIGH). The proportion of words associated with each copy of a given constraint determines the likelihood of associating new words with that copy. If the ratio is *different* for the clones of the specific constraint and the general constraint, novel words satisfying the phonological condition for the specific constraint will be more or less likely to take the default form than others. In the Hebrew example, if a *greater* percentage of nouns with [o] take *-ot* than others, then the *higher-ranked* copy of LOCAL(o) will be associated with a *greater* percentage of nouns than $\acute{\sigma}$/HIGH, so the *higher-ranked* copy of LOCAL(o) will be *more attractive* to novel words than the higher-ranked copy of $\acute{\sigma}$/HIGH, which means that novel nouns with [o] are *more likely* to take *-ot* than other novel nouns.

The cloned constraint model of learning generalizations from lexical patterns of allomorphy does not necessarily require that these patterns be phonologically optimizing, for two reasons. First, both the specific and general constraints involved in the Hebrew plural analysis prefer *-ot*, so both are in a sense optimizing, just for different reasons. It is possible in most cases to find *some* constraint that would favor one possible allomorph over another, even if this constraint is a normally very low-ranked constraint penalizing individual segments, like *m. Since the more popular clone of this constraint need not be highly ranked throughout the language, its application can be fairly narrowly tailored and divorced from language-wide calculations of phonological optimality. Second, the ratios of the specific constraint can go in either direction. The only role of LOCAL(o) is to single out nouns with [o] in the last syllable from all others; however, the propensity of such nouns to take *-ot* could have gone in either direction. If nouns with [o] took *-ot less* than other nouns, this pattern could be learned as well, even though this is the *opposite* of the expected pattern, since more phonological constraints prefer *-ot* with stems with [o] than with other stems.

Thus, this model requires that phonological categories productively taking one allomorph at a different rate than the baseline be defined in a way that is phonologically *grounded*; that is, there must be some constraint that implicates this phonological category in relation to the allomorphs in question. This point is driven home by Becker et al. (2011) (also discussed in Section 2.1.3), who find that ungrounded generalizations in Turkish lexically variable consonant voicing alternations are not learned. In particular, participants in a nonce word study failed to productively extend correlations in the lexicon between the quality of a noun's stem vowel and whether or not its stem-final consonant undergoes a voicing alternation; this is predicted under the cloned constraint model, because there is no plausibly universal phonological constraint linking the quality of a vowel and voicing of an adjacent consonant.

The cloned constraint model is thus designed to be unable to capture "ungrounded" phonological generalizations on allomorph selection. However, other nonce word studies have found that speakers are able to learn ungrounded preferences. For example, Gouskova et al. (2015) found that Russian speakers learned a dispreference for the diminutive suffix *-ók* towards stems with hiatus (see Chapter 6 for more about Russian diminutives). This effect is non-local and unrelated to the content of the diminutive suffix: there is no phonological connection between stem-internal hiatus and the suffix *-ók*.

The cloned constraint model has other restrictions as well. The most important is that the phonological categories must be *product-oriented*: since they are grounded in markedness, they apply to the output inflected form, not to the base. However, as Bybee (1995) and Gouskova et al. (2015) point out (and as I discussed briefly in Section 2.2.2.1), some generalizations learned are also *source-oriented*, applying to the shape of the base rather than the inflected form (see Bybee, 1995; Kapatsinski, 2010). One such pattern, which I replicate in Section 6.4, again involves the Russian diminutive *-ók*: nouns ending in velars like [krʲuk] 'hook' usually take *-ók* (rather than the other options, *-ʲik* and *-tʃʲik*), but *-ók* itself causes velars [k g x] to mutate to [tʃʲ ʐ ʂ], so the diminutive is [krʲutʃʲ-ók]. However, nouns with surface [tʃʲ ʐ ʂ], whether underlying or alternating with other consonants, do not prefer *-ók* as strongly: for example, [palʲets] 'finger' undergoes its own consonant mutation to get the diminutive [palʲtʃʲ-ik], while [ʂalaʂ] 'shack' has the diminutive [ʂalaʂ-ɨk]. Thus, this really is a source-oriented generalization linking stem-final velars and *-ók*, and that is what speakers have learned.

Another restriction of the cloned constraint model is that, as described by Becker (2009), speakers can attend to one phonological category at a time—for example, in Hebrew, nouns with [o] in the last syllable are associated with a constraint that targets such words. However, in my nonce word studies, speakers are capable of weighing multiple factors at once. In my Russian nonce word study in Section 6.4, I show that speakers' choice of diminutive suffix can be sensitive to properties of the final consonant and vowel: nouns ending in dorsals greatly prefer *-ók*, while nouns with low or rounded vowels in the last syllable disprefer *-ók*. These countervailing tendencies are operative at the same time. Under the cloned constraint model, nouns ending in velars would be placed into the velar class and be assigned a diminutive in accordance with the behavior of velars in the lexicon; at that point, the choice of vowel should have no additional effect.

The takeaway from this section is that theories of phonological productivity requiring categorical or variable generalizations to be phonologically natural, grounded, or optimizing are too restrictive: not all allomorphy selection is phonologically optimizing, and not all learned generalizations are phonologically natural. In addition, the generalizations can be both source- and product-oriented, grouping words together by their base or their inflected form. However, since hard-coded generalizations actively apply to surface forms in derivations, they must be product-oriented. In Chapter 3, I present the sublexicon model of learning generalizations (Allen & Becker, 2015; Gouskova et al., 2015), which allows for a wide range of generalizations without any learning biases. In this model, generalizations are neither hard-coded into the grammar nor necessarily phonologically grounded; instead, speakers may learn arbitrary patterns from their lexicon and apply them to new forms. In the next section, I argue that this model is best suited for capturing morphological generalizations, just as it is for capturing phonological ones.

2.3.2 Morphological generalizations

In the previous section, I discussed theories of encoding phonological generalizations. I focused on two criteria: whether these generalizations were *hard-coded* into the derivational process and whether they were *grounded* in language-wide or universal phonological principles. This dissertation is primarily about *morphological* generalizations rather than phonological ones, so I will now discuss theories of encoding such generalizations using the same criteria. In some cases, it is not obvious how, or even whether, it makes sense to talk about morphological patterns being grounded or hard-coded, so I first define how I apply these terms morphologically before discussing the benefits and drawbacks of each.

2.3.2.1 Grounding

First, I will talk about grounding. What does it mean for a morphological pattern to be grounded? What would it be grounded in? There are two common approaches: first, morphological patterns that appear arbitrary can be placed into the syntax or phonolgy and derived from properties of abstract structure. As discussed in Section 2.1.5.1, the syntactic approach is often taken to account for *syncretism*, where exponents are shared by multiple paradigm cells. In fact, morphosyntactic grounding is often used to define syncretism: for example, Embick (2003) contrasts systematic syncretism with accidental homophony, in which two paradigm cells share a phonological form but not a morphosyntactic feature set.

The phonological instantiation of morphological patterns is used by Trommer (2021) to account for the German plural, as described in Section 2.2.2.4. In German, the plural can be realized by a number of suffixes and/or a vowel shift in the root known as umlaut; while some plural suffixes can cooccur with umlaut, the suffix *-n* generally does not. Trommer (2021) derives this complementary distribution from highly phonologically abstract underlying forms for the plural that can combine with abstract representations of noun stems to be realized as either umlaut or *-n* but not both. As discussed in Section 2.2.2.4, this phonologizing approach to encoding morphological generalizations is hard-coded and, in a sense, grounded, but requires the use of phonological buliding blocks that combine in a non-obvious way.

A second approach to grounding morphological patterns is paradigm uniformity. In Chapter 4, I discuss the distribution of possessive forms in Hungarian. Table 2.11 shows the possessive paradigms for four Hungarian nouns, which inflect for both noun number and the person and number of the possessor. The possessive paradigms have two points of variation. The first involves the "linking vowel" that appears in the 1SG, 2SG, and 2PL for singular nouns ('my X', 'your X', 'youse's X'). For most consonant-final nouns, like [dɒl] 'song' and [tʃont] 'bone', this is the mid vowel [o], but a small number of "lowering stems" like [vaːlː] 'shoulder' and [hold] 'moon' instead have a low linking vowel [ɒ]. The second point of variation is the presence or absence of [j] in singular forms with third person possessors and all possessed plural forms. Some words, like [tʃont] and [hold] (e.g. [tʃont-jɒ] 'her bone'), have [j] in their paradigms, while others, like [dɒl] and [vaːlː], do not (e.g. [dɒl-ɒ] 'her song').

noun	dɒl	tʃont	vaːlː	hold
gloss	‘song’	‘bone’	‘shoulder’	‘moon’
possessor	*singular noun*			
1SG	dɒlom	tʃontom	vaːlːɒm	holdɒm
2SG	dɒlod	tʃontod	vaːlːɒd	holdɒd
3SG	dɒlɒ	tʃont**j**ɒ	vaːlːɒ	hold**j**ɒ
1PL	dɒlunk	tʃontunk	vaːlːunk	holdunk
2PL	dɒlotok	tʃontotok	vaːlːɒtok	holdɒtok
3PL	dɒluk	tʃont**j**uk	vaːlːuk	hold**j**uk
possessor	*plural noun*			
1SG	dɒlɒim	tʃont**j**ɒim	vaːlːɒim	hold**j**ɒim
2SG	dɒlɒid	tʃont**j**ɒid	vaːlːɒid	hold**j**ɒid
3SG	dɒlɒi	tʃont**j**ɒi	vaːlːɒi	hold**j**ɒi
1PL	dɒlɒink	tʃont**j**ɒink	vaːlːɒink	hold**j**ɒink
2PL	dɒlɒitok	tʃont**j**ɒitok	vaːlːɒitok	hold**j**ɒitok
3PL	dɒlɒik	tʃont**j**ɒik	vaːlːɒik	hold**j**ɒik

Table 2.11: Sample Hungarian possessive paradigms

As seen in Table 2.11, all four combinations of linking vowel and [j] presence are possible. However, among the small class of lowering stems with linking vowel [ɒ], the 3SG *-ɒ* is much more common than *-jɒ*—that is, there are many more words like [vaːlː] and only a few like [hold], and as I show in Section 4.4, speakers productively apply this correlation when assigning 3SG possessive forms to nonce words. Rebrus (2013) and Rebrus et al. (2017) derive this correlation from a strong preference for paradigm uniformity—in particular, for suffix-initial vowels to agree throughout a stem's inflectional paradigm. In the case of [vaːlː], the stem is followed by the linking vowel [ɒ] in forms like singular 1SG and the affix [ɒ] in plural and singular third-person possessed forms, so

10 of the 12 cells have suffix-initial [ɒ]. However, for the case of [hold], most of the forms have an intervening [j], which leads to a less uniform paradigm. For regular nouns with linking vowel [o], the [j] is inconsequential for paradigm uniformity, since [o] is different from both possible suffix-initial segments, [ɒ] and [j].

The paradigm uniformity preference proposed by Rebrus et al. (2017) is grounded: it links the correlation learned between two paradigm cells to broader paradigmatic principles. However, this and similar constraints are not proposed as morphological universals and are not necessarily hard-coded: "The constraints in the paper are not meant as OT constraints (which are part of UG), but as generalisations over (sets of) surface forms that speakers have memorised" (Rebrus et al., 2017, p. 167); "The forms that occur (with the greatest probability) are the ones that are facilitated by the constraints and the forms not facilitated are dispreferred (or do not occur)" (Rebrus et al., 2017, p. 176). It is not clear how strong a bias these constraints are intended to impose on the learning process. For example, consider an alternate version of Hungarian in which lowering stems instead *preferred* the forms with [j] (opposite to reality). In this alternative Hungarian, there is a morphological dependency between lowering stems and possessive [j], but one that *reduces* analogical uniformity rather than improves it. Presumably, Rebrus et al. (2017) would predict that the dependency in alternative Hungarian would go unlearned, because it cannot be captured through the paradigm uniformity constraints—this would constitute a learning bias in favor of paradigmatically "grounded" morphological dependencies.

The Hungarian analysis is an invocation of paradigm uniformity constraints in explaining allomorph selection. In other cases, output–output correspondence between two or more paradigm cells is used to drive or block *phonological* processes. For example, Transderivational Correspondence Theory of Benua (1997) explains under- or overapplication of phonological processes in a complex derived form due to faithfulness to the base form. This is asymmetrical output–output correspondence: the derived form must correspond with a base, but not vice versa. In Sundanese, for example, nasality spreads from nasal consonants to following vowels until blocked by an oral consonant: in [ŋũliat] 'stretch', nasality spreads from [ŋ] to the following vowel [u], but not across [l] to later vowels. Verbs pluralize with an infixed *-ar-* after the first consonant; if the first consonant is nasal, nasality spreads not just to the infix vowel but also to the vowel(s) after the infix. Thus, the plural of [ɲĩãr] 'seek' is [ɲ-ãl-ĩãr] rather than *[ɲ-ãl-iar], which would be expected given that [l] typically blocks nasality spreading. Benua (1997, p. 70) attributes this overapplication of nasal spreading to an output–output faithfulness constraint requiring segments in the plural to agree in nasality with their singular base.

Paradigmatic faithfulness constraints require identity, so they are only applicable to morphological dependencies where the selected allomorphs have some phonological material in common. This is a possible explanation for some of the cases discussed in this dissertation, like the Hungarian possessive described above. For others, however, the relationship is not one of identity: for example, in Section 6.4 I show that speakers have learned the lexical correlation that words with plural *-a* are more likely to take *-ók* as a diminutive. There is some similarity: both plural *-a* and diminutive *-ók* are obligatorily stressed. However, the correlation between the two suffixes holds *even once stress is taken into account*: speakers are more likely to assign diminutive *-ók* to words with stressed plural *-a* than to other words with stressed plurals. This sort of effect cannot be captured by paradigm uniformity factors. Thus, in my model of morphological dependencies presented in Chapter 3, I allow for speakers to learn *any* arbitrary correlation between two inflected forms, not just those that are morphologically grounded in paradigm uniformity effects.

2.3.2.2 Hard-coding

Let us now turn to morphological hard-coding. By this I mean that the relationship between two morphological patterns (for example, the selection of allomorphs in two different cases) follows from the structure of morphosyntactic and diacritic features used in different rules of realization. As discussed in Section 2.1.5.2, for example, morphological dependencies not involving paradigm uniformity are very often hard-coded into morphological analyses in the form of *inflection classes*. An inflection class can be defined as a group of words "whose members each select for the same set of inflectional realizations" (Aronoff, 1994, p. 64). One example of typical inflection classes, in Russian, is repeated in Table 2.12. According to this analysis, Russian nouns divide into four classes that fully determine their suffixes across all cases. Thus, for example, the four-class presentation of Russian assumes that nouns that have *-i* in the dative like [kost^j] 'bone' also have *-ju* in the instrumental, null in the nominative, and so on.

class	I	II	III	IV
example	'law'	'school'	'bone'	'wine'
nominative	zakon	ş̆kol-a	kost^j	vin-o
accusative	zakon	ş̆kol-u	kost^j	vin-o
dative	zakon-u	ş̆kol-e	kost^j-i	vin-u
genitive	zakon-a	ş̆kol-i	kost^j-i	vin-a
instrumental	zakon-om	ş̆kol-oj	kost^j-ju	vin-om
locative	zakon-e	ş̆kol-e	kost^j-i	vin-e

Table 2.12: Inflection classes with singular case forms for Russian nouns (from Corbett, 1982)

As discussed in Section 2.1.5.1, under the analysis of Müller (2004), class II nouns are marked $[-\alpha, +\beta]$, while class II nouns are $[-\alpha, -\beta]$. The vocabulary items spelling out various case suffixes reference these features, which fully determine a noun's paradigm: $[-\alpha, +\beta]$ nouns all have -*a* in the nominative, -*e* in the dative, -*oj* in the instrumental, and so on. This is morphological *hard-coding*: morphological dependencies are built into the structure of the inflection class features. That is, the morphological dependency between nominative -*a* and dative -*e* is captured inherently by the set of rules of realization used actively in derivations.

This hard-coding is quite efficient for "canonical" inflection classes like those in Table 2.12, in which a noun's macroclass fully determines its inflection and knowing one or two forms of a word generally allows one to infer its entire paradigm (Corbett, 2009; Corbett & Baerman, 2006). It does not work so well for cases like the Hungarian described above, in which a noun's linking vowel ([o] vs. [ɒ]) can vary separately from the presence vs. absence of [j] in some possessive forms. If we tried to group these together into two inflection classes—for example, [o]/[ɒ] vs. [j]/no [j]—we would leave large numbers of nouns unaccounted for and have to mark them separately anyway. Instead, I argue that the correlation between lowering stems (with linking vowel [ɒ]) and possessives without [j] should be learned as a gradient generalization between two more narrowly targeted inflectional features: lowering stems, which have a feature [lower], also usually have a feature [−j] indexing a possessive paradigm without [j] (see Section 4.2 for a full analysis). Thus, the underlying form of [vaːlː] 'shoulder' would be /$\text{vaːlː}_{\text{[lower,}-\text{j]}}$/, with two separate features indexing its 1SG form [vaːlː-ɒm] and its 3SG possessive [vaːlː-ɒ].

The Russian analysis of Müller (2004) is thus categorically different from the proposed analyses of Hungarian: the hard-coded feature structure of the former allows inference of new inflected forms from known ones, whereas the latter does not. However, linguists like Ackerman et al. (2009) and Baerman et al. (2017) argue that the difference between the two languages is one of *degree*, rather than kind: Russian has a greater degree of inflectional cohesion than Hungarian (at least in the corners discussed in this paper), in that forms are more informative of one another than in Hungarian. Indeed, as Parker and Sims (2020) point out, the presentation of Russian inflection in Table 2.12 hides various exceptions, subclasses, and other complexities. For example, they describe nine subclasses of nouns that do not fit neatly into a single class—in one, nouns like [vremʲa] 'time' have *-i* in the dative like class III ([vrem-enʲ-i], with a stem extension), but instrumental *-em* ([vrem-enʲ-em]), which is the usual ending for class I and IV. As I describe in more detail in Chapter 6, Russian nouns can vary in their inflectional stress patterns, and while certain stress patterns occur more often with certain inflectional classes (Brown et al., 1996), these are variant tendencies. In general, a noun's stress pattern must be stored separately from its inflection class.

In Section 2.3.1, I suggested that models of phonological generalizations that require phonological grounding and hard-coding are too restrictive, and that we should instead allow for analogical learning of arbitrary preferences of allomorphs. I argue that the same should be done for morphological dependencies in both Hungarian and Russian. In Section 4.4, I show that Hungarian speakers have learned a correlation between linking vowel and presence of [j]—that is, a correlation between two lexical diacritic features (the other studies in this dissertation establish similar patterns). Thus, an account of speakers' productive language faculty must include some mechanism for linking these two features without hard-coding, such as what I propose in Chapter 3. Since these correlations are necessary anyway, we can dispense with hard-coded inflection classes altogether.

What would this look like in Russian? In (26), we see rules spelling out the same case features as in (4) from Section 2.1.5.1, but instead of having umbrella features [$\pm\alpha$] and [$\pm\beta$], each suffix allomorph is indexed by its own narrowly targeted feature, as in Hungarian. For the purposes of illustration, I retain the case decompositions of Müller (2004)—for example, the rule in (26d) is satisfied by the genitive, dative, locative, and instrumental, all of which are [+obl]. It is never invoked in the instrumental, which satisfies the more specific rules in (26a) and (26b). Nouns that have the [DL:e] feature have their dative and locative (both [−subj]) spelled out through (26c), while nouns with the [G(DL):i] feature have their genitive and, if they have no [DL:e] feature, their dative and locative spelled out through (26d).

(26) *Rules of realization for syncretic Russian suffixes in class II and III (without inflection classes)*

a. [−subj, −gov, +obl] $\leftrightarrow$ oj / [I:oj] ___

b. [−subj, −gov, +obl] $\leftrightarrow$ ju / [I:ju] ___

c. [−subj, +obl] $\leftrightarrow$ e / [DL:e] ___

d. [+obl] $\leftrightarrow$ i / [G(DL):i] ___

Rather than having overarching inflection class features that determine full paradigms, each lexical entry now has one feature for each of the rules invoked in (26). For example, a class II noun like [ʂkola] 'school' has the underlying form /ʂkol$_{\text{[G(DL):i,DL:e,I:oj]}}$/, while a class II noun like [kost^{j}] is underlyingly /ʂkol$_{\text{[G(DL):i,I:oj]}}$/. In this analysis, what defines a class II noun is that it has the [G(DL):i], [DL:e], and [I:oj] features; likewise, class III nouns share the features [G(DL):i], and [I:ju]. Under this approach, there is no symbolic unit of representation corresponding to an inflection class. Instead, inflection classes are *emergent*: since these features tend to cooccur on the same words, the learner's sublexical grammars for each feature would have clusters of extremely strong morphological constraints enforcing inflection class cohesion. For example, the [DL:e] sublexical grammar would have a very heavily weighted constraint *[I:ju].

For the purposes of illustration, I retain the syncretisms in the analysis of Müller (2004): nouns with both [DL:e] and [G(DL):i] (class II nouns) spell out the dative and locative ([−subj, +obl] cases) as *-e* by (26c) and the genitive ([+subj, +obl]) as *-i* by (26d); nouns with [G(DL):i] but not [DL:e] (class III nouns) spell out all three cases as *-i* by (26d). We could also shift the burden of inflectional cohesion further to morphological constraints by governing syncretism here as well. Each rule could spell out exactly one case, with features to match: class II nouns would have features [G:i], [D:e], and [L:e], while class III nouns would have [G:i] (the same genitive feature as class II nouns), [D:i], and [L:i]. In this case, though, the fact that the dative and locative case suffixes are homophonous in both classes (intra-paradigmatic syncretism) becomes a coincidence.

The lexical items necessitated by the rules in (26) contain more redundancies than one with inflectional class features II and III, but offer two advantages mentioned above. First, the approach using these lexical items can more readily handle the intricacies and exceptions described by Parker and Sims (2020) for Russian: nouns that cut across classes like [vrem^{j}a] 'time' do not require any new features, only an unusual combination of already existing features (like [D:i,I:em]). Second, this analysis provides a path for learning inflection classes: they emerge from features tracking the distribution of individual affixes—which, as described above, are independently necessary. In contrast, an analysis with features like [$\pm\alpha$] and [$\pm\beta$] requires speakers to formally arrive at the right combinations of rules and endings. Thus, an emergent class analysis requires fewer pieces of theoretical machinery and enables a formal unification of relatively canonical inflection class systems like Russian (cf. Corbett & Baerman, 2006) and lexically variable cases of feature cooccurrence like the Hungarian linking vowel and possessive [j].

2.3.3 Summary

In Chapter 3, I present a model for encoding generalizations in lexically variable patterns of allomorphy. To understand the architectural choices that go into this model, let us review the findings of this chapter. Speakers are able to learn and productively apply all kinds of gradient generalizations to lexically determined patterns of allomorphy. Much of the past experimental work has explored *phonological* generalizations; the purpose of the current dissertation is to test and model *morphological* dependencies as well: speakers use known forms of a word to infer unknown forms. In order to model how speakers store morphological dependencies, we must first choose a representation for the lexically exceptional forms over which they are generalizing. I argued that lexically specific behavior is best indexed using symbolic diacritic features. These can appear in a wide range of grammatical uses across phonology and morphology, and are properly located in the phonological underlying forms of exponents. The encoding of morphological patterns through diacritic features on underlying forms allows the model in Chapter 3 to treat phonological and morphological generalizations in the same way. Finally, I evaluated theories of storing gradient generalizations along two criteria: hard-coding and grounding. Hard-coding and grounding are too restrictive to capture the full range of generalizations that speakers use productively, so in my model, speakers are able to learn arbitrary patterns from the lexicon and store them in a separate pattern-matching module not actively involved in derivations.

The *sublexicon* model (Allen & Becker, 2015; Becker & Gouskova, 2016; Gouskova et al., 2015) satisfies the criteria argued for in this chapter. In this model, speakers learn such patterns as arbitrary statistical preferences over partitions of the lexicon defined by lexical diacritic features. Previous applications of the sublexicon model account for *phonological* generalizations, but I use the properties of diacritic features to extend the sublexicon to capture *morphological* dependencies as well. Motivated by the similarity between phonological and morphological generalizations described in this section, my extended sublexicon model brings both types of generalizations under the purview of a single pattern-matching mechanism.

3 Sublexicon models of morphological learning

The main task I am testing in this dissertation is the affixation of nonce words. I use this task to show that speakers have learned and productively apply phonological and morphological generalizations over patterns of allomorphy in their lexicon like those discussed in Section 2.1. In this chapter, I describe a theoretical proposal for how speakers perform nonce word tasks (and, more generally, productively extend lexical patterns to new words). This has two main architectural features corresponding to the overviews in Section 2.2 and Section 2.3: first, lexically specific patterns are encoded through diacritic features that are attached to the underlying forms of lexical items and index their behavior by being invoked in the context of rules of realization for the appropriate allomorph. Second, generalizations between lexically specific patterns (morphological dependencies) are encoded as correspondence relations between two diacritic features. These relations are not *hard-coded*, meaning that they are not invoked in the active derivation of items; instead, they are stored in a separate pattern-matching module that encodes generalizations learned in the lexicon, but only comes into use as a backup for unfamiliar words that do not already have their behavior encoded directly in their lexical entry. These relations are also not necessarily phonologically or morphologically *grounded*: any arbitrary pattern can be learned and stored in the grammars, not just those within a limited set of "natural" patterns. This work's novel empirical and theoretical contribution is that I test how one complex form of a word influences another in the nonce word task, so this chapter extends a previous theoretical proposal—the sublexicon model of morphological learning (Allen & Becker, 2015; Becker & Gouskova, 2016; Gouskova et al., 2015)—to cover the new domain of morphological dependencies. I also provide a new extension of Distributed Morphology to account for cases of variation in which individual lexical items can variably take one of multiple allomorphs, where this choice can also depend on the syntactic context in which they appear. I then adapt the sublexicon model to account for learning and productively extending patterns where variable lexical items have underlying distributions of allomorphs encoded in their lexical entry through a feature that can take variable strength, not just a binary feature (present/absent) as assumed in earlier versions of the sublexicon model.

In Section 3.1, I describe the sublexicon model, which encodes phonological generalizations over lexically specific variation as phonotactic constraints in grammars describing partitions of the lexicon. These grammars comprise a separate pattern-matching module of the grammar that learns generalizations over the lexicon but only comes into play when needed to determine of the behavior of new words that do not have their behavior already stored. New words, then, are assigned a feature probabilistically according to their performance on these sublexical grammars. To demonstrate this model, I use the example of the Hungarian possessive, which I study in detail in Chapter 4. In Section 3.2, I present my novel extension of the sublexicon model to account for implicative structure between complex forms built from the same root. Here, I take advantage of the fact that I store selectional restrictions as essentially phonological features: correlations between morphological behaviors can be cast as constraints penalizing cooccurrence of features on the same lexical item. Again, I use the Hungarian possessive as an example here.

In my model, speakers have multiple grammars used for diferent purposes: each feature defines a sublexical grammar that is invoked when needed to determine the behavior of a novel form, but otherwise lies dormant. In Section 3.3, I compare this model with a simpler one with only a single, language-wide phonotactic grammar, in which generalizations over lexically specific patterns are stored as indexed constraints ranked above the general phonotactic constraint that applies acloss all forms in the leanguage. The multiple grammar model and the single grammar model differ in their predictions about the behavior of nonce words: the multiple grammar (sublexicon) model predicts that speakers will assign a feature to a nonce word once, meaning that repeated trials with the same nonce word should yield consistent responses. In contrast, the single grammar model predicts that speakers should not store lexical properties of nonce words unless they receive positive evidence about its behavior, so repeated trials of the same nonce word should yield variable responses. As I explain, however, these predictions are very difficult to test against one another in practice.

In the basic sublexicon model, and in general in my case studies, it is assumed that individual lexical items behave categorically, meaning that their behavior can be indexed by a binary feature. This is built into the architecture of the sublexicon model, which forms generalizations over the set of words that share a feature. However, this is a simplifying assumption that at times ignores real and interesting variability within lexical items. In particular, allomorphy in the Czech locative, studied in Chapter 5, involves pervasive variability for individual lexical items, and any analysis of the phenomenon is incomplete without accounting for the effect of *syntactic context* on locative allomorphy for variable lexical items. In the remainder of this chapter, I describe proposals for how to extend my theoretical mechanisms to account for variable lexical items. In Section 3.4, I first present a model that brings lexical and syntactic variation within a generative framework (assuming Distributed Morphology), using the Czech locative as an example. As before, I assume that selectional restrictions are stored as lexical diacritic features. My model of variation relies on *variable, weighted* features that set a lexical item's baseline distribution of allomorphs. These features trigger splits in derivations, which are weighted accordingly; at the end of the derivation, one valid path is chosen probabilistically. Next, the basic sublexicon model must be modified to account for features that are weighted rather than binary. In Section 3.4.5, I describe an extension of the sublexicon model in which new words are probabilistically assigned a feature weight.

3.1 The basic sublexicon model

3.1.1 Sublexicons and features

The sublexicon model (Allen & Becker, 2015; Becker & Gouskova, 2016; Gouskova et al., 2015) encodes phonological generalizations in lexically specific variation. This allows learners to pick up on the partial phonological predictability determining a given lexical item's choice of allomorph. As such, it follows in the path of previous models, like the Minimal Generalization Learner (Albright & Hayes, 2003), that use structured, formalized comparison to extract generalizations over the set of lexical items to which a given morphophonological rule applies (see Guzmán Naranjo (2019) for an overview of proposals—in both the generative tradition and others, such as usage-based linguistics—using analogical modelling to predict morphological class affiliation).

In the sublexicon model, the learner divides the lexicon into *sublexicons* that pattern together. These sublexicons correspond with morphological features that track inflectional patterns. Consider the example of the Hungarian plural and possessive, discussed in Section 2.3.2.1 and studied at length in Chapter 4. In Table 3.1, I show two allomorphs for each (modulo vowel harmony). Most Hungarian nouns have plural *-ok*, while *-ɒk* appears with a small class called "lowering stems". In the possessive, nouns may take *-ɒ* or *-jɒ* (which I call *-V* and *-jV*, abstracting over vowel harmony), and both are very frequent. The plural and possessive suffixes may vary independently; all four combinations are possible.

noun	dɒl	tʃont	vaːlː	hold
gloss	'song'	'bone'	'shoulder'	'moon'
plural	dɒl-ok	tʃont-ok	vaːlː-ɒk	hold-ɒk
possessive	dɒl-ɒ	tʃont-**j**ɒ	vaːlː-ɒ	hold-**j**ɒ

Table 3.1: Possible combinations of Hungarian plural and possessive suffixes

Lowering stems cannot be identified by their phonological form, nor, in many cases, can a noun's possessive (Rácz & Rebrus, 2012; Siptár & Törkenczy, 2000). Thus, both require lexical marking on underlying forms (at least some of the time). I assume that lowering stems are marked with [lower]; non-lowering stems, which are the large majority, are unmarked here. An alternative would be to mark *all* nouns: lowering stems are [+lower] and others are [−lower]. I use this binary approach for the more balanced possessive allomorphy: nouns taking -*jV* and -*V* are marked with [+j] and [−j], respectively. This means that the lexical entries for Hungarian nouns can have one or multiple features indexing their inflectional patterning, as shown in (27). Here, I group all nouns that share a feature together into a list, which I refer to as a *sublexicon*. Each feature defines a sublexicon comprising lexical entries containing that feature.

(27) *Lexical entries for Hungarian nouns*

a. [lower]: /vaːlː$_{[\textbf{lower},-\text{j}]}$/ 'shoulder', /hold$_{[\textbf{lower},+\text{j}]}$/ 'moon', /ɟaːr$_{[\textbf{lower},-\text{j}]}$/ 'factory', /ɲaːr$_{[\textbf{lower},+\text{j}]}$/ 'poplar', . . .

b. [+j]: /tʃont$_{[\mathbf{+j}]}$/ 'bone', /hold$_{[\text{lower},\mathbf{+j}]}$/ 'moon', /paːr$_{[\mathbf{+j}]}$/ 'pair', /ɲaːr$_{[\text{lower},\mathbf{+j}]}$/ 'poplar', . . .

c. [−j]: /dɒl$_{[\mathbf{-j}]}$/ 'song', /vaːlː$_{[\text{lower},\mathbf{-j}]}$/ 'shoulder', /kaːr$_{[\mathbf{-j}]}$/ 'damage', /ɟaːr$_{[\text{lower},\mathbf{-j}]}$/ 'factory', . . .

Each feature is associated with rules in (28) to get the right output forms. These rules will suffice for the purposes of this chapter, but ignore some complexities of Hungarian morphophonology. For a full analysis, see Section 4.2.

(28) *Rules of realization for the Hungarian plural and possessive (simplified)*

a. PL ↔ ɒk / [lower] ___   c. POSS ↔ jɒ / [+j] ___

b. PL ↔ ok /   d. POSS ↔ ɒ / [−j] ___

In the process of learning, Hungarian speakers append these features to lexical entries to mark their behavior. However, when Hungarian speakers wish to form the possessive of a new lexical item, they must determine which of the rules of realization in (28) applies to this particular item. Since a new lexical item, by definition, lacks a [±j] feature, to create a possessive, speakers must figure out whether this word gets a [+j] or [−j] diacritic—that is, which sublexicon in (27) it belongs to. For previously unseen words, there is no lexically listed possessive form, so they must place the word into a sublexicon on the basis of its phonology. They do so by using a *sublexical grammar*, described in the next section.

3.1.2 Sublexical phonotactic grammars

The sublexicon model learns generalizations by extending the concept of phonotactic grammars, which describe speakers' knowledge of what segments and sequences are preferred or dispreferred in a language (that is, what makes a good word of a language). Hayes and Wilson (2008) present a model of phonotactic learning in which a learner captures generalizations over a language's surface forms through a constraint-based phonotactic grammar. In their proposal, the learner keeps track of sounds or sequences of sounds (defined in terms of features) that are rare or absent in the lexicon and proposes constraints against them, weighting them in accordance with the strength of the generalization. For example, in Hungarian, adjacent obstruents generally agree in voicing (Siptár & Törkenczy, 2000), so words with a voiced obstruent followed by an unvoiced obstruent will appear far less than expected from the frequency of voiced and unvoiced obstruents on the whole. Accordingly, the learner should generate a phonotactic constraint, *[−son,+voice][−voice], which penalizes voiced obstruents before voiceless consonants, and weight it heavily.[1] This is how the

[1]Hayes and Wilson (2008) released an implementation of their learning model, the UCLA Phonotactic Learner. In practice, it does capture many strong phonotactic tendencies, but also learns many constraints that strike linguists as phonologically unnatural and do not correspond to the phonotactic knowledge of real speakers (Hayes & White, 2013). When applied to the Hungarian data, the Phonotactic Learner also failed to learn many moderate tendencies to which speakers displayed sensitivity. In this section, I focus on the conceptual framework of sublexical phonotactic grammars rather than any particular model of how they are learned. The Sublexical Learner, a freely available implementation of the sublexicon model described by Allen and Becker (2015), does not induce new constraints, but rather weights given ones.

speaker knows that surface forms with mixed-voice obstruent clusters are unlikely in Hungarian.

The sublexicon model uses phonotactic learning to capture generalizations over words that pattern together by treating them as discrete subsets of the lexicon that have their own generalizations about what makes a good word—that is, sublexicons. When the learner encounters a case of lexically specific allomorphy, she marks words that take the various allomorphs with lexical diacritic features; words sharing a feature are grouped into a sublexicon. The learner induces a phonotactic grammar not just for the entire lexicon, but also for each sublexicon (that is, for each diacritic feature), capturing patterns specific to that sublexicon. As I will show in Chapter 4, Hungarian nouns ending in sibilants and palatals categorically take possessive *-V* (for example, the possessive of [haːz] 'house' is [haːz-ɒ], not *[haːz-jɒ]), while nouns ending in vowels always take *-jV*. The sublexical grammar for the [+j] sublexicon should include heavily weighted constraints penalizing final sibilants and palatals, and the [−j] sublexicon should penalize word-final vowels. I do not include an explicit learning process for how these weights are determined: this process requires negative evidence (words that are *not* in the sublexicon), and there are two possibilities for where this negative evidence comes from. First, Hayes and Wilson (2008) assume that all of the words in the lexicon are compared with a generated set of non-words, whose properties are contrasted with the set of real worlds. Alternately, the set of words in a given sublexicon is compared with all the other words in the language that are not in the sublexicon—that is, words that could be in the sublexicon, but are not. In both cases, the goal is to isolate properties that appear more often in the words in the sublexicon than those outside of it. As I mention in Section 3.4.5.4, the latter approach is *more informative* than the former, since it compares words in a sublexicon with other actual words of the language, providing a much more narrowly targeted set of forms used as negative evidence.

These sublexical grammars are then reflected in speakers' behavior. When a speaker wishes to form the possessive of a novel word, she evaluates the stem against each sublexicon's grammar, where each sublexical grammar yields a score for that word. The better a word fares on the [+j] sublexicon relative to the [−j] sublexicon, the more likely it is to be placed into this sublexicon, and thus take *-jV*.

In Figure 3.1 and Figure 3.2, we see two nonce words from Chapter 4, [ruɲɒs] (orthographically *runyasz*) and [fuːzaːt] (*fúzát*), tested on toy sublexical grammars with the constraints described above. Here, [ruɲɒs] is penalized by *[+strident]#, penalizing word-final sibilants, in the [+j] sublexicon, but not by the constraint against word-final vowels in the [−j] sublexicon; [fuːzaːt] accrues no penalities. For the purposes of illustration, I assume that all three constraints have a weight of 5.

constraint / *weight*	*[+strident]# / 5	*[+palatal]# / 5	total
ruɲɒs	−5	0	−5
fuːzaːt	0	0	0

Figure 3.1: Evaluation of nonce words *runyasz* and *fúzát* on the [+j] sublexical grammar

constraint / *weight*	*[+syllabic]# / 5	total
ruɲɒs	0	0
fuːzaːt	0	0

Figure 3.2: Evaluation of nonce words *runyasz* and *fúzát* on the [−j] sublexical grammar

Since [ruɲɒs] has a better score on the [−j] sublexical grammar than the [+j] sublexical grammar, it is much more likely to be placed into the former and form its possessive with *-V*. Specifically, this is a maximum entropy model (Hayes & Wilson, 2008): a word's likelihood of being placed into a sublexicon is proportional to its (negative) score raised to the power of *e*. Here, the probability of [ruɲɒs] being assigned to the $[+j]$ sublexicon is $\frac{e^{-5}}{e^{0}+e^{-5}} = .0067 = .67\%$. On the other hand, since [fuːzaːt] has the same score on both sublexicons, it has a 50% chance of being assigned to each.

The sublexicon model is designed to capture generalizations over the phonological shape of each sublexicon's members. However, as I discussed in Section 2.3.2.1, there is also a morphological generalization in the lexicon, which speakers observe (see Chapter 4): lowering stems are more likely to have possessive -*V*. In a feature-based analysis, this means that [lower] and [−j] are likely to cooccur on lexical items. A feature-based analysis casts this as a cooccurrence relation between features: if a lexical entry has a [lower] feature, it is also likely to have a [−j] feature. In the next section, I extend the sublexicon model to accommodate these relations.

3.2 A sublexicon model with morphology

In my proposal, diacritic features appear in phonological underlying forms (see Section 2.2.2.3). If the constraints of the sublexical grammars operate over phonological features, then these can include the diacritic features as well. Accordingly, I propose that each sublexicon's grammar has constraints penalizing diacritic features alongside those penalizing phonological features. For example, every member of the [+j] sublexicon has [+j] (by definition), but very few have [lower], since lowering stems rarely take -*jV*. Since [lower] is underrepresented in the [+j] sublexicon, the [+j] sublexical grammar should contain a heavily weighted constraint *[lower] penalizing nouns with both [lower] and [+j]. The [−j] sublexicon, comprising words that take -*V*, will also have a *[lower] constraint, but it will not be as strong, since lowering stems are better represented among -*V* words (though still uncommon).

Why have I set up these constraints to *penalize* feature cooccurrence (e.g. between [lower] and [+j]) rather than *reward* it (e.g. between [lower] and [−j])? The choice is a theoretical one: in my proposal, lexical features are placed alongside phonological features in underlying forms, so they should be treated like phonological features. In the phonotactic grammars of Hayes and Wilson

(2008), constraints penalize uncommon phonological sequences, so I treat the morphological features the same way. However, there is no difference in the resulting grammar: penalizing [lower] in the [+j] sublexicon is equivalent to rewarding it in the [−j] sublexicon. The two grammars will output the same probabilities of assigning allomorphs to words, as I discuss in Section 3.4.5.2.

Figure 3.3 and Figure 3.4 show the evaluation of our two nonce words on the toy grammars, now containing *[lower]. This constraint has a heavier weight in the [+j] grammar than in the [−j] grammar (2 and 1, respectively). Here, the speaker knows that the plurals of these words are [ruɲɒs-ɒk] and [fuːzaːt-ɒk], so she has marked both with [lower].

constraint	*[+strident]#	*[+palatal]#	*[lower]	total
weight	5	5	2	
$\text{ruɲɒs}_{\text{[lower]}}$	−5	0	−2	−7
$\text{fuːzaːt}_{\text{[lower]}}$	0	0	−2	−2

Figure 3.3: Evaluation of nonce lowering stems *runyasz* and *fúzát* on the [+j] sublexical grammar with *[lower]

constraint	*[+syllabic]#	*[lower]	total
weight	5	1	
$\text{ruɲɒs}_{\text{[lower]}}$	0	−1	−1
$\text{fuːzaːt}_{\text{[lower]}}$	0	−1	−1

Figure 3.4: Evaluation of nonce lowering stems *runyasz* and *fúzát* on the [−j] sublexical grammar with *[lower]

The *[lower] constraint slightly reduces the likelihood of [+j] (and possessive *-jV*) being assigned to /$\text{ruɲɒs}_{\text{[lower]}}$/, from .67% to .25%. For /$\text{fuːzaːt}_{\text{[lower]}}$/, the effect is more visible: the likelihood of [+j] goes from 50% to $\frac{e^{-2}}{e^{-1}+e^{-2}} = .269 = 26.9\%$. This shows how the sublexicon model can accommodate the effects found in the nonce word experiment, both phonological and morphological: nonce words ending in sibilants are less likely to be assigned *-jV* (that is, be placed in the [+j] sublexicon), as are words shown as lowering stems. These effects can all be assessed in a single calculation, correctly allowing them to compound or cancel out. This can be contrasted with the cloned constraint approach in Becker (2009) described in Section 2.3.1.2, which only allows for forms to "count" towards one generalization at a time.

3.3 A single grammar alternative

In this section, I compare the sublexicon model described in this chapter to a minimally different alternative that generalizes the language-wide phonotactic grammar of Hayes and Wilson (2008) in a different way. Hayes and Wilson (2008) developed the phonotactic grammars used in the sublexicon model to capture speakers' knowledge of phonotactics operationalized as phonotactic well-formedness judgements—for example, which clusters are good onsets in English? In their theory, a single grammar is trained on the entire lexicon describing typical words in the language. The output of the grammar is a score that determines whether a candidate word in a language is a good word for that language. There are two things I wish to emphasize about this model. First, the generalizations in the phonotactic grammar are not *hard-coded* in the sense described in Section 2.3: the grammar describes speakers' knowledge about the language, but is not (necessarily) used to actively choose between candidate forms in a derivation. Instead, the output of the evaluation of a phonotactic grammar is a *grammaticality judgement* for a single candidate. In this sense, the phonotactic grammar is (or at least can be) distinct from the grammar that evaluates candidates within a generative phonological derivation. Second, the constraint space is *unrestricted*: candidate constraints are generated freely from the data and are not necessarily grounded in the sense described in Section 2.3—that is, they need not penalize what are traditionally seen as marked structures. As Hayes and Wilson (2008, p. 281) explain, this freedom of constraint generation sets an inductive baseline that can be compared to stronger theories with biases suggesting that certain phonotactic generalizations are unlearnable because they are phonologically unnatural.

In the sublexicon model, phonotactic grammars are generated over portions of the lexicon defined by their inflectional patterning. I assume that the constraint space is similarly unrestricted: speakers can learn any phonological generalizations on allomorphy, not just phonologically "natural" ones (but see Becker et al., 2011). The outputs of the sublexical grammar are the same as with the lexical phonotactic grammar: well-formedness scores. However, the sublexicon model includes an extra step: a candidate's scores on the different sublexical grammars are compared, and the candidate is placed into one of the sublexicons probabilistically in accordance with its scores on the sublexicons' respective grammars. These grammars do not themselves evaluate candidates in the course of a phonological derivation; rather, they fill in the lexicon so it is sufficiently well-defined for the derivation to subsequently be carried out. However, this is not the only way to extend the language-wide phonotactic grammar to account for generalizations over nouns sharing a feature. In the remainder of this section, I compare it with an alternative, in which there is still only one language-wide phonotactic grammar, but constraints can be indexed to apply only to morphemes with particular diacritic features. As I describe, the two models make slightly different predictions for nonce word studies: the multiple grammar model predicts that individuals should treat nonce words consistently over repeated trials, while the single grammar model does not. The multiple grammar model also allows for a wider range of relations between individual sublexicons and language-wide phonotactics.

3.3.1 Properties of a single grammar model

In this section and beyond, I refer to the sublexicon model as the *multiple grammar* model, since it uses sublexical grammars defined for each feature to probabilistically outputs a form for nonce words based on its phonological (and morphological) characteristics. The alternative has only one language-wide grammar, so I call it the *single grammar* model. As before, this grammar comprises weighted constraints, which assign scores to candidates and chooses one of the candidates in accordance with their relative scores. In this, it resembles a Harmonic Grammar (e.g. Legendre et al., 1990, 2006; Pater, 2016; Potts et al., 2010) and other stochastic constraint-based grammars (e.g. Becker, 2009; Hayes & Londe, 2006). However, there is one key difference between these grammars, as well as theories of phonological allomorph selection like Wolf (2008), and my single grammar model: the typical grammars have the goal of penalizing phonological markedness, so that the candidate that best balances markedness and faithfulness will win. However, my multiple grammars approach explicitly allows for speakers to learn and apply non-optimizing patterns of allomorphy. The constraints in the single grammar model proposed here are similarly unrestricted, in that they need not be grounded. Under the current assumption that individual words behave categorically, they are also not hard-coded: as in the multiple grammar model, this grammar is only invoked to productively extend patterns from the lexicon when an unfamiliar word lacks a sufficiently defined lexical entry and its behavior is not determined by an existing diacritic feature on its lexical entry.

The sublexical grammars shown in Figure 3.3 and Figure 3.4 are *source-oriented* (see the discussion in Section 2.2.2.1): they evaluate the base form of the word, not the resulting possessive. I made this choice for the sake of illustration, although sublexical grammars can also capture *product-oriented* generalizations by evaluating the potential inflected form; in Gouskova et al. (2015), each sublexicon is associated with two grammars, a source-oriented one and a product-oriented one. Indeed, as they discuss, both types of generalizations are common in cases of lexically variable allomorphy. However, if the grammar is choosing between inflected forms, the candidates must be product forms, and the generalizations thus product-oriented. All of the phonological constraints in the illustration above can be recast in product-oriented terms (though this is not necessarily true for all possible constraints—several generalizations describing the Russian diminutive studied in Chapter 6 must be source-oriented, as discussed in Section 2.3.1.2 and Section 6.2). For example, the prohibition against stem-final stridents and palatals with *-jV* can be recast as penalizing these sounds before [j]: *[+strident]j and *[+palatal]j. These sequences are quite rare in Hungarian overall, so it is plausible that a general phonotactic grammar would account for them. However, the fact that vowel-final nouns consistently take *-jV* may be problematic: hiatus is relatively tolerated in Hungarian (and can be quite profligate, as in [indiɒiɒk] 'Indians'), so any language-wide phonotactic constraint penalizing hiatus would probably not be strong enough to create the consistent choice for *-jV* over *-V* for vowel-final words.

This brings us to another feature of this single grammar model: the constraints can have different weights for specific constructions, as in Lexical MaxEnt (Coetzee & Pater, 2011; Gouskova & Linzen, 2015). That is, the anti-hiatus constraint (*[+syllabic][+syllabic]) should be weighted *more heavily* when one of the vowels is part of a possessive allomorph. This constraint may have a weight of 2 in general but 5 in a possessive context. Below, I notate this, taking after (Gouskova & Linzen, 2015), by marking its weight as $2+3_p$: 2 in general, with an extra 3 if one of the vowels is in the possessive.

3.3.2 Examples

Let us recast the example of the nonce words [ruɲɒs] and [fuːzaːt] from Figure 3.1 and Figure 3.2 in our new single grammar with phonological constraints. In Figure 3.5, the grammar evaluates the candidate forms created by attaching *-ɒ* and *-jɒ* to [ruɲɒs]. As before, the stem-final strident causes a violation next to [j] for *-jV*, while *-V* violates none of the listed constraints. Thus, the candidate [ruɲɒsjɒ] has a score of -5, while [ruɲɒsɒ] is unpenalized with a score of 0. None of the constraints violate either candidate possessive for [fuːzaːt], so both have a score of 0, as shown in Figure 3.6.

/ruɲɒs/-{ɒ, jɒ}$_p$	*constraint* / *weight*	*[+strident]j / 5	*[+palatal]j / 5	*[+syllabic][+syllabic] / $2+3_p$	score
ruɲɒsjɒ		-5	0	0	-5
ruɲɒsɒ		0	0	0	0

Figure 3.5: Evaluation of nonce word *runyasz* on the single grammar

/fuːzaːt/-{ɒ, jɒ}$_p$	*constraint* / *weight*	*[+strident]j / 5	*[+palatal]j / 5	*[+syllabic][+syllabic] / $2+3_p$	score
fuːzaːtjɒ		0	0	0	0
fuːzaːtɒ		0	0	0	0

Figure 3.6: Evaluation of nonce word *fúzát* on the single grammar

As with the multiple grammar model, I assume that candidates are selected using the maximum entropy principle: the exponential of a candidate's score is proportionate to its likelihood of being chosen. In this case we are comparing candidates evaluated by the *same* grammar. The probabilities, however, are the same: the likelihood of [ruɲɒsjɒ] is $\frac{e^{-5}}{e^{0}+e^{-5}} = .0067 = .67\%$, while [fuːzaːt] is equally likely to take *-V* and *-jV*, since the scores of the two candidates on the single grammar are the same.

To demonstrate the differential effect of hiatus, let us consider another nonce word (not used in my study), [piɒduː], which has both internal hiatus and a stem-final vowel. Both candidates violate *[+syllabic][+syllabic] once because of the stem-internal hiatus and are penalized 2 points. However, [piɒduːɒ] has an additional violation between the stem-final vowel and the possessive suffix. Since this hiatus involves a possessive suffix (here marked with *p*), the violation is more egregious, costing this candidate 5 points. This increased severity makes *-V* very unlikely for this word, since the candidate with this suffix has a score 5 lower than [piɒduːjɒ]. This corresponds to a probability of $\frac{e^{-7}}{e^{-2}+e^{-7}} = .0067 = .67\%$ (as I show in Section 3.4.5.2, the distribution of probabilities depends only on the difference between the candidate scores, not their individual values).

/piɒduː/-{ɒ, jɒ}$_p$	*constraint*	*[+strident]j	*[+palatal]j	*[+syllabic][+syllabic]	score
	weight	5	5	$2+3_p$	
piɒduːjɒ		0	0	-2	-2
piɒduːɒ		0	0	$-2, -(2+3)$	-7

Figure 3.7: Evaluation of nonce word *piadú* on the single grammar

3.3.3 Real and nonce words

In the previous section, I showed how the single grammar stochastically chooses a possessive suffix for Hungarian nonce words. By definition, such words do not have a listed form for the possessive, and speakers must resort to defaults or inference in forming their possessives. However, existing words like [bɒraːt] 'friend' have a set possessive form (in this case, [bɒraːt-jɒ]). The grammar in Figure 3.6 should not apply to them. As with the multiple grammar model, we can assume that such words are lexically marked with [±j], at which case the rules in (28) can apply to spell out the possessive in accordance with a noun's lexical feature. Thus, under this model, the grammar only applies to assign a possessive when the regular rules fail.

This assumption works under the current assumption that existing words in the lexicon show categorical behavior. In Section 3.4.5.3, I extend the multiple grammar model to account for variable lexical items—the adjustment is fairly minimal. However, the single grammar model most easily accounts for variable lexical items with a more substantive change: hard-coding. That is, if lexical items can be variable, the single grammar model is active in the course of derivations by setting a baseline distribution for the different allomorphs that is *adjusted* by the weighted features of variable lexical items. I compare the two models with the assumption of variability at length in

Section 5.4.6.3.

3.3.4 Comparison with the multiple grammar model

The single grammar model makes at least two predictions that are different from the multiple grammar model. Although the studies in this dissertation cannot, on their own, distinguish between the two, I describe the predictions here for future work.

The first predicted difference involves the behavior of nonce words in repeated trials. In the multiple grammar model, the stochastic process is one of assigning a lexical feature to a new word. Once this feature is assigned, it should be there for good: we would predict a given speaker to be consistent in their inflection if prompted repeatedly. In contrast, the single grammar model predicts within-speaker variation for nonce words. In the single grammar model, the grammar outputs an inflected form for words with no lexical features; since the nonce word will continue to have no features, the suffix for each trial should be chosen stochastically.[2] To test this, we would need a nonce word study in which the same speakers are tested on the same nonce word multiple times.

Secondly, the single grammar model requires that individual suffixes should have *more restrictive* phonotactic tendencies than the language as a whole, while the multiple grammar model allows for grammars of individual sublexicons to be either *more or less restrictive* than the language as a whole. In the single grammar model, each candidate is subject to all of the constraints in the language's general phonotactic grammar. Certain candidates can be especially penalized with *heavier* weights, but they cannot be *weaker*. However, in the multiple grammar model, sublexical grammars do not necessarily come with all of the constraints active in the phonotactic grammar of the language as a whole. The multiple grammar model, but not the single grammar model, predicts the possibility of a sublexicon that accepts phonotactically deviant forms: if one sublexicon has a lenient grammar in general, it will rate most words highly, including words that are poor fits for

[2] This prediction does not hold if we assume that speakers treat the nonce word task as a feedback loop: they treat their own output as meaningful input and store the word's properties learned from this self-generated input.

the language as a whole (see also the discussion of Kapatsinski (2005) in Section 2.2.2.2); if other sublexicons are more picky, the leftovers will fall into the sublexicon with the lenient grammar. This is another sense in which the multiple grammar model is maximally *unrestrictive*: it is able to learn a broader range of patterns than a single grammar model with many of the same architectural assumptions.

While I do not compare this version of the single and multiple grammar models further, I argue that their variable versions can be distinguished with a particular combination of corpus studies, one of which is shown in Section 5.6. I compare the two models further in Section 5.4.6.3.

3.4 A model of syntactic and lexical variation

In the original implementation of the sublexicon model, sublexicon membership is binary: a word either has a feature or it does not. However, as Becker and Gouskova (2016) note, many cases of morphological patterning—including all of those in this dissertation—have variable lexical items, which appear with more than one possible suffix allomorph. For example, the Hungarian possessive, discussed at length in Chapter 4, has allomorphs both with and without [j] (*-ɒ* and *-jɒ*). While most words take one suffix or the other, words like [ɒblɒk] 'window' can take either [ɒblɒk-ɒ] or [ɒblɒk-jɒ]. Similarly, the Czech locative, discussed in Chapter 5, has allomorphs *-u* and *-ɛ*. In this case, many stems accept both forms, so the locative of [tiːl] 'back of the head' can be [tiːl-u] or [tiːl-ɛ]. The distinction between variable and non-variable forms is not binary; rather, variable forms show off the full range of allomorph distributions: some Czech words almost exclusively take *-u*, others almost exclusively take *-ɛ*, and everything in between. Thus, any proper account of lexical variation must allow for individual lexical items to have different baseline rates of allomorph distribution. Becker and Gouskova (2016) account for variation by assuming that sublexicon membership can be gradient: for example, if a Czech noun takes locative *-ɛ* 60% of the time, it is 60% in the *-ɛ* sublexicon and 40% in the *-u* sublexicon. As I show in this section, this assumption alone is not enough to fully account for Czech and Hungarian variation. In this section, I propose a more complex model that allows regular derivations in a Distributed Morphology framework to exhibit variable outcomes by splitting the derivations in accordance with variable (weighted) features. In particular, this model accounts for effects of *syntactic* context on the distribution of allomorphs.

Variation in both the Czech locative and the Hungarian possessive shows sensitivity to the syntactic context in which a given form is used. For example, Kiefer (1985, p. 108) notes that for some variable forms, the type of possession can affect the possessive allomorph: *-V* prefers inalienable possession, while *-jV* prefers alienable possession.

> In general, the suffix *-ja*/*-je* can be used to render conspicuous the relation of real possession whereas the other *habeo*-relations are indicated by means of the suffix *-a*/*-e*. For example, *anyag-a* 'its material' as in 'the material of the costume' and *anyag-ja* 'his material' as in the merchant's material' [sic]. Similarly, one can express that the house has no windows with *ablak-a* or that the carpenter has no windows with *ablak-ja*. [...] The distinction is gaining ground and it can often be encountered in everyday talk. Of course, only words which can take either possessive suffix may develop such a semantic differentation. [...] This difference is not fully grammaticalized as yet. That is, one can use invariably *-a*/*-e* or, alternatively, *-ja*/*-je* without thereby affecting meaning. This is a distinction *in statu nascendi* and if nothing unforeseen happens it may very well become grammaticalized in the not too remote future.

Many languages distinguish between alienable and inalienable possession morphosyntactically, and indeed, den Dikken (2015) argues that the morphological split for words like [ɒblɒk] reflects an alienability distinction encoded in the syntax. I follow him in assuming that the Hungarian distinction lies in the *syntax*, not just in the semantics as assumed by Ritter (2023).

The choice between the Czech locative allomorphs *-u* and *-ɛ* also shows some syntactic conditioning, also described in Section 5.1.2: -ε is more frequent in "canonical" uses of the case, namely marking the object of prepositions expressing location like [v] 'in', and less frequent in less "canonical" uses, such as with prepositions like [o] 'about' that do not express location (Bermel & Knittl, 2012). I likewise assume that this distinction is not exclusively semantic, but has a syntactic reflex as well. This assumption can find support in the suffixes' cognates in a related Slavic language, Russian: Jakobson (1984), agreeing with Kuznecov (1953), argues that there are two locative (and, similarly, genitive) cases in Russian, which are only distinguished in a small number of words. He contrasts the two locatives of [sn^{j}eg] 'snow', [sn^{j}ég^{j}e] (which he labels L1) and [sn^{j}egú] (L2):

(29) *L1:* Artists look for something in [sn^{j}ég^{j}e], but there is nothing picturesque in [sn^{j}ég^{j}e].

(30) *L2:* The crows were looking for something in the [sn^{j}egú], but there was no food in the [sn^{j}egú].

(Jakobson, 1984, p. 125)

The difference between L1 and L2 is one of "directionality of the action onto the object": the former "ascribe[s] to the object a property, or a state resulting from the action directed towards the object, and consequently may be called ascriptive (directional)" (Jakobson, 1984, p. 125). This is a fairly close fit to the above description of Czech, suggesting that a syntactic split for variable items in Czech is a reasonable hypothesis. Thus, my model of formal variation must also allow for circumstances like the Hungarian possessive and Czech locative, which I call *syntactically conditioned variation*.

3.4.1 Desiderata for a model of variation

The model of variation laid out in this section is designed to satisfy a number of theoretical and architectural criteria, listed below. In particular, this model is designed as a way to get around modularity concerns: if the syntactic component of a derivation strictly precedes vocabulary insertion, then the latter should only operate within the bounds of the former. That is, if the Czech locative -ε (for example) is used in one syntactic context 80% of the time, standard assumptions of modularity would make it that a given lexical item cannot push the rate of -ε much above 80%. However, in reality, individual lexical items can cover the full range of probabilities.

- **The model can account for both lexically and syntactically conditioned variation.** Each type of variation is handled in its appropriate module, following standard assumptions of modularity and locality in Distributed Morphology (Harley, 2014; Harley & Noyer, 1999). That is, while lexically conditioned variation is handled at vocabulary insertion, syntactically conditioned variation is handled earlier, in the syntax. Vocabulary insertion does not have access to arbitrary amounts of syntactic information, but rather operates on the (perhaps variable) output of the syntactic module. In addition, analyses of syntactically conditioned variation should make use of typologically principled syntactic distinctions: for example, alienable vs. inalienable possession for the Hungarian possessive, or two locative cases in Czech as in Russian.

- **Lexically and syntactically conditioned variation should compound.** Any given instance of variation may be lexically conditioned, syntactically conditioned, or both. However, if certain variable items show syntactically conditioned variation, then all variable items should. The model should not predict (or, at least, not necessarily predict) a situation where two lexical items show the same *overall* rates of allomorph usage but only one has that allomorphy *dependent on syntactic context*.

- **There is no categorical distinction between variable and non-variable behavior.** Non-variable items should be analytically equivalent to variable items with an extremely high probability of putting out one form. This is the opposite assumption from that of Guzmán Naranjo and Bonami (2021): in their analysis of the Czech locative, they take syntactic conditioning as *prima facie* evidence of a categorical split between variable items (whose variation is syntactically conditioned) and non-variable items (which are, by definition, invariant regardless of syntactic context). One possible exception—that is, one allowable discontinuity between variable and non-variable items—is that a given lexical entry may lack a variable feature (see next bullet point) entirely, in which case it exhibits categorical default behavior.

- **Variability is associated with features.** Variability in the grammar is driven by morphological and morphosyntactic *features* that may have different strengths (that is, probability of outcomes) for each lexical item. For example, lexically conditioned variation may be handled through the variable application of a particular rule of vocabulary insertion. However, the variability in application is due to the variable presence or absence of a feature in its context. Once the feature is either selected or not, the standard vocabulary insertion process of Distributed Morphology is unchanged: given a set of features, a rule is selected to expone them according to the Subset Principle. This is a different architectural assumption from approaches that assume that post-syntactic rules, including rules of realization, themselves apply probabilistically (Adger & Smith, 2010; Labov, 1969; Nevins & Parrott, 2010), and is more in line with the overall spirit of the theoretical proposals in this chapter.

3.4.2 A split-derivation model of variation

Here I present a model of variation in which variable features cause the morphosyntactic derivation to split into multiple derivations, which are computed in parallel and assigned a score. At the end of the derivational process, one derivation is chosen: the greater a derivation's score, the higher its likelihood of being chosen. This split crucially allows variable vocabulary insertion—which, in a modular theory like Distributed Morphology, strictly follows syntax—to be influenced by variable syntactic processes without being limited by them at the extremes.

A derivation is initialized with a score of 0 and proceeds as usual until it hits a feature that is marked in the grammar as being variable. For example, in the case of lexically conditioned variation, a rule of vocabulary insertion may have a variable feature in its context; for syntactically conditioned variation, a functional head may assign a variable case feature to its object. At this point, the derivation splits into two derivations: one in which the variable feature enters the derivation, and the other in which it does not. The score of the branch without the feature remains unchanged. (If the derivation had not previously been split, the score for this branch remains at its initial value of 0.) The other branch of the derivation, which contains the variable feature, updates the score from before the split, s, using a linear function, $f(s) = as + b$. Since the derivation starts with a score of 0, after the first split, the branch with the variable feature will have a score of b. The values of a and b are specified in the lexicon. Each variable feature has a single value a that is the same for all lexical items with that feature. On the other hand, each lexical entry with the feature associates the feature with a distinct value b. These b weights are similar to the additive lexical scaling factors of Linzen et al. (2013) and Gouskova and Linzen (2015). Thus, in a given derivation, the value of the parameter b in the scoring function depends on the lexical item responsible for introducing the variable feature into the derivation.

As the derivation continues, it may encounter other variable features, in which case the derivation splits again and scores are updated. At the end of the derivation, each derivation's score is converted into a probability using the maximum entropy algorithm (Goldwater & Johnson, 2003; Hayes & Wilson, 2008): a derivation's probability is proportional to the exponential function of its score (that is, e to the power of its score), scaled so that the sum of all probabilities adds up to 1.

I now provide a formal mathematical description of how derivations are scored before providing examples of variable derivations in Section 3.4.3. Given a derivation d involving variable features $f_1, f_2, \ldots, f_n$, its final score $s(d) = s_n(d)$ is calculated in (31), where $s_i(d)$ is its score after the application of f_i, and the initial score $s_0(d) = 0$. Here, a_{f_i} is the multiplicative parameter specified for feature f_i, and b_{f_i,w_i} is the additive parameter for feature f_i in the lexical entry of word w_i.

(31) *Calculating the score of a derivation*

$$
\begin{aligned}
&s_0(d) = 0 \\
&s_1(d) = a_{f_1} s_0(d) + b_{f_1,w_1} = b_{f_1,w_1} \text{ if } f_1 \in d, \qquad s_1(d) = s_0(d) = 0 \text{ if } f_1 \notin d \\
&s_2(d) = a_{f_2} s_1(d) + b_{f_2,w_2} \qquad \text{if } f_2 \in d, \qquad s_2(d) = s_1(d) \qquad \text{if } f_2 \notin d \\
&\vdots \\
s(d) = {} &s_n(d) = a_{f_n} s_{n-1}(d) + b_{f_n,w_n} \qquad \text{if } f_n \in d, \qquad s_n(d) = s_{n-1}(d) \quad \text{if } f_n \notin d
\end{aligned}
$$

Of course, d is just one possible derivation in the set of all possible derivations D. Its probability $P(d)$ is then calculated from its score $s(d)$ as shown in (32). This formula shows that the probability of a given derivation d is proportional to the exponential of its score, divided by the sum of all probabilities $P(d_i)$ such that their sum adds up to 1.

(32) *Calculating the probability of a derivation from its score*

$$P(d) = \frac{e^{s(d)}}{\sum\limits_{d_i \in D} e^{s(d_i)}}$$

3.4.3 Examples

To show how the model works, I present three cases of variation, based on the Czech locative: one with lexically conditioned variation, one with syntactically conditioned variation, and one with both (that is, the real situation).

3.4.3.1 Example 1: Lexically conditioned variation

In the first example, nouns can show variation with different rates of -*ɛ* (or have categorical -*u*), but the choice between -*u* and -*ɛ* is dependent only on the noun to which it is attaching, not the syntactic context. This case is analogous to the one described by Becker and Gouskova (2016): roots belong to a sublexicon variably, and the choice is made to include them with or without the variable feature in each derivation.

This situation has the two rules of realization in (33), which invoke a variable diacritic feature, [+lvar].

(33) *Rules of realization for Example 1 (lexically conditioned variation)*

a. [LOC] ↔ ɛ / [+lvar]___

b. [LOC] ↔ u

The rules in (33) spell out locative case as -*ɛ* in the context of a [+lvar] feature (attached to the head noun) and -*u* otherwise.

In Table 3.2, we see the derivational process of four nonce words, which I have listed together to facilitate easy comparison. The first, [tʃal], does not have [+lvar] in its lexical entry at all. Thus, there is no variation and the derivation never splits—this word displays categorical default behavior, always taking -*u* in the locative. The other words have the variable feature [+lvar] with different additive parameters: the parameter of [+lvar] in the lexical entry for [bɛːl], $b_{+\text{lvar},\text{bɛːl}}$, is -1; for [moʊl], $b_{+\text{lvar},\text{moʊl}} = 1$. Accordingly, they have different rates of -*ɛ* locatives: about 27% and 73%, respectively. Finally, [+lvar] has a much higher parameter (10) in the lexical entry of [ciːl], meaning that its locative is [ciːl-ɛ] essentially categorically: [ciːl-u] is only predicted for about one out of every 20,000 locative tokens of [ciːl]. Of course, if the parameter of [+lvar] is even higher, the behavior will be even more lopsided. With $b_{+\text{lvar},\text{ciːl}} = 20$, locative -*u* is predicted only once out of every 500,000,000 locative tokens—that is, never within a human lifetime. Thus, a sufficiently high strength for a variable feature yields categorical behavior as an extreme of grammatically encoded variation.

Here I set $a_{+\text{lvar}} = 2$. However, since [+lvar] always causes the first and only split in this example, this parameter is irrelevant, because the previous score is always 0, so $2 \cdot 0 = 0$. To get a sense for how a derivation proceeds, let us follow the derivational paths of the locative of [bɛːl] in Table 3.2 more closely. The derivation splits at vocabulary insertion of the root, before insertion of the locative: when we are inserting the root, we can either insert it with or without the [+lvar] feature.

In Derivation 1, [+lvar] is inserted with the root. Since the b weight of [+lvar] in the lexical entry of [bɛːl] is -1, Derivation 1 gets a score of -1. Derivation 2, in which [+lvar] is not inserted, retains its prior score of 0. These scores are then used to calculate the likelihood of Derivation 1 getting chosen: 26.9%. Since the locative is spelled out as -ɛ if and only if the derivation contains [+lvar], the likelihood of locative -ɛ is 26.9%.

output of syntax:	[K n K[LOC]]
noun:	tʃal
[+lvar] parameters:	not in lexical entry
	Derivation 1 [+lvar] absent
rule of realization:	(33b)
output:	tʃal-u
final score s:	0
probability:	$\frac{e^0}{e^0} = 1$

output of syntax:	[K n K[LOC]]	
noun:	bɛːl	
[+lvar] parameters:	$a = 2, b = -1$	
	Derivation 1 [+lvar] present	Derivation 2 [+lvar] absent
rule of realization:	(33a)	(33b)
score s_1:	$2 \cdot 0 - 1$	0
output:	bɛːl-ɛ	bɛːl-u
final score s:	-1	0
probability:	$\frac{e^{-1}}{e^{-1}+e^0} \approx .269$	$\frac{e^0}{e^{-1}+e^0} \approx .731$

output of syntax:	[K n K[LOC]]	
noun:	moʊl	
[+lvar] parameters:	$a = 2, b = 1$	
	Derivation 1 [+lvar] present	Derivation 2 [+lvar] absent
rule of realization:	(33a)	(33b)
score s_1:	$2 \cdot 0 + 1$	0
output:	moʊl-ɛ	moʊl-u
final score s:	1	0
probability:	$\frac{e^1}{e^1+e^0} \approx .731$	$\frac{e^0}{e^1+e^0} \approx .269$

output of syntax:	[K n K[LOC]]	
noun:	ciːl	
[+lvar] parameters:	$a = 2, b = 10$	
	Derivation 1 [+lvar] present	Derivation 2 [+lvar] absent
rule of realization:	(33a)	(33b)
score s_1:	$2 \cdot 0 + 10$	0
output:	ciːl-ɛ	ciːl-u
final score s:	10	0
probability:	$\frac{e^{10}}{e^{10}+e^0} \approx .99995$	$\frac{e^0}{e^{10}+e^0} \approx .00005$

Table 3.2: Derivations for Example 1 (lexically conditioned variation)

3.4.3.2 Example 2: Syntactically conditioned variation

The second example shows a situation where words can behave either variably or categorically, but all variable items behave the same way, purely conditioned by the syntax. That is, variable items all have the same baseline rate of *-ɛ* locatives; any difference in their observed rates is due solely to the fact that words may tend to appear in different syntactic environments. (This is an idealized scenario, since variation is likely always conditioned to some extent by individual lexical items.) One way to accomplish this is with the rules of realization in (34). Here, variable items are marked with [+varl], which is a regular (non-variable) feature; it *always* appears on words with which it is marked, and does not cause a derivational split. However, there is another feature, [+core]—equivalent to the [+directionality] of Jakobson (1984) (see Chvany, 1986)—that is variably assigned in the syntax. (I discuss this feature further in Section 5.2.1.) Prepositions that assign locative case to their objects can assign the case either with or without the [+core] feature in a given derivation. Each locative-assigning preposition has the [+core] feature in its syntactic lexicon entry with a particular *b* parameter weight. I set $a_{+\text{core}} = 1$; however, as in the previous example, this is irrelevant because it is always multiplied by the initial score of 0. According to the rules of realization in (34), the locative is spelled out as *-ɛ* when a [+core] locative is assigned to a noun that allows *-ɛ*, as lexically marked with its [+varl] feature. If either [+core] or [+varl] (or both) is absent, the locative is spelled out as *-u*.

(34) *Rules of realization for Example 2 (syntactically conditioned variation)*

a. [LOC, +core] ↔ ɛ / [+varl]___

b. [LOC] ↔ u

In Table 3.3, I show the derivational process of two nouns, [tʃal] and [bɛːl], each attaching to two prepositions, [v] 'in' and [o] 'about'. As in the previous example, [tʃal] is non-variable and lacks the [+varl] feature, while [bɛːl] allows variability and thus has the [+varl] feature. However, unlike in the previous example, this feature is *categorical*: it is not associated with a parameter, and does not cause a split in the derivation. On the other hand, the [+core] feature causes a split in the derivation, and the two prepositions assign it with different parameters: [v] has [+core] with $b_{+\text{core,v}} = 2$, while [o] has it with $b_{+\text{core,o}} = -2$. The result is that both prepositions assign *-u* to the non-variable item [tʃal], while variable items like [bɛːl] have differing behaviors with the two prepositions: [v bɛːl-ɛ] is much more common than [v bɛːl-u] (88.1%), while with [o], the situation is reversed: [o bɛːl-ɛ] appears only around 11.9% of the time, while [o bɛːl-u] is much more common.

preposition:	v	
[+core] parameters:	$a = 1, b = 2$	
	Derivation 1 [+core] present	Derivation 2 [+core] absent
output of syntax:	[$_K$ n K $\begin{bmatrix}\text{LOC}\\ +\text{core}\end{bmatrix}$]	[$_K$ n K [LOC]]
score s_1:	$1 \cdot 0 + 2$	0
noun:	tʃal	tʃal
[+varl] presence:	absent	absent
rule of realization:	(34b)	(34b)
output:	v tʃal-u	v tʃal-u
final score s:	2	0
probability:	$\frac{e^2}{e^2+e^0} \approx .881$	$\frac{e^0}{e^2+e^0} \approx .119$

preposition:	o	
[+core] parameters:	$a = 1, b = -2$	
	Derivation 1 [+core] present	Derivation 2 [+core] absent
output of syntax:	[$_K$ n K $\begin{bmatrix}\text{LOC}\\ +\text{core}\end{bmatrix}$]	[$_K$ n K [LOC]]
score s_1:	$1 \cdot 0 - 2$	0
noun:	tʃal	tʃal
[+varl] presence:	absent	absent
rule of realization:	(34b)	(34b)
output:	o tʃal-u	o tʃal-u
final score s:	-2	0
probability:	$\frac{e^{-2}}{e^{-2}+e^0} \approx .119$	$\frac{e^0}{e^{-2}+e^0} \approx .881$

preposition:	v	
[+core] parameters:	$a = 1, b = 2$	
	Derivation 1 [+core] present	Derivation 2 [+core] absent
output of syntax:	[$_K$ n K $\begin{bmatrix}\text{LOC}\\ +\text{core}\end{bmatrix}$]	[$_K$ n K [LOC]]
score s_1:	$1 \cdot 0 + 2$	0
noun:	bɛːl	bɛːl
[+varl] presence:	present	present
rule of realization:	(34a)	(34b)
output:	v bɛːl-ɛ	v bɛːl-u
final score s:	2	0
probability:	$\frac{e^2}{e^2+e^0} \approx .881$	$\frac{e^0}{e^2+e^0} \approx .119$

preposition:	o	
[+core] parameters:	$a = 1, b = -2$	
	Derivation 1 [+core] present	Derivation 2 [+core] absent
output of syntax:	[$_K$ n K $\begin{bmatrix}\text{LOC}\\ +\text{core}\end{bmatrix}$]	[$_K$ n K [LOC]]
score s_1:	$1 \cdot 0 - 2$	0
noun:	bɛːl	bɛːl
[+varl] presence:	present	present
rule of realization:	(34a)	(34b)
output:	o bɛːl-ɛ	o bɛːl-u
final score s:	-2	0
probability:	$\frac{e^{-2}}{e^{-2}+e^0} \approx .119$	$\frac{e^0}{e^{-2}+e^0} \approx .881$

Table 3.3: Derivations for Example 2 (syntactically conditioned variation)

Note that the rules in (34) do not allow for forms that categorically take -ɛ in the locative. In order to capture these forms, we would need an additional categorical feature, [+loc]. This feature would then be referenced by an additional rule, shown in (35b), which always assigns -ɛ to [+loc] nouns regardless of whether [+core] is present in the syntax.

(35) *Rules of realization for Example 2 (syntactically conditioned variation), with an additional rule for non-variable -ɛ locative nouns*

a. [LOC, +core] ↔ ɛ / [+varl]___

b. [LOC] ↔ ɛ / [+loc]___

c. [LOC] ↔ u

In this case, the difference between categorical nouns and variable nouns is one of kind, not degree: variable nouns have a different feature ([+varl]) from categorical -ɛ nouns, which have [+loc]. Similarly, in this case categorical behavior cannot be an extreme of variable *morphological* behavior, since individual variable nouns all have the same rate of locative -ɛ. However, categorical behavior could still be an extreme of variable *syntactic* behavior: for example, if a particular preposition always assigned locative -*u* even to variable nouns, this could be achieved by having a very low [+core] parameter like −20.

3.4.3.3 Example 3: Lexically and syntactically conditioned variation

The last example, which is more realistic, combines the two types of variation: prepositions assign locative with the variable morphosyntactic feature [+core], and nouns can be associated with the variable morphological feature [+lvar]. Of course, both of these features can split the derivation, meaning that variable items can follow four different derivational paths.

The rules of realization are the same as those of (34) from the previous example; however, in this case both [+core] and [+lvar] are variable features, whereas previously, in Example 2, [+varl] was not.

(36) *Rules of realization for Example 3 (lexically and syntactically conditioned variation)*

a. [LOC, +core] ↔ ε / [+lvar]___

b. [LOC] ↔ u

This example highlights the need for splitting and scoring derivations, rather than treating syntactically and lexically conditioned variation as a series of successive coin flips. Suppose a given syntactic context assigns [+core] 11.9% of the time, as in Table 3.3. If this syntactic choice is made definitively, before the choice to insert the [+lvar] feature, then -*ε* can be inserted no more than 11.9% of the time for a given word, no matter how high its [+lvar] weight. Instead, as I discuss in more detail in Section 3.4.4, splitting and scoring derivations allows a highly weighted [+lvar] feature to push rates of -*ε* higher than the baseline rate for a given syntactic context, which is in fact what we see in Czech: some variable nouns appear with -*ε* almost all the time, even in syntactic contexts that are much more favorable for -*u*.

Since [+core] and [+lvar] are both variable features, each are associated with lexical items with particular b parameters: prepositions (r) are lexically specified to assign [+core] locatives with a given $b_{+\text{core},r}$, while nouns (n) may include the [+lvar] feature with a listed parameter $b_{+\text{lvar},n}$. The tables below show the derivations of the four nouns from before, each combining with two prepositions. This example is different from the previous ones in a number of ways. First of all, this time, the a parameter of [+lvar] matters: since the syntactic variable feature of [+core] can make the score non-zero before vocabulary insertion, the [+lvar] functions can operate on this non-zero score. As before, $a_{+\text{lvar}} = 2$ (and $a_{+\text{core}} = 1$, though this does not play a role). Second, the b parameters are different: the preposition [v] assigns [+core] with $b_{+\text{core},\text{v}} = 5$, while for the preposition [o], $b_{+\text{core},\text{o}} = 3$. As in Example 1, [tʃal] lacks [+lvar] entirely; as shown in Table 3.4 the derivation only splits once, on the syntactic variable feature, though both derivations output -*u*.

preposition:	v	
[+core] parameters:	$a = 1, b = 5$	
	Derivation 1	Derivation 2
	[+core] present	[+core] absent
output of syntax:	[K n [K [LOC, +core]]]	[K n [K [LOC]]]
score s_1:	$1 \cdot 0 + 5$	0
noun:	tʃal	tʃal
[+lvar] parameters:	not in lexical entry	not in lexical entry
rule of realization:	(36b)	(36b)
output:	v tʃal-u	v tʃal-u
final score s:	5	0
probability:	$\frac{e^5}{e^5+e^0} \approx .993$	$\frac{e^0}{e^5+e^0} \approx .007$

preposition:	o	
[+core] parameters:	$a = 1, b = 3$	
	Derivation 1	Derivation 2
	[+core] present	[+core] absent
output of syntax:	[K n [K [LOC, +core]]]	[K n [K [LOC]]]
score s_1:	$1 \cdot 0 + 3$	0
noun:	tʃal	tʃal
[+lvar] parameters:	not in lexical entry	not in lexical entry
rule of realization:	(36b)	(36b)
output:	o tʃal-u	o tʃal-u
final score s:	3	0
probability:	$\frac{e^3}{e^3+e^0} \approx .953$	$\frac{e^0}{e^3+e^0} \approx .047$

Table 3.4: Derivations for the word [tʃal] in Example 3 (lexically and syntactically conditioned variation)

The other three words have different [+lvar] parameters in this example: [bɛːl] in Table 3.5 has $b_{+\text{lvar},\text{bɛːl}} = -5$, [moʊl] in Table 3.6 has it with $b_{+\text{lvar},\text{moʊl}} = -3$, and [ciːl] in Table 3.7 has [+lvar] with a parameter of 3.

<table>
<tr><td>preposition:
[+core] parameters:</td><td colspan="4">v
$a = 1, b = 5$</td></tr>
<tr><td></td><td colspan="2">Derivation 1
[+core] present</td><td colspan="2">Derivation 2
[+core] absent</td></tr>
<tr><td>output of syntax:</td><td colspan="2">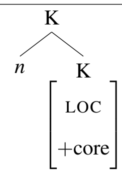</td><td colspan="2">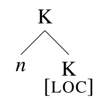</td></tr>
<tr><td>score s_1:</td><td colspan="2">$1 \cdot 0 + 5$</td><td colspan="2">0</td></tr>
<tr><td>noun:
[+lvar] parameters:</td><td colspan="2">bɛːl
$a = 2, b = -5$</td><td colspan="2">bɛːl
$a = 2, b = -5$</td></tr>
<tr><td></td><td>Derivation 1a
[+lvar] present</td><td>Derivation 1b
[+lvar] absent</td><td>Derivation 2a
[+lvar] present</td><td>Derivation 2b
[+lvar] absent</td></tr>
<tr><td>rule of realization:
score s_2:</td><td>(36a)
$2 \cdot 5 - 5$</td><td>(36b)
5</td><td>(36b)
$2 \cdot 0 - 5$</td><td>(36b)
0</td></tr>
<tr><td>output:
final score s:
probability:</td><td>v bɛːl-ɛ
5
$\frac{e^5}{e^5+e^5+e^{-5}+e^0} \approx .498$</td><td>v bɛːl-u
5
$\frac{e^5}{e^5+e^5+e^{-5}+e^0} \approx .498$</td><td>v bɛːl-u
-5
$\frac{e^{-5}}{e^5+e^5+e^{-5}+e^0} \approx .000$</td><td>v bɛːl-u
0
$\frac{e^0}{e^5+e^5+e^{-5}+e^0} \approx .003$</td></tr>
</table>

<table>
<tr><td>preposition:
[+core] parameters:</td><td colspan="4">o
$a = 1, b = 3$</td></tr>
<tr><td></td><td colspan="2">Derivation 1
[+core] present</td><td colspan="2">Derivation 2
[+core] absent</td></tr>
<tr><td>output of syntax:</td><td colspan="2">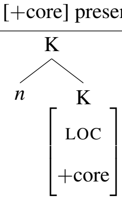</td><td colspan="2">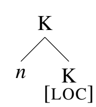</td></tr>
<tr><td>score s_1:</td><td colspan="2">$1 \cdot 0 + 3$</td><td colspan="2">0</td></tr>
<tr><td>noun:
[+lvar] parameters:</td><td colspan="2">bɛːl
$a = 2, b = -5$</td><td colspan="2">bɛːl
$a = 2, b = -5$</td></tr>
<tr><td></td><td>Derivation 1a
[+lvar] present</td><td>Derivation 1b
[+lvar] absent</td><td>Derivation 2a
[+lvar] present</td><td>Derivation 2b
[+lvar] absent</td></tr>
<tr><td>rule of realization:
score s_2:</td><td>(36a)
$2 \cdot 3 - 5$</td><td>(36b)
3</td><td>(36b)
$2 \cdot 0 - 5$</td><td>(36b)
0</td></tr>
<tr><td>output:
final score s:
probability:</td><td>o bɛːl-ɛ
1
$\frac{e^1}{e^1+e^3+e^{-5}+e^0} \approx .114$</td><td>o bɛːl-u
3
$\frac{e^3}{e^1+e^3+e^{-5}+e^0} \approx .844$</td><td>o bɛːl-u
-5
$\frac{e^{-5}}{e^1+e^3+e^{-5}+e^0} \approx .000$</td><td>o bɛːl-u
0
$\frac{e^0}{e^1+e^3+e^{-5}+e^0} \approx .042$</td></tr>
</table>

Table 3.5: Derivations for the word [bɛːl] in Example 3 (lexically and syntactically conditioned variation)

preposition:	v			
[+core] parameters:	$a=1$, $b=5$			
	Derivation 1 [+core] present		Derivation 2 [+core] absent	
output of syntax:	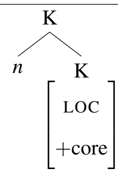		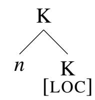	
score s_1:	$1 \cdot 0 + 5$		0	
noun:	moʊl		moʊl	
[+lvar] parameters:	$a=2$, $b=-3$		$a=2$, $b=-3$	
	Derivation 1a [+lvar] present	Derivation 1b [+lvar] absent	Derivation 2a [+lvar] present	Derivation 2b [+lvar] absent
rule of realization:	(36a)	(36b)	(36b)	(36b)
score s_2:	$2 \cdot 5 - 3$	5	$2 \cdot 0 - 3$	0
output:	v moʊl-ɛ	v moʊl-u	v moʊl-u	v moʊl-u
final score s:	7	5	-3	0
probability:	$\frac{e^7}{e^7+e^5+e^{-3}+e^0} \approx .880$	$\frac{e^5}{e^7+e^5+e^{-3}+e^0} \approx .119$	$\frac{e^{-3}}{e^7+e^5+e^{-3}+e^0} \approx .000$	$\frac{e^0}{e^7+e^5+e^{-3}+e^0} \approx .001$

preposition:	o			
[+core] parameters:	$a=1$, $b=3$			
	Derivation 1 [+core] present		Derivation 2 [+core] absent	
output of syntax:	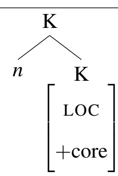		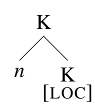	
score s_1:	$1 \cdot 0 + 3$		0	
noun:	moʊl		moʊl	
[+lvar] parameters:	$a=2$, $b=-3$		$a=2$, $b=-3$	
	Derivation 1a [+lvar] present	Derivation 1b [+lvar] absent	Derivation 2a [+lvar] present	Derivation 2b [+lvar] absent
rule of realization:	(36a)	(36b)	(36b)	(36b)
score s_2:	$2 \cdot 3 - 3$	3	$2 \cdot 0 - 3$	0
output:	o moʊl-ɛ	o moʊl-u	o moʊl-u	o moʊl-u
final score s:	3	3	-3	0
probability:	$\frac{e^3}{e^3+e^3+e^{-3}+e^0} \approx .487$	$\frac{e^3}{e^3+e^3+e^{-3}+e^0} \approx .487$	$\frac{e^{-3}}{e^3+e^3+e^{-3}+e^0} \approx .001$	$\frac{e^0}{e^3+e^3+e^{-3}+e^0} \approx .024$

Table 3.6: Derivations for the word [moʊl] in Example 3 (lexically and syntactically conditioned variation)

preposition:	v			
[+core] parameters:	$a = 1, b = 5$			
	Derivation 1		Derivation 2	
	[+core] present		[+core] absent	
output of syntax:	[K n [K [LOC, +core]]]		[K n [K [LOC]]]	
score s_1:	$1 \cdot 0 + 5$		0	
noun:	ciːl		ciːl	
[+lvar] parameters:	$a = 2, b = 3$		$a = 2, b = 3$	
	Derivation 1a	Derivation 1b	Derivation 2a	Derivation 2b
	[+lvar] present	[+lvar] absent	[+lvar] present	[+lvar] absent
rule of realization:	(36a)	(36b)	(36b)	(36b)
score s_2:	$2 \cdot 5 + 3$	5	$2 \cdot 0 + 3$	0
output:	v ciːl-ɛ	v ciːl-u	v ciːl-u	v ciːl-u
final score s:	13	5	3	0
probability:	$\frac{e^{13}}{e^{13}+e^{5}+e^{3}+e^{0}} \approx .9996$	$\frac{e^{5}}{e^{13}+e^{5}+e^{3}+e^{0}} \approx .0003$	$\frac{e^{3}}{e^{13}+e^{5}+e^{3}+e^{0}} \approx .0000$	$\frac{e^{0}}{e^{13}+e^{5}+e^{3}+e^{0}} \approx .0000$

preposition:	o			
[+core] parameters:	$a = 1, b = 3$			
	Derivation 1		Derivation 2	
	[+core] present		[+core] absent	
output of syntax:	[K n [K [LOC, +core]]]		[K n [K [LOC]]]	
score s_1:	$1 \cdot 0 + 3$		0	
noun:	ciːl		ciːl	
[+lvar] parameters:	$a = 2, b = 3$		$a = 2, b = 3$	
	Derivation 1a	Derivation 1b	Derivation 2a	Derivation 2b
	[+lvar] present	[+lvar] absent	[+lvar] present	[+lvar] absent
rule of realization:	(36a)	(36b)	(36b)	(36b)
score s_2:	$2 \cdot 3 + 3$	3	$2 \cdot 0 + 3$	0
output:	o ciːl-ɛ	o ciːl-u	o ciːl-u	o ciːl-u
final score s:	9	3	3	0
probability:	$\frac{e^{9}}{e^{9}+e^{3}+e^{3}+e^{0}} \approx .995$	$\frac{e^{3}}{e^{9}+e^{3}+e^{3}+e^{0}} \approx .002$	$\frac{e^{3}}{e^{9}+e^{3}+e^{3}+e^{0}} \approx .002$	$\frac{e^{0}}{e^{9}+e^{3}+e^{3}+e^{0}} \approx .000$

Table 3.7: Derivations for the word [ciːl] in Example 3 (lexically and syntactically conditioned variation)

3.4.4 General behavior of the variable feature model

The rules in (36) spell out the locative as -*ε* when both [+core] and [+lvar] are in the derivation, and -*u* otherwise. That is, in the derivations above, Derivation 1a yields -*ε*, while the other three derivations yield -*u*. In this section, I describe the general behavior of this model mathematically, including how it operates at the extremes.

Increasing the lexical parameter of either [+core] or [+lvar] in the derivation increases the overall probability of locative -*ε*. For example, [o] assigns [+core] with $b = 3$ and [bɛːl] has [+lvar] with $b = -5$, so the derivation of [o bɛːl-ɛ] has a score of $2(1 \cdot 0 + 3) - 5 = 1$, which converts to a probability of 11.4%. On the other hand [v] has [+core] with $b = 5$, so [v bɛːl-ɛ] has a score of $2(1 \cdot 0 + 5) - 5 = 5$ and a probability of 49.8%. Similarly, the noun [moʊl] has [+lvar] with $b = -3$, so [o moʊlɛ] has a higher score than [o bɛːl-ɛ] (in this case, $2(1 \cdot 0 + 3) - 3 = 3$, which corresponds to a higher probability, 48.7%).

In Example 1 (Table 3.2), with only lexically conditioned variation, a noun with $b = 10$ for [+lvar] behaved nearly categorically, showing locative -*ε* 99.995% of the time. Adding syntactic conditioning, we can get a similar effect, although the parameters are different: here, when [ciːl] has [+lvar] with $b = 3$, that gives us behavior very close to categorical: [v ciːl-ɛ] appears 99.96% of the time, while [o ciːl-ɛ] has 99.5% probability. In fact, as can be clearly seen in Figure 3.8 below, the asymptote of the probability of locative -*ε* is slightly less than 1: as $b_{+\text{lvar},n}$ increases, the probability approaches not 1 but a smaller number, $\frac{1}{1+e^{-a_{+\text{lvar}} b_{+\text{core},n}}} = \frac{1}{1+e^{-2b_{+\text{core},n}}}$ for this derivation, since $a_{+\text{lvar}}$ is fixed at 2. The proof of this asymptotic relationship is presented in (37), where $s(d)$ is the score of derivation d, as in (31) and (32); note that $\lim_{x \to \infty} e^{-x} = 0$:

(37) *Deriving the asymptote of P(locative -ε)*

$$\begin{aligned}
\lim_{b_{+\text{lvar},n}\to\infty} P(\text{locative -}\varepsilon) &= \lim_{b_{+\text{lvar},n}\to\infty} P(\text{Deriv. 1a}) \\
&= \lim_{b_{+\text{lvar},n}\to\infty} \frac{e^{s(\text{Deriv. 1a})}}{\sum\limits_{d_i\in D} e^{s(d_i)}} \\
&= \lim_{b_{+\text{lvar},n}\to\infty} \frac{e^{s(\text{Deriv. 1a})}}{e^{s(\text{Deriv. 1a})}+e^{s(\text{Deriv. 1b})}+e^{s(\text{Deriv. 2a})}+e^{s(\text{Deriv. 2b})}} \\
&= \lim_{b_{+\text{lvar},n}\to\infty} \frac{e^{a_{+\text{lvar}}b_{+\text{core},r}+b_{+\text{lvar},n}}}{e^{a_{+\text{lvar}}b_{+\text{core},r}+b_{+\text{lvar},n}}+e^{b_{+\text{core},r}}+e^{b_{+\text{lvar},n}}+e^{0}} \cdot \frac{e^{-b_{+\text{lvar},n}}}{e^{-b_{+\text{lvar},n}}} \\
&= \lim_{b_{+\text{lvar},n}\to\infty} \frac{e^{a_{+\text{lvar}}b_{+\text{core},r}}}{e^{a_{+\text{lvar}}b_{+\text{core},r}}+e^{b_{+\text{core},r}-b_{+\text{lvar},n}}+1+e^{-b_{+\text{lvar},n}}} \\
&= \frac{e^{a_{+\text{lvar}}b_{+\text{core},r}}}{e^{a_{+\text{lvar}}b_{+\text{core},r}}+0+1+0} \cdot \frac{e^{-a_{+\text{lvar}}b_{+\text{core},r}}}{e^{-a_{+\text{lvar}}b_{+\text{core},r}}} \\
&= \frac{1}{1+e^{-a_{+\text{lvar}}b_{+\text{core},r}}}
\end{aligned}$$

The relationship between [+core] and [+lvar] parameters is shown in Figure 3.8. Each of the curves in Figure 3.8 shows how the probability of locative -ε changes with the [+lvar] weight of a noun *for a preposition with a given [+core] weight*. That is, each curve represents the probability given a set value for $b_{+\text{core},r}$ and different values for $b_{+\text{lvar},n}$. For example, when $b_{+\text{core},r} = 0$, the probability of locative -ε given $b_{+\text{lvar},n}$ is a sigmoid curve with asymptotes at $P = 0$ and $P = .5$. That is, in this syntactic context, a noun with [+lvar] parameter of 0 would take locative -ε 25% of the time; increasing or decreasing the strength brings this probability closer to 50% and 0, respectively.

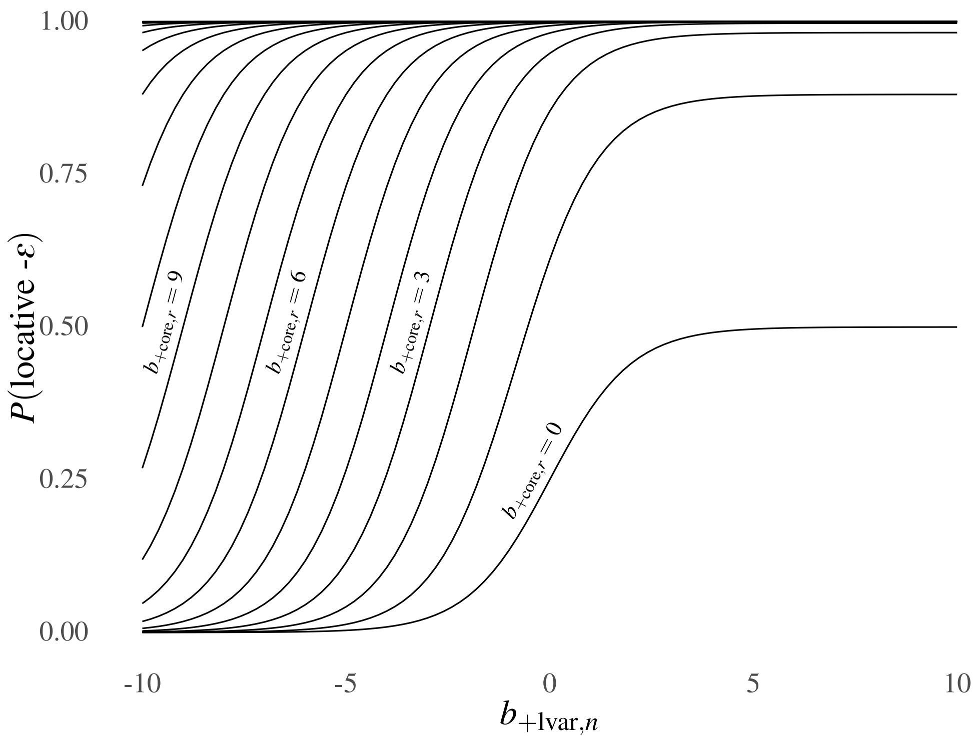

Figure 3.8: The probability of locative -ε given $b_{+\text{lvar},n}$, for non-negative values of $b_{+\text{core},r}$ at intervals of 1 with fixed $a_{+\text{lvar}} = 2$

When $b_{+\text{core},r}$ is sufficiently high for a preposition r, the asymptote of the sigmoid curve is close enough to 1 to get the full range of lexical and syntactic conditioning and the curves become regular:[3]

- For a given syntactic context, nouns appear at different points on that context's sigmoid probability curve
- For a given noun, changing its syntactic context shifts the curve horizontally at a set rate
- Nouns with sufficiently high [+lvar] parameters show categorical behavior in all syntactic contexts

[3]When $a_{+\text{lvar}} = 1$, the curves do not regularize nicely in the same way, meaning that we can have syntactic conditioning or categorical behavior but not both. This is the motivating factor for a parameters in the first place, rather than having all scores be additive.

How can this model account for a case like the Czech locative, where a noun can range from 0% -ε to 100% -ε with everything in between (cf. Guzmán Naranjo & Bonami, 2021)? This is possible as long as all syntactic contexts have a sufficiently high [+core] parameter, perhaps greater than 4 or 5. Accordingly, most variable nouns would have a strongly negative [+lvar] parameter, ranging between approximately 0 and -15.

3.4.5 Assigning feature weights to nonce words

So far in this section, I have taken a theoretical construct used for lexical indexation of categorical behavior—diacritic features—and extended it to account for variability, by replacing binary features with variable, weighted features. The primary theoretical proposal of this dissertation is an extension of the sublexicon model for assigning diacritic features to unfamiliar words. In the sublexicon model, features are associated with sublexical grammars that assign scores to potential members, and a word is more likely to be assigned a feature whose grammar gives it a high score. As described above, this model was initially developed to assign categorical features to new lexical items. The existence of variable lexical items thus represents a puzzle: how can sublexical grammars describe weighted features, and even assign feature weights to nonce words? In this section, I show how to adapt the sublexicon model to feature weights: the scores produced by a sublexical grammar can be treated as centers of a probability distribution used to assign weights to new lexical items. Much of this section involves technical issues of implementation, including mathematical derivations.

3.4.5.1 Categorical sublexicons

Let us recall the basic sublexicon model for categorical features (Allen & Becker, 2015; Gouskova et al., 2015), as applied to Hungarian in Section 3.1.2. In Hungarian, nouns can take either -*V* or -*jV* in the possessive. This is indexed to a lexical diacritic feature, [±j]: nouns that take -*V* have a [−j] feature on their lexical entry, while nouns that take -*jV* are decorated with a [+j] feature.

In order to form the possessive of a novel form, speakers must first assign it a [±j] feature. Each feature is associated with a *sublexical grammar* that captures generalizations over lexical entries that share that feature. That is, the [+j] grammar scores stereotypical [+j] words well and unlikely [+j] words poorly. These grammars comprise a set of weighted constraints. For example, nouns that end in sibilants and palatals categorically take *-V*, while nouns ending in vowels categorically take *-jV*. Thus, the [+j] sublexical grammar contains strong constraints against word-final palatals and sibilants (ensuring that such words are very unlikely to be assigned [+j]), while the [−j] grammar has a strong constraint against word-final vowels. In the example in Figure 3.9 and Figure 3.10, repeated from Figure 3.1 and Figure 3.2, these constraints all have a weight of 5.

I also encoded correlations between inflection classes into the sublexicon model. For example, there is a small class of nouns called “lowering stems”, which exceptionally take *-ɒk* in the plural instead of the usual *-ok*. I mark these with a [lower] feature. These nouns are also more likely to take *-V* in the possessive (see Chapter 4). This preference is encoded as a constraint against the [lower] feature active in the [+j] sublexical grammar, with a more moderate weight of 2 in the example. Since lowering stems are quite rare overall, they are also poorly represented among [−j] words. I thus also include a *[lower] constraint in the [−j] sublexical grammar in Figure 3.10, albeit with a weaker weight of 1: the penalty should not be as strong as it in the [+j] grammar.

The examples in Figure 3.9 and Figure 3.10 show two nonce words, [ruɲɒs] and [fuːzaːt], being evaluated on the two sublexical grammars. In this case, both words have *-ɒk* in the plural: [ruɲɒs-ɒk] and [fuːzaːt-ɒk]. Thus, both already have a [lower] feature on their lexical entry. To make the following discussion easier, I present constraint weights here as *negative*, so that the weight of a constraint is assessed to words that violate it. This is the opposite of how I presented weights in Figure 3.3 and Figure 3.4—that is, these tableaux are identical to those earlier ones, but the sign of the weight is switched from positive to negative. This change, which I make for expositional reasons, reverses the normal convention that constraint weights are positive but violations—and thus scores—are negative.

constraint *weight*	*[+strident]# −5	*[+palatal]# −5	*[lower] −2	total
ruɲɒs$_{[lower]}$	−5	0	−2	−7
fuːzaːt$_{[lower]}$	0	0	−2	−2

Figure 3.9: Evaluation of nonce lowering stems *runyasz* and *fúzát* on the [+j] sublexical grammar with *[lower]

constraint *weight*	*[+syllabic]# −5	*[lower] −1	total
ruɲɒs$_{[lower]}$	0	−1	−1
fuːzaːt$_{[lower]}$	0	−1	−1

Figure 3.10: Evaluation of nonce lowering stems *runyasz* and *fúzát* on the [−j] sublexical grammar with *[lower]

Both words have a better score on the [−j] sublexical grammar, so they are both more likely to be given a [−j] feature. This is a maximum entropy model (Goldwater & Johnson, 2003; Hayes & Wilson, 2008), so a word's likelihood of being assigned a feature is proportional to its score raised to the power of *e*. For example, since /fuːzaːt$_{\text{lower}}$/ has a score of −2 on the [+j] grammar and −1 on the [−j] grammar, it will be assigned a [+j] feature $\frac{e^{-1}}{e^{-1}+e^{-2}} = .731 = 73.1\%$ of the time.

3.4.5.2 For binary features, one sublexical grammar is enough

If there are only two possible outcomes (that is, a single binary feature with two sublexical grammars), the distribution of outcomes for a word is dependent only on the difference between the two scores, not their absolute value. (I discuss how to extend the sublexical model in cases with more than two outcomes in Section 6.2.1.) I derive this mathematically as follows: Suppose we have a binary feature with values f_1 and f_2, and let s_1 and s_2 be the scores assigned to a word by the sublexical grammars for these two features respectively. The probability $P(f_1)$ of assigning f_1 to the word is a function of $s_1 - s_2$, as seen in (38).

(38) *Deriving the $P(f_1)$ as the difference between s_1 and s_2*

$$\begin{aligned} P(f_1) &= \frac{e^{s_1}}{e^{s_1} + e^{s_2}} \cdot \frac{e^{-s_2}}{e^{-s_2}} \\ &= \frac{e^{s_1 - s_2}}{e^{s_1 - s_2} + e^{s_2 - s_2}} \\ &= \frac{e^{s_1 - s_2}}{e^{s_1 - s_2} + e^{0}} \\ &= \frac{e^{s_1 - s_2}}{e^{s_1 - s_2} + 1} \end{aligned}$$

Of course, since there are two outcomes, the probabilities of both outcomes should add up to 1 in all cases, so the other outcome, of assigning f_2 to this word, should likewise depend only on the difference between s_1 and s_2, and also be equal to $1 - P(f_1)$. This is in fact the case:

(39) *Deriving $P(f_2)$ as $1-P(f_1)$*

$$\begin{aligned}
P(f_2) &= \frac{e^{s_2}}{e^{s_1}+e^{s_2}} \cdot \frac{e^{-s_1}}{e^{-s_1}} \\
&= \frac{e^{s_2-s_1}}{e^{s_1-s_1}+e^{s_2-s_1}} \\
&= \frac{e^{s_2-s_1}}{1+e^{s_2-s_1}} \\
&= \frac{e^{s_2-s_1}}{1+e^{s_2-s_1}} + 1 - \frac{1+e^{s_2-s_1}}{1+e^{s_2-s_1}} \\
&= 1 + \frac{e^{s_2-s_1}}{1+e^{s_2-s_1}} - \frac{1+e^{s_2-s_1}}{1+e^{s_2-s_1}} \\
&= 1 + \frac{e^{s_2-s_1} - 1 - e^{s_2-s_1}}{1+e^{s_2-s_1}} \\
&= 1 + \frac{-1}{1+e^{s_2-s_1}} \\
&= 1 - \frac{1}{1+e^{s_2-s_1}} \cdot \frac{1+e^{s_1-s_2}}{1+e^{s_1-s_2}} \\
&= 1 - \frac{e^{s_1-s_2}}{e^{s_1-s_2}+e^{s_2-s_1+s_1-s_2}} \\
&= 1 - \frac{e^{s_1-s_2}}{e^{s_1-s_2}+1} \\
&= 1 - P(f_1)
\end{aligned}$$

A word's score on a grammar is the sum of the weights of the constraints it violates. Thus, the effect of each constraint depends on the difference of its weights in the two sublexicons. For example, *[lower] has a weight of -2 in the [+j] grammar and -1 in the [−j] grammar, but the results would be equivalent if these scores were -3 and -2, or -15 and -14, or -1 and 0, so long as its weight for [+j] is 1 lower than its weight for [−j].

We can take this equivalence to its extreme: if a constraint c has weight w_1 in the [+j] grammar and w_2 in the [−j] grammar, this is equivalent to c having a weight of $w_1 - w_2$ in the [+j] grammar and 0 in the [−j] grammar. If we do this for every constraint, then the [−j] grammar is inert: every constraint has a weight of 0. However, some of the weights in the [+j] grammar may now be positive: if w_2 is lower (that is, more heavily negatively weighted) than w_1, then $w_1 - w_2 > 0$. This is a bit odd, since now candidates can be *rewarded* for violating constraints. If we accept this, then all we need is a single [±j] grammar with all of the constraints. We can see how this grammar evaluates our two nonce words in Figure 3.11. The first two constraints, against word-final sibilants and palatals, were only in the [+j] grammar and remain unchanged. The next, against word-final vowels, was only in the [−j] grammar, so its sign is reversed: words that end in vowels are now *rewarded* 5 on their score by this grammar. Finally, the constraint against lowering stems was in both sublexical grammars, so its score is the difference of the two. Since it was stronger in the [+j] grammar, its sign is still negative.

constraint / *weight*	*[+strident]#	*[+palatal]#	*[+syllabic]#	*[lower]	total
	−5	−5	5	−1	
ruɲɒs$_{\text{[lower]}}$	−5	0	0	−1	−6
fuːzaːt$_{\text{[lower]}}$	0	0	0	−1	−1

Figure 3.11: Evaluation of nonce lowering stems *runyasz* and *fúzát* on the combined [±j] sublexical grammar with *[lower]

This grammar puts out the same probability of assigning [+j] as before. For example, *fúzát* has a score of −1, so its probability of getting [+j] is $\frac{e^{-1}}{1+e^{-1}} = .731 = 73.1\%$. (The 1 in the denominator is derived from the fact that the score on the empty [−j] grammar is always 0, and $e^0 = 1$.)

Once we manipulate the math in this way, sublexical grammars are identical to *logistic models* (see Hayes & Wilson, 2008; Potts et al., 2010). In a logistic model, we predict a binary outcome by summing up weighted predictors—equivalent to the weighted constraints of our sublexical grammar. How do these predictors correspond to a binary outcome? By predicting the *probability* that a given trial will fall into one of the outcomes. Accordingly, the sum of the predictors, s, is converted into a probability P of outcome using the logistic function $P = \frac{e^s}{e^s+1}$, which is the formula derived for the sublexical grammars in (38). This transformation is used because it converts the range from 0 to 1—the range of probabilities—into a continuous function from negative infinity to positive infinity. This eliminates the potential issue of breaking the model by adding up predictors to a sum greater than 1 or less than 0 (James et al., 2013, pp. 129–134).

3.4.5.3 Extending to gradient features

Let us consider the new model: instead of a binary feature, we have a single weighted feature, and the goal of the sublexical grammar is to assign a weight to new words. As before, we wish for the sublexical grammars to stochastically assign a feature to novel lexical items that then determines the new words' behavior. Here, the sublexical grammars already output a score: the sum of the weights of violated constraints. The simplest option, then, would be to say that this score *is* the weight assigned to the word: if our single sublexical grammar assigns a score of -2 to a word, that word gets assigned a feature weight of -2. This is equivalent to a *linear model*, which predicts a continuous variable as the sum of weighted predictors. In fact, a logistic model is mathematically similar to a linear model: probabilities are converted to real numbers with the logistic function, and the model's fit is linear with respect to the transformed probabilities.

The issue with this proposed method of assigning feature weights to new words is that it is *deterministic*: a word's weight can be fully predicted from its score on the sublexical grammar. (This does not mean that the *output* on a given token is deterministic: recall that the feature weight determines the distribution of allomorphs for variable words.) By contrast, in the original sublexical model with binary features, feature assignment is *non-deterministic*: a word's score determined the likelihood that it would be assigned a feature, which would never be 100% (though it might be very, very close). By hypothesis, I would like the feature assignment process to remain stochastic (noisy), so that a given speaker may assign different weights—and thus output different distributions of allomorphs—to words with the same score. To do this, we must introduce some randomness: rather than the score fully determining the assigned weight, the score should instead be the *most likely* assigned weight.

We can do this by assuming that the score is the center of a normal distribution of probabilities:

the likelihood of a word being assigned a particular feature weight is dependent on the weight's proximity to the score, so that a word is most likely to be assigned a weight close to its score and will only rarely receive a more distant weight. This approach is very similar to that of noisy Harmonic Grammar (Boersma & Pater, 2016), in which a grammar's constraint weights in any given derivation are subject to noise via a Gaussian (that is, normally distributed) random variable.

The distribution of scores described is shown in Figure 3.12, which plots distance from the score *s* on the x-axis and the density of probability of assigning that feature weight to the word on the y-axis. Since this is a continuous function defined over all real numbers, the y-axis technically does not represent the actual probability of assigning any value (since there are infinitely many real numbers), but this is roughly what is being shown: the higher the curve at a weight on the x-axis, the higher the probability that that weight will be chosen.

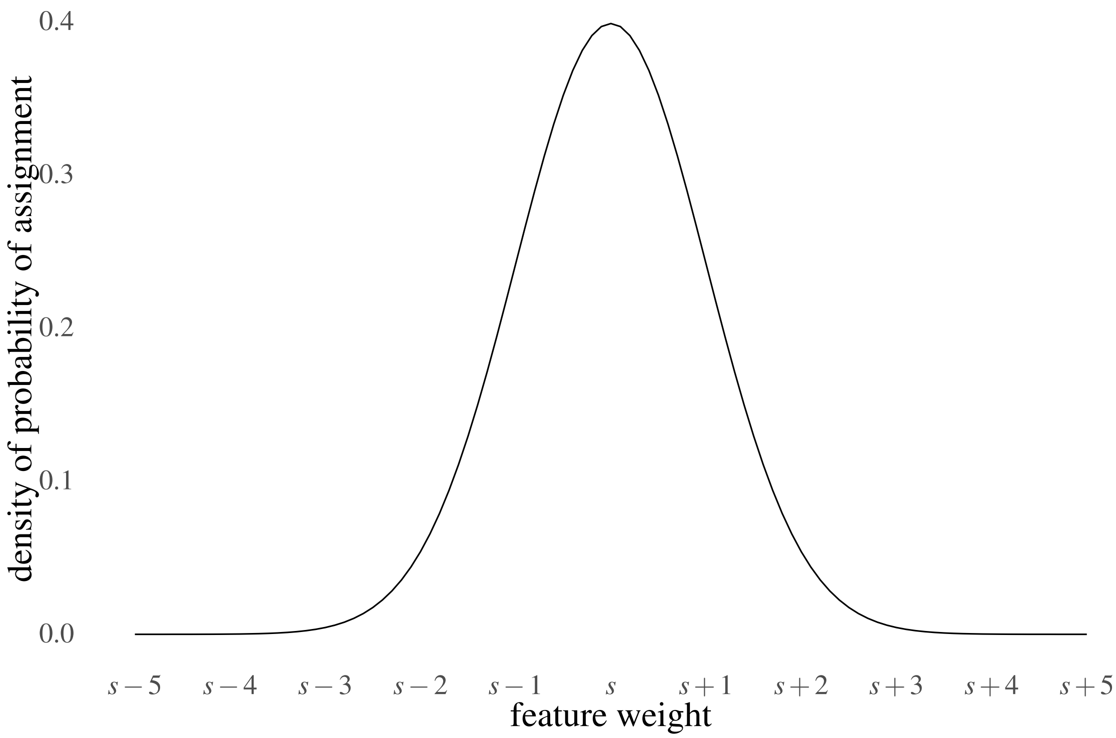

Figure 3.12: The probability of assigning a given feature weight to a word that gets a score of *s* on that feature's sublexical grammar, according to a normal distribution with standard deviation 1

The area under the curve in Figure 3.12, which represents the total probability across all weights, is equal to 1. However, this is not the only such normal distribution centered around *s*. We have one free parameter: the standard deviation, represented by σ, which governs the steepness of the curve. Figure 3.13 shows several such distributions with different standard deviations (.5, 1, 2, and 4). In each case, the area under this curve (that is, the total probability) is equal to 1. For lower standard deviations, this probability is more heavily concentrated in points close to *s*—that is, the resulting weight assigned is likely to be quite close to *s*, making weight assignment less noisy. As the standard deviation gets higher, the distribution is flatter, meaning that the chosen weight is more likely to be further away from *s* and the assignment process is noisier.

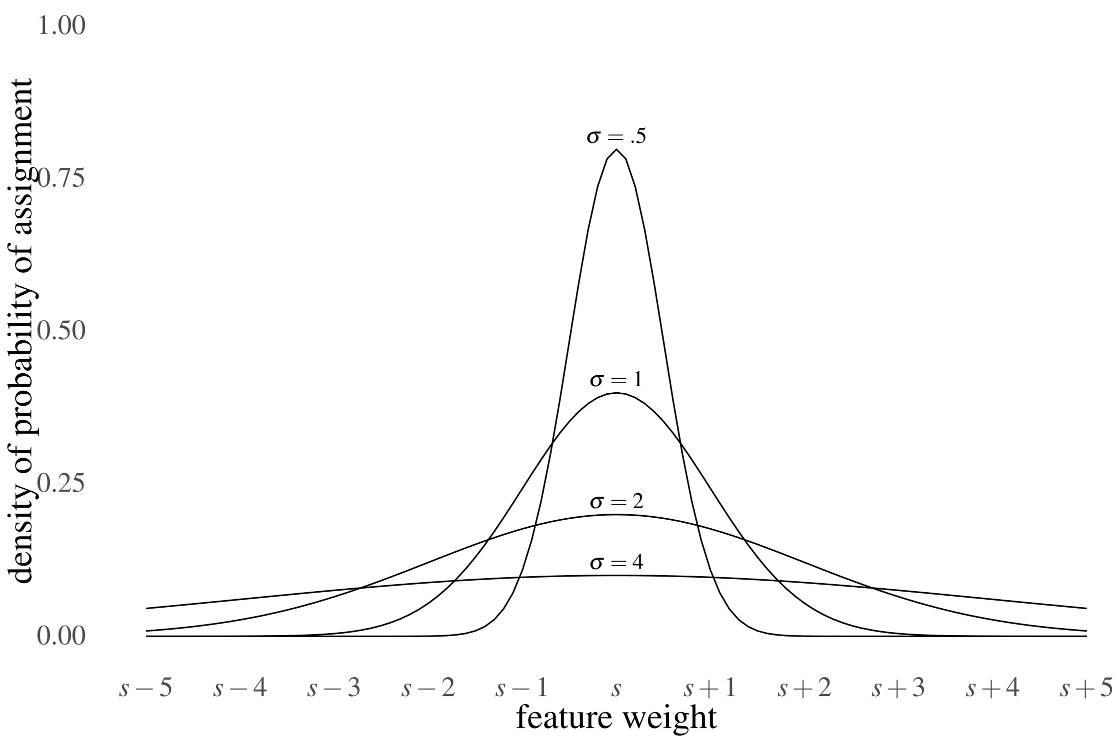

Figure 3.13: Normal distributions centered at *s* with standard deviations (σ) of .5, 1, 2, and 4

I have no principled way of eliminating this parameter, and for now, at least, I assume that it is set (by the linguist) empirically. This leaves open the possibility that each speaker has a different σ, or even a different σ for each feature. One possibility is that this parameter can be derived empirically from the distribution of forms: the better the grammar accounts for the distribution of feature weights (that is, variable rates) in the lexicon, the less noisy feature weight assignment should be.

To demonstrate how the σ parameter might be derived from patterns in the lexicon, let us consider two cases of allomorphy in languages A and B. In language A—which resembles the case of the Hungarian possessive in Chapter 4—the distribution of allomorphs is largely predicted by its phonology and morphology. In this case, the sublexical grammar is a very good fit to the lexicon, and speakers should thus have a very good sense of what the behavior of a novel word "should" be. In this case, we might expect less noisiness in the assignment of feature weights (a lower standard deviation); the weights assigned to nonce words should hew close to their scores from the sublexical grammar. By contrast, in language B (like the Czech locative in Chapter 5), the distribution of allomorphs among words is more random, and a word's phonological and morphological profile is less predictive. Here the sublexical grammar, while better than nothing, is not a great fit to the lexicon, and speakers should expect greater variability in how a nonce word will behave. In this case, nonce words should be assigned feature weights with less certainty and more noise—that is, a higher standard deviation.

This hypothesis for setting the standard deviation is, in theory, testable: if true, then the lexicon grammar of language A, which is a better fit to the distribution of allomorphs, should also be a better predictor for a wug test on language A. The sublexical grammar for language B, which is a worse fit, should accordingly be worse at predicting experimental results. However, we must be careful in testing it. There are two reasons why a linguist's model of the lexicon may not be a very good fit for the results of a given experiment. The first is that the model excludes relevant factors. In this case, we would expect the model to be as poor a fit to the experimental results as it is to the lexicon, since presumably speakers are also using these omitted factors. The other possibility is that the linguist's model includes all relevant factors, but that these factors are simply not very predictive. In this case, speakers may nonetheless choose to apply them strongly, which would lead to the lexicon model being a good fit to the experiment results (and disconfirm the hypothesis). Thus, in order to test the hypothesis that speakers apply better models more strongly than worse models, we must be sure that we are capturing all relevant factors, and that there is no model that would provide a better description of the lexicon.

3.4.5.4 Learning a sublexical grammar for gradient features

The basic model of a sublexical grammar for variable, gradient features laid out in this section is quite similar to the original sublexicon model for categorical features: the sublexical grammar evaluates new words and outputs a score. Once the grammar is in place, it does not really matter what happens after the score is generated; the core mechanism is the same. However, the process of training the sublexical grammar in the first place must be modified as well. In the original conception of the sublexicon model (Becker & Gouskova, 2016; Gouskova et al., 2015), inherited from the phonotactic grammars of Hayes and Wilson (2008), the learner induces constraints based on sequences of segments or natural classes that are underrepresented among members of the sublexicon and weights them accordingly. In the variable model, however, *every* word contains the feature, and the point of the grammar is to predict each word's weight. Thus, the method of constraint induction proposed in Hayes and Wilson (2008) would not work well for variable features. However, this method of constraint induction is limited in its ability to create human-like grammars anyway (Hayes & White, 2013), so here, as elsewhere, I set the problem of constraint induction aside and focus on constraint weighting.

In the categorical sublexicon model, constraints are weighted in a sublexical grammar according to how often they are violated by its members. Technically, the grammar is only being trained on words that are in the sublexicon; the set of words used as negative evidence is randomly generated, as described in Section 3.1.2. If anything, this makes the learning process *more difficult*, since we are lacking negative evidence in the form of *real* words that could be in the sublexicon but are not. If, as described above, we recast sublexical grammars of a binary feature as a single sublexical grammar whose constraints can have both positive and negative weights, then the sublexicon grammar can be trained on words that take a given feature value and words that take the opposite value. At this point, we have a logistic model, and given a set of constraints (or factors), a logistic model can be fitted using logistic regression. Similarly, learning a sublexical grammar for assigning variable feature weights is a process of fitting a linear model with constraints (or factors) by a linear regression. The two processes are equally difficult. Thus, learning a sublexical grammar for weighted features is no more difficult than learning one for categorical binary features.

3.4.6 Summary

The baseline theoretical proposals assumed in this dissertation—Distributed Morphology and sublexical grammars—are not designed to handle variable lexical items, especially if these are also subject to syntactically conditioned variation as is seen in the Czech locative (discussed at length in Chapter 5). To account theoretically for this sort of variation, in this section I developed an account of variation dependent on variable feature weights. In order for the sublexicon model to be able to extract generalizations over weighted features and assign them to new lexical items, we must change the way its sublexical grammars are structured. Under the extended sublexicon model that I proposed to account for these variable lexical items, the grammar introduces two points of variation when a speaker wishes to generate a novel form of a nonce word. First, the grammar stochastically assigns the word a weight for the relevant feature according to the word's evaluation on that feature's sublexical grammar; next, when the grammar needs to actually generate an inflected form, it stochastically chooses one based on the newly assigned feature weight.

This section is largely theoretical in nature, because the precise predictions of its architectural choices are relatively difficult to test. In Section 5.4.6.3, I discuss how my wug tests are rather blunt instruments that cannot distinguish between this two-layer model (with multiple grammars) and a simpler one with only one point of random variation (the variable extension of the single grammar model described in Section 3.3). Indeed, in analyzing the results of the Czech wug test in Section 5.4, I group words together categorically, since I cannot really dig into the subtleties of words with different rates of variability. However, I argue that my corpus study of how individual authors produce variable words in Czech (Section 5.6) may provide tentative evidence for my two-stage model over a simpler one.

4 Hungarian lowering stems and possessive allomorphy

The first case study of this dissertation looks at factors influencing the distribution of the two allomorphs of the possessive suffix in Hungarian, *-V* (*-ɒ* or *-ɛ*, depending on vowel harmony) and *-jV* (*-jɒ* or *-jɛ*). In particular, one class of irregular nouns, known as lowering stems, shows a bias for the *-V* allomorph in the lexicon, as I show in a corpus study. In a nonce word study, I show that speakers have learned and apply both phonological tendencies in the distribution of the possessive and the morphological correlation: lowering stems are assigned *-V* more frequently than non-lowering stems.

I begin with this case study for several reasons. First, the distribution of allomorphs is relatively simple and shows an active mix of interlocking phonological and morphological effects, some variable and some categorical. Second of all, this is a case where a dependency between two features is needed to explain the phenomenon: as described in Section 2.3.2.2, lowering stems and possessives cannot be unified under a single inflection class diacritic. Finally, this case study contributes empirical evidence for a case of a morphological dependency that has been well-described in the literature on Hungarian, the psychological reality of which has not previously been demonstrated experimentally.

4.1 Background

4.1.1 Basic outline

In Section 3.1, I presented an outline of the Hungarian plural and possessive (also described, in sligthly different terms, in Section 2.3.2.2), which I repeat here. As mentioned above, the possessive has two basic allomorphs, *-V* and *-jV*, both of which are very frequent. In the case of the plural, most nouns have *-ok*, while a small class called "lowering stems" instead takes *-ɒk*. In Table 4.1, we see that all four combinations of plural and possessive are possible.

noun	dɒl	tʃont	vaːlː	hold
gloss	'song'	'bone'	'shoulder'	'moon'
plural	dɒl-o̱k	tʃont-o̱k	vaːlː-ɒ̱k	hold-ɒ̱k
possessive	dɒl-ɒ	tʃont-**j**ɒ	vaːlː-ɒ	hold-**j**ɒ

Table 4.1: Possible combinations of Hungarian plural and possessive suffixes

A noun's possessive is sometimes fully predictable from its phonological form—for example, nouns ending in palatals and sibilants always take *-V*. In other cases, the phonology provides a clue, but does not predict a noun's possessive with full reliability. I discuss these and other phonological factors predicting possessive allomorphy in the corpus study in Section 4.3. In addition to the phonological factors, the possessive observes a morphological correlation: most lowering stems take *-V*. This correlation is the main object of my study. Before presenting the studies, I discuss more details about the possessive and lowering stems.

4.1.2 Vowel harmony alternations

Hungarian words have either back or front harmony, and suffix vowels alternate accordingly. The mid suffixes also show rounding harmony: front-harmonizing suffixes with mid vowels have rounded and unrounded variants to match the last vowel of the stem. These alternations, for short vowels, are shown in Table 4.2; see Siptár and Törkenczy (2000, pp. 63–73) for more details.

height	*back*	*front*		*example suffix*			*example words*		
		rounded	*unrounded*				haːz 'house'	føld 'land'	kɛrt 'garden'
high	u	y		*-unk/-ynk*	1PL	'our'	haːz-unk	føld-ynk	kɛrt-ynk
mid	o	ø	ɛ	*-hoz/-høz/-hɛz*	ALL	'to'	haːz-hoz	føld-høz	kɛrt-hɛz
low	ɒ	ɛ		*-bɒn/-bɛn*	INESS	'in'	haːz-bɒn	føld-bɛn	kɛrt-bɛn

Table 4.2: Vowel harmony alternations for Hungarian suffixes (from Siptár & Törkenczy, 2000, p. 65)

Examples in this chapter have back harmony. This chart can be used to find the front-harmonizing version of each suffix. Thus, the front-harmonizing equivalents of possessive *-ɒ* and *-jɒ* are *-ɛ* and *-jɛ*. Regular-stem plural *-ok* (from the mid vowel set) has two front-harmonizing variants, depending on rounding, *-øk* and *ɛk*, while the lowering stem plural *-ɒk* (from the low vowel set) only has one front-harmonizing variant, *-ɛk*. Words with front unrounded harmony can only have plural *-ɛk* and thus cannot be distinguished on the surface as lowering stems. Siptár and Törkenczy (2000, p. 225) nonetheless mark some nouns with front unrounded harmony as lowering stems on the basis of other properties that correlate (more or less reliably) with lowering stem status. Since this difference is not marked in my corpus and cannot be entirely inferred, I assume that all words with front unrounded harmony are undetermined for stem class. In the nonce word experiment (Section 4.4), I treat stimuli with front unrounded harmony as fillers.

A stem's harmony class is usually but not always predictable from its vowels (Hayes & Londe, 2006; Hayes et al., 2009; Rebrus et al., 2012, 2019; Siptár & Törkenczy, 2000)—thus, some nouns must be explicitly marked for harmony class, as discussed in Section 2.2.2.4. I assume that vowel harmony is handled *in the phonology proper*: *-ɒ* and *-ɛ* are surface variants of a single underlying form inserted in the context of [−j], and words taking *-V* are marked with a unified [−j] feature. Likewise, [+j] marks words taking *-jɒ* or *-jɛ*.

4.1.3 The possessive

4.1.3.1 Morphosyntactic details

The full paradigm of possessives for the four words in Table 4.1 are shown in Table 4.3 (see Rounds, 2008, pp. 135–137). Hungarian distinguishes between the person and number of possessors, as well as the number of the possessed noun, so [dɒl-om] means 'my song', while [dɒl-ɒ-i-m] means 'my songs', and so on.

noun	dɒl	tʃont	vaːlː	hold
gloss	‘song’	‘bone’	‘shoulder’	‘moon’
possessor	*singular noun*			
1SG	dɒlom	tʃontom	vaːlːɒm	holdɒm
2SG	dɒlod	tʃontod	vaːlːɒd	holdɒd
3SG	dɒlɒ	tʃont**j**ɒ	vaːlːɒ	hold**j**ɒ
1PL	dɒlunk	tʃontunk	vaːlːunk	holdunk
2PL	dɒlotok	tʃontotok	vaːlːɒtok	holdɒtok
3PL	dɒluk	tʃont**j**uk	vaːlːuk	hold**j**uk
possessor	*plural noun*			
1SG	dɒlɒim	tʃont**j**ɒim	vaːlːɒim	hold**j**ɒim
2SG	dɒlɒid	tʃont**j**ɒid	vaːlːɒid	hold**j**ɒid
3SG	dɒlɒi	tʃont**j**ɒi	vaːlːɒi	hold**j**ɒi
1PL	dɒlɒink	tʃont**j**ɒink	vaːlːɒink	hold**j**ɒink
2PL	dɒlɒitok	tʃont**j**ɒitok	vaːlːɒitok	hold**j**ɒitok
3PL	dɒlɒik	tʃont**j**ɒik	vaːlːɒik	hold**j**ɒik

Table 4.3: Hungarian possessive paradigms for some back-harmonizing words

There are two main points of variation among these paradigms. The first is the alternation between [o] and [ɒ] (underlined in Table 4.3) in the 1SG, 2SG, and 2PL singular. This is the same lowering stem alternation as in the plural, and will be addressed in Section 4.2.2. The second is the variable presence of [j] (bolded in Table 4.3) in singular nouns with 3SG and 3PL possessors and plural nouns with all possessors.[1] This is the possessive morpheme, with allomorphs, *-V* and *-jV*. Its vowel deletes before 3PL *-uk*.

[1]In general, [j] is either present or absent throughout the paradigm. One very rare exception is [bɒraːt] ‘friend’, which takes *-jV* in the singular ([bɒraːt-jɒ] ‘her friend’) and *-V* in the plural ([bɒraːt-ɒ-i] ‘her friends’).

Under the standard syntactic analysis (cf. Bartos, 1999; Dékány, 2018; É. Kiss, 2002), *-V* and *-jV* are realizations of a Poss head, which has a zero allomorph when adjacent to a first- or second-person possessor marker. (The marker for third singular possessors is null.) Thus, I gloss *-V* and *-jV* as POSS (not 3SG), while *-(V)m*, *-(V)d*, etc. mark 1SG, 2SG, and so on.

4.1.4 Lowering stems beyond the plural

In Table 4.3, we see that the 1SG, 2SG, and 2PL possessor markers undergo the same lowering stem alternation as the plural. I repeat these possessive forms and the plural of the regular stem [dɒl] ‘song’ and the lowering stem [vaːlː] ‘shoulder’ here, also including forms of [kɒpu] ‘gate’. The non-possessed plural and the three possessive markers show the same pattern. The suffix is a bare consonant (or *-tok*) after a vowel, including the vowel-final noun [kɒpu] and the possessed plural *-i*. Otherwise, this consonant is preceded by a “linking vowel” (cf. Siptár & Törkenczy, 2000, p. 219), which is mid after [dɒl] and low after [vaːlː] (see Table 4.2 above).

noun	dɒl		vaːlː		kɒpu	
gloss	‘song’		‘shoulder’		‘gate’	
possessor	*singular*	*plural*	*singular*	*plural*	*singular*	*plural*
none	dɒl	dɒlo̲k	vaːlː	vaːlːɒ̲k	kɒpu	kɒpuk
1SG	dɒlo̲m	dɒlɒim	vaːlːɒ̲m	vaːlːɒim	kɒpum	kɒpuim
2SG	dɒlo̲d	dɒlɒid	vaːlːɒ̲d	vaːlːɒid	kɒpud	kɒpuid
2PL	dɒlo̲tok	dɒlɒitok	vaːlːɒ̲tok	vaːlːɒitok	kɒputok	kɒpuitok

Table 4.4: Lowering stem alternations in the plural and possessive markers

This suggests that the naïve analysis presented in Section 3.1.1, in which *-ok* and *-ɒk* are contextually dependent allomorphs of the plural, is not the right approach. In the next section, I address this and other representational questions.

4.2 Formal analysis

In Section 3.1.1, I presented basic rules of realization, in (28), for the Hungarian possessive and plural, repeated here:

(40) *Rules of realization for the Hungarian plural and possessive (simplified)*

a.	PL ↔ ɒk / [lower] ___	c.	POSS ↔ jɒ / [+j] ___
b.	PL ↔ ok /	d.	POSS ↔ ɒ / [−j] ___

In this section, I discuss my representational choices and develop a more nuanced (though still quite basic) analysis of lowering.

4.2.1 Marked and default allomorphs

In (40), I assume that both *-V* and *-jV* are marked and neither is a default. Although Prasada and Pinker (1993), Marcus et al. (1995), and Yang (2016) assume that productive default rules drive language acquisition and usage, there are cases of lexical variation where children fail to form a productive rule, memorizing each word individually (Dąbrowska, 2001; Schuler et al., 2021), as discussed in Section 2.2.2.2. If the Hungarian possessive is one such case, it would yield the rules in (28). Rácz and Rebrus (2012), however, argue that *-jV* is a productive default: it is becoming more frequent (Rounds, 2008) and is used for most (Kiefer, 1985; Rebrus et al., 2017), if not all (Rácz & Rebrus, 2012), recent loanwords that do not end in sibilants or palatals. My experimental results do not support this claim: as in previous nonce word studies (e.g. Gouskova et al., 2015), participants use *-V* and *-jV* for the same nonce words and do not treat *-jV* as a default (see Section 4.4.7).

Thus, I keep the assumption that *every* word is marked for its possessive. In Section 3.4, I proposed that generalizations over the distribution of *-V* and *-jV* are learned over the sets of lexical entries that have [−j] and [+j], respectively. This falls out naturally if all words have one or the other (see Gouskova et al., 2015, pp. 44–46).

4.2.2 Lowering stems and the plural

In Section 3.4, I treated lexical variation in the plural as allomorph selection, in which a class of nouns called “lowering stems”, marked with a [lower] feature, selects *-ɒk* instead of the usual *-ok*. This analysis is sufficient for the purposes of this paper, but it obscures the actual nature of lowering stems and its complicated interaction with vowel harmony. In Table 4.4, I showed that the plural and a number of other suffixes (1SG, 2SG, 2PL) actually have a three-way distinction between *-C*, *-oC*, and *-ɒC*. The analysis in (28) would require each of these suffixes to have three allomorphs: with no vowel, [o], and [ɒ]. A more parsimonious analysis would insert the bare form of each suffix with a feature, [LV],[2] indicating that it undergoes linking vowel alternations:

(41) *Rules of realization for linking vowel suffixes (better version)*

a. PL $\leftrightarrow$ $\text{k}_{\text{[LV]}}$

b. 1SG $\leftrightarrow$ $\text{m}_{\text{[LV]}}$

...

Readjustment rules can then insert the appropriate linking vowel after consonants:

(42) *Readjustment rules for linking vowels*

a. Ø → ɒ / [lower] ___ [LV]

b. Ø → o / C ___ [LV]

This analysis is a more elegant approach to lowering stems than the rules in (28), and correctly predicts that a noun that has a low linking vowel in one suffix (e.g. the plural) will have a low linking vowel in all suffixes. In both analyses, nouns that take *-ɒk* in the plural are marked with a [lower] diacritic.[3] Thus, for my purposes, they make equivalent predictions: speakers should be

[2]The linking vowel in these suffixes is not predictable from phonotactics, and so must be marked (Siptár & Törkenczy, 2000, p. 219).

[3]One possibility is that lowering stems are phonologically predictable, showing “height harmony”. This is not the case: most noun stems with low vowels are not lowering stems, and many lowering stems, like [hold] ‘moon’ (plural [hold-ɒk]), have non-low vowels.

able to learn a dependency between [lower] and the [−j] feature marking possessives that take *-V*, using the sublexical grammars described in Section 3.2.

4.2.3 The representation of lowering stems

The analysis in the previous section assumes that the lowering alternation is encoded *morphologically*: lowering stems are marked with [lower], and suffixes with linking vowels have an [LV] feature. Siptár and Törkenczy (2000, s. 8.1.4) instead propose an abstract *phonological* analysis: lowering stems have a floating low feature [+open_1] and linking vowel suffixes have an underlying vowel unspecified for height. This vowel surfaces as low in the presence of [+open_1], otherwise it surfaces as mid after consonants and deletes after vowels.

These analyses represent two approaches to morphophonologically exceptional morphemes, compared in Section 2.2.2. In my analysis, exceptional lexical items are marked with a *diacritic* indexing a morpheme-specific rule or constraint (e.g. Gouskova, 2012; Inkelas et al., 1997; Pater, 2010; Rysling, 2016). Siptár and Törkenczy (2000) instead use *abstract structure*: defective segments and subsegmental units that cannot surface in their underlying form, but behave differently from full segments (e.g. Lightner, 1965; Rubach, 2013; Trommer, 2021).

The two approaches are not mutually exclusive (for example, Lightner (1965) uses both, as shown in Section 2.2.2), and the choice between them is often one of elegance and coverage. Moreover, both have been criticized on similar grounds: Pater (2006) and Gouskova (2012) argue that underspecification accounts can overgenerate and be hard to learn, while Bermúdez-Otero (2012, 2013), Haugen (2016), and Caha (2021) argue that arbitrary lexical marking and readjustment rules are unrestrained and weaken our theory of grammar. In this case, the two analyses are largely equivalent: for Siptár and Törkenczy (2000), the floating feature has no phonological effect beyond producing a low linking vowel, making it akin to what Kiparsky (1982) calls "purely diacritic use of phonological features" (see Section 2.2.2.4).[4]

The present chapter argues that Hungarian speakers learn generalizations over the [lower] feature on lowering stems. In the analysis of Siptár and Törkenczy (2000), the floating [+open$_1$] feature is unique to lowering stems. This is compatible with my main hypothesis that speakers learn generalizations over features that index unpredictable morphophonological behavior.

The purpose of this section was to lay some formal groundwork for the correlation studied in the rest of this chapter between lowering stems (in my analysis, lexical items with a [lower] feature) and nouns that take *-V* in the possessive (in my analysis, lexical items with a [−j] feature). I now turn to the corpus and nonce word studies.

4.3 Hungarian synchronic corpus study

This study determines whether being a lowering stem is a reliable predictor for a word taking *-V* as its possessive allomorph. I have two reasons for doing this: first, this corpus study establishes the baseline distribution of the lexicon, which I use as a model for what speakers have learned from the lexicon in the nonce word study in Section 4.4. Second, I also compare the results of my study to a previous corpus study of the Hungarian possessive, in particular Rácz and Rebrus (2012). I find some differences between their results and mine, generally attributable to differences in our corpora.

4.3.1 Data

In this section I discuss the corpus that I use to represent the Hungarian lexicon. I discuss the consequences of my corpus construction in Section 4.3.4 and Section 4.4.7.

[4] "Self-lowering" (Siptár & Törkenczy, 2000, pp. 228–229) verbal suffixes, which show a vowel–zero alternation whose vowel is always low, could potentially distinguish the two analyses. However, Siptár and Törkenczy (2000) argue that the "self-lowering" alternation is morphological (allomorph selection) rather than phonological (underspecification), converging with my analysis for these cases.

My source of data is Papp (1969), a morphological dictionary of Hungarian which I transcribed manually. I use Papp (1969) for its comprehensive tagging of derivational morphology, but it has disadvantages: it is over 50 years old and reflects lexicographic work rather than pure corpus data. A comparison with the Hungarian National Corpus (Oravecz et al., 2014) shows that the two sources are closely correlated in their distribution of possessive allomorphs, so I conclude that the benefits of this dictionary outweigh any potential negatives.

This morphological tagging allows me to make an informed choice about my corpus construction: whether to include monomorphemic words only, or to allow complex (derived) forms as well. Under standard assumptions in Distributed Morphology, lexical information like allomorph selection is stored for roots and affixes, not complex stems (Embick & Marantz, 2008). Thus, if speakers are generalizing over the frequency of types in the lexicon (cf. Albright & Hayes, 2003; Bybee, 1995, 2001; Hayes & Wilson, 2008; Hayes et al., 2009; Pierrehumbert, 2001), derived words and compounds with the same head (rightmost affix or root) should not count as separate types. Root-based storage predicts that words ending in the same suffix should take the same possessive, which is largely true in Hungarian (Rácz & Rebrus, 2012). I adopt the assumption of root-based storage by limiting my corpus to monomorphemic nouns. In Section 4.4.7, I argue that this corpus accurately reflects the behavior of Hungarian speakers and compare my results to the predictions of stem-based storage. For the other studies in this dissertation, I do not have reliable morphological tagging within stems, so I include all words, complex or not; I believe the negative consequences of this choice are limited. However, without these technical limitations, I would ideally use only monomorphemic stems for all of my corpus studies.

Although adjectives can also take possessive suffixes, I limit my corpus to nouns. Unlike nouns, most adjectives are lowering stems (Siptár & Törkenczy, 2000, pp. 229–230), and some forms behave differently depending on the syntactic environment (Rebrus & Szigetvári, 2018), so including adjectives would complicate the relationship between lowering stems and possessive allomorphy. I excluded vowel-final words, since these categorically take *-jV* and would be undefined for a number of the factors in my regression. I also removed the few words ending in orthographic *h*, which is phonologically complicated (see Siptár & Törkenczy, 2000, pp. 274–276). Finally, I excluded nouns with variable or unknown possessive to allow for binary coding of the possessive variable (*-V* vs. *-jV*). This leaves 2,427 noun types.

My data set, which is based on a structured dictionary, differs significantly in its design from that of Rácz and Rebrus (2012), who use a web corpus. Their search includes all consonant-final nominals that appear in the possessive, including morphologically complex stems and stems listed by Papp (1969) as including adjectival uses, for a total of 22.8 thousand consonant-final types and about 10 million tokens. Indeed, they explain some of their results with reference to the selectional restrictions of individual suffixes, including the comparative. To the extent that my results in Section 4.3.3 differ from theirs, much of this difference can be reduced to the effect of excluding morphologically complex words, which reflects different assumptions about the lexicon.

In using Papp (1969), the question arises of whether a hand-compiled dictionary several decades old is appropriately representative of modern Hungarian. I selected this corpus instead of the Hungarian National Corpus (Oravecz et al., 2014), which draws from a large body of contemporary texts, primarily because the better morphological tagging of the former allowed me to easily filter out compounds and derived forms (as discussed above). In addition, Papp (1969) can distinguish between homonyms with different morphological behavior and requires less processing of loanwords that do not conform to regular Hungarian orthography.

Comparing the set of monomorphemic nouns from the dictionary to those labelled as monomorphemic by the corpus, the latter again overpredicts *-V* relative to the experimental results. There are two potential explanations for this. The first is that, for some reason, the monomorphemic words in the web corpus simply do have a higher rate of *-V* than those in the dictionary. The alternative is that the difference is driven by words that are incorrectly labelled in the web corpus as being monomorphemic. Indeed, this seems to be the case: 53.2% of nouns (1189 of 2236) marked as monomorphemic in both data sets take *-V* (compared to 50.7% in the dictionary data set), while the proportion of nouns marked as monomorphemic in the web corpus but complex in the dictionary that take *-V* is much higher, 87.0% (2843 of 3266). Inspection of this latter class reveals scientific and Greco-Latin words with borrowed suffixes (like [nɒtsionɒlizmuʃ] 'nationalism'), compounds (like [uːjhold] 'new moon', from [uːj] 'new' and [hold] 'moon'), and derived forms like [viʒgaːlɒt] 'examination', from the verb [viʒgaːl] 'examine'. Accordingly, I conclude that the chief relevant difference between the dictionary and the web corpus is that the latter is not as good at reliably marking morphologically complex stems, and that correct segmentation is essential for accurately capturing the empirical results. Thus, I use the dictionary rather than the corpus—the benefits of its segmentation outweigh any effects of age or reliance on introspection.

4.3.2 Methods and analysis

As described above, the purpose of this corpus study is to provide an accurate representation of the lexicon to use as a baseline to predict how Hungarian speakers will respond to the phonological and morphological properties of nonce words. Accordingly, I use a number of phonologial predictors that could plausibly affect the choice of possessive, without necessarily having a convincing reason for *why* each effect may be relevant: as I explain in Section 2.3, my model allows for speakers to learn arbitrary generalizations, even if they are not phonologically grounded.

I fitted two linear regressions on my corpus with possessive suffix as the dependent variable. The first regression includes predictors representing a stem's phonological form: the place and manner of its final consonant, the height and length of its final syllable's vowel, its vowel harmony class, the complexity of its final coda, and whether it is monosyllabic.[5] Word-final consonant quality affects possessive allomorphy in Hungarian (Rácz & Rebrus, 2012). Becker et al. (2011) showed

[5]Hungarian has fixed word-initial stress, so this factor also marks whether the suffix is attaching to the stressed syllable.

that the quality of the last vowel is not relevant for Turkish lexically marked consonant voicing alternations, but given that harmony class affects Hungarian possessive selection (Rácz & Rebrus, 2012), I included other factors of vowel quality as well.

The second regression includes the same set of phonological predictors, plus the morphological factor of stem class. Stems were classified as lowering, non-lowering, variable, or indeterminate (nouns with front unrounding harmony, see Section 4.1.2). The models were assembled by forward stepwise comparison using the `buildmer` function in R from the package of the same name (R Core Team, 2022; Voeten, 2022). This function adds factors to the model one at a time such that each additional factor improves the model's Akaike Information Criterion (AIC), which measures how well the model fits the data while penalizing model complexity (that is, number of factors). One additional factor, the roundedness of the final syllable's vowel, did not significantly improve the model and was left out.

In the second regression, lowering stems should be significantly less likely to take *-jV* than non-lowering stems. The second regression should also provide a better model fit—as, indeed, it does. This shows that the dependency between lowering stems and *-V* is robust even when phonological factors are taken into account.

4.3.3 Results

4.3.3.1 Phonology

Table 4.5 contains the full model with phonological factors listed in the order in which they were added to the model (which roughly corresponds to their importance). Most of the examined factors are significant. The most influential are the place and manner of the final consonant. This effect strength is probably driven by the categorical effects of sibilants and palatals, which have the strongest negative effect size (favoring *-V*): words ending in these nouns, like [haːz] 'house', always take *-V*, with very rare exceptions. However, other places and manners have significant effects as well. Other phonological factors are also significant, e.g. front-harmonizing words like [køb] 'cube' (possesive [køb-ɛ] take *-jV* less than back-harmonizing words like [dob] 'drum' (possessive [dob-jɒ]), and nouns ending in geminates prefer *-jV* relative to nouns ending in singleton consonants. The model predicts a word's possessive quite well ($R^2 = .68$).

	β coef	*SE*	*Wald z*	*p*
Intercept	**3.02**	**.32**	**9.55**	**<.0001**
C Manner (default: plosive)				
fricative	**–1.44**	**.39**	**–3.73**	**.0002**
sibilant	**–10.69**	**.80**	**–13.36**	**<.0001**
nasal	**–1.95**	**.27**	**–7.16**	**<.0001**
approximant	**–4.08**	**.30**	**–13.47**	**<.0001**
C Place (default: alveolar)				
labial	**–2.02**	**.26**	**–7.94**	**<.0001**
palatal	**–8.88**	**1.10**	**–8.06**	**<.0001**
velar	**–3.26**	**.29**	**–11.19**	**<.0001**
Harmony (default: back)				
front	**–2.03**	**.18**	**–10.96**	**<.0001**
variable	**2.26**	**.97**	**2.33**	**.0197**
V Height (default: mid)				
high	**1.73**	**.22**	**7.89**	**<.0001**
low	0.28	.19	1.50	.1342
V Length (default: short)				
long	**1.40**	**.17**	**7.98**	**<.0001**
Coda (default: singleton)				
geminate	**2.47**	**.40**	**6.25**	**<.0001**
cluster	0.04	.21	0.18	.8602
Syllables (default: monosyllabic)				
polysyllabic	**1.15**	**.17**	**6.67**	**<.0001**

Table 4.5: Regression model with phonological predictors of possessive *-jV*, with significant effects bolded

This regression is intended as an approximation of the sublexical grammar for the [+j] sublexicon described in Section 3.1.2, so I now describe how it works mathematically. The model takes an input word and calculates a coefficient x which measures the predicted probability P that that word takes *-jV*, $P = \frac{e^x}{1+e^x}$. This coefficient is the sum of the β coefficients of the intercept and a word's value for each factor when it differs from the default. The model can predict the possessive of nonce words as well. For example, the nonce word [lufɒn] has a coefficient of $\beta_{\text{Intercept}} + \beta_{\text{C place: nasal}} + \beta_{\text{V height: low}} + \beta_{\text{Syllables: polysyllabic}} = 3.02 - 1.95 + 0.28 + 1.15 = 2.50$, corresponding to a probability of $\frac{e^{2.50}}{1+e^{2.50}} = .924 = 92.4\%$: if this were a real word, its possessive would likely be [lufɒn-jɒ]. I refer to these coefficients as *phon_odds* and use them as predictors of the nonce word experiment in Section 4.4.6. These odds represent the odds of a speaker assigning [+j] to a nonce word according to the scores assigned to that word by the [+j] and [−j] sublexical grammars. As I show in Section 3.4.5.2, these odds are determined by the difference between the two scores, and can thus be represented by the single *phon_odds* coefficient.

4.3.3.2 Phonology and morphology

Adding stem class to the model significantly improves it ($\chi^2 = 112.9$, $p < .0001$), raising the correlation to $R^2 = .71$. Stem class is significant and the most important factor after final C manner and place. Otherwise, the new model, shown in Table 4.6, is very similar to the phonological model in Table 4.5: the same phonological factors are added to the model (though in a slightly different order), and the effect sizes are quite similar. The effect of lowering stems is strongly negative: independent of their phonology, lowering stems are more likely to take *-V*. The effect of undetermined stem class is smaller and not significant. As discussed in Section 4.1.2, this class comprises nouns with front unrounded harmony like [ɲɛlv] 'language', so its effect should be masked by the factor of harmony.

	β coef	*SE*	*Wald z*	*p*
Intercept	**3.53**	**.33**	**10.60**	**<.0001**
C Manner (default: plosive)				
fricative	**–1.03**	**.44**	**–2.37**	**.0179**
sibilant	**–11.07**	**.80**	**–13.86**	**<.0001**
nasal	**–2.07**	**.28**	**–7.39**	**<.0001**
approximant	**–4.06**	**.31**	**–13.10**	**<.0001**
C Place (default: alveolar)				
labial	**–2.22**	**.27**	**–8.35**	**<.0001**
palatal	**–9.25**	**1.13**	**–8.22**	**<.0001**
velar	**–3.54**	**.31**	**–11.55**	**<.0001**
Stem class (default: non-lowering)				
lowering	**–3.71**	**.44**	**–8.44**	**<.0001**
undetermined	–0.25	.25	–0.98	.3278
variable	**–2.76**	**.69**	**–4.00**	**<.0001**
V Height (default: mid)				
high	**1.85**	**.23**	**8.09**	**<.0001**
low	**0.77**	**.21**	**3.66**	**.0003**
Harmony (default: back)				
front	**–1.98**	**.27**	**–7.41**	**<.0001**
variable	**2.25**	**1.04**	**2.17**	**.0297**
Coda (default: singleton)				
geminate	**2.43**	**.41**	**5.97**	**<.0001**
cluster	–0.08	.22	–0.36	.7161
V Length (default: short)				
long	**1.30**	**.19**	**6.97**	**<.0001**
Syllables (default: monosyllabic)				
polysyllabic	**0.79**	**.18**	**4.31**	**<.0001**

Table 4.6: Regression model with phonological and morphological predictors of possessive *-jV*, with significant effects bolded

I confirmed that stem class is indepedent of the phonological effects by testing its variance inflation factor (VIF) using the `check_collinearity` function from R's `performance` package (Lüdecke et al., 2021). This measures whether different factors are describing the same effect. Stem class had a low correlation with the other factors (see James et al., 2013), meaning that its effect on possessive allomorphy cannot be reduced to some combination of phonological factors.

4.3.4 Discussion

The corpus study shows that a number of phonological factors are good predictors of a noun's choice of possessive, as is lowering stem class—participants are expected to productively extend these phonological and morphological generalizations to nonce words as well in the experiment in Section 4.4. However, some of these results differ from those of Rácz and Rebrus (2012), who address a subset of the phonological factors in my analysis. I compare their main findings with mine in Table 4.7, highlighting the differences.

	rate of -jV according to…	
factor	Rácz and Rebrus (2012)	Table 4.5
harmony class	back > front	back > front
final coda	vowel = 100% > **complex** > **singleton**	geminate > **singleton** ≈ **cluster**
final C place	**labial** > **alveolar** > velar ≫ palatal = 0%	**alveolar** > **labial**, velar, palatal

Table 4.7: Comparison of phonological effects on possessive allomorphy from Rácz and Rebrus (2012) and my analysis

There are two main differences: first, Rácz and Rebrus (2012) find that nouns ending in complex codas take *-jV* at a higher rate than those ending in singleton consonants. They refer to these as "clusters", but they group geminates and clusters together, and attribute the effect to derivational suffixes that end in geminates. I make a finer-grained distinction and find that the difference is driven specifically by geminates. Second, Rácz and Rebrus (2012) find that nouns ending in labial stops take *-jV* more than nouns ending in alveolar stops, while I find the opposite place effect. The discrepancy is due to choice of corpus discussed in Section 4.3.1: Rácz and Rebrus (2012) attribute the high rate of labial *-jV* to the comparative suffix *-bː*, which prefers *-jV*. My corpus includes neither adjectives nor derived words, so it lacks these comparative adjectives. All in all, the differences between my analysis and that of Rácz and Rebrus (2012) can be attributed to different choices in coding and corpus construction. In Section 4.4, I use the phonological model in Table 4.5 as the representation of the scores from the [±j] sublexical grammars, and discuss its accuracy as a representation of the lexicon in Section 4.4.6.

4.4 Hungarian nonce word study

In Section 4.3, I described the gradient phonological and morphological effects on the distribution of possessive allomorphs. In this section, I present a novel experimental paradigm testing whether Hungarian speakers productively apply these generalizations. While previous nonce word studies have focused on phonological generalizations (e.g. Becker et al., 2011; Gouskova et al., 2015; Hayes et al., 2009), I show that speakers apply a morphological generalization as well: nonce words are assigned *-V* more often when presented as lowering stems (with plural *-ɒk*).

4.4.1 Predictions

I hypothesize that speakers form the possessives of novel words by taking both their phonology and their morphology (specifically, their plural) into account. I first show that speakers observe phonological effects, then show the effect of stem class.

Hayes et al. (2009) propose the “law of frequency matching” (see Section 2.1.2): when adult speakers are asked to extend variable lexical patterns, they usually do so by choosing stochastically in a way that roughly matches the frequency of each variant in the lexicon. This pattern is also found in artificial language studies (e.g. Hudson Kam & Newport, 2005), and I expect to see it in the present experiment (see Jarosz (2022) for discussion). Since my primary concern is the morphological dependency, I focus on phonological frequency matching in the aggregate and only discuss particular phonological effects where necessary to explain the effects of morphological sensitivity.

As discussed in Section 4.1.3, Rácz and Rebrus (2012) argue that *-jV* is the productive default for most words. If this is true, speakers should categorically assign *-jV* to most words rather than showing frequency matching.

4.4.2 Participants

Subjects were recruited through Prolific and had to be born in Hungary and raised as monolingual Hungarian speakers. I recruited 30 participants for the stimulus norming study and 91 for the stimulus testing study. One additional subject was rejected for poor quality, and an additional 48 subjects were recruited for earlier versions of the stimulus testing study; their data are not presented here.

4.4.3 Stimuli

I trained the UCLA Phonotactic Learner (Hayes & Wilson, 2008) on the corpus of Hungarian nouns used in Section 4.3. Part of the program's output included a "sample salad" of 1,968 nonce words. Of these, I selected all words with the shape (C)VC(C) or (CV)CVC(C) that were not coincidentally real words. I also removed words with mixed backness harmony, i.e. disyllabic words with one front vowel and one back vowel. This left a final set of 317 nonce word stimuli, some of which had rather unlikely clusters (like [ɲɒsm]). Each word was presented in the singular and plural, in the latter case with either a regular or lowering plural suffix.

4.4.4 Procedure

This experiment was split into two studies. First, subjects rated the 317 nonce word stimuli for plausibility as Hungarian words. The ratings obtained in this study were used to select a smaller set of stimuli to be used for testing in the main experiment, which probed the morphological dependency between lowering stems and *-V* possessives.

4.4.4.1 Stimulus norming and selection

Participants completed 50 trials, each with a different stimulus. Each trial had a frame sentence containing the target stimulus twice. In its first occurrence, the stimulus appeared in bare nominative form; the second time, the stimulus had a plural suffix (and sometimes additional suffixes as well). Most stimuli were shown with regular plurals (e.g. *-ok*), but 8 randomly chosen trials instead showed stimuli as lowering stems (e.g. with plural *-ɒk*). Participants rated each stimulus as a potential Hungarian word on a scale of 1 to 5.

These ratings were used as inputs to a Python script that selected a set of stimuli with a high average rating and a phonological distribution similar to the base corpus. I examined high-ranking sets manually and selected a set with 81 stimuli to use for the main testing phase.

4.4.4.2 Morphological dependency testing

Participants each completed 35–50 trials, which had the format shown in Figure 4.1. First, the stimuli were presented in the same frame sentences as in the stimulus norming experiment. In Figure 4.1, the nonce word [lufɒn] has a regular plural *-ok*, but in 8–12 trials, the stimuli were presented as lowering stems, e.g. plural [lufɒn-ɒk]. As an attention check, participants had to correctly select the plural form appearing in the first sentence. Next, a second frame sentence appeared, in which participants had to select 1SG and possessive forms. As discussed in Section 4.2.2, the 1SG suffix has the same regular and lowering stem variants as the plural, so the linking vowel should match that of the plural: in this case, [lufɒn-om].[6] The choices included both back and front variants; the possessive should have the same harmony class as the plural (in this case, [lufɒn-ɒ] or [lufɒn-jɒ]). Trials in which speakers chose a discordant 1SG or antiharmonic possessive were discarded.

[6]The possessive morpheme *-V/-jV* does not appear in first singular possessive forms, only the possessor marker (see Section 4.1.3).

A good **lufɒn** is one who knows how to make other **lufɒnok** laugh.

back regular

Please select the word's plural form: [lufɒnøk / lufɒnɒk / lufɒnɛk / **lufɒnok**]

front regular — back lowering — front lowering — back regular

That's correct! Now select the word in the appropriately inflected form according to you.

My [lufɒnɒm / lufɒnɛm / lufɒnøm / **lufɒnom**] couldn't sing well, however my husband's

back lowering — front lowering — front regular — back regular

[lufɒnɛ / lufɒnjɛ / **lufɒnɒ** / **lufɒnjɒ**] sang brilliantly.

front -*V* — front -*jV* — back -*V* — back -*jV*

Figure 4.1: Trial for Hungarian stimulus testing study, with forms annotated for harmony and stem class and acceptable answers bolded

4.4.5 Analysis

I discarded discordant trials (as described above) and trials with filler stimuli with front unrounded harmony, which do not show the lowering stem alternation (see Section 4.2.2 for discussion). This left a total of 2,398 trials with 57 stimuli.

As in the corpus study, I fitted two mixed logistic regressions whose dependent variable is the possessive suffix selected by the participant (-*V* vs. -*jV*). The first regression describes how participants used a nonce word's phonology to assign its possessive. If speakers are matching the distribution of the lexicon, then the experimental results should correlate with the *phon_odds* coefficients for nonce words showing their likelihood of -*jV* according to the phonological model of the lexicon in Table 4.5 (see Section 4.3.3 for details), which I use to estimate speakers' sublexical grammars for the possessive. Thus, the first regression includes the *phon_odds* coefficients for the nonce words and random effects: a random intercept for participant and item, and a by-participant random slope for participant. These effects measure how participants differ in their rate of *-jV* and their sensitivity to the fixed effect of phonology.

The second model includes these two factors as well as stem class and a by-participant random intercept for stem class: whether a nonce word was presented with a regular plural *-ok* or a lowering stem plural *-ɒk* in a given trial. This tests whether, and to what extent, participants show sensitivity to the morphological dependency.

4.4.6 Results

4.4.6.1 Phonology

In Table 4.8, we see the effects of the mixed logistic regression predicting participant responses given random intercepts for participant and item, as well as a fixed effect of *phon_odds* calculated from the phonological model of the lexicon in Table 4.5 and a by-participant random slope for *phon_odds*. (The raw distribution of possessive responsives according to several phonological factors is shown in Section 4.4.7.3 below.) The mean of *phon_odds* is fairly close to 0 and the coefficients were not centered around the mean; a model with centered *phon_odds* has a lower but still positive intercept and other effect sizes more or less unchanged.

Random effects	*variance*	*SD*		
Participant				
Intercept	0.62	.79		
Phon_odds	0.01	.12		
Item	0.50	.71		
Fixed effects	*β coef*	*SE*	*Wald z*	*p*
Intercept	**0.68**	**.00**	**1031.20**	**<.0001**
Phon_odds	**0.38**	**.00**	**572.13**	**<.0001**

Table 4.8: Effects of mixed logistic model with predictions of the phonological model of the lexicon (Table 4.5) for experimental use of possessive *-jV*, with significant effects bolded

This model shows an overall bias towards *-jV* (since the intercept is positive): there were 1,031 responses of *-V* and 1,367 of *-jV*. The results also show a correspondence between predicted rates and actual rates: the coefficient for *phon_odds* is positive. However, the effect size is much smaller than 1: an increase of 1 in *phon_odds* corresponds to a predicted increase in the experimental likelihood of a *-jV* response of only .38, although the two operate on the same scale. This means that the overall range of likelihood predicted by the experimental model is much narrower than that predicted by the model of the lexicon.

The random intercept for item shows that different words were given fairly divergent rates of *-jV* even once phonology (in the form of *phon_odds*) is taken into account. The random intercept for participant shows that different participants had different baseline rates of *-jV*, but the by-participant random slope for *phon_odds* has a very low variance, suggesting that participants treated the phonological effect in roughly the same way. The primary effect of the random slope is to make the standard error of the fixed effects very small and make the overall model nearly unidentifiable (with a very large eigenvalue). The model also requires a greater tolerance in order to converge. Rescaling *phon_odds* does not remove these warnings, and a similar model without a random slope has no such issue and is otherwise very similar. Nonetheless, I include this random slope because it significantly improves the model ($\chi^2 = 15.62$, $p = .0004$) and because the effect of stem class discussed in Section 4.4.6.2 is clearer when compared to this fuller phonological model.

We see the difference between the lexical and experimental models in Table 4.9, which shows the two words predicted to be most ([oluːnt]) and least ([jøs]) likely to take *-jV* and the word with a predicted rate closest to 50% ([ʃokːol]). These predictions follow from the words' phonology: for example, words ending in sibilants like [jøs] categorically take *-V* in the lexicon, and other properties of this word push its prediction further towards *-V*, such as its front harmony. By contrast, [oluːnt] has numerous properties that strongly prefer *-jV*, especially the fact that it ends in a cluster and an alveolar stop.

The two extremes, [jøs] and [oluːnt], are predicted by Table 4.5 to be essentially categorical in the lexicon, but showed mixed responses in the experiment. Correspondingly, the model trained on the experimental results (Table 4.8) predicts that one variant should be dominant, but not effectively categorical. Table 4.9 also shows the effect of the random intercept for stimulus, which is to account for variance that the fixed effects alone cannot by bringing observations closer to the line of best fit. In the case of [jøʃ], the base model substantially underestimates the likelihood of *-jV*: 4.8% instead of the observed 17.4%. This word thus has a substantial positive random intercept, which adjusts the predicted rate up to 9.2%—closer to the observed rate. For the other two words, the model with fixed intercepts for item already does a very good job at matching the experimental rate: for example, [ʃokːol] is predicted to have a rate of 67.1% and has an actual rate of 67.5%. For this word, the random intercept overcompensates slightly and brings the predicted rate up to 69.6%; for [oluːnt], the random intercept yields a slight improvement, reaching 96.3%, very close to the true rate of 94.7%.

nonce word	*predicted likelihood of -jV in lexicon model*	*experimental rate of -jV*	*predicted likelihood of -jV in experimental model*	
			fixed intercept	*random intercept*
jøʃ	0.006%	17.4% (8/46)	4.8%	9.2%
ʃokːol	52.097%	67.5% (27/40)	67.1%	69.6%
oluːnt	99.934%	94.7% (36/38)	96.9%	96.3%

Table 4.9: Predicted likelihood of *-jV* for nonce words according to models trained on lexicon (Table 4.5) and experimental results (Table 4.8), including the adjustment of the random intercept for item

Figure 4.2 shows the relationship between the predicted likelihood (according to *phon_odds* from the model trained on the lexicon) of each nonce word taking *-jV* and its experimental rate of *-jV*. Both axes are shown in terms of log odds (that is, the coefficients) in order to make the relationship linear. A rate of 0 corresponds to a log odds of negative infinity, so the nonce word *fátyúsz* [faːcuːs], which speakers assigned *-V* in every trial, should be at negative infinity. It is included at the bottom edge of the graph in Figure 4.2. The graph shows each nonce word twice: in black, the word's position on the x-axis assumes a fixed intercept, so each word's position is solely a function of its *phon_odds*. The lighter gray includes the adjustment of the random intercept for word. As discussed above, the effect of the random intercept is to bring each word's predicted rate closer to the line of best fit, so the gray words (including the random intercept) are closer to the line than the unadjusted words: stimuli above the line have a random intercept shifting them to the right, while those below the line move left with the random intercept. Figure 4.3 shows the same data plotted on scales of raw probability. The graphs also include a line corresponding to the fit of the model in Table 4.8.

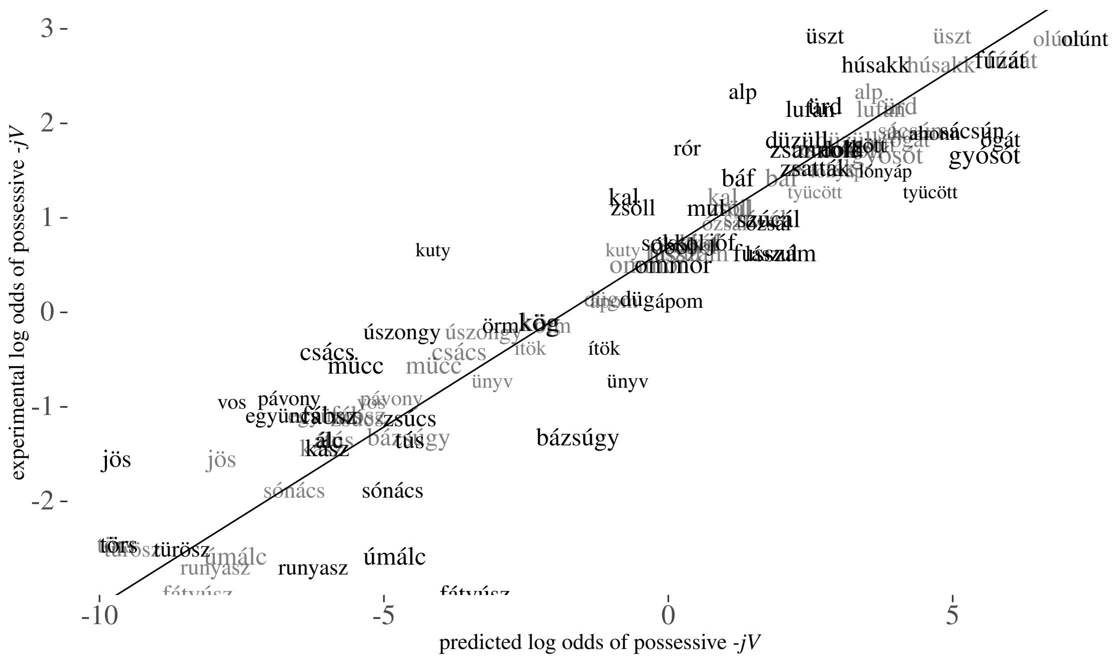

Figure 4.2: The relationship between predicted and experimental log odds of possessive *-jV* for individual nonce words with (gray) and without (black) the random intercept, sized according to number of trials, with a line showing the fit of the experimental model in Table 4.8

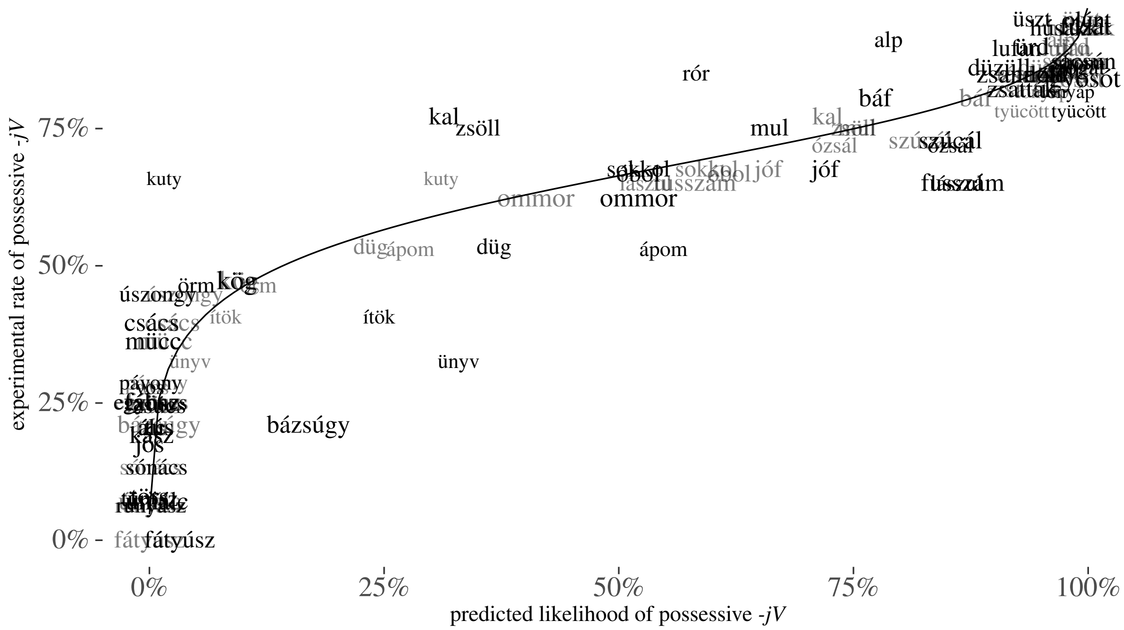

Figure 4.3: The relationship between predicted likelihood and experimental rate of possessive *-jV* for individual nonce words with (gray) and without (black) the random intercept, sized according to number of trials, with a line showing the fit of the experimental model in Table 4.8

A nonce word with a higher predicted likelihood of *-jV* generally has a higher experimental rate of *-jV*, and the relationship is linear (Figure 4.2): the phonological model of the lexicon fits the experimental results well Figure 4.3 shows that the experimental results are less extreme than the predicted likelihood, especially on the low end: nouns ending in palatals and sibilants, which categorically take *-V* in the lexicon and thus had a near-zero predicted likelihood of *-jV*, were assigned *-jV* in the experiment up to nearly 50% of the time. Nonce words with a very high predicted likelihood of *-jV* had a high experimental rate of *-jV*.

4.4.6.2 Phonology and stem class

Table 4.10 shows the effects of the regression including the factor of stem class (that is, the plural shown on the nonce word in a given trial) alongside *phon_odds* and the random intercept for participant. As before, this model requires a larger tolerance to converge, although otherwise it has no issues.

Random effects	*variance*	*SD*		
Participant				
Intercept	0.62	.79		
Phon_odds	0.01	.12		
Stem class	0.47	.69		
Item	0.51	.72		
Fixed effects	*β coef*	*SE*	*Wald z*	*p*
Intercept	**0.78**	**.15**	**5.23**	**<.0001**
Phon_odds	**0.38**	**.03**	**12.19**	**<.0001**
Stem class (default: non-lowering)				
lowering	**−0.39**	**.17**	**−2.34**	**.0192**

Table 4.10: Effects of mixed logistic model with predictions of the phonological model of the lexicon (Table 4.5) and stem class for experimental use of possessive *-jV*, with significant effects bolded

The factor of stem class appears to make a slight difference, as expected from the lexicon, given the distribution of responses: stimuli (like [huːsɒkː]) presented as regular stems (with plural [huːsɒkː-ok]) were assigned *-jV* ([huːsɒkː-jɒ]) 58.1% of the time (1,090 out of 1,876 trials), while participants assigned *-jV* to stimuli slightly less often when they were presented as lowering stems (with plural [huːsɒkː-ɒk]), 53.1% of the time (277 of 522 trials). This is a rather small difference, but it is significant in the model in Table 4.10. This model performs significantly better than the model without the morphological factor shown in Table 4.8 ($\chi^2 = 13.12$, $p = .011$). As before, the random slope for *phon_odds* is not doing much, with a variance of .01. However, the random slope for stem class has a much higher variance, .47. This suggests that some speakers observed the correlation between lowering stems and *-V* more than others.

To get a better sense of the data, let us look at the behavior of the individual nonce words. Figure 4.4 and Figure 4.5 show the same data as Figure 4.2 and Figure 4.3, but each nonce word is now split between trials when it was presented as a regular stem (in black), and a lowering stem (in gray). The lowering stem words are always smaller than the regular words because each word had fewer lowering stem trials. A line connects the two conditions for each word, going leftward from the regular word to the lowering stem word because the model of the experimental results predicts a lower likelihood of *-jV* for lowering stems. Probabilities of 0 and 1 correspond to log odds of (negative) infinity, so words with categorical behavior in one condition are shown at the bottom and top edges of Figure 4.4 and connected with a dashed line. The graphs show lines corresponding to the fit of the model in Table 4.10 for regular (black) and lowering stem (gray) conditions.

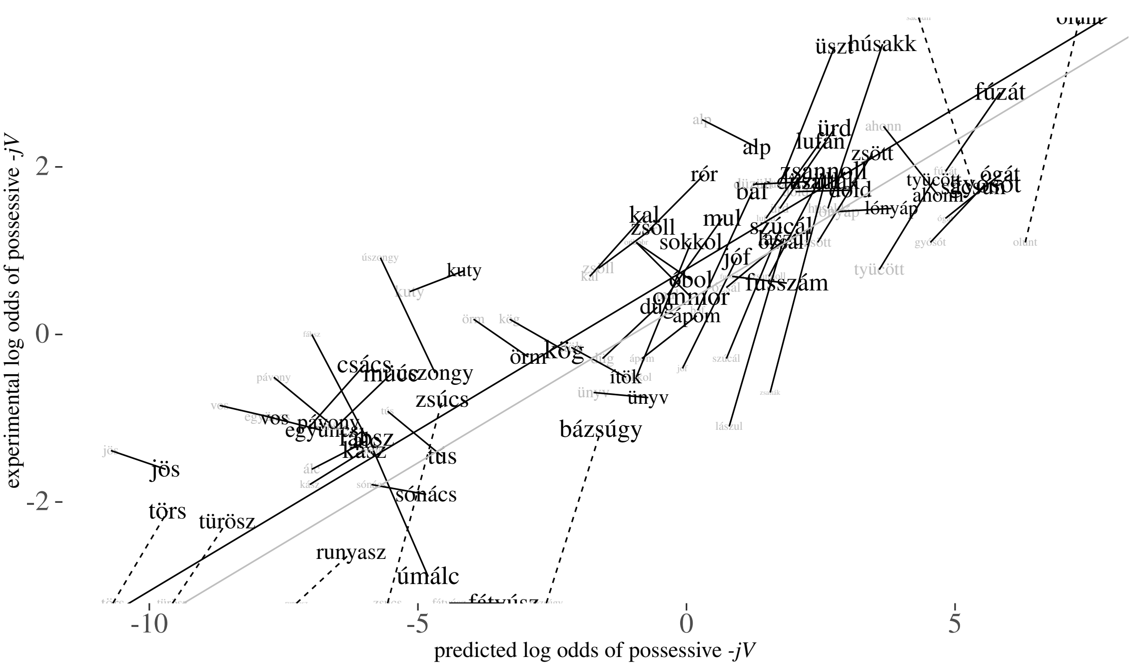

Figure 4.4: The relationship between predicted and experimental log odds of possessive *-jV* for individual nonce words presented with regular (black) and lowering (gray) plurals, sized according to number of trials, with a line showing the fit of the experimental model in Table 4.8

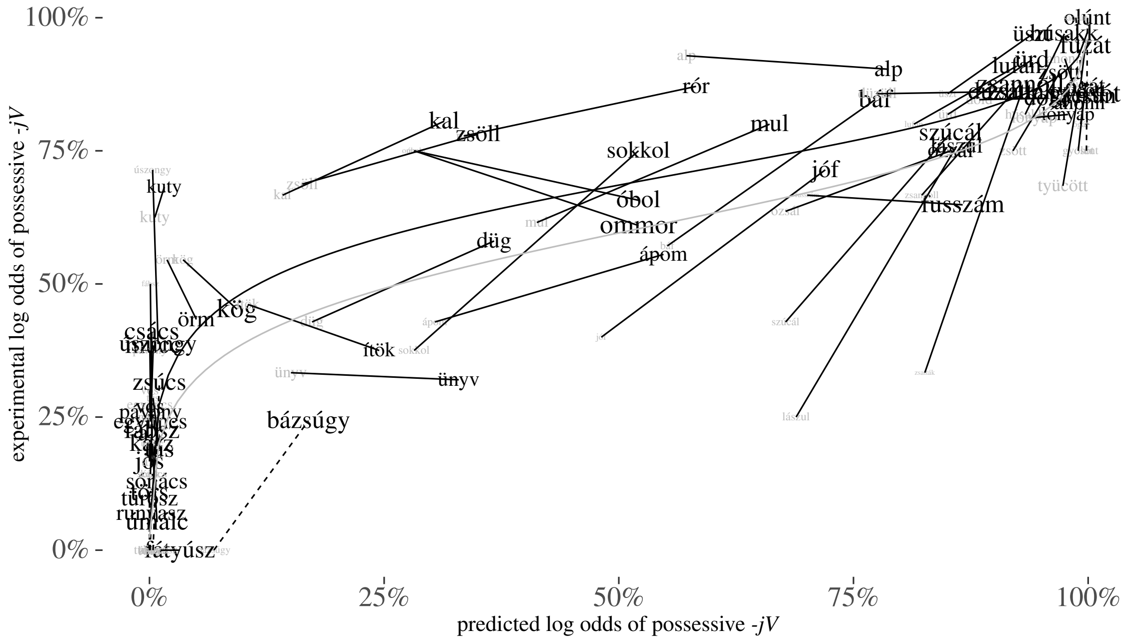

Figure 4.5: The relationship between predicted likelihood and experimental rate of possessive *-jV* for individual nonce words presented with regular (black) and lowering (gray) plurals, sized according to number of trials, with a line showing the fit of the experimental model in Table 4.8

Most of the lines in these figures slope downward and to the left, indicating that most nonce words had a lower experimental rate of *-jV* when presented as a lowering stem. This is especially true in the top right section of the graph—that is, nonce words with a higher rate of *-jV*. On the other hand, nouns with a lower expected rate of *-jV* are more likely to have a *higher* rate of *-jV* when presented as a lowering stem. These words, which end in palatals and sibilants, behaved unexpectedly: they categorically take *-V* in the lexicon but showed moderate rates of *-jV* in the experiment. Indeed, if we remove stimuli ending in sibilants or palatals, the model does better (Table 4.11). The effect size of *phon_odds* is very slightly smaller, and the variance of the random intercept for item is much smaller, suggesting that these words are better behaved with respect to their phonology. The effect of stem class is substantially stronger, at $-.56$ compared to $-.39$ in Table 4.8. The intercept is also higher, as expected given that this data set excludes nouns with lower rates of *-jV*. Participants also apply the effect of stem class more uniformly to these words: the by-participant random slope for stem class has a variance of .22, compared to .47 in Table 4.10.

Random effects	*variance*	*SD*		
Participant				
Intercept	0.62	.78		
Phon_odds	0.03	.16		
Stem class	0.22	.47		
Item	0.21	.46		
Fixed effects	*β coef*	*SE*	*Wald z*	*p*
Intercept	**0.95**	**.16**	**5.95**	**<.0001**
Phon_odds	**0.35**	**.05**	**6.37**	**<.0001**
Stem class (default: non-lowering)				
lowering	**−0.56**	**.19**	**−2.93**	**.0034**

Table 4.11: Effects of mixed logistic model with predictions of the phonological model of the lexicon (Table 4.5) and stem class for experimental use of possessive *-jV* for stimuli not ending in palatals or sibilants, with significant effects bolded

As before, the model including stem class is a significantly better fit than the model whose only fixed effect is *phon_odds*, whose effects I do not show here ($\chi^2 = 15.72$, $p = .003$).

4.4.7 Discussion

4.4.7.1 Speakers show sensitivity to stem class

The experiment supports my main hypothesis: on the whole, Hungarian speakers observed the morphological dependency in the lexicon, assigning possessive *-V* more often to nonce words when they were presented as lowering stems (with plural *-ɒk*). Thus, my experiment serves as a proof of concept for nonce word studies manipulating the inflection of novel forms. Lowering stem nouns generally comprise a closed class, but speakers nonetheless extended the lexical generalization that lowering stems prefer *-V*. It is unsurprising that speakers can apply patterns for unproductive stem classes like lowering stems, because there exist rare lowering stems, like [maːl] 'belly fur', that a speaker might first encounter as an adult.

The by-participant random slope for stem class suggests that different participants applied the correlation between lowering stems and *-V* more strongly than others. This is certainly possible: inspection shows a wide range of possessive distributions for lowering stems, although removing stimuli ending in palatals and sibilants substantially reduced the variance. These individual differences require further study: each participant had at most 12 lowering stem trials, and a good number of them were discarded for discordant responses, as described in Section 4.4.4.2. While the large variance in this random slope is interesting and suggestive of individual differences, given the structure of the task I am tentative in my conclusion and aim to confirm this result with future experiments.

4.4.7.2 How did speakers match the frequency of the lexicon?

In general, Hungarian speakers match the phonological distribution of *-V* and *-jV* in the lexicon quite closely and consistently with novel words. However, my results showed a less extreme distribution than the lexicon—in particular, nouns ending in sibilants and palatals categorically take *-V* in the lexicon but were sometimes assigned *-jV*. This pattern holds true for all of the experiments in this dissertation. One likely explanation is that the task, in which participants had to select from several options, boosted the salience of usually unlikely alternatives. Another possibility is some amount of noise: I had trouble getting participants to pay attention to the relevant details, suggesting that the task was fairly difficult, and they may have sometimes chosen completely at random.

It is clear, though, that subjects applied gradient patterns from the lexicon, counter to the claim by Rácz and Rebrus (2012) that novel words categorically take one possessive suffix (see Section 4.1.3): *-V* for nouns ending in palatals and sibilants, *-jV* otherwise. One possible reason for this is that the experimental task may be different from what speakers do in real life. Even if so, I have shown that speakers store and can apply generalizations about possessive allomorphy that are both phonological and morphological in nature. The primary purpose of my study is to probe these generalizations, and in this, it is successful. It also aligns with most nonce word studies, which find matching of lexically variable patterns, as discussed in Section 2.1.2.

Alternately, the claim of productive defaults in the literature may be overstated. A search of the Hungarian National Corpus (Oravecz et al., 2014) reveals occasional uses of *-jV* with sibilant-final loanwords: one speaker forms the possessive of the Britney Spears song "Sometimes" as [saːmtaːjmz-jɒ]. Thus, speakers sometimes generate the "non-productive" forms in spontaneous text. Perhaps speakers show frequency matching when inflecting totally novel words, but these quickly stabilize to the default as the more likely form dominates and spreads. More research is required to reconcile my results with the lexical behavior of nouns newly entering the Hungarian lexicon.

4.4.7.3 Are speakers generalizing over roots or stems?

Finally, I return to one ancillary question implicated in this study. In Section 4.3.1, I argued for a corpus of roots (i.e. monomorphemic words) over a stem-based corpus that counts every derived form and compound separately. Phonological models trained on the two corpora yield similar phonological effects, but very different baseline rates of *-jV* (that is, model intercepts). Most derived forms take *-V* (Rácz & Rebrus, 2012), so the stem-based corpus has a much lower overall rate of *-jV*. Although the phonological effects from the stem-based corpus are somewhat better predictors of the experimental results, comparing the rates of *-V* and *-jV* makes it clear that participants are matching the distribution of roots, not stems. We see this in Table 4.12, which compares the experimental rates of *-jV* across final consonant place and manner (the most important phonological effects). With the exception of sibilants and palatals, which are equally unexpected for all accounts, the rates across monomorphemic words are a much better fit.[7]

[7] The substantially higher rate of *-jV* for approximants in the experiment is because no nonce words ended in the palatal approximant [j], which always takes *-V* in the lexicon.

	experiment (responses)			*lexicon (types)* *monomorphemic nouns*			*all nominals*		
	-V	*-jV*	*% -jV*	*-V*	*-jV*	*% -jV*	*-V*	*-jV*	*% -jV*
labial	112	212	65.4%	117	156	57.1%	1781	565	24.1%
alveolar	745	957	56.2%	712	816	53.4%	12 041	2214	15.5%
palatal	93	61	39.6%	275	1	0.4%	2035	3	0.1%
velar	81	137	62.8%	126	224	64.0%	2796	667	19.2%
plosive	207	599	74.3%	200	660	76.7%	4993	2318	31.7%
fricative	49	77	61.1%	36	30	45.5%	567	54	8.7%
sibilant	550	127	18.8%	496	2	0.4%	7886	2	0.0%
nasal	91	174	65.7%	180	267	59.7%	2136	503	19.1%
approximant	134	390	74.4%	318	238	42.8%	3072	572	15.7%
all	1031	1367	57.0%	1230	1197	49.3%	18 653	3449	15.6%

Table 4.12: Type frequency of *-jV* in two versions of the lexicon and experimental frequency of *-jV* responses, by final C place and manner

I take Table 4.12 as empirical support for my assumption, discussed in Section 4.3.1, that Hungarian speakers are counting over roots and affixes, not over complex stems. This is required by root-based storage, which is used in Distributed Morphology (e.g. Embick & Marantz, 2008; Halle & Marantz, 1993), fitting with the theory of morphological dependencies I presented in Section 3.1. The predictions of stem-based storage are less clear: to capture the results in Table 4.12, speakers would need a way to mark forms that share a head as duplicating one another and not count them. That is, the analogical mechanism must be able to recognize that stems sharing certain structures in common (namely, a head) should only count as one type, but that stems sharing other structures (for example, a final consonant) should each be counted individually. This would require a more complicated analogical mechanism than root-based storage.

4.5 General discussion and summary

In this study, I looked at a tendency for a small class of Hungarian nouns, known as lowering stems, to take the possessive allomorph *-V*. In the lexicon, the correlation between the two is strong: only a few lowering stems take the other option, *-jV*. In a nonce word study, participants (at least some) followed this pattern: nonce words that were presented as lowering stems, with plural *-ɒk* instead of the regular *-ok*, were more likely to be assigned *-ɒ* in the possessive. Much stronger were the phonological correlations: a phonological model of the lexicon was a very good predictor of participants' possessive. These results thus show that Hungarian speakers learn and can apply gradient tendencies from the lexicon that are both phonological and morphological in nature, and consider the effects together.

This result is neatly explained by the sublexicon model I described in Chapter 3: Hungarian nouns are marked with a lexical diacritic feature [±j] that encodes the possessive they take, and speakers develop a grammar that evaluates candidates for how well they fit the sublexicon of nouns that take a given feature. The relevant factors can be phonological, but also morphological: if lowering stems are likewise marked with a [lower] feature, then possession of this feature can also make a word a better or worse fit for one possessive sublexicon. In particular, my results suggest that features are properties of roots, as predicted by Distributed Morphology, not stems.

The next two studies, in Czech and Russian, differ in their details, but both affirm the same basic principle demonstrated here: speakers learn correlations between morphological properties of a word and willingly extend these correlations to new words, as shown through nonce word studies.

5 Czech genitive and locative

In this case study, I look at variation within one of the main inflectional classes in Czech, known as hard-stem masculine inanimate nouns. This class shows variation in the genitive and locative singular, and I study the extent to which a given noun's realization of the two cases is aligned. This study differs from the Hungarian study in Chapter 4 in two ways. In Hungarian, the morphological dependency was not very strong compared to some of the phonological generalizations. By contrast, in Czech the phonological factors influencing locative allomorphy are relatively weak, and the morphological dependency is much stronger. This holds true in the lexicon and in how speakers apply patterns in the nonce word study. The second difference is that this case study is marked by the widespread presence of *variable lexical items*: a given noun that allows the minority suffix for one of the cases usually allows the more common one as well. I discussed this case study in my model of variation in Section 3.4 and look at the lexicon extensively through both a variable and a categorical lens in the corpus study in Section 5.3. The nonce word study does not have the resolution to test whether speakers assign variable behavior to nonce words, so I set the question of variation aside in Section 5.4. However, the prevalence of variable items does allow me to conduct two further studies that shed more light on the morphological dependency. First, in Section 5.5, I conduct a similar experiment using *real* variable words. Speakers show a correlation between genitive and locative for nonce words but not real words—this means that the morphological dependency shown in the nonce word study really is the productive extension of a lexical pattern to novel words, as opposed to some other effect like priming. Next, in Section 5.6, I look at variable words in the work of individual authors. I show that authors have a correlation between genitive and locative for individual words. This correlation can come from two sources: if this bias similarly appears in their *input* for individual words, Czech speakers may have faithfully learned a biased input. However, if this bias is *not* present in their input, it must instead come from a bias in their grammar. Specifically, as I argue in Section 5.4.6.3, this bias would be facilitated by a variable version of the multiple grammar model presented in Section 3.4.5.3, but not a variable version of the single grammar model described in Section 3.3. At present, I do not have any way to determine speakers' input, so I do not know which of these two explanations is correct. Nonetheless, the author study further affirms the correlation between genitive and locative in speakers' lexicons.

5.1 Background

5.1.1 Czech inflectional suffixes

Nouns in Czech inflect for seven cases and two numbers; this study focuses on two case endings (genitive and locative) in the singular for one large subset of nouns. Morphophonologically, there is a split in Czech between soft-stem nouns that take primarily front-vowel suffixes and end in palatals [c ɟ ɲ j], post-alveolars [ʃ ʒ tʃ], selected alveolars [ts r̝], and sometimes alveolars and labials [s z l p b f v m]; and hard-stem nouns that take a mix of back-vowel and front-vowel suffixes, the latter sometimes triggering alternations of stem-final consonants. An example is shown in Table 5.1 showing the singular case forms of the neuter nouns [kolo] 'wheel' and [pole] 'field', a hard-stem and soft-stem neuter noun, respectively. When a word ends in an ambiguous consonant like [l], it may be either hard or soft.

class	hard-stem	soft-stem
example	‘wheel’	‘field’
nominative	kol-o	pol-ɛ
genitive	kol-a	pol-ɛ
dative	kol-u	pol-i
accusative	kol-o	pol-ɛ
vocative	kol-o	pol-ɛ
locative	kol-ɛ	pol-i
instrumental	kol-ɛm	pol-ɛm

Table 5.1: Singular case endings for hard-stem and soft-stem neuter nouns

Many of the vowel correspondences between the two sets of endings are fairly regular: paradigm cells that have [o] or [a] for hard-stem nouns generally have [ɛ] for soft-stem nouns, and hard-stem [ɛ] and [u] usually correspond to soft-stem [i]. This has led some linguists (e.g. Janků, 2022; Williams et al., 2020) to propose that the distinction between soft- and hard-stem nouns is largely phonological, and that many of the suffixes are underlyingly the same. The variation studied in this chapter appears only in hard-stem nouns and has no soft-stem analogue, so I look only at hard-stem nouns that have suffixes broadly similar to [kolo].

Morphosyntactically, Czech nouns distinguish three genders (masculine, feminine, and neuter), and masculine nouns further show the same animacy distinction described in Section 2.1.5.2 for Polish: the accusative singular is morphosyntactically identical to the nominative for inanimate masculine nouns and to the genitive for animate masculine nouns.[1] For hard-stem masculine inanimate nouns, the default ending for both the genitive and the locative, is *-u*. This form is a historical innovation, and a number of words retain the original endings: *-a* in the genitive and *-ɛ*

[1]I describe this identity as *morphosyntactic* rather than *morphological* because it appears in adjectival agreement as well as inflection.

in the locative, the latter of which triggers a palatalization alternation in the last consonant of the stem. Examples of these forms can be seen in Table 5.2, which also contains one animate noun. I exclude animates from my data set, because they do not show variation: they exclusively have *-a* in the genitive, and their locatives are formed with a third suffix, *-ovi*, with occasional *-u* as well (but never *-ɛ*).

noun	'problem'	'match'	'evening'	'church'	'student'
nominative	problɛːm	zaːpas	vɛtʃɛr	kostɛl	studɛnt
genitive	problɛːm-u	zaːpas-u	vɛtʃɛr-a	kostɛl-a	studɛnt-a
locative	problɛːm-u	zaːpas-ɛ	vɛtʃɛr-u	kostɛl-ɛ	studɛnt-ovi

Table 5.2: Czech hard-stem masculine genitive and locative singular forms (mostly inanimate, with one animate)

One of the suffixes, *-ɛ*, triggers regular alternations of the last consonant of the stem, as shown in Table 5.3 below. The alternations of stem-final labials and alveolars are not very salient, and this "palatalization" is marked in both cases on the vowel ([pjɛ] is orthographic *pě*, [cɛ] is *tě*). The alternation between [r] and the fricated trill [r̝] is somewhat salient and marked orthographically on the consonant: *r* vs. *ř*. The velar alternations, on the other hand, are much more salient and are marked orthographically by completely different letters: *k* vs. *c*, *h* vs. *z*, and *x* vs. *š* ([g] only appears in loan words). This salience likely explains the fact that nouns ending in dorsals are much more resistant to *-ɛ* than other final consonants: speakers seem to avoid a suffix that triggers a salient alternation. Interestingly, this contrasts with a pattern in the Russian diminutive, shown in Section 6.1.2, where certain nouns actually *prefer* a suffix that triggers a stem alternation.

consonant	*alternations*	example	
labials	C∼Cj	ostrov ∼ ostrovjɛ	‘island’
alveolar stops	t∼c, d∼ɟ, n∼ɲ	svjɛt ∼ svjɛcɛ	‘world’
alveolar trill	r∼r̝	papiːr ∼ papiːr̝ɛ	‘paper’
dorsals	k∼ts, g/ɦ∼z, x∼ʃ	jazɪk ∼ jazɪtsɛ	‘language’

Table 5.3: Consonant alternations triggered by Czech locative suffix *-ɛ*

Most masculine inanimate hard-stem nouns only allow the standard *-u* endings in both cases. Of those that do allow the minority forms, most do so variably: alternating between *-a* and *-u* in the genitive, or *-ɛ* and *-u* in the locative. Variable lexical items are not uniform in their behavior: some lexemes overwhelmingly prefer one form or the other, while others have a more even split, and everything in between.

5.1.2 Syntactic effects on Czech case allomorphy

The variation also depends in part on *syntactic context*, as discussed in Section 3.4. Bermel and Knittl (2012) note that, in both cases, variable items occur with *-u* more often in less “canonical” uses of the case. For the genitive, this means that *-a* is more preferable when marking possessors and after prepositions of motion like [z] ‘out of’, and worse when following non-motion prepositions like [ftʃɛtɲɛ] ‘including’ or as the object of a verb that inherently selects for a genitive object, such as [dosaːɦnoʊt] ‘reach’. Similarly, the locative *-ɛ* ending is more acceptable as the object of a preposition expressing location, like [v] ‘in’, than one not expressing location (such as [o] ‘about’) or in a construction where a verb selects for a prepositional phrase with a locative noun, such as [zaːlɛʒɛt] ‘depend’ which, like its English counterpart, takes as its object a phrase headed by the preposition [na] ‘on’. These reported syntactic effects have been confirmed in a corpus study (Guzmán Naranjo & Bonami, 2021) and in acceptability judgements Bermel and Knittl (2012).

The syntactic effect of preposition is demonstrated in Table 5.4, which shows the distribution of inflected forms for the noun [rɪbɲiːk] 'pond' combining with the above genitive and locative prepositions. This noun has different baseline rates for *-u* in the two cases: *-u* is much more common in the locative than the genitive; *-ɛ* (which triggers a consonant alternation) is relatively uncommon. Within each case, we see that the preposition has a substantial effect on the case suffix: for example, [o rɪbɲiːts-ɛ] 'about the pond' is quite rare, while [v rɪbɲiːts-ɛ] 'in the pond' is more robustly attested, though still less common than [v rɪbɲiːk-u].

		genitive					*locative*		
preposition		*rɪbɲiːk-u*	*rɪbɲiːk-a*	*% -u*	preposition		*rɪbɲiːk-u*	*rɪbɲiːts-ɛ*	*% -u*
z	'out of'	1816	6612	21.5%	v	'in'	14 227	6510	68.6%
ftʃɛtɲɛ	'including'	37	65	36.3%	o	'about'	237	15	94.0%

Table 5.4: Distribution of suffixes for the doubly variable noun [rɪbɲiːk] 'pond' after the genitive prepositions [z] 'out of' and [ftʃɛtɲɛ] 'including' and the locative prepositions [v] 'in' and [o] 'about'

5.2 Formal analysis

The basic analysis of the Czech genitive and locative is very simple and parallel for the genitive and locative: nouns that take *-a* in the genitive are marked with one diacritic feature, nouns that take *-ɛ* in the locative are marked with another, and speakers learn a morphological dependency between the two features with a constraint in their sublexical grammars, as in the previous cases. I spend the remainder of this section discussing the theoretically more complex issue of variation.

5.2.1 Handling variation

In Section 3.4, I laid out a model of lexical and syntactic variation using the example of the Czech locative. In this model, a Czech masculine noun that takes locative *-ɛ* (at any frequency) is associated with a variable feature [+lvar], which has a parameter strength *b* corresponding to its rate of locative -*ε*. This accounts for the lexical variation. Similarly, I proposed that Czech has two locatives distinguished by a variable feature [+core], and that each locative-assigning preposition has the [+core] feature in its syntactic lexical entry also associated with a weight *b* corresponding to its propensity to take complements with locative -*ε*. Each time a speaker wishes to form a locative, she chooses a derivation that either contains or lacks the [+core] and [+lvar] features; the likelihood of these features entering the derivation corresponds to the weight of these features on the lexical entry of the preposition and noun, respectively. (See Section 3.4.3 for full details.) These features then correspond to locative realization using the rules in (36), repeated here.

(36) *Rules of realization for Czech locative with lexically and syntactically conditioned variation*

a. [LOC, +core] ↔ ε / [+lvar]___

b. [LOC] ↔ u

When the selected derivation includes both [+core] and [+lvar], the locative is spelled out as -*ε*; otherwise, it is spelled out as -*u*. Thus, the higher the weight of these features in the respective lexical entry of the preposition and the noun in a locative prepositional phrase, the higher the likelihood that that noun will end up with locative suffix -*ε*.

The genitive shows the same patterns of lexical and syntactic variation as the locative, as described in Section 5.1.2, so it can be handled in the same way: with a variable feature [+gvar] whose lexical parameter controls the rate of a given word's production with genitive -*a*. The syntactic variation is handled in the same way as the locative: the Czech genitive is split into cases with and without [+core], and prepositions assign [+core] to their genitive complements at different rates.

(43) *Rules of realization for Czech genitive with lexically and syntactically conditioned variation*

a. [GEN, $+$core] $\leftrightarrow$ a / [$+$gvar]___

b. [GEN] $\leftrightarrow$ u

While the distribution of the two locatives is biased towards particular semantic categories and contexts as described above, I assume that the distinction is at heart morphosyntactic, and does not correspond to any categorical difference in semantics.

My proposal for the Czech genitive and locative mirrors that of Jakobson (1984) and Chvany (1986) for Russian, similar to that used by Müller (2004) as discussed in Section 2.1.5.1: both "cases" are really two largely overlapping cases that differ in a syntactic feature, which I call [$+$core]. This [$+$core] is a purely formal syntactic feature, part of a project of decomposing cases into a number of binary features. For example, Müller (2004) decomposes the Russian locative as [$-$subject, $-$governed, $+$oblique]. This approach to Russian has three binary features, yielding $2^3 = 8$ logical combinations, but only six cases. The two missing cases are the [$-$oblique] counterparts of the genitive and locative, which is why, as mentioned in Section 3.4, Jakobson (1984) proposes that there are really two genitives and two locatives. The Czech vocative, which Russian does not have, causes problems for this neat feature decomposition, but if we were to pursue it further, we could identify [$\pm$core] with [$\pm$oblique]: the [$+$core] genitive and locative are the "standard" cases, while the [$-$core] versions are the "extra" cases. These features do not necessarily correspond to any semantic properties, certainly not obligatory formal semantic units, since cases can appear with a wide range of meanings. There is a rough correspondence, in that prepositions with certain meanings (for example, prepositions of location and movement, in both cases) are more likely to assign the [$+$core] case feature to their objects. In any given derivation, however, there can be a triple dissocation between meaning, presence of [$+$core] in the syntax (which is not directly detectable), and case allomorph.

5.2.2 Are all words lexically variable?

I now turn to another theoretical issue about the structure of variable features in the lexicon. The studies in this dissertation concern how speakers extend lexical patterns to new words, and I generally assume that they do so by deciding whether to assign features to newly formed lexical entries using sublexical grammars. When this feature is itself variable, like [+lvar], there is an additional question: what weight to assign to [+lvar], as discussed in Section 3.4.5.3.

However, this still leaves one question unanswered: are categorical nouns treated as extremes of variable nouns, or as categorically distinct from them? There are two analyses compatible with the patterns of variation seen in Czech, where the majority of nouns categorically take *-u* in the locative. The difference between them is the lexical representation of categorical *-u* nouns. In the discussion, I assume a *gradient extremes* analysis: all nouns have [+lvar], and nouns that show categorical behavior have assigned such a low weight to the feature that the probability of even a single token of locative *-ɛ* being recorded is vanishingly small. In this analysis, the nonce word task is modelled as assigning a lexical weight to a novel noun's [+lvar] feature. The higher the assigned weight, the more likely a response of *-ɛ* in the observed trial. However, the nonce word study in Section 5.4 cannot distinguish between different baseline rates of nouns for a given speaker, since each speaker only sees each nonce word once.

The alternative is a *categorical* analysis: nouns that categorically take *-u* lack the [+lvar] feature entirely. That is, there is a categorical split between words that *never* take *-ɛ* (no [+lvar]) and words that take *-ɛ* sometimes, even very rarely ([+lvar] with low weight). In determining a new word's behavior, the sublexical grammar decides to either attach [+lvar] or not. Then, secondarily, if a noun is given the [+lvar] feature, the sublexical grammar determines its weight as described in Section 3.4.5.4. Thus, this analysis requires an extra step in determining the behavior of a nonce word: presence vs. absence of [+lvar], and if so, then [+lvar] weight.

The gradient extremes analysis and the categorical analysis make different predictions (although the studies in this chapter cannot distinguish between them). If some nouns lack [+lvar], as in the categorical analysis, the factors influencing presence or absence of [+lvar] may be different from those influencing the parameter value of [+lvar]. In fact, the factors determining feature *presence* and the factors determining feature *weight* can even pull in *opposite* directions: for example, one group of nouns (those ending in labials, say) may have a relatively high rate of variable nouns (labials are good predictors for the *presence* of [+lvar]), but also more correlated with low rates of locative -ɛ (labials are good predictors of *low* [+lvar] parameter values). In this example, the effect of labials is to raise the probability of a word taking locative -ɛ *some of the time* (by raising its probablity of having [+lvar] at all), but also to lower the *frequency* with which it takes -ɛ (by lowering its [+lvar] parameter). On the other hand, if all nouns are given a [+lvar] parameter as in the gradient extremes analysis, such effects are impossible: if a given factor predicts a higher parameter value, it must predict both a higher *rate* of locative -ɛ for variable nouns and a higher probability of *being a variable noun at all*, since in this approach, categorical behavior is simply the extreme of variable behavior and all nouns are placed along a single spectrum.

What would evidence for the two analyses look like? In Section 5.3.3, I tentatively conclude that we see a split pattern: categorically, nouns with final fricatives and/or long vowels in the last syllable are *more likely* to take -ɛ some of the time (than nouns ending in alveolars with final-syllable short vowels), but given that they are variable, they take -ɛ *less often* than the baseline. If speakers have learned this pattern, it would provide evidence for the categorical analysis in which categorical -*u* nouns lack [+lvar] entirely. However, if they fail to learn it, this would support the gradient extremes analysis in which categorical -*u* nouns have [+lvar] at a very low weight. This would constitute a *surfeit of the stimulus* (Becker et al., 2011): speakers fail to learn certain patterns present in the input because of the architecture of their grammar. To test this difference, we would need to establish each speaker's baseline rate of locative -ɛ for each nonce word, which requires many repeated trials. In my nonce word study, each speaker only sees each nonce word once, so we cannot establish a speaker's underlying rate of -*ɛ*. (I find similar issues with other proposed model comparisons, for example in Section 5.4.6.4.) Thus, the difference between these hypotheses is not directly tested here. In the nonce word study in Section 5.4, I assume the categorical analysis, grouping words categoricallly into those that always take -*u* in the locative and those that sometimes or always take -*ɛ*. This allows for the cleanest analysis of the experimental data at the resolution available to me.

5.3 Czech synchronic corpus study

In this study, I look at the factors predicting the locative realization of Czech masculine inanimate hard-stem nouns. As with previous studies, I look at two sources of information: the phonological shape of a noun and its inflectional patterns, in particular whether nouns whose genitives can be formed with -*a* are also more likely to allow locative forms with -*ɛ*. Similarly to the Hungarian corpus study in Section 4.3, the primary goal of this study is to create a baseline expectation of the lexical patterns that participants are expected to apply in the experiments in Section 5.4 and Section 5.5. For the nonce word study in Section 5.4, I use the effects of the categorical model of the lexicon in Section 5.4.6.4 as a representation of the generalizations speakers have learned; for the variable word study in Section 5.5, I use the coefficients calculated in Section 5.3.3.1 as representations of the baseline rate of -*u* for existing lexical items adjusted for the syntactic contexts in which they appear.

5.3.1 Data

As a corpus, I used nouns from SYNv11 (Křen et al., 2022), a corpus from the Czech National Corpus comprising a large number of texts, mostly from the past two decades. This corpus is morphologically annotated, so I searched for masculine inanimate nouns in the genitive and locative preceded by a preposition (for a full list of prepositions used, see Table **??**) and any number of adjectives in the same case. For my main corpus study, I discarded words that included uppercase or numbers or had uncommon spelling or stem changes (mostly loan words), as well as any nouns that did not appear at least five times in the two cases combined, with at least one token in each case. I also excluded soft-stem nouns; nouns that can take either hard-stem or soft-stem endings are counted when they take hard-stem endings (see Section 5.1.1). This left a total of 10,839 nouns in the locative and 10,851 in the genitive.

In Section 5.3.3.1, I calculate the effects of individual prepositions in the genitive and locative. This requires a corpus of tokens, in which I included all nouns, not just those with five tokens or more. This data set includes 20,062,111 genitive tokens and 42,312,097 locative tokens from 21,354 types.

The majority of both types and tokens take *-u* in the locative. The totals are shown in Table 5.5, along with counts broken down by several phonological characteristics (final C place, final coda, stem length) and, for lemmas, the correlation between genitive and locative. For type counts, lemmas are classified as taking *-u* or *-ɛ* (or *-a*, for the genitive) if they do so in at least 99% of tokens; nouns that take *-u* between 1% and 99% of the time are classified as variable. The distribution is much more extreme for types (94.1%) than tokens (70.0%), suggesting that *-ɛ* is more common among more frequent words. Beyond this, we see a substantial difference according to the final consonant of the stem: nouns ending in alveolars, like [svjɛt] 'world', are much more likely to take *-ɛ* (e.g. [svjɛcɛ]) than those ending in labials (like [fstup] 'entry', locative [fstupu]) and especially velars. The avoidance of *-ɛ* with velars, in particular, is likely due to the fact that *-ɛ* triggers a salient consonant alternation: [jazɪk] 'language' has the locative [jazɪtsɛ]. However, the velar-final nouns that take *-ɛ* do so very frequently: velar-final nouns have a greater proportion of *tokens* with *-ɛ* than alveolar- and labial-final nouns. We see a similar difference if we look at complexity of stem codas: nouns ending in clusters are less likely to take *-ɛ* than nouns ending in singleton consonants (one rare exception is [most] 'bridge', whose locative is usually [moscɛ]). Here, though, the difference is much more dramatic when looking at individual tokens. Likewise, monosyllabic lemmas are slightly more likely to allow -*ɛ* than longer words, but by token count, -*ɛ* attaches *much* more often to monosyllabic stems. (The seven stems without vowels are monosyllabic stems with vowel–zero alternations in suffixed forms: [rɛt] 'lip' has locative [rtu].) Finally, the morphological effect is by far the strongest of those pictured: nouns that variably and, especially, categorically take -*a* in the genitive—like [svjɛt] 'world', whose genitive is [svjɛta]—are much more likely to at least variably take -*ɛ* in the locative as well.

	types				*tokens*		
	-u	*variable*	*-ɛ*	*% -u*	*-u*	*-ɛ*	*% -u*
total	10 201	574	64	94.1%	29 604 076	12 708 021	70.0%
labial	1113	25	14	96.6%	3 722 940	870 273	81.1%
alveolar	5552	550	49	90.3%	16 876 324	7 099 948	70.3%
velar	3536	9	1	99.7%	9 004 745	4 737 023	65.5%
singleton	6930	505	62	92.4%	20 072 574	12 551 809	61.2%
cluster	3271	69	2	97.9%	9 531 502	156 212	98.4%
no vowels	7	0	0	100.0%	54 667	0	100.0%
monosyllabic	1466	96	9	93.3%	10 994 512	7 994 984	57.9%
polysyllabic	8728	478	55	94.2%	18 554 897	4 713 037	79.7%
genitive -*u*	9686	523	21	94.7%			
genitive variable	145	18	3	87.3%			
genitive -*a*	32	18	31	39.5%			

Table 5.5: Type and token frequency of locative allomorphs -*u* and -*ɛ*, by final C place, final coda complexity, stem length, and genitive

The main takeaway from the type counts in Table 5.5 is that phonology does not have much of an impact, unlike the Hungarian possessive discussed in Chapter 4. Instead, by far the greatest effect is the morphological one: the small number of genitives that allow *-a* are also much more likely to allow locative *-ɛ*. The velar avoidance of *-ɛ* has a clear functional explanation: avoidance of a salient stem change. What about the other effects shown, where shorter words ending in singletons are more likely to take *-ɛ*? This likely has the same historical explanation as the relative abundance of *-ɛ* in the token counts: *-ɛ* is the older suffix, and is more likely to appear with frequent native words. These are more likely to be shorter, and also, probably, more likely to end in singletons: inspection of the data suggests that many of the words ending in clusters are loans from English, German, and French.

Indeed, more frequent words are more likely to take *-ɛ*, as shown in Table 5.6. Interestingly, this is because there are fewer categorical *-u* words as frequency increases: the raw number of words that take *-ɛ* is remarkably steady across the frequency spectrum. This is due in large part to the presence of low-frequency compounds built off of high-frequency locative *-ɛ* words, like [ɛʊrosvjɛt] 'Euro-world', which appears six times with the locative [ɛʊrosvjɛcɛ] and is transparently built from the very common word [svjɛt] 'world'. In addition, less frequent words are more likely to be classified as variable due to a small number of stray *-ɛ* tokens: [dukaːt] 'ducat' takes locative *-ɛ* only once in its 28 locative tokens, but this pushes the proportion of *-ɛ* tokens below the 99% threshold needed to be classed as categorical. Since these factors inflate the frequency of variable or categorical *-ɛ* among less common words, the frequency effect becomes even stronger once we take them into account.

lemma frequency	*-u*	*variable*	*-ɛ*	*% -u*
5–50	5154	119	38	97.0%
50–500	2749	149	5	94.7%
500–5,000	1512	144	2	91.2%
5,000–50,000	611	121	5	82.9%
50,000–500,000	168	37	11	77.8%
500,000–5,000,000	7	4	2	53.8%
5,000,000–	0	0	1	0.0%

Table 5.6: Type counts of locative allomorphs -*u* and -*ɛ*, by lemma frequency

As will be shown in Section 5.3.3, there are other phonological tendencies that emerge as significant. However, these do not have an obvious functional or historical explanation—as far as I can tell, they are simply patterns that have emerged in the data. Speakers are nonetheless expected to learn and productively apply them, so I include them in my model.

5.3.2 Methods and analysis

As with the other corpus studies, the main goal of this study is to model speakers' representation of lexical patterns, so I test a collection of assorted phonological and morphological properties of nouns to determine what speakers are expected to have learned. As discussed in Section 5.2.2, for this purpose I make the simplifying assumption that nouns are categorically grouped as taking either -*u* in the locative (categorically) or -*ɛ* (categorically or variably), and use the categorical phonological model in Section 5.3.3.3 to predict the experimental results in Section 5.4. However, since lexically variable items are of theoretical importance, I analyze the lexicon extensively with these in mind as well.

In the analysis presented in Section 5.2, speakers handle lexical and syntactic variation in the locative through weighted parameters: each noun (or at least, each variable noun) has a [+lvar] feature whose weight indicates its baseline rate of -*ɛ*, and each preposition has a [+core] feature that indicates its propensity to select -*ɛ*. If speakers are generalizing over their mental lexicon, they should be generalizing over the presence and weight of [+lvar] features in lexical entries, *not* over a noun's actual rate of -*u* in the language. Accordingly, in this corpus study I look for phonological and morphological predictors of a noun's "preposition-adjusted locative coefficient", a derived measure which I describe below, rather than its actual rate of locative -*u* in the corpus. Since I wish to establish a correlation between a word's genitive and locative, I use the same process to find preposition-adjusted genitive coefficients.

The first step of the analysis is to calculate preposition-adjusted coefficients for the genitive and locative. For each case, I fitted a mixed regression model whose dependent variable was the case suffix taken by a given token of a noun following a preposition. The regressions included a fixed factor of preposition and a random intercept of noun lemma. In this model, preposition and lemma each contribute to the predicted likelihood of -*u* in a given token. The effect sizes for the prepositions correspond to the [+core] feature weights, and the random intercepts represent the [+lvar] weights—what I call the preposition-adjusted coefficients.

As with the other corpus studies, I then fitted pairs of logistic regressions. The first included hand-picked phonological predictors of the stem (final coda size, number of syllables, final C place, final C manner, final syllable nucleus type, last V length, last V backness, last V height), while the second included all of the phonological variables as well as preposition-adjusted *genitive* coefficient, which was calculated in the same way. (One of my regressions operationalized genitive realization differently, as described below.) I built up the regressions one factor at a time using a forward stepwise algorithm using the `buildmer` function in R from the package of the same name (R Core Team, 2022; Voeten, 2022). This function adds factors to the model one at a time such that each additional factor improves the model's Akaike Information Criterion (AIC), which measures how well the model fits the data while penalizing model complexity (that is, number of factors).

I fitted these paired phonological and morphological regressions using different dependent variables reflecting different operationalizations of locative realization. First, I ran linear regressions whose dependent variable is preposition-adjusted locative coefficients for all nouns, both categorical and variable. This corresponds to the gradient extremes analysis in Section 5.2.2 that *all* nouns have a [+lvar] feature, and categorical -*u* nouns have it with a very low weight. Next, I ran binomial regressions whose dependent variable was locative realization coded as a binary: categorical -*u* ($< 1\%$ -*ε*) vs. other ($> 1\%$ -*ε*). For the morphological regression in this pair, I also coded genitive categorically, with three levels: categorical ($> 99\%$) -*u*, categorical ($> 99\%$) -*ε*, and variable. This pair of regressions reflects the categorical analysis in Section 5.2.2 that categorical -*u* nouns lack the [+lvar] feature entirely, splitting nouns into those that have and lack [+lvar]. As explained above, I use these regressions as representations of the lexicon in my analysis of the nonce word study. In this analysis, only variable nouns have [+lvar]. Thus, I ran a third pair of regressions. These were similar to the first in that the dependent variable was preposition-adjusted locative coefficient; however, this latter pair only looked at *variable* locative nouns.

The second and third pairs of regressions are intended to isolate any effects that are located solely in the categorical or variable nouns but not both. This is relevant for the discussion in Section 5.2.2, where I showed that the categorical analysis of variable nouns predicts that there may be differences between categorical and variable effects, while the gradient extremes analysis does not. There are many fewer variable words, so the third pair of regressions include much less input data. Phonological effects that are significant in the categorical regression may not be significant in the variable regression. However, this may be due to lack of statistical power in the latter. In testing for a difference between categorical and variable effects, I will take as positive evidence for the gradient extremes analysis *significant effects* moving in *opposite directions* between the categorical and variable regressions—mere absence of variable effect is not enough evidence that the two should be treated differentially.

5.3.3 Results

5.3.3.1 Calculating preposition-adjusted coefficients

I first calculate preposition-adjusted locative and genitive coefficients for each noun. These are the random intercepts for nouns in a logistic regression plotting the locative of a given token whose fixed effect is preposition, listed in decreasing order of token frequency; the baseline is the most common preposition, [v]. The effects of the locative regression are shown in Table 5.7.[2] This table shows that the preposition [na] 'on' is slightly but significantly more likely to condition locative -*ɛ* than [v] 'in', while the other prepositions ([o] 'about', [po] 'along, after', and [pr̩i] 'during') cooccur with -*ɛ* much less often than [v] does. As expected, this model is skewed heavily negative, since the vast majority of tokens (and lemmas) take -*u*. The vast majority of the variance in this model is explained by the random intercepts for noun lemma: the conditional $R^2 = .985$, while the marginal $R^2 = .008$. This makes sense, as nouns can cover the full range from categorical -*u* to (near-)categorical -*ɛ*, and the variance in random intercepts is very large accordingly.

Random effects	*variance*	*SD*		
Lemma	221.65	14.89		
Fixed effects	*β coef*	*SE*	*Wald z*	*p*
Intercept	**–16.03**	**.23**	**–69.16**	**<.0001**
Preposition (default: v)				
na	**0.35**	**.01**	**67.74**	**<.0001**
o	**–3.31**	**.01**	**–433.94**	**<.0001**
po	**–1.11**	**.01**	**–169.89**	**<.0001**
pr̩̊i	**–1.74**	**.01**	**–185.07**	**<.0001**

Table 5.7: Effects of mixed logistic model predicting locative suffix for tokens in the corpus used to calculate preposition-adjusted locative coefficients, with significant effects bolded

[2]This model includes the five prepositions taking locative: [na] 'on', [o] 'about', [po] 'along, after', [pr̩i] 'during', and [v] 'in'.

These results are similar but not identical to the findings of Guzmán Naranjo and Bonami (2021, p. 24), who found that variable nouns took -*ɛ* most often with the preposition [v], followed by [na] and [po], with [o] and [pr̥̝i] showing highest rates of locative -*u*. There are two differences between their analysis and mine: first, I include all nouns, not just ones with variability in the locative. However, this is accounted for by the random intercepts, and limiting my analysis to variable nouns produces identical fixed effects for the prepositions, so this is not the source of the discrepancy. Second, their analysis is based on raw token counts, not a statistical residualization. Thus, their analysis conflates the propensity for a given preposition to change a given noun's baseline rate of -*ɛ* with the propensity for a given noun to combine with a given preposition in the first place. That is, the discrepancy is expected if nouns with higher rates of -*ɛ* combine more often with [v] than with [na] (however, this is does not seem to be the case in my data). My analysis separates out these two factors by comparing the rate of -*ɛ* for each noun with each preposition.

As we see in Figure 5.1, a noun's preposition-adjusted coefficient (plotted on the y-axis) corresponds quite well to its actual rate of locative -*u* (plotted on the x-axis in a log odds scale, with higher rates of -*u* at the left). That is, adjusting a noun's baseline distribution of locative allomorphs to take into account the prepositions it combines with usually does not substantially affect its baseline predicted rate. On the left edge are the vast majority of words that categorically take -*u* (technically, these are at negative infinity, since a probability of 0 corresponds to a log odds of $-\infty$). Their preposition-adjusted coefficients are clustered at a very low number, with slight differences corresponding to the nouns' frequency and the distribution of prepositions with which they appear (the effects of which will be explained in the next paragraph). Likewise, the right edge (technically, at infinity) shows words that categorically appear in the corpus with -*ɛ*. In between are words with variable locative realization. In this and subsequent figures, each dot represents a noun, and the size of the noun corresponds to its token frequency in the corpus: larger nouns are more common.

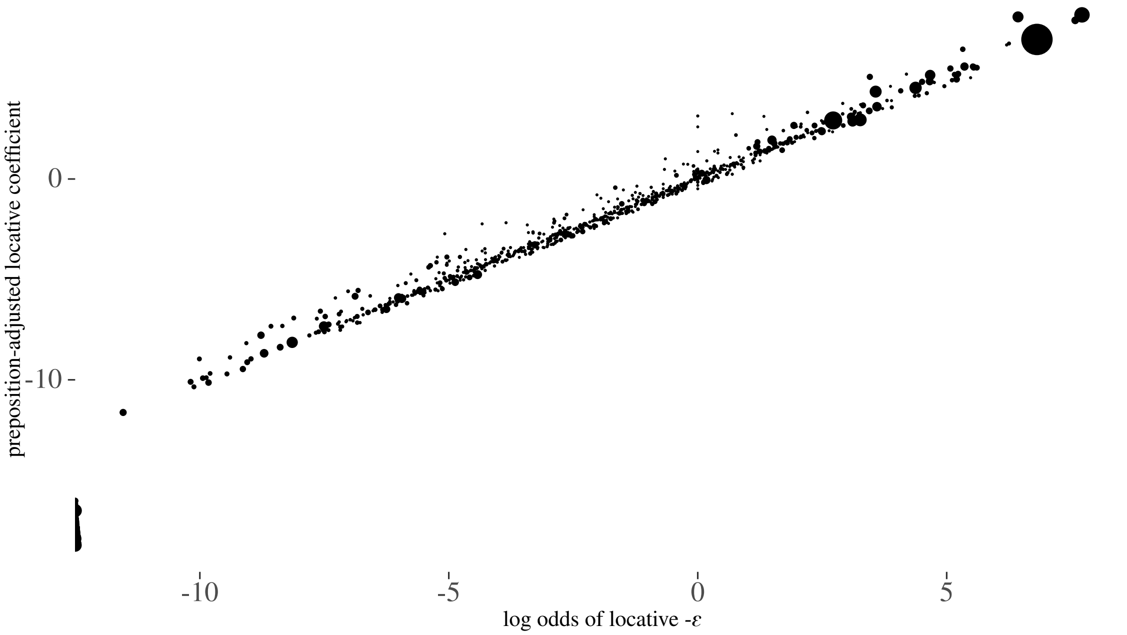

Figure 5.1: The relation between preposition-adjusted locative coefficient and locative -ɛ for individual nouns, sized according to their frequency (note: the *vast majority of nouns* are in the cluster at the bottom left, indicating higher rates of -*u*)

There are two main deviations from the straight line in Figure 5.1. The first is that some words have a higher preposition-adjusted coefficient than expected from their actual rate of locative -ɛ. This is due to the distribution of prepositions with which they appear: the vast majority of locative tokens occur with the preposition [v] 'in', which prefers the ending -ɛ. Words with high rates of prepositions that prefer -*u*, like [o] 'about, during' and [po] 'after', will thus appear above the main line in Figure 5.1. This is the case with [polotʃas] 'halftime', for example. Next, some nouns that categorically take -ɛ have a *lower* preposition-adjusted coefficient than nouns with rates of -ɛ slightly below 100%. The latter tend to be very high-frequency words like [rok] 'year' and [svjɛt] 'world', while categorical -ɛ nouns generally have much lower frequency. (Indeed, many of them are compounds built off these nouns, like [psɛudosvjɛt] 'pseudo-world'.) Thus, the lower preposition-adjusted coefficients reflect some amount of uncertainty about whether these words are "truly" categorical or merely undersampled: if they appeared as often as [svjɛt], they would probably get a few stray tokens with -*u*, as all sufficiently high-frequency words seem to.

This overlap between the preposition-adjusted coefficients of categorical -*ɛ* words and very high probability -*ɛ* words is thus desirable.

Using the same process, I calculate preposition-adjusted coefficients for the *genitive*. The fixed effects for the various prepositions are shown in Table 5.8, listed in order of decreasing frequency,[3] along with the variance and standard deviation of the random effect of lemma. The baseline is [z] ‘out of’, the most frequent preposition. Some of these prepositions are not very common, and have small effect sizes, so I only discuss several effects. Two of the most common prepositions, [do] ‘to’ and [u] ‘at’, take -*a* much more often than [z], while another common one, [od] ‘from’, takes -*a* less often than [z]. Three less common prepositions—[podlɛ] ‘according to’, [uprostr̝ɛd] ‘in the middle of’, and [za] ‘during’—also cooccur with genitive -*a* more often than [z] does. This model is skewed even more heavily negative than the locative preposition model in Table 5.7, which is expected, because -*u* is even more common for the genitive than the locative. As for the locative, most of the variance in this model is explained by the intercepts for lemma: the marginal R^2 is .003, while the conditional R^2, taking the random effects into account, is .994.

[3]This model includes a fairly large number of the most common genitive-taking prepositions: [bjɛfiɛm] ‘during’, [bɛz] ‘without’, [bliːsko] ‘near’, [do] ‘to’, [kolɛm] ‘around’, [kromɲɛ] ‘except’, [miːsto] ‘instead of’, [od] ‘from’, [ofilɛdɲɛ] ‘regarding’, [okolo] ‘around’, [podlɛ] ‘according to’, [pomotsiː] ‘by means of’, [prostr̝ɛdɲitsviːm] ‘by means of’, [u] ‘at’, [uprostr̝ɛd] ‘in the middle of’, [uvɲitr̝] ‘inside’, [ftʃɛtɲɛ] ‘including’, [z] ‘out of’, and [za] ‘during’.

Random effects	*variance*	*SD*		
Lemma	535.15	23.13		
Fixed effects	*β coef*	*SE*	*Wald z*	*p*
Intercept	**–20.72**	**.38**	**–54.47**	**<.0001**
Preposition (default: z)				
do	**2.70**	**.01**	**219.83**	**<.0001**
od	**–0.94**	**.02**	**–56.69**	**<.0001**
u	**0.78**	**.02**	**39.74**	**<.0001**
bɛz	–0.10	.08	–1.30	.1932
kromɲɛ	**–0.55**	**.07**	**–7.54**	**<.0001**
ftʃɛtɲɛ	**–0.82**	**.08**	**–9.81**	**<.0001**
miːsto	**–0.24**	**.09**	**–2.67**	**.0077**
vɛdlɛ	0.03	.08	0.43	.6650
pomotsiː	**–1.48**	**.15**	**–10.13**	**<.0001**
kolɛm	**0.09**	**.03**	**3.11**	**.0019**
podlɛ	**1.41**	**.09**	**16.40**	**<.0001**
prostr̝̊ɛdɲitsviːm	**–0.60**	**.22**	**–2.78**	**.0055**
okolo	**0.13**	**.06**	**2.34**	**.0193**
uvɲitr̝̊	**–0.31**	**.13**	**–2.31**	**.0211**
uprostr̝̊ɛd	**0.84**	**.09**	**9.44**	**<.0001**
za	**2.63**	**.20**	**13.48**	**<.0001**
oɦlɛdɲɛ	**–0.90**	**.31**	**–2.92**	**.0036**
bjɛɦɛm	**–1.40**	**.04**	**–31.55**	**<.0001**
bliːsko	–0.07	.15	–0.46	.6459

Table 5.8: Effects of mixed logistic model predicting genitive suffix for tokens in the corpus used to calculate preposition-adjusted genitive coefficients, with significant effects bolded

The expectation from the literature (see Bermel & Knittl, 2012) is a cline of prepositions going from friendliest to *-a* to friendliest to *-u*: prepositions of motion like [z] and [do] should take *-a* more than prepositions of location like [u] and [uprostr̝̊ɛd], which should in turn (presumably; they do not explicitly state this) take *-a* more than other prepositions like [ftʃɛtɲɛ] 'including'. This is not exactly what we see: although [z] 'out of' should be in the highest *-a* category, we see higher rates of *-a* for prepositions in all three categories, e.g. [do] 'to' (motion), [u] 'at' (location), and [podlɛ] 'according to' (other).

I do not have a full explanation for this discrepancy, but one possibility is an animacy effect: masculine animate hard-stem nouns always take *-a* in the genitive, so prepositions that appear more frequently with animate nouns should appear with *-a* more frequently across their entire distribution. Indeed, [podlɛ] 'according to' often has rational objects and [u] is used to say you are visiting someone or at someone's house (e.g. [u jan-a] 'at Jan's'). However, [do] and [z] should have similarly low rates of animate objects (expressing motion into and out of the middle of a space, respectively; in Czech bodies are generally not spaces), while [od] should take animate objects more often, as it is used in constructions like [od jan-a] 'from Jan'. In fact, [od] takes *-a* much less than [z] does. Thus, the animacy explanation cannot explain all the deviations from the expected distribution of genitive allomorphs with various prepositions.

The correlation between a noun's actual rate of genitive *-u* (on the x-axis in a log odds scale, with higher rates of *-u* at the left) and its preposition-adjusted genitive coefficient (on the y-axis) is shown in Figure 5.2. Once again, the graph somewhat obscures the actual distribution of the lexicon: the vast majority of nouns categorically take *-u* and have a preposition-adjusted coefficient around -20, appearing in a tight cluster in the bottom-left corner of the graph. Here the relationship between a noun's rate of genitive *-a* and its preposition-adjusted coefficient is not quite as tight as in the locative: adjusting for preposition has more of an effect for individual nouns. The greater spread may be due to a more even distribution of genitive prepositions for each lemma, or simply the larger number of prepositions. The interpretation of the bar on the right edge of the graph is the same as for the locative: nouns that never appear with genitive *-u* often have somewhat low frequency, so the model does not have the confidence to assign them a high coefficient. However, the genitive does have more high-frequency words that categorically take *-a*, leading to a clump of words in the top right that are high-confidence genitive *-a* words, such as [lɛs] 'forest'.

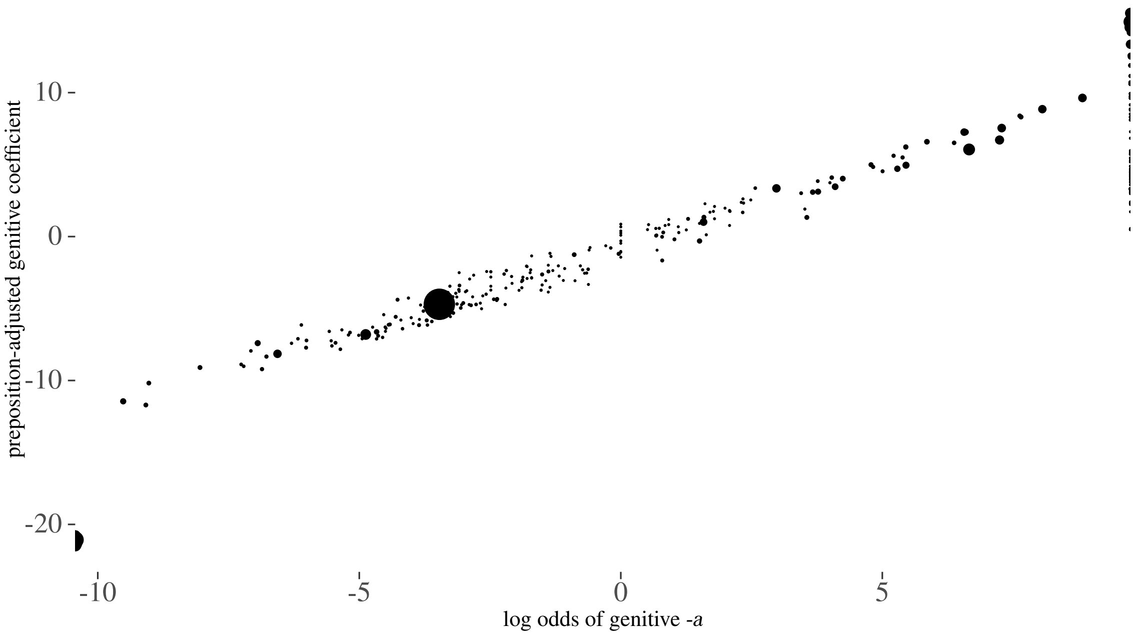

Figure 5.2: The relation between preposition-adjusted genitive coefficient and locative -*a* for individual nouns, sized according to their frequency (note: the *vast majority of nouns* are in the cluster at the bottom left, indicating higher rates of -*u*)

5.3.3.2 Preposition-adjusted locative coefficients

Now that I have shown the models that calculated preposition-adjusted coefficients, I present the phonological and morphological regressions with preposition-adjusted locative coefficients as the dependent variable, representing the gradient extremes analysis assuming that all nouns have a variable feature weight and that categorical -*u* nouns have a very low weight. These models include all nouns, categorical and variable. (I omitted nouns with no genitive tokens, and thus no preposition-adjusted genitive coefficients, so that this model could be compared to the morphological model below, which uses these genitive coefficients as a predictor.) The effects of the phonological regression are shown in Table 5.9. As in the other corpus studies, a number of the phonological factors are significant predictors of locative coefficients, and the largest effects involve the last consonant. Nouns ending in dorsals like [jazɪk] ‘language’ are less likely to take -ɛ than those ending in alveolars, like [svjɛt] ‘tongue’.

	β coef	*SE*	*Wald z*	*p*
Intercept	**–13.64**	**.14**	**–97.44**	**<.0001**
C Place (default: alveolar)				
labial	**–1.42**	**.13**	**–11.05**	**<.0001**
dorsal	**–2.18**	**.10**	**–21.13**	**<.0001**
C Manner (default: plosive)				
affricate	0.76	.69	1.10	.2706
fricative	**0.96**	**.12**	**8.07**	**<.0001**
nasal	**–1.04**	**.13**	**–8.26**	**<.0001**
liquid	**–0.99**	**.11**	**–8.61**	**<.0001**
V Length (default: short)				
long	**0.57**	**.10**	**5.82**	**<.0001**
none	–0.67	1.58	–0.42	.6730
Coda (default: singleton)				
cluster	**–0.65**	**.09**	**–7.31**	**<.0001**
V Backness (default: back)				
front	**–0.35**	**.10**	**–3.69**	**.0002**
none	–0.59	.40	–1.49	.1364
V Height (default: mid)				
high	0.18	.10	1.84	.0663
low	**0.34**	**.11**	**3.18**	**.0015**
Syllables (default: monosyllabic)				
polysyllabic	**–0.26**	**.11**	**–2.50**	**.0126**

Table 5.9: Regression model with phonological predictors of preposition-adjusted locative coefficients (lower represents more -*u*), with significant effects bolded

One striking feature of the model in Table 5.9 is its extremely low intercept: since the vast majority of nouns—in a tight cluster at the bottom left of Figure 5.1—categorically take -*u*, the regression is very skewed in favor of highly negative coefficients. Although the model does find significant effects, it still predicts that every word should have effectively categorical -*u*: the predicted coefficients range from -12.03 (corresponding to a .0006% rate of -*ɛ*) to -19.29 (corresponding to a .0000004% rate of -*ɛ*). For context, nouns that categorically take -*u* have a locative coefficient between -18.24 and -16.03, while nouns with a single locative -*ɛ* token bottom out at -11.62 (for the noun [tɛrɛːn] 'terrain', which has 102,874 tokens with -*u* and 1 with -*ɛ*). Thus, interpreting the coefficients on their own terms and not as probabilities, we see that this regression predicts many nouns to have a coefficient higher than those of categorical -*u* words, but does not quite reach the coefficient associated with variable -*ɛ* words. Perhaps because it puts many nouns in this no man's land, this model is quite a poor fit, with $R^2 = .08$.

We will now add the morphological factor: the preposition-adjusted *genitive* coefficients. The relationship between nouns' genitive and locative coefficients is shown in Figure 5.3. This chart contains four groups of nouns: the vast majority of nouns sit in a tight cluster on the bottom left representing categorical -*u* in both cases. The line running along the left edge represents nouns that are categorical -*u* in the genitive but variable in the locative, and vice versa along the bottom, which shows nouns that are categorically -*u* in the locative but variable in the genitive. Finally, the cloud of nouns in the top and middle are variable (or never take -*u*) in both cases. This graph clearly shows a positive correlation among variable nouns: in the variable cloud, nouns with higher genitive coefficients also tend to have higher locative coefficients. However, from this figure, it is difficult to see whether any categorical effect holds (that is, whether nouns that sometimes take -*a* in the genitive, at any frequency, are also likely to allow -*ɛ* in the locative. I return to this question when I look at categorical effects below.

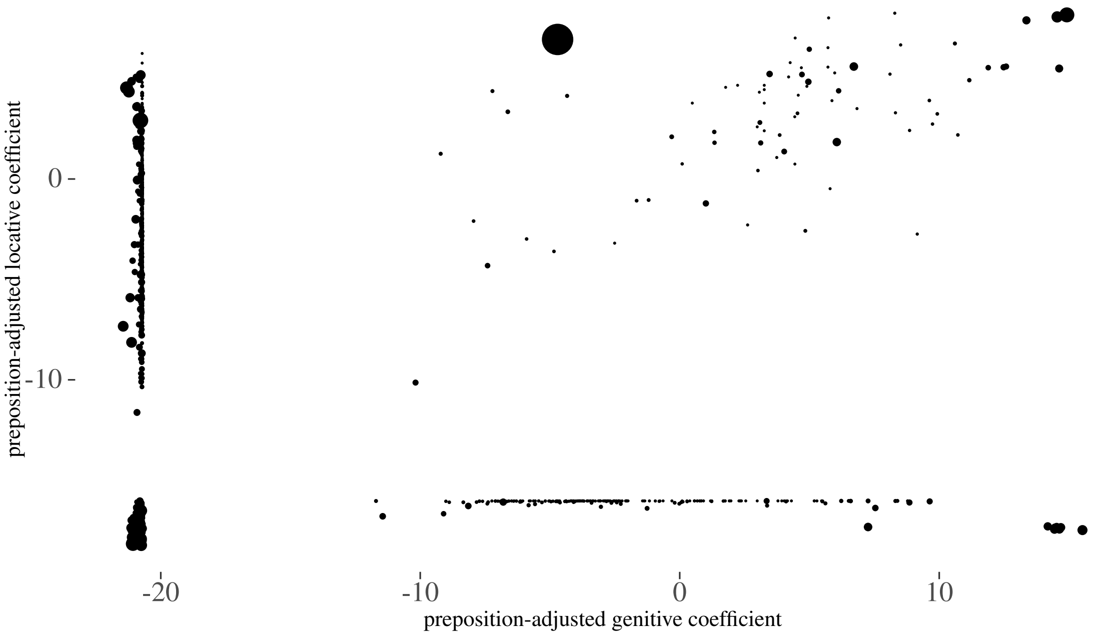

Figure 5.3: The relation between preposition-adjusted genitive and locative coefficients for individual nouns, sized according to their frequency

We can now add preposition-adjusted genitive coefficient to our model to see whether the correlation between genitive and locative coefficients is robust in the context of the phonological effects. The morphological model built up by stepwise comparison includes genitive coefficient and the same phonological factors as the phonological model in Table 5.9. Adding genitive coefficient significantly improves the model ($F = 602.3$, $p < .0001$). In fact, genitive coefficient is a better predictor of locative coefficient than any of the phonological factors: it is added first to the model, and brings the fit up to $R^2 = .13$. Thus, this model shows that nouns that are more likely to take *-a* in the genitive are also more likely to take *-ɛ* in the locative, which was the predicted morphological dependency. The full model is seen in Table 5.10.

	β coef	*SE*	*Wald z*	*p*
Intercept	**–8.62**	**.25**	**–35.08**	**<.0001**
Genitive coefficient	**0.25**	**.01**	**24.54**	**<.0001**
C Place (default: alveolar)				
labial	**–1.43**	**.13**	**–11.43**	**<.0001**
dorsal	**–2.18**	**.10**	**–21.80**	**<.0001**
C Manner (default: plosive)				
affricate	0.88	.67	1.31	.1906
fricative	**0.94**	**.12**	**8.15**	**<.0001**
nasal	**–1.12**	**.12**	**–9.18**	**<.0001**
liquid	**–1.03**	**.11**	**–9.21**	**<.0001**
V Height (default: mid)				
high	**0.26**	**.09**	**2.72**	**.0065**
low	**0.49**	**.11**	**4.61**	**<.0001**
none	–0.67	.39	–1.73	**.0841**
Coda (default: singleton)				
cluster	**–0.66**	**.09**	**–7.67**	**<.0001**
V Length (default: short)				
long	**0.46**	**.10**	**4.88**	**<.0001**
none	–1.14	1.54	–0.75	**.4566**
V Backness (default: back)				
front	**–0.43**	**.09**	**–4.58**	**<.0001**
Syllables (default: monosyllabic)				
polysyllabic	–0.15	.10	–1.44	.1492

Table 5.10: Regression model with phonological and morphological predictors of preposition-adjusted locative coefficients (lower represents more *-u*), with significant effects bolded

Adding genitive coefficient to the model barely changes the size of the phonological effects. However, the intercept is now much higher: -8.62 compared to -13.64 in Table 5.9. This is due largely to the distribution of genitive coefficients: nouns that categorically take *-u* in the genitive (that is, the vast majority) have a genitive coefficient between -21.45 and -20.72. Thus, these nouns have a predicted locative coefficient of at most $-8.62 + .25 \cdot (-20.72) = -13.80$, which is very similar to the intercept in Table 5.9. This means that adding genitive coefficient to the model does not substantially change the baseline predicted locative coefficient. However, the range of predicted locative coefficients goes much higher than in the phonological model in Table 5.9, whose predicted coefficients reached a maximum of -12.03. The highest genitive coefficient, for categorical genitive -*a* words, is 15.52, which corresponds to a baseline predicted coefficient (before phonological factors are taken into account) of $-8.62 + .25 \cdot 15.52 = -4.74$. This is still a very low predicted likelihood of locative -*ɛ* (0.86%), but would not yield categorical behavior for sufficiently frequent words. Thus, this model also underpredicts the existence of locative -*ɛ*, but not as severely as the phonological model in Table 5.9: words with high genitive coefficients are predicted to have locative coefficients corresponding to very low but observable rates of locative -*ɛ*. The overall pattern here, typical for the Czech locative, is that the morphological dependency between genitive and locative coefficient is the most powerful effect in the data.

5.3.3.3 Categorical effects

As discussed in Section 5.3.2, the regressions above—whose dependent variable is the locative coefficient for all words, categorical and variable—conflate categorical and variable effects. In addition, while these regressions do give a sense of various phonological effects, they are quite skewed at predicting the actual rate of locative -*ɛ* for any given noun. To address some of these issues, I now look at the effects on locative realization viewed from a categorical lens. I use the regression in Table 5.11 as my phonological representation of the lexicon in the nonce word study in Section 5.4.

First, Table 5.11 shows the phonological effects on locative realization, where nouns are coded as categorical -*u* ($>$ 99% -*u*) vs. variable/categorical -*ɛ* ($>$ 1% -*ɛ*). This is a statistical analysis of the patterns shown in Table 5.5. A higher coefficient in this model represents a prediction that a word is more likely to take locative -*ɛ* at least some of the time. As before, I removed nouns with no genitive tokens to better allow for comparison with the second model including genitive realization as a factor. The effect sizes here are very, very close to those in the previous phonological regression, Table 5.9. This makes sense given that the two regressions are plotted on the same scale (a linear increase in effect size corresponds to a linear increase in the log odds of taking locative -ɛ) and most nouns are categorical anyway. The most notable difference is that here, the effect size for *none* under final V length is much larger. This factor only includes the very small number of words with no vowels in their locative stems (that is, words with vowel–zero alternations like [rɛt] 'lip', whose locative is [rt-u]). Thus, this is not a meaningful difference, as shown by its enormous standard error and lack of significance.

	β coef	*SE*	*Wald z*	*p*
Intercept	**–1.79**	**.15**	**–11.71**	**<.0001**
C Place (default: alveolar)				
labial	**–1.42**	**.18**	**–7.97**	**<.0001**
dorsal	**–3.96**	**.33**	**–12.13**	**<.0001**
C Manner (default: plosive)				
affricate	0.71	.48	1.46	.1439
fricative	**0.78**	**.11**	**6.89**	**<.0001**
nasal	**–0.97**	**.15**	**–6.29**	**<.0001**
liquid	**–0.64**	**.12**	**–5.40**	**<.0001**
Coda (default: singleton)				
cluster	**–1.09**	**.14**	**–7.86**	**<.0001**
V Length (default: short)				
long	**0.69**	**.11**	**6.06**	**<.0001**
none	–9.42	215.90	–0.04	.9652
V Backness (default: back)				
front	**–0.36**	**.12**	**–3.10**	**.0020**
none	–0.79	.75	–1.06	.2914
V Height (default: mid)				
high	0.19	.12	1.56	.1187
low	**0.38**	**.12**	**3.15**	**.0017**
Syllables (default: monosyllabic)				
polysyllabic	**–0.32**	**.12**	**–2.69**	**.0071**

Table 5.11: Regression model with phonological predictors of categorical locative realization, with significant effects bolded

The intercept of this regression is much more reasonable: most nouns are predicted to be somewhat unlikely to allow locative -ɛ, but not astronomically so. The distribution is still skewed in that no nouns are predicted to be more than 50% likely to allow -ɛ, but at the same time, individual factors can have a substantial effect. For example, nouns ending in alveolars, like [pot] 'sweat', have a baseline prediction of -1.79 (i.e., the intercept), which corresponds to a 14.3% likelihood of allowing -*ε*. However, the baseline prediction for nouns ending in dorsals, like [rok] 'year', is $-1.79+(-3.96)=-5.75$, which is only a 0.3% likelihood of allowing -*ε*. (In fact, [rok] is one of the few exceptions, with locative [rotsɛ].) Thus, this regression shows that nouns ending in dorsal consonants are much more likely to uniformly take locative -*u*. The more moderate intercept only does so much to improve the fit relative to the locative coefficient model in Table 5.9: it is still quite poor, with $R^2 = .09$.

We can now return to the question of whether the morphological factor of genitive realization shows a categorical effect. The relationship between the categories is shown in Figure 5.3, a visualization of the morphological effect shown in Table 5.5. The primary result is that the vast majority of nouns have -*u* in both cases. However, of the small number of nouns that are variable in the genitive, a slightly greater perecentage allow -*ε* in the locative as well. For nouns that categorically take genitive -*a*, the effect is much more pronounced: the majority of these allow -*ε* in the locative as well. This suggests a relationship in the lexicon between genitive and locative realization—as was the case when comparing genitive and locative coefficients in Figure 5.3.

Figure 5.4: The relation between genitive and locative realizations, treated categorically

This correlation between genitive and locative stands out when we add genitive realization to the phonological model in Table 5.11. Adding this morphological factor significantly improves the model ($\chi^2 = 240.7$, $p < .0001$) and leads to a better, though still modest, fit ($R^2 = .15$). In this regression, genitive realization is the second factor added to the model after final C place and has a strong effect: nouns with variable genitives are more likely to allow locative -*ε*, and nouns with categorical -*a* genitive are *much* more likely to allow -*ε* in the locative. The phonological effect sizes are slightly greater in magnitude than those in Table 5.11, suggesting that adding the morphological factor to the model is constructive, allowing the phonological effects to stand out more clearly.

	β coef	*SE*	*Wald z*	*p*
Intercept	**–1.97**	**.16**	**–12.32**	**<.0001**
C Place (default: alveolar)				
labial	**–1.68**	**.20**	**–8.46**	**<.0001**
dorsal	**–4.03**	**.33**	**–12.25**	**<.0001**
Genitive (default: -u)				
variable	**1.62**	**.27**	**6.12**	**<.0001**
-a	**4.14**	**.29**	**14.20**	**<.0001**
C Manner (default: plosive)				
affricate	0.76	.49	1.57	.1170
fricative	**0.78**	**.12**	**6.65**	**<.0001**
nasal	**–1.15**	**.16**	**–6.99**	**<.0001**
liquid	**–0.68**	**.12**	**–5.61**	**<.0001**
Coda (default: singleton)				
cluster	**–1.12**	**.14**	**–7.86**	**<.0001**
V Height (default: mid)				
high	**0.36**	**.13**	**2.85**	**.0045**
low	**0.56**	**.13**	**4.51**	**<.0001**
none	–0.86	.76	–1.13	.2581
V Length (default: short)				
long	**0.63**	**.12**	**5.38**	**<.0001**
none	–9.36	214.13	–0.04	.9651
V Backness (default: back)				
front	**–0.49**	**.12**	**–3.97**	**<.0001**
Syllables (default: monosyllabic)				
polysyllabic	**–0.28**	**.13**	**–2.23**	**.0258**

Table 5.12: Regression model with phonological and morphological predictors of categorical locative realization, with significant effects bolded

5.3.3.4 Variable effects

Finally, I look at the phonological and morphological effects of genitive coefficient for variable words only. This allows me to look for possible effects to test the gradient extremes analysis of variable lexical items discussed in Section 5.2.2, which predicts that patterns that go in opposite directions for variable nouns compared to the categorical binning of nouns should not be learned. The dependent variable is preposition-adjusted locative coefficient, and I only look include nouns that take locative -*ε* at least 1% of the time. There are 615 of these, compared to 10,478 words used in the previous regressions. (As before, words that never appear in the genitive are filtered out as well, to make comparison easier once the morphological factor is added.) Despite this small sample size, we do see some significant effects.

The phonological model built by stepwise comparison is shown in Table 5.13. This model is a similarly poor fit to the previous ones ($R^2 = .12$), though for different reasons: in this case, there is no mass of categorical nouns whose coeffients are clumped closely together. Here, the coefficients are better behaved (in that they are spread out more evenly), but the fit is simply not very good. Why might this be? First of all, fewer phonological effects emerged as significant and improved the model—in particular, final C manner is absent from the model in Table 5.13 despite being one of the more significant factors in the previous regressions. More surprisingly, the effects are mostly *opposite* from those in the previous regressions: for example, the effect of final labials and dorsals is positive in Table 5.13 (suggesting a higher rate of -*ε* relative to alveolars), but negative elsewhere. If this pattern proves robust, it would indicate that the categorical and variable effects on locative realization really are different in Czech. If speakers learn this pattern, it is evidence in favor of the categorical analysis, in that it shows that speakers do not treat categorical locative -*u* words cannot be treated as extremes of variable words. However, as I explained in Section 5.2.2, this prediction is difficult to test in practice.

	β coef	*SE*	*Wald z*	*p*
Intercept	**0.82**	**.33**	**2.48**	**.0132**
V Height (default: mid)				
high	**–1.30**	**.32**	**–4.07**	**<.0001**
low	**–1.05**	**.31**	**–3.38**	**.0008**
none	–1.37	1.87	–0.74	.4628
Coda (default: singleton)				
cluster	**–1.55**	**.36**	**–4.32**	**<.0001**
C Place (default: alveolar)				
labial	**1.27**	**.47**	**2.71**	**.0070**
dorsal	1.06	.84	1.27	.2059
V Backness (default: back)				
front	**0.82**	**.31**	**2.65**	**.0083**
V Length (default: short)				
long	**–0.65**	**.27**	**–2.45**	**.0148**
Syllables (default: monosyllabic)				
polysyllabic	**–0.62**	**.29**	**–2.13**	**.0333**

Table 5.13: Regression model with phonological predictors of preposition-adjusted locative coefficients (lower represents more -*u*) for variable locative nouns, with significant effects bolded

Before we reach any strong conclusions about Table 5.13, let us reintroduce the morphological factor of genitive coefficient. Figure 5.3 above shows a tight positive correlation between genitive and locative coefficients when *both* are variable. However, these words are quite rare: only 71 nouns take -*u* less than 99% of the time in both the genitive and the locative, compared to 544 nouns that allow -*ɛ* in the locative but only have -*u* in the genitive. The density plot in Figure 5.5 shows the distribution of genitive and locative coefficients broken up into different regions, where a darker color indicates that more nouns lie in that region (the color scale is log-transformed, to allow regions with fewer nouns to still be somewhat visible). This figure includes all variable locative nouns, and is thus an alternate way of visualizing the data in Figure 5.1 with all the nouns at the bottom (which categorically take -*u* in the locative) removed. As expected, the bar on the left is much darker than the region further to the right, reflecting the fact that most variable locative nouns categorically take -*u* in the genitive. However most of these nouns greatly prefer -*u* in the locative as well: the nouns are somewhat darker (more concentrated) towards the bottom of this bar, indicating that nouns with extremely low genitive coefficients also tend to have lower locative coefficients. The doubly variable nouns are arranged in a roughly diagonal line, similarly indicating that locative coefficients are positively correlated with genitive coefficients when both are variable.

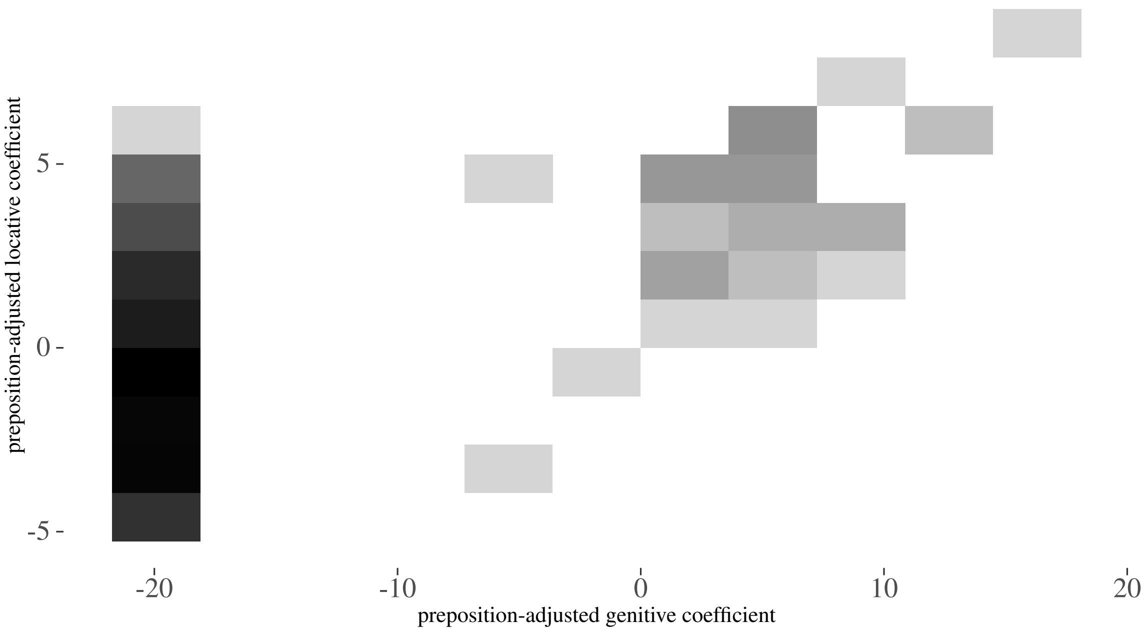

Figure 5.5: The relation between preposition-adjusted genitive and locative coefficients for variable locative nouns, colored by the log of the number of nouns in a particular region

When genitive coefficient is added to the model in Table 5.13, it becomes a much better fit ($R^2 =$ .34). Thus, genitive coefficient is by far the most important factor in this model. This model also has different phonological factors from Table 5.13: it includes final C manner but not final V backness and height. However, most of the effects still move in the same direction as in the phonological model of variable nouns Table 5.13—which is the opposite direction from the effects of the categorical model in Table 5.12. The one difference between this model and the phonological model of variable nouns is that final dorsals now have a negative effect size (that is, variable nouns ending in dorsals are less likely to take -ε), matching the effect in the other regressions (which has been discussed at length). However, this effect is not quite significant.

	β coef	*SE*	*Wald z*	*p*
Intercept	**3.94**	**.35**	**11.22**	**<.0001**
Genitive coefficient	**0.18**	**.01**	**15.24**	**<.0001**
V Length (default: short)				
long	**−1.25**	**.24**	**−5.27**	**<.0001**
Coda (default: singleton)				
cluster	**−1.32**	**.32**	**−4.12**	**<.0001**
Syllables (default: monosyllabic)				
polysyllabic	**−0.53**	**.25**	**−2.13**	**.0338**
C Manner (default: plosive)				
affricate	−1.26	.94	−1.34	.1803
fricative	**−0.66**	**.25**	**−2.67**	**.0078**
nasal	**−0.75**	**.34**	**−2.17**	**.0301**
liquid	−0.23	.26	−0.88	.3803
C Place (default: alveolar)				
labial	0.21	.41	0.52	.6017
dorsal	−1.47	.75	−1.96	.0500

Table 5.14: Regression model with phonological and morphological predictors of preposition-adjusted locative coefficients (lower represents more *-u*) for variable locative nouns, with significant effects bolded

5.3.4 Discussion

The prevalence of variable locative nouns (and near-total absence of categorical locative *-ɛ* nouns) makes statistical analysis of the distribution of locatives in the Czech lexicon rather complicated. The positive correlation between genitive and locative realization is robust across all ways of measuring. That is, nouns that allow *-a* in the genitive are more likely to allow *-ɛ* in the locative (a categorical effect), and variable locative nouns that have a higher rate of genitive *-a* are likely to have a higher rate of locative *-ɛ* as well (a variable effect).

The phonological effects, however, are less uniform, and not very strong in general. One relatively consistent effect is that nouns ending in dorsals [k ɦ x] are more likely to take *-u*. As described in Section 5.1, this is likely because the *-ɛ* suffix triggers a salient consonant alternation for dorsals:

those consonants alternate with [ts z ʃ], as in [jazɪk] 'language' with locative [jazɪtsɛ]. Thus, this effect can be seen as an avoidance of a stem change.

As discussed in Section 5.3.1, other phonological effects may reflect the historical distribution of locative suffixes: *-ɛ* is an older remnant of a previously productive suffix, and we would expect it to occur more often with older, likely more frequent native words. This may explain the fact that monosyllables are more likely to take *-ɛ* (older and native words are likely shorter), and perhaps also the fact that nouns ending in clusters are more likely to take *-u*: while Czech has profligate word-initial clusters, word-final clusters are somewhat less common, and manual inspection suggests that many of the words in my corpus ending in clusters are loans, often not particularly well assimilated; such nouns are exceedingly likely to take *-u* in general. Interestingly, these effects appear in both the categorical and variable models, which is not the case for some of the phonological effects that lack obvious historical or functional explanations.

One such effect, which goes one way in the categorical model and the other for variable nouns, involves the manner of the stem-final consonant: among the small number of variable locative nouns, those that end in a fricative and/or have a long vowel in the last syllable, like [paːs] 'belt', have a *lower* rate of *-ɛ*; viewed categorically, though, such nouns are *more likely* to allow *-ɛ*. Let us accept this as a true representation of the Czech lexicon, despite the caveat that the sample of variable nouns is quite small and the fit of phonological factors not very strong.

If speakers fail to learn these two countervailing tendencies, it would suggest a unified representation of locative variation, where categorical *-u* nouns are stored as one extreme of variable nouns; that is, with a very low value of the [+lvar] feature in Section 5.2 (the gradient extremes analysis in Section 5.2.2). This would constitute a *surfeit of the stimulus* (Becker et al., 2011): speakers fail to learn patterns in the data due to learning biases. In this case, the bias would be limitations on reasonable feature structures, such that there is pressure for lexically specific variation to be handled by a single feature.

However, as I discuss in Section 5.2.2, it is difficult to test precisely what patterns speakers have learned using a wug test. A single token of a nonce word cannot definitively show whether subjects have placed a word into a categorical category or a variable one; likewise, there is no way to ascertain a word's precise rate of variability for a word without multiple trials for the same word.

Thus, in the nonce word study to follow, I focus on the morphological effect of genitive: the phonological effects are not very strong (especially for variable nouns) and less consistent in their patterning.

5.4 Czech nonce word study

After I found a morphological dependency between genitive *-a* and locative *-ɛ* in the lexicon, I tested whether speakers extend this correlation to novel forms. Like Section 4.4 in Hungarian, this study is divided up into two sections: a stimulus norming study, where subjects evaluate novel forms for plausibility as Czech words, and the main stimulus testing study, in which subjects select novel locative forms for the stimuli given nominative and genitive forms. In the lexicon, the correlation between genitive and locative is stronger than any phonological effects, and participants show a similar pattern: although certain phonological properties of a nonce word make participants more or less likely to assign *-ɛ* to it, they assigned locative *-ɛ* much often to words to which they also assigned genitive *-a*. That is, they have learned a strong morphological dependency and only relatively weak phonological dependencies.

5.4.1 Participants

Subjects were recruited through Prolific and had to be located in the Czech Republic and raised as monolingual Czech speakers. I recruited 30 participants for the stimulus norming study and 90 for the stimulus testing study. The first three participants completed a slightly different form of the study, described in Section 5.4.3.2. One participant was rejected for poor performance on this more complex version, after which the study design was changed for the remaining 87 participants. One participant was also rejected for listing their birth place as Slovakia. Two more were rejected for anomalous behavior as described below. Thus, the results include 86 participants.

5.4.2 Stimuli

Unlike in the Hungarian study in Section 4.4, the UCLA Phonotactic Learner (Hayes & Wilson, 2008) trained on the corpus of Czech nouns used in Section 5.3 did not produce realistic stimuli. Instead, I generated 2000 stimuli using a bigram-based Markov chain, which randomly created stimuli segment-by-segment such that bigrams (that is, two-segment sequences) that were more likely in the corpus were more likely to be generated. I selected monosyllables and disyllables of the form (CV)(C)CV(C)C, where V could be either a vowel or a syllabic liquid [r l]. I removed all nonce words that are coincidentally real words according to Hajič et al. (2020), as well as any words that ended in unambiguous phonologically soft consonants [ts tʃ ʃ ʒ r̝ c ɟ ɲ j], which take a different set of endings (see Section 5.1). I also took out words with [f ɡ aʊ ɛʊ], which (almost) exclusively appear in loan words. This left a total of 427 nonce word stimuli. Each word was presented in the nominative and genitive singular, the latter of which had the suffix *-u* or *-a*.

5.4.3 Procedure

In the first study (stimulus norming), subjects rated the nonce word stimuli for plausibilty as Czech words. The ratings obtained in this study were used to select a smaller set of stimuli to be used for testing in the main experiment, which tested the morphological dependency between genitive *-a* and locative *-ɛ*.

5.4.3.1 Stimulus norming and selection

Participants were shown 50 stimuli, each presented in a frame sentence containing the target nonce word twice (as shown in Figure 5.6 below). The first occurrence of the nonce word was in the nominative singular (which is the bare stem). In its second occurrence, the stimulus had a genitive singular suffix, which was usually *-u.* For 12 randomly chosen stimuli, the genitive form instead had the suffix *-a.* Participants rated each stimulus as a potential Czech word on a scale of 1 to 5, where 1 is the most likely to be a Czech word and 5 is the least likely (this scale, where lower numbers are better, is familiar to Czechs as the grading system at school).

These ratings were then used as inputs to a Python script that selected a set of stimuli with a high average rating and a phonological distribution similar to the base corpus in the categories used as predictors in Section 5.3. I examined high-ranking sets manually and selected one with 83 stimuli to use for the morphological dependency testing phase. This set contained one real word not previously filtered out, [bar] 'bar' (a borrowing from English), so I removed this for a total of 82 stimuli.

5.4.3.2 Morphological dependency testing

Participants were shown 50 stimuli, of which 12 were presented with the genitive suffix *-a.* These were presented in the same frame sentences as in the stimulus norming study. For each trial, the stimulus (first) sentence and the target (second) sentence were chosen randomly and independently. The nouns were highlighted together with the prepositions preceding them, and subjects had to select preposition–noun collocations. As an attention check, participants had to correctly select the genitive form appearing in the first sentence. This was to ensure that participants were actually reading and internalizing the genitive form as it appeared. Once participants selected the matching genitive form, a second frame sentence appeared, in which participants had to select a genitive singular and locative singular form for the stimulus. In all cases, the stimulus was listed with two possible suffixes: -*u* and -*a* in the genitive, and -*u* and -*ɛ* in the locative. When stimuli beginning with consonants were paired up with the prepositions [v] 'in' or [z] 'out of', subjects were also given the choice of vocalized variants of the prepositions ([vɛ] and [zɛ]), which appear before some words, especially those beginning with clusters (see Dickins, 1998; Lundová, 2018). Thus, in some cases, subjects had four options to choose from rather than two. A sample trial is shown in Figure 5.6 with the prepositions [z] 'out of', which takes the genitive, and [v], which takes the locative. Other trials used frame sentences with different prepositions for the genitive and locative, or the genitive suffix *a*, e.g. [ze ʃpoda].

In our area we usually grill **ʃpod** or products made **zɛ ʃpodu**.

Please select the word's genitive form: [z lufana / **zɛ ʃpodu**]

That's correct! Now select the word in the appropriately inflected form according to you.

After removing the drill [**z ʃpoda / zɛ ʃpodu / z ʃpodu / zɛ ʃpoda**], a clean, circular hole will remain [**v ʃpoɟɛ / vɛ ʃpoɟɛ / vɛ ʃpodu / v ʃpodu**].

Figure 5.6: Trial for Czech stimulus testing study containing prepositions [z] 'out of' (takes genitive) and [v] 'in' (takes locative)

The first three participants saw a slightly different version of the study, identical to Figure 5.6 except that the drop-down menu for the locative also included forms with the locative suffix -*ovi*, which is used for masculine *animate* nouns (see Table 5.2). One participant selected -*ovi* forms about half the time, while the other two never did. The first participant was thus rejected. To prevent this from happening again, and also to make the task simpler since it was more time-consuming than expected, the design was changed so that locative forms ending in -*ovi* were not available, generally removing the possibility of animate readings for the nonce words.[4] This is particularly relevant because the genitive suffix *-a* is used for a small number of inanimates and all animates of the masculine hard stem inflection class, so offering *-ovi* as a possibility perhaps nudged certain participants towards animate readings (and, thus, genitive *-a*).

5.4.4 Analysis

For each case, both suffixes in Figure 5.6 are consistent with information given, so all trials were kept, even when the genitive selected did not match the genitive presented in the first frame sentence. This is because both genitive and locative are subject to wide-scale variation partially dependent on syntactic context, so a participant may select a genitive with *-u* for a stimulus presented with genitive *-a* for reasons other than inattention. Thus, all trials were kept, although two speakers whose genitive responses matched the presented genitive in fewer than 29 trials out of 50 were discarded.

In total, there were 4297 trials. (Three trials were discarded due to a data error.) I fitted two mixed logistic regressions whose dependent variable is the locative suffix selected by the participant (*-ɛ* or *-u*) with a random intercept for participant and item and by-participant random slopes. The first regression describes how participants used a nonce word's phonology to assign its locative. If speakers are matching the distribution of the lexicon, then the experimental results should correlate with the likelihood of taking *-ɛ* assigned to nonce words by the models of the lexicon presented in Section 5.3.3 that predict a word's locative realization given its phonological characteristics. As in the case of Chapter 4, I call these coefficients *phon_odds*; see Section 4.3.3 for an explanation of how these coefficients are calculated.

[4] In fact, many of the frame sentences preclude animate readings anyway. In Czech, the accusative for masculine nouns is syncretic with the nominative (usually null) for inanimate nouns and with the genitive (for these words, *-a*) for animate nouns. Since many of the frame sentences presented the word with a bare accusative, this should have eliminated an animate reading. However, the discarded participant sometimes selected locative *-ovi* even for nouns presented with bare accusatives, suggesting a lack of attention.

Unlike in the Hungarian experiment, in Section 5.3.3 I fitted several phonological models of the lexicon to account for widespread variability in the genitive and especially the locative. In one instance, I grouped words into two locative bins: one that always took *-u*, and one that took *-ε* at least some of the time. As an alternative, I calculated baseline rates of *-u* for each word and treated them as a continuous dependent variable for modelling. While the latter approach gives a more precise view of the lexicon, it is a precision that is not captured by the study at hand, in which each speaker forms the locative of a given nonce word exactly once. Thus, I calculate the *phon_odds* coefficients from the categorical (binned) model in Table 5.11 (whose effects are, in any case, broadly similar to the continuous-variable regression in Table 5.9): the binomial regression modelling the lexicon (always *-u* vs. sometimes *-ε*) maps most closely onto the experimental task of selecting between *-u* and *-ε* in the locative for a given trial. However, we should keep in mind that the fit is not perfect, even theoretically: if a participant selects locative *-u* for a nonce word in a given trial, that does not necessarily mean that she has assigned it categorical *-u* behavior. Another option is that she has assigned it variable behavior, then selected *-u* for that trial. However, if she selects *-ε*, that means that she has *not* assigned it categorical *-u* behavior.

The second regression predicting experimental results includes as predictors both the *phon_odds* coefficient and a morphological factor indicating a word's genitive. If speakers have learned the correlation between genitive and locative present in the lexicon, then a nonce word with *-a* in the genitive should be more likely to have *-ε* in the locative. The morphological variable also requires explanation. Unlike in the Hungarian study in Chapter 4, trials in which speakers selected a different genitive from what was presented were not discarded: because variable genitive nouns are common, speakers may select a different genitive from what was presented for a nonce word either out of inattention or because they have assigned it variable behavior. Thus, when predicting selected locative from genitive in a given trial, we can include either the presented genitive or the selected genitive, or both. I look at three variants of this second regression, which use presented genitive, selected genitive, and the intersection, respectively. Because the genitive selected by a participant for a word should be a better reflection of their mental representation of that word than presented genitive, selected genitive should be a better predictor of selected locative than presented genitive. Indeed, this is what we see: while presented genitive is clearly predictive of selected locative, this effect is mediated by selected genitive.

Both the phonological regression and the morphophonological regression also include a factor accounting for the context in which the locative is presented, namely the preposition. This is to test whether speakers apply the effect of syntactic context (discussed in Section 5.1.2) in a nonce word task environment.

Finally, as a check, I fit a regression predicting speakers' selected genitive. This includes both phonological factors (*phon_odds* calculated from a regression predicting *genitive* realization in the lexicon given phonology), presented genitive, and syntactic context (that is, preposition) for both genitives. This is to measure the extent to which presented genitive affected speakers' formation of the same genitive, and to what extent this depends on external factors. If speakers are simply ignoring the presented genitive, this factor should not be very predictive.

5.4.5 Results

5.4.5.1 Descriptive summary

The primary factor that influenced participants' choice of locative suffix for a nonce word is its genitive. Unlike in the Hungarian study in Section 4.4, phonology had only a very small effect, even smaller than the already small phonological discrepancies from the lexicon shown in Table 5.5. Those lexicon rates can be compared with the experimental rates broken down by various phonological factors in Table 5.15 below. The most salient phonological effect, as in the lexicon, was that of final C place: nouns ending in alveolars took -*ɛ* more than nouns ending in labials, and nouns ending in velars showed the highest preference for -*u*. This accords with the type frequencies in the lexicon, though not token frequencies—as shown in Table 5.5, the few velar-final nouns that take -*ɛ*, especially [rok] 'year' (locative [rotsɛ]), are so frequent that the token frequency of locative -*ɛ* for velar-final nouns is actually *higher* than for other nouns. The slight preference for -*u* among cluster-final words in the lexicon is not repeated experimentally, and the extremely small difference in the lexicon based on word length is totally erased experimentally. In general, the relationship between the distribution of locative suffixes in the lexicon and the experiment is one of flattening: the effects are less extreme than they are in the lexicon. This goes for the overall distribution of the suffixes as well: -*u* is used for 94.1% of all types in the lexicon, but only 74.4% of experimental results. As with the other studies, there seems to be greater use of minority variants; this is likely a task effect, since the explicit option to choose -*ɛ* makes it more salient than it might be otherwise.

	locative		
	-*u*	-*ɛ*	% -*u*
total	3199	1098	74.4%
labial	316	108	74.5%
alveolar	1828	719	71.8%
velar	1055	271	79.6%
singleton	2268	772	74.9%
cluster	931	326	74.1%
monosyllabic	1642	563	74.5%
disyllabic	1557	535	74.4%

Table 5.15: Experimental frequency of selected locative allomorphs -*u* and -*ɛ*, by final C place, final coda complexity, and stem length

In the lexicon, the preposition [v] 'in' conditions locative -*ɛ* on its object quite a bit more frequently than the preposition [o] 'about', but less than [na] 'on' (see Table 5.7). As Table 5.16 shows, participants in the nonce word study showed the same effect.

		locative		
preposition		*-u*	*-ɛ*	*% -u*
v	‘in’	1188	483	71.1%
na	‘on’	323	165	66.2%
o	‘about’	1688	450	79.0%

Table 5.16: Experimental frequency of selected locative allomorphs *-u* and *-ɛ*, by preposition in frame sentence

The most dramatic effect in the lexicon is the morphological dependency between the genitive and the locative, and this appears in the experimental results as well—though, as with the other effects, the discrepancy is smaller than in the lexicon. Here we must be more sophisticated in our tallying: Table 5.17 shows both the genitive presented in the first frame sentence and the genitive selected in the target sentence. Let us look at how presented genitive influences selected genitive, and how they both influence selected locative. First of all, most of the time, participants assigned the same genitive to a nonce word that they were shown. Participants assigned genitive *-a* to a word presented with genitive *-u* particularly rarely (399 out of 3267 trials, or 12.2%). Second of all, the genitive suffix has a significant effect on the locative suffix: participants selected locative *-u* much more often when they saw and chose genitive *-u* than then they saw and chose genitive *-a*. Finally, when participants chose the opposite genitive from what they were shown, the effect on the locative aligned with the genitive they selected, not the genitive they were shown.

genitive			*locative*		
presented	*selected*	*total*	*-u*	*-ɛ*	% *-u*
-*u*	-*u*	2868	2273	595	79.3%
-*u*	-*a*	399	223	176	55.9%
-*u*	total	3267	2496	771	76.4%
-*a*	-*u*	336	259	77	77.1%
-*a*	-*a*	694	444	250	64.0%
-*a*	total	1030	703	327	68.3%

Table 5.17: Experimental frequency of selected locative allomorphs -*u* and -*ɛ*, by presented and selected genitive

The overall picture presented by Table 5.17 is as follows: participants usually selected the genitive for a nonce word to match the genitive they were shown. Their choice of locative, in turn, is influenced by their choice of genitive. In cases where the participants ignored the presented genitive, their choice of locative was still influenced by their choice of genitive. Thus, the choice of locative is affected by the presented genitive, but this relation appears to be indirect, mediated by the selected genitive.

5.4.5.2 Phonology

Table 5.18 shows the effects of the mixed logistic regression predicting locative responses given random intercepts for participant and item and a fixed effect and by-participant random slope for *phon_odds* calculated from the model of the lexicon in Table 5.11, which represents the phonological patterns that speakers are expected to have learned for locative allomorphy. This model also includes a fixed effect and a by-participant random slope for locative preposition; adding the latter requires a greater tolerance for the model to converge.

Random effects	*variance*	*SD*		
Participant				
Intercept	2.16	1.47		
Phon_odds	0.02	.14		
Preposition (default: v)				
na	0.03	.18		
o	0.12	.35		
Item	0.14	.38		
Fixed effects	*β coef*	*SE*	*Wald z*	*p*
Intercept	**–0.72**	**.22**	**–3.23**	**.0012**
Phon_odds	**0.16**	**.04**	**3.83**	**.0001**
Preposition (default: v)				
na	**0.30**	**.14**	**2.13**	**.0332**
o	**–0.55**	**.11**	**–4.99**	**<.0001**

Table 5.18: Effects of mixed logistic model with predictions of the phonological model of the lexicon (Table 5.11) for experimental use of locative -ɛ, with significant effects bolded

This model shows an overall bias towards -*u*, as indicated by the negative intercept. In fact, this effect size understates the bias towards -*u*, since the *phon_odds* factors themselves are strongly negative, ranging from -7.21 for the nonce word [mɛlk] (which is predicted as extremely likely to take -*u* by the phonological model) to $-.49$ for [viːz] (predicted as only somewhat likely to take -*u*). However, speakers were much more willing to select locative -ɛ than was predicted from corpus type frequencies. Only 638 of the 10,839 nouns in my corpus (5.9%) have locative -ɛ at least 1% of the time, and even these nouns were heavily skewed towards -*u*: variable words appeared with -*u* an average of 97.6% of the time; the frequency-weighted average is 70.0%, since more frequent words appear show a greater preference for -ɛ, even among variable words. In contrast, my participants assigned -ɛ about 26% of the time (1098 out of 4297). The variance of the random effect in Table 5.18 is quite high, suggesting that participants had very different baseline rates of -ɛ usage. Indeed, six participants never selected -ɛ at all for the locative, while twelve participants chose -ɛ at least half the time, with a maximum of 36/50 -ɛ tokens. The random effect accounts for this variation; the fixed effects reveal that there are nonetheless general patterns in the behavior of individual participants.

Although a word's phonology is predictive of the experimental results, the effect is quite weak: the effect size of *phon_odds* is .15, meaning that the range of likelihoods predicted for nonce words according to the corpus model is substantially compressed in describing the experimental results. This may reflect the fact that Czech locative allomorphy is only weakly correlated with phonological factors in the lexicon (see discussion in Section 3.4.5.3). The strongest effect found in the lexicon is that nouns ending in dorsals strongly prefer -*u* even beyond the already strong baseline preference. We find traces of this effect in the experimental results as well: in general, nouns ending in [k] cluster together with low experimental rates of locative -*ɛ*. The random intercept for item indicates moderate variance, suggesting that there are some word-specific effects beyond phonology, but not a huge amount. In addition, the variance of the by-participant random slope for *phon_odds* is very small, indicating that participants applied the phonological effects more or less uniformly.

The phonological effect can be seen in Figure 5.7, which shows the relationship between the predicted likelihood of each nonce word allowing locative -*ɛ* and its experimental rate of -*ɛ*. Both axes are shown in terms of log odds, making the relationship linear: the *phon_odds* coefficients for nonce words are plotted on the x-axis, while their experimental rate of -*ɛ* is on the y-axis. Figure 5.8 shows the same plot, but the axes are plotted on untransformed scales. The two figures also include a line corresponding to the fit of the model in Table 5.11. From the latter figure, it is clear that all stimuli have quite a low predicted likelihood of allowing -*ɛ* in the lexicon, and that the experimental results are similar: no word was assigned -*ɛ* in a majority of trials. These graphs also show the effect of the random intercept for item: each word appears twice. In black we see the predicted rate of -*ɛ* based entirely on the *phon_odds* (with a fixed intercept), while the gray shows the addition of the random intercept. The effect of the random intercept is to bring the predicted rates of individual words closer to the line of best fit, encapsulating what the phonological factor cannot. Thus, the random intercept brings words closer to the line: to the left if it is below the line and to the right if it is above it. In most cases, the predicted rate of -ɛ is very low, so many of the nouns have a substantial shift to the right from their random intercept. This is clear from the graph with raw probabilities in Figure 5.8.

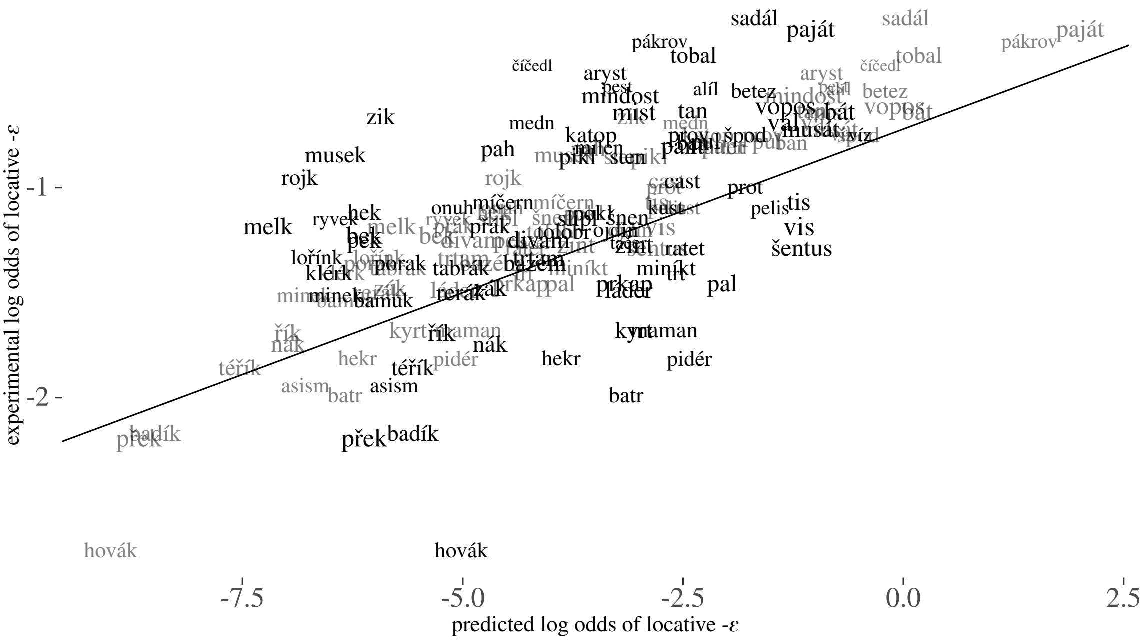

Figure 5.7: The relationship between predicted and experimental log odds of locative -ɛ for individual nonce words with (gray) and without (black) the random intercept, sized according to number of trials, with a line showing the fit of the experimental model in Table 5.11

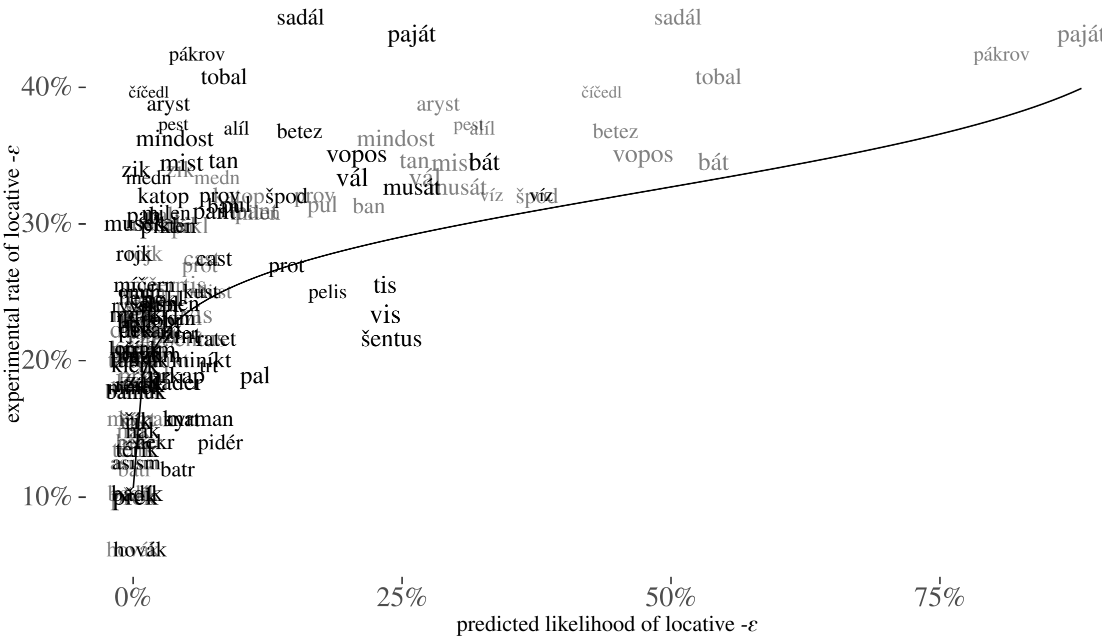

Figure 5.8: The relationship between predicted likelihood and experimental rate of locative -ɛ for individual nonce words with (gray) and without (black) the random intercept, sized according to number of trials, with a line showing the fit of the experimental model in Table 5.11

Figure 5.7 shows the fairly loose fit of *phon_odds*: the line's slope is not very steep, and while the cloud of words does follow it upward as it goes right, the cluster is not very tight.

We also see a significant effect of syntactic context: -ɛ is much more common with [v] 'in' (a "canonical" locative preposition expressing location, see Section 5.1.2) than [o] 'about' (a "non-canonical" locative preposition). In addition, -ɛ appears significantly more frequently with [na] 'on' than with [v]. This pattern of locative -ɛ distribution ([na] $>$ [v] $\gg$ [o]) is the same as that found by the mixed model that produced preposition-adjusted locative coefficients in Section 5.3.3.1. The by-participant random slopes for preposition are fairly low in variance, meaning that speakers do not vary much in their sensitivity to syntactic context.

5.4.5.3 Phonology and morphology

We can now see whether the correlation between genitive and locative is statistically significant. Recall that speakers were shown a genitive for a nonce word and had to select a genitive for that same nonce word, but the two did not have to match. Thus, we can include either or both as factors (as well as an interaction term between them). The results clearly show that selected genitive is a better predictor than presented genitive—as discussed in Section 5.4.5.1, presented genitive does predict the locative, but indirectly, mediated by selected genitive. Accordingly, I use selected genitive as the morphological factor in Table 5.19. Adding presented genitive on top of this (with or without an interaction term with selected genitive) slightly improves the fit of the model to the experimental results, but not significantly so, and not enough to overcome the penalty of adding a factor according to the Akaike Information Criterion (AIC). Thus, I leave this feature out and only include selected genitive.

Table 5.19 shows the effects of the mixed logistic regression with selected genitive added to the model in Table 5.18. As before, this model also includes random intercepts for participant and item, fixed effects for locative preposition and *phon_odds* calculated from the phonological model of the lexicon, and by-participant random slopes for all fixed effects.

Random effects	*variance*	*SD*
Participant		
Intercept	3.45	1.86
Phon_odds	0.02	.14
Selected genitive (default: -u)		
-a	1.74	1.32
Preposition (default: v)		
na	0.11	.34
o	0.17	.41
Item	0.13	.35

Fixed effects	*β coef*	*SE*	*Wald z*	*p*
Intercept	**–1.16**	**.22**	**–3.23**	**.0012**
Phon_odds	**0.18**	**.04**	**3.83**	**<.0001**
Selected genitive (default: -u)				
-a	**1.40**	**.18**	**7.69**	**<.0001**
Preposition (default: v)				
na	0.29	.15	1.89	.0596
o	**–0.64**	**.12**	**–5.40**	**<.0001**

Table 5.19: Effects of mixed logistic model with predictions of the phonological model of the lexicon (Table 5.11), preposition, and selected genitive for experimental use of locative -ɛ, with significant effects bolded

The effect of selected genitive is very large: when participants chose *-a* for the genitive of a nonce word, they were also much more likely to choose -ɛ as its locative. In fact, the morphological effect is much stronger than the phonological effect. The random slope shows that speakers vary considerably in how much they correlate the genitive and locative. This model also has a substantially greater variance for the by-participant random intercept, indicating that this more complex model emphasizes the differences in participants' baseline usage of -ɛ.

The difference in magnitude of the phonological and morphological effects can be seen in Figure 5.9 and Figure 5.10. These show the same data as Figure 5.7 and Figure 5.8, but each nonce word is split into trials according to which genitive was selected: *-u* in black, and *-a* in gray. The size of the words represents the number of trials in which they were shown; the gray words are smaller, since *-u* was selected as a genitive more often than *-a.* The two tokens of each word are connected by a line; this line is dashed for the small number of words that were always assigned locative *-u* together with genitive *-a*, since these words (whose gray tokens are at the bottom edge of the graph) are technically at negative infinity in Figure 5.9, which corresponds to a probability of 0. These graphs also show lines corresponding to the fit of the model in Table 5.19 for nonce words with selected genitive *-u* (black) and *-a* (gray). These lines are very light to make the graphs as legible as possible.

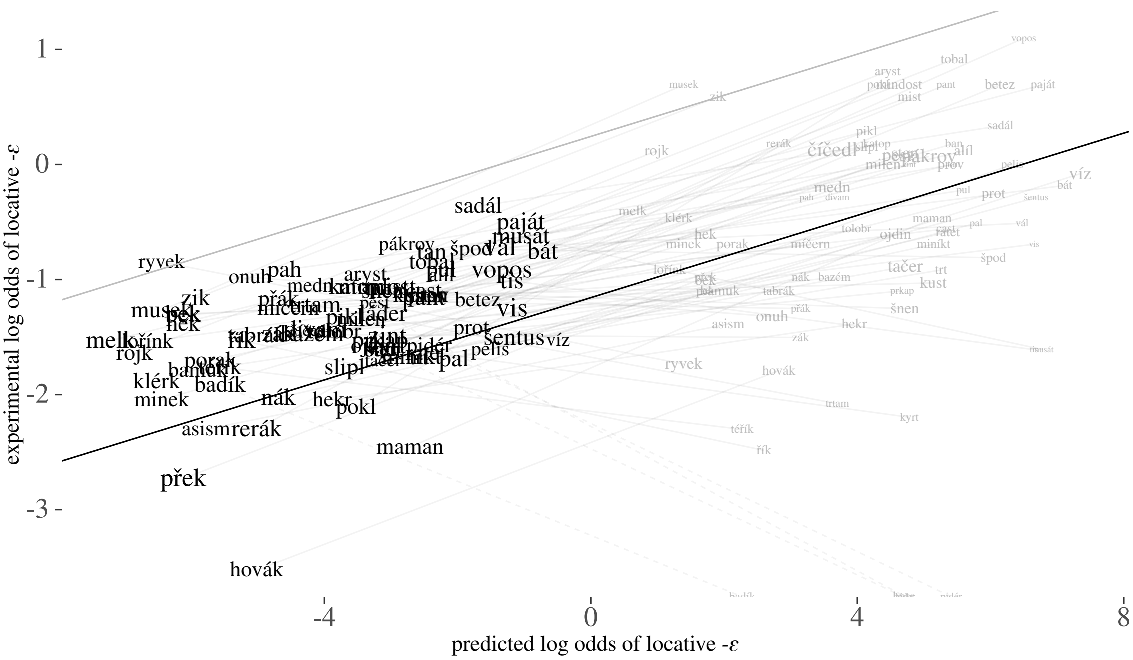

Figure 5.9: The relationship between predicted and experimental log odds of locative -ε for individual nonce words with selected *-u* (black) and *-a* (gray) genitives, sized according to number of trials, with a line showing the fit of the experimental model in Table 5.11

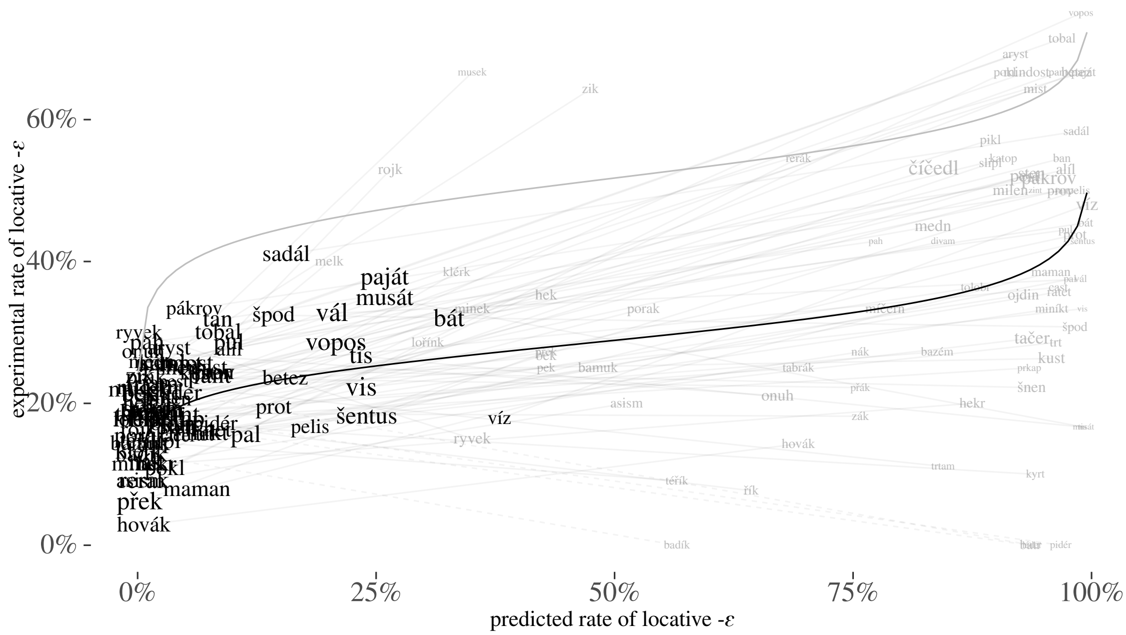

Figure 5.10: The relationship between predicted likelihood and experimental rate of locative -ɛ for individual nonce words with selected -*u* (black) and -*a* (gray) genitives, sized according to number of trials, with a line showing the fit of the experimental model in Table 5.11

In Figure 5.9, the words in the two genitive conditions have very little overlap. This means that the words with genitive -*a* are almost all more likely to take locative -ɛ than the words with genitive -*u*, regardless of their phonology. This can be seen in Table 5.19: the effect size of selected genitive is 1.40, almost 8 times greater than the effect size of *phon_odds* (.18). As mentioned above, the *phon_odds* coefficients range from −7.21 (for [mɛlk], the word least likely to take -ɛ) to −.49 (for [viːz], the most likely [-ɛ] word), a difference of 6.72. This means that going from genitive -*u* to -*a* has a slightly larger effect on locative response than going from one end of the range of phonological effects to the other.

Most of the lines in Figure 5.9 and Figure 5.10 slope upward and to the right, again indicating that most nonce words were assigned locative -ɛ more often when they were also assigned genitive -*a*. This is the strongest factor predicting locative suffix, with locative preposition and nonce word phonology also playing roles. In this model, the preposition [o] conditions -*u* more than [v]; this time, the effect of the preposition [na] does not quite reach significance. This corresponds with the fact that the difference between [v] and [o] in the lexicon (see Section 5.3.3.1) and previous descriptions (see Section 5.1.2) is much more substantial than that between [v] and [na], which are expected to pattern similarly. For all tested effects, speakers' behavior matches the distribution of locative suffixes in the lexicon.

5.4.5.4 Selecting the genitive

In the descriptive summary in Section 5.4.5.1, I suggested that locative is influenced by the genitive presented to participants, as mediated by the genitive they selected. Having shown the latter half (selected genitive predicts selected locative), I now turn to the factors influencing participants' choice of genitive. As with the locative, I look at three factors influencing choice of genitive: a nonce word's phonology, the preposition with which it appears (in both the stimulus and target sentences), and its presented genitive in the frame sentence. As expected, presented genitive is far and away the most important factor predicting selected genitive—in fact, the phonological factors predicting genitive are even weaker than they are for locative. This finding completes our understanding of the task, linking presented genitive to selected genitive to selected locative.

The first step in determining how speakers choose a genitive suffix for nonce words is to study the effect of phonology. I have not previously looked at the phonological properties predicting genitive realization in the corpus, so I do so now. Let us begin with the lexicon. Phonology is even less predictive of a noun's genitive suffix in the lexicon than of the locative: the phonological model of the lexicon predicting genitive (binned in the same way as locative, categorical -*u* vs. variable and categorical -*a*), shown in Table 5.20, has just $R^2 = .01$. In part, this low correlation may be due to the fact that even fewer nouns in my data set allow genitive -*a* (255) than locative -*ɛ* (615).[5] The

[5]There is one complicating factor: for *animate* nouns in this class, not in my data set, -*a* is the only genitive suffix. There are some nouns that have ambiguous animacy, as well as inanimate nouns (like [ʃlofiːtʃɛk] 'nap (diminutive)' and car names like 'Fiat') that can behave as syntactically animate and thus take genitive -*a*. Thus, there is a slight confound with animacy in the genitive data; however, inspection suggests that this is confined to a few stray nouns and not a systematic issue.

small number of nouns with genitive *-a* is likely not the only reason for the poor fit, though: the model itself also shows a weak effect of phonology. Strikingly, the features of the final consonant, which are directly adjacent to the suffix and generally the most important factor both within Czech and across languages, do not improve the model and are not added. Instead, the most salient effects are those of the last syllable's vowel, especially its height: stems with mid vowels [ɛ ɛː ɛʊ o oː oʊ] in their last syllables, like [lɛs] 'forest' allow genitive *-a* (e.g. [lɛsa]) more frequently than stems with low or high vowels. It is surprising that these are the most important factors, and I do not have a good explanation for this fact. One possibility is that this effect is attributable to specific suffixes that prefer *-u* or *-a*. Since my data set does not mark derivational morphemes, it is hard to measure this precisely.

	β coef	*SE*	*Wald z*	*p*
Intercept	**–3.29**	**.17**	**–19.34**	**<.0001**
V Height (default: mid)				
high	**–0.55**	**.15**	**–3.59**	**.0003**
low	**–1.28**	**.24**	**–5.41**	**<.0001**
none	0.21	.60	0.34	.7311
V Length (default: short)				
long	**0.78**	**.15**	**5.26**	**<.0001**
none	–1.47	1.25	1.18	.2365
Syllables (default: monosyllabic)				
polysyllabic	**–0.58**	**.16**	**–3.77**	**.0002**
V Backness (default: back)				
front	**0.46**	**.16**	**2.93**	**.0034**
none	–0.79	.75	–1.06	.2914

Table 5.20: Regression model with phonological predictors of categorical genitive realization in the lexicon, with significant effects bolded

The effect of preposition is more distinctive: of the 20 genitive prepositions measured in my data set, all but three have rates of *-a* significantly different from my selected baseline, [z] 'out of', in the model used to calculate preposition-adjusted genitive coefficients. (That is, with a random intercept for lemma; see Section 5.3.3.2.) The distribution of prepositions in the stimulus sentences (the presented frame sentences showing the genitive) and the target sentences (in which participants selected genitives and locatives), along with their effect sizes in the lexicon model in Table 5.8 (which are all significantly different from the baseline of [z]), are shown in Table 5.21.

preposition		*number of trials*		*effect size*
		stimulus sentence	*target sentence*	*in lexicon model*
z	'out of'	2130	2399	(baseline)
do	'into'	1614	1056	2.70
od	'from'	553	241	–0.94
u	'by'	—	365	0.78
kolɛm	'around'	—	121	0.09
ofilɛdɲɛ	'regarding'	—	115	–0.90

Table 5.21: Number of trials for genitive prepositions in stimulus and target sentences and their effect sizes in mixed model of tokens in lexicon (positive → higher rate of *-a*)

The effect sizes presented in Table 5.21 should be taken with a grain of salt. First of all, the mixed model of the lexicon requires a higher tolerance to converge: perhaps the large number of prepositions or the general skewed distribution of genitive suffixes (usually *-u*) is problematic. Second of all, these results clash somewhat with previous findings arguing that prepositions of motion like [z], [do], and [od] have higher rates of *-a* than prepositions of location like [u] and [kolɛm], which should in turn have higher rates than non-locational prepositions like [ofilɛdɲɛ] (see Bermel & Knittl, 2012). In this light, [od] has an unexpectedly negative effect size, while [u] and [kolɛm] should be lower than [z]. On the other hand, [z] and [do] should be similar, leaving the large discrepancy between them unexplained.

Using the established literature as our guide, we would expect higher rates of *-a* selected with [z], [do], and [od], though not necessarily any difference between them. However, the other prepositions, which should have higher rates of *-u*, may be too infrequent in the experiment to yield a measurable difference. Given that these first three prepositions are the only ones that appear in the stimulus sentence, I do not expect the difference between them to be substantial.

Now that we have looked at the factors predictive of genitive in the lexicon, we can see how these factors predict the experimental results. Table 5.22 shows the results of a mixed model predicting genitive responses selected in the experiment. The model is based on one built by stepwise comparison using the `buildmer` function from the R package of the same name (Voeten, 2022), where factors were added if they improved the model's Akaike Information Criterion (AIC), which balances model fit and penalizes model complexity (additional factors). Candidate factors were presented genitive, preposition in the stimulus and target sentences, plus interactions of all three, and predictions of the phonological model of the lexicon for genitive suffix, as well as dummy variables of trial number, average score of stimulus, and the order in which choices were presented. The model also included random intercepts for participant and nonce word. Only three factors made it into the model: presented genitive, the phonological coefficient *phon_odds*, and the preposition in the *stimulus* sentence (that is, the first sentence which initially showed the genitive, not the target sentence). The last factor improved the model, but none of its effects were significant. I then added by-participant random slopes for the fixed effects; doing so requires a greater tolerance for the model to converge. I manually checked the factor of stimulus preposition in the models again: once random slopes were added, this factor (in fixed effect and random slope) no longer improved the model significantly and did not improve the AIC, so I removed it. This left the model in Table 5.22 with only two fixed effects: presented genitive (very strong, as expected) and *phon_odds* (much weaker, not significant).

Random effects	*variance*	*SD*		
Participant				
Intercept	4.75	2.18		
Presented genitive (default: -u)				
-a	4.74	2.18		
Phon_odds	0.09	.57		
Item	0.33	.57		
Fixed effects	*β coef*	*SE*	*Wald z*	*p*
Intercept	**–1.86**	**.50**	**–3.72**	**.0002**
Presented genitive (default: -u)				
-a	**3.67**	**.27**	**13.75**	**<.0001**
Phon_odds	0.22	.12	1.82	.0693

Table 5.22: Effects of mixed logistic model with predictions of the phonological model of the lexicon (Table 5.20), stimulus preposition, and presented (stimulus) genitive for experimental use of genitive *-a*, with significant effects bolded

Although it is not significant, the *phon_odds* factor is stronger than it is in the model predicting selected locative in Table 5.19, .22 vs. .17. Does this mean that phonology has *more* of an effect on genitive than locative? No: since there are fewer factors in the lexicon model predicting genitive, there is much less differentiation in the genitive *phon_odds* of individual words, and their overall range is much narrower, ranging from -5.14 to -2.04 (a total range of 3.10). The effect of presented genitive on selected genitive, 3.67, is 5.4 times greater than the total range of phonology ($3.10 \cdot .22$). By comparison, the effect of selected genitive on selected locative was only slightly larger than the range of phonology, as described above. Thus, while phonology does have an effect, it pales in comparison to that of selected genitive. Speakers' choice of genitive is mostly informed by the genitive already presented, and only slightly by their weak priors about the phonology of a word (in addition to some amount of random noise).

It is perhaps not surprising that speakers differ in the correlation between the genitive they chose and the genitive they were presented with. The random intercept for participant and by-participant random slope for presented genitive have very large variances. Presumably some speakers paid more attention to the presented genitive than others, or were more susceptible to having their priors changed by the presented form. However, this does not mean that participants who leaned on their priors had heavy phonological input: the by-participant random slope for *phon_odds* does not have a very high variance, indicating that individual participants gave similar weight to the influence of phonology. Thus, the difference seems to be that some participants were selecting more randomly than others.

5.4.6 Discussion

5.4.6.1 Patterns from the lexicon

The main finding of the Czech nonce word experiment is that participants show a strong correlation between the genitive and locative that they assign to stimuli: genitive -*u* is correlated with locative -*u*, and genitive -*a* with locative -*ε*. Specifically, participants responded to the experimental manipulation in a mediated fashion: the genitive suffix with which a stimulus was shown influenced the genitive that participants assigned to it, and the genitive they selected was tightly correlated with the locative they selected. This correlation between the two cases also exists in the lexicon, indicating that Czech speakers are learning correlations between inflected forms from the lexicon and applying them productively in new situations. (We would expect the same result if the experiment were conducted in the opposite direction: locative should likewise predict genitive.) This is the primary finding of this experiment, and supports the general hypothesis of this work. However, participants are not all applying the correlation between the cases equally: there is much individual difference in this effect size. This difference may be due to differing internal grammars: some speakers have a grammatical correlation between the two cases, while others do not. It may also be due to different approaches in the task: for whatever reason, the artificiality of the task may have led certain speakers to choose one case without regard for the other, whereas this might not be the case for more naturally learned items.

The results show that speakers are also applying other patterns from the lexicon. Although the effect of phonology on locative realization is much weaker than it was in Hungarian (see Chapter 4), speakers nonetheless applied these phonological patterns in the production of new forms. In particular, the strongest phonological generalization in the lexicon is that nouns ending in dorsals are more likely to take -*u*, avoiding the salient consonant alternation that dorsals undergo before locative -*ɛ* (e.g. [jazɪk] 'language' and its locative [jazɪtsɛ]); speakers likewise assigned -*ɛ* less often to nouns ending in dorsals.

In addition, participants productively apply the syntactically conditioned variability present in the input (discussed in Section 5.1.2). In Czech, variable locative nouns tend to take -*ɛ* more often with prepositions [v] 'in' and [na] 'on', and less often with prepositions like [o] 'about' (cf. Bermel & Knittl, 2012; Guzmán Naranjo & Bonami, 2021). My model of the lexicon calculating preposition-adjusted locative coefficients, shown in Table 5.7, found that [na] occurs with -*ɛ* somewhat more often than [v], and [o] occurs with -*ɛ* much less often. Participants in this study mirrored this pattern precisely. As discussed in Section 5.3.3.1, my corpus study disagreed slightly with that of Guzmán Naranjo and Bonami (2021), who found a slightly *higher* rate of -*ɛ* for [v] than for [na]. The fact that participants used locative -*ɛ* more often with [na] than with [v] suggests that my corpus study provides a more accurate representation of the Czech lexicon than theirs, although I do not have a good explanation for what the relevant difference is.

In my analysis, syntactically conditioned variability is not a case in which speakers are productively extending patterns from the lexicon to new words. In Section 3.4, I presented a theory of syntactically conditioned variation, exemplified by Czech, in which the lexical entries for prepositions contain features whose strength conditions the surface probability of locative suffixes on their complements. Under this approach, participants in the nonce word task are stochastically choosing a locative suffix conditioned in part by their previously existing lexical entries for prepositions, as they do when forming the locative of *real* words. In this sense, it is reassuring that participants applied the syntactic patterns present in the written language: if they had failed to do so, it would indicate that the task was in some way unrealistic.

5.4.6.2 The genitive as the morphological factor

In this study, I chose not to discard trials in which the genitive suffix selected by speakers diverged from the genitive that they were shown. This is because many nouns vary in their genitive (just like the locative), so speakers who choose a divergent genitive may be exhibiting grammatically grounded variation rather than simply ignoring their input for the task. I treated discordant trials differently from the Hungarian experiment in Chapter 4, where such trials were discarded because lexical variability in the relevant case was unexpected. If discordant genitive trials are not due solely to inattention, we would expect presented genitive to create a very strong prior on selected genitive and for other effects in the input (phonological and syntactic factors) to be relevant as well, alongside some randomness to account for the stochastic selection of genitive suffixes given nouns that have been lexically specified as variable (see Section 3.4). This is largely what we find: presented genitive is by far the strongest predictor of selected genitive—though some speakers are much more sensitive to it than others—but speakers observe the (weak, rather unexpected) phonological patterns in the lexicon as well. I did not replicate the syntactic conditioning effect reported in the literature, but this may be due in part to the choice of prepositions in my frame sentences: most trials had [z] ‘out of’, [od] ‘from’, and [do] ‘into’, which should have similar rates of genitive *-a*. Other prepositions like [kolɛm] ‘around’, which should condition higher rates of *-u* according to the literature, may have been too infrequent in the test trials to have an effect (see Section 5.4.5.4).

Although I did not find an effect of syntactic conditioning in the genitive, participants’ choice for a word’s genitive was largely as predicted: it was heavily influenced by the genitive they were shown and to its phonological properties; they were not simply choosing randomly. Accordingly, presented genitive serves as a meaningful predictor of selected locative suffix, though one mediated by selected genitive—further affirming my study's main result.

5.4.6.3 Matching the theoretical model

As described elsewhere, the current study does not have the resolution to seriously probe how speakers deal with variable lexical items, so I assume for the purposes of this study that genitive and locative behavior are categorical. However, it is a point of theoretical interest, so in the remainder of this section, I discuss the models I compared in Section 3.3 in light of widespread variation, and how they might be distinguished in future studies.

In Section 3.3, I presented a single grammar model as an alternative to my proposed multiple grammar model. In the multiple grammar model, constraint-based grammars describe phonotactics for the language as a whole and for each sublexicon, defined by the presence of lexical diacritic features. The "multiple grammars" refer to the language-wide phonotactic grammar and one or more sublexical phonotactic grammars. The process of outputting a nonce word involves assigning a feature to its lexical entry. In the single grammar model, similar to stochastic constraint-based Harmonic Grammars (e.g. Legendre et al., 2006; Pater, 2016; Potts et al., 2010), there is only the language-wide phonotactic grammar, which evaluates candidates for nonce words. This grammar, however, can contain constraints indexed to particular morphemes, allowing it to similarly capture arbitrary generalizations over words that take a given suffix allomorph. In the single grammar model, nonce words have underspecified lexical entries, so their inflected froms are determined stochastically by the language-wide phonotactic grammar.

In that section, I briefly mentioned how the two models could be adapted to include gradient weighted features that index variation for individual lexical items. The multiple grammar model, presented in Section 3.4.5.3, assumes a two-step process for the experimental task in this study. First, as before, the sublexical grammar must fully determine the nonce word's lexical entry by assigning it a weight for the variable diacritic feature [+lvar]. Then, to actually generate a form, the weight of [+lvar] is used to stochastically choose a locative suffix. The strength on a given lexical item corresponds to that lexical item's likelihood of being assigned -ε in a given derivation. How, in the first step, does the sublexical grammar assign a weight for [+lvar] to a nonce word? First, the word's phonological and morphological characteristics are evaluated in the grammar and given a score. Next, some random noise is added to the score and assigned as the feature weight. That is, words are assigned a weight such that the expected weight is the most likely outcome, and outcomes close to the expected weight are more likely than those further away. The single grammar model works as before: a nonce word is evaluated on the phonotactic grammar and a form is stochastically chosen accordingly.

How do the single grammar and multiple grammar models differ in their predictions for behavioral experiments? For the current nonce word study, they cannot be distinguished. However, there are possible study designs in which they could theoretically be tested. These require participants to be repeatedly tested on the same nonce word. In Table 5.23, we see an example of the process used by the two models to generate locative tokens (ignoring the effect of syntactic context, which is the same for both). In both models, the speaker calculates a score based on phonological characteristics of the word and morphological factors (that is, the genitive form, if known). In this example, the nonce words [stod] and [blod] are very similar phonologically, so the speaker gives them both a score of -2 (where higher numbers lead to a higher likelihood of -ε). Here the two paths diverge. In the single grammar model, these scores are directly used to calculate the probability of locative -ε, using the principle that an outcome's probability is proportional to its score raised to the power of e, with the assumption that the score for locative -*u* is 0. Thus, since [stod] and [blod] have the same score, they should have the same distribution of locative allomorphs, getting -*u* 88.1% of the time. In the multiple grammar model, the speaker adds an [+lvar] feature to the lexical entry for [stod] and [blod] whose weight is determined by the scores. Since both words have a score of -2, the speaker assigns a [+lvar] weight stochastically using a probability distribution centered around -2—that is, -2 is the most likely weight, and the probability of a weight being selected is proportional to its distance from -2. In this instance, the speaker assigns [stod] a [+lvar] weight of -1 and [blod] a [+lvar] weight of -2.7. These weights are then used to calculate the probability of locative -*ɛ*, using the same principle as in the single grammar model. This time, since the two words have different lexical entries, their distribution of locative allomorphs will differ: [stod] will get -*ɛ* 26.9% of the time, but [blod] will get -*ɛ* only 6.3% of the time.

	single grammar model		*multiple grammar model*	
• Get nonce word	stod	blod	stod	blod
• Calculate score *s* from phonological and morphological factors	-2	-2	-2	-2
• Stochastically assign [+lvar] weight ($b_{+\text{lvar}}$) to lexical entry using probability distribution centered around *s*	⋮	⋮	-1	-2.7
• Calculate probability *p* of locative -*ɛ*…	…based on *s*		…based on $b_{+\text{lvar}}$	
	$\frac{e^{-2}}{e^{-2}+e^{0}} = .119$	$\frac{e^{-2}}{e^{-2}+e^{0}} = .119$	$\frac{e^{-1}}{e^{-1}+e^{0}} = .269$	$\frac{e^{-2.7}}{e^{-2.7}+e^{0}} = .063$
• Stochastically choose locative suffix based on *p*	stodu (88.1%), stoɟɛ (11.9%)	blodu (88.1%), bloɟɛ (11.9%)	stodu (73.1%), stoɟɛ (26.9%)	blodu (93.7%), bloɟɛ (6.3%)

Table 5.23: Comparison of single grammar and multiple grammar models of locative assignment (ignoring syntactic context, which is the same for both models)

This example shows how the predictions between the two models differ: in the single grammar model, a nonce word's proportions of locative -*u* and -*ɛ* depend entirely on its phonological and morphological properties, such that two words with similar phonology (and genitive) should have similar rates of locative -*ɛ* for a given speaker. However, the multiple grammar model predicts some divergence between a given speaker's rates of locative -*ɛ* for words with similar phonology and genitive due to a noise factor in assigning a word's rate (through the feature weight). One subcase of this is that, in both models, some words may always take -*u*, but the single grammar model predicts that only typical -*u* words (phonologically speaking) should exhibit categorical behavior, while the multiple grammar model predicts that a given speaker may produce a larger range of words with categorical locative -*u*.

In order to distinguish between the two models, we must establish a baseline rate of locative -*ɛ* for each speaker and nonce word; this requires many trials for each nonce word. A study testing this hypothesis would either need to be much longer or contain many fewer distinct stimuli. If we were to perform this experiment, we should do a classical wug test, without the genitive factor. In the context of repeated trials, exposure to the genitive could cause dynamic update of a noun's lexical entry, which could be a confound. Instead, we would like a noun's lexical entry in the multiple grammar model to be as static as possible after initial exposure to ensure that we are sampling from a single distribution of locative forms.

Although the two models cannot be distinguished by the sort of wug test I have used, there are, potentially, other ways to compare the models. To explore these, we need to be more explicit about the representation of existing variable nouns in the single grammar model. As mentioned in Section 3.3, the single grammar requires a lexical "fudge factor" for variable real words. Let us look at one example of why a lexical weight is necessary. The word [fxod] 'entrance' should have the same score (-2, in the example in Table 5.23) as the nonce words [stod] and [blod]. That is, if [fxod] were a nonce word with no existing lexical entry, we should expect only 11.9% of locative tokens for this word to be [fxoɟ-ɛ]. However, [fxod] actually appears more frequently with -*ɛ*, about 77% of the time in my corpus. For a speaker to correctly match this input, in the single grammar model, their lexical entry for the word must include a factor to tip the scales of the phonotactic grammar more towards [vxoɟ-ɛ]. In other words, the single grammar model also requires the use of [+lvar], or something very much like it. However, the weight of the lexical [+lvar] feature on individual nouns should derive *entirely* from the input: the feature directly tracks the discrepancy between the predicted and observed locative distribution.

In the multiple grammar model, every new word must pass through the gatekeeper of the sublexical grammar and be assigned a weight for [+lvar] on its new lexical entry before actually being placed into a derivation where it receives a locative suffix. Every noun must have its baseline rate of locative -ɛ encoded in the weight of the [+lvar] feature on its lexical entry. Thus, during the learning process, words are assigned an [+lvar] weight when they are first encountered (or when the speaker first forms their locative), with some amount of randomness as described above. These weights can be adjusted in accordance with the input, so that the impact of the random initialization decreases as the learner gets more confident, especially for very frequent words. However, the randomness in initialization could lead to slight variation in adult grammars that compound the slightly different inputs that individuals have.

Thus, the single grammar model predicts that an individual's locative suffix distribution for individual existing words should be *solely* a product of their input, and to the extent that individuals differ, it is due to their input. The multiple grammar model predicts some *additional variation* due to the randomness in initially assigning [+lvar] weights to nouns, even if this randomness is largely washed out by feedback from the input.

5.4.6.4 How to test the single grammar and multiple grammar models

We can tentatively use this difference in the predicted behavior of established words to test the two models: if we find a systematic difference between the behavior of variable words in the *input* for individual speakers and how these speakers *produce* variable words, this difference must be due to a grammatical bias introduced in going from input to output. Moreover, the grammar must be structured in such a way that enables such biases to enter into learning—that is, the multiple grammar model. I will explain this using a relevant example. The Czech nonce word study finds that nonce words with -*a* in the genitive for a given speaker are more likely to be assigned locative

-ɛ by that speaker—that is, this is a bias in the grammar. I argue that, if we see a similar effect in individuals' baseline rates of *real words*, this would constitute a bias in the output. Of course, this evidence is crucially dependent on an unbiased input, a claim for which I have no evidence. However, if it can be shown at a later time that there is no similar bias in the input, the bias must come from the learning process, which would provide evidence for the multiple grammar model over the single grammar model.

Let us consider two relatively common words: [rɪbɲiːk] 'pond' and [komiːn] 'chimney'. These are one of just a handful of nouns that are *doubly variable* in both the genitive and the locative (see Section 5.5.2). As shown in Table 5.24, the two nouns each occur frequently in both cases, and while they have nearly identical rates of *-u* in the genitive (both are weighted towards *-a*), their distribution in the locative is quite different: [rɪbɲiːk] has a majority of locative tokens with *-u*, while most locative tokens for [komiːn] use *-ɛ*.

		genitive			*locative*		
noun		*-u*	*-a*	*% -u*	*-u*	*-ɛ*	*% -u*
rɪbɲiːk	'pond'	9692	47 156	17.0%	29 837	10 178	74.6%
komiːn	'chimney'	1716	8406	17.0%	1171	7794	13.1%

Table 5.24: The distribution of genitive and locative suffixes of two doubly variable nouns, by token count

This variable distribution is an average across an entire corpus, and does not necessarily reflect the behavior of individual speakers: some presumably always say [komiːna] as the genitive for 'chimney', some say [komiːnu] categorically, and some have variable usage at different rates. However, the input that individual learners receive would, by hypothesis (to be tested at a later date), be somewhat more uniform and closer to the numbers in Table 5.24. This is because each speaker's input is an aggregate of the output of several individuals, who each have their own genitive and locative usage; this should roughly average out to the rates in Table 5.24.

As previously stated, both of these nouns are quite frequent in the locative. Thus, in the single grammar model, speakers can be expected to have a lexical [+lvar] weight reflecting the exact rate of locative -*ɛ* for these words in their input. The rates of genitive -*a* and locative -*ɛ* are not expected to be correlated in the *input* for individual speakers (which includes tokens from a number of different speakers, as argued above). In the single grammar model, then, they are not expected to be correlated in the output.

By contrast, the multiple grammar model predicts a correlation in the output even if there is none in the input. In the nonce word study, I showed that Czech speakers have learned a correlation between genitive -*a* and locative -*ɛ* that influences their assignment of locative to nonce words. In the multiple grammar model, we should be able to see this effect in the way individual speakers produce locatives of *real* words as well. The multiple grammar model requires the assignment of a [+lvar] weight to all words even before there is sufficient data in the input for the weight to reflect the true rate of locative -*ɛ* in the input. Thus, the initial scores must be assigned using other cues, including a bias correlating the genitive and locative, so that learners are expected to assign higher [+lvar] weights to nouns that have appeared more frequently with genitive -*a*.[6] Although these weights are later adjusted in accordance with the input, some residue of the learning bias may remain. So, given two speakers *with the same input*, the multiple grammar model predicts that they should have different rates of genitive and locative -*u* for both [rɪbɲiːk] and [komiːn]. In comparing the genitive and locative forms of speakers with the same input, there should be a bias: a speaker that frequently says [komiːna] in the genitive should also say [komiːɲɛ] frequently in the locative, and vice versa; a speaker that says [rɪbɲiːku] in the genitive should be more likely to say [rɪbɲiːku] in the locative. This would show both that a grammatical correlation exists between genitive and locative exponents (which is the main hypothesis of this experiment), and that this correlation can be detected in both real words and nonce words (which is possible in the multiple grammar model

[6] This explication assumes that speakers see the genitive of a given word before the locative, which is an oversimplification.

but not the single grammar model—if it can be shown by future research that the correlation does not exist for individual learners in the input).

I provide one half of the argument needed to test this prediction by looking at the output of individual authors in a second corpus study, presented in Section 5.6. We do see a correlation between a given speaker's genitive and locative for [rɪbɲiːk] and [komiːn]. (The full corpus study in Section 5.6 confirms that this correlation is statistically robust across a larger range of words.) The data are quite sparse, but if we look at authors who have used the words at least twice in both cases, we can track which suffix they use in the genitive and locative, or whether they have used both. In Table 5.24, we saw that both nouns have fairly high rates of genitive -*a*, and this is reflected in Table 5.25: few authors have genitive -*u* in every token; most use -*a* at least some of the time. The two nouns differ in their locative patterns: [rɪbɲiːk] usually has locative -*u*, and we see that most authors have -*u* in at least some locative tokens. On the one hand, authors that always have -*a* in the genitive are more likely to always have -*ɛ* in the locative (15 out of $64+79+15=158$, or 9.5%) than authors that use both genitives (5 out of $31+67+5=103$, or 4.9%). On the other hand, authors that always have -*a* in the genitive are also more likely to always have -*u* in the locative (64 out of 158, or 40.5%) than authors that use both genitives (31 out of 103, or 30.1%). Under the assumption that this correlation is not in the speakers' input (which, again, is to be studied at a later date), the former is in line with the predictions of the multiple grammar model, while the latter conflicts with it. For [komiːn], the pattern is more clearly in line with the predictions of the multiple grammar model (assuming unbiased input). This word usually takes -*ɛ* in the locative, and all authors used -*ɛ* for this word at least some of the time. However, authors that variably used -*u* in the genitive were also more likely to do so in the locative (13 out of $13+17=30$, or 43.3%) relative to authors that always used -*a* in the genitive, who had variable locatives much less commonly (12 out of $12+43=55$, or 21.8%).

	rɪbɲiːk			komiːn		
	LOC-*u*	LOC-*u*/-*ɛ*	LOC-*ɛ*	LOC-*u*	LOC-*u*/-*ɛ*	LOC-*ɛ*
GEN -*u*	3	4	0	0	1	1
GEN -*u*/-*a*	31	67	5	0	13	17
GEN -*a*	64	79	15	0	12	43

Table 5.25: Authors by the genitive and locative suffixes they use for [rɪbɲiːk] ‘pond’ and [komiːn] ‘chimney’

The results in Table 5.25 show that genitive and locative suffixes for at least some words are correlated for *individual speakers* when both can vary. If the pattern for words like [komiːn] holds across the language and proves statistically robust (and it is shown that this correlation is plausibly absent in the input), it would provide evidence for the multiple grammar model as well as further confirmation of the main result of the wug test, which is that Czech speakers’ grammar contains a correlation between genitive and locative. The full results of the author study, which looks at all doubly variable words in Czech, can be found in Section 5.6, and they do indicate the correlation between genitive -*a* and locative -*ɛ* shown for [komiːn] in Table 5.25.

Before presenting this corpus study, I present the results of a variant of the wug test study whose stimuli are *real* words that are variable in one or both cases, which clarifies the results of the nonce word study.

5.5 Czech variable word study

Under the theory of the nonce word task assumed in this dissertation, speakers use phonological and morphological cues to bias the stochastic assignment of lexical diacritic features to nonce words (see Section 3.1.2, Section 3.2, and Section 5.4.6.3 for discussion). In the nonce word study in Section 5.4, I showed that speakers applied a morphological dependency between the genitive and locative: speakers were more likely to assign locative *-ɛ* to stimuli to which they also assigned genitive *-a*, and they were in turn more likely to assign genitive *-a* to nonce words that were shown with genitive *-a*. Thus, the experimental manipulation (the genitive with which the nonce word was presented in the frame sentence) influenced the choice of locative indirectly, mediated by speakers' own choice of the genitive.

The results of the nonce word study in Section 5.4 are compatible with the hypothesis that Czech speakers apply a correlation between the genitive and the locative. However, it does not specifically locate that application at the point of determining the behavior of nonce words. For example, another possibility is a much more immediate effect: whenever a speaker hears, sees, or uses *-u* in one form (the genitive), she is also more primed to use *-u* in other available forms (e.g. the locative). We would expect to see this sort of identity-based priming effect for both nonce words and real words. In contrast, the morphological dependency bias should only apply to nonce words, not to real words. This is because speakers already have representations for the genitives and locatives of real words, and a single genitive token should not influence the process of selecting a locative token from a distribution based on that locative entry (as described in Section 3.4. Thus, if we do *not* find a correlation between genitive and locative with *real words*, this confirms that the correlation is really a pattern being productively extended to unknown words.

To test the hypothesis, I conduct a study with a similar design to that of the nonce word study, except the stimuli are *real* Czech words that are variable in the genitive and (usually) the locative. While participants do show a correlation between the genitive suffix presented in the frame sentence and the genitive they select (a genitive–genitive priming effect), there is no significant effect on the locative of either presented or selected genitive (no genitive–locative priming effect). This study thus confirms that the effect found in the nonce word study in Section 5.4 is a true learned morphological dependency between the Czech genitive and locative.

5.5.1 Participants

Subjects were recruited through Prolific and had to be Czech nationals born in the Czech Republic and raised as monolingual Czech speakers. In addition, participants who had taken the stimulus testing study for nonce words (Section 5.4) were excluded from this one. I recruited 90 participants and discarded one for listing that they have not spoken Czech since childhood. Thus, the results include 89 participants.

5.5.2 Stimuli

For the stimuli in this study, I used 20 words that are variable in the genitive, most of which are also variable in the locative. Following Bermel and Knittl (2012), I considered a noun variable if it took -*u* between 1% and 99% of the time in my corpus with at least 10 tokens of each noun in each case. This left a total of 18 nouns, of which I used 15, shown in Table 5.26 below. Two of the remaining three, [apriːl] 'April Fool's joke'[7] and [koscitr̝̊as] 'boneshaker',[8] were removed for being very uncommon, while the third, [sokol] 'Sokol movement club', was removed for being homophonous (indeed, named after) with the word for 'falcon', which is animate—this is problematic, because masculine animate nouns in this Czech class always have -*a* in the genitive (see Section 5.1.1). To this I added five nouns that are variable in the genitive but not the locative. Participants should be comfortable seeing these stimuli (along with the others) with either -*u* or -*a* in the genitive, but are predicted to show no variation for these words in the locative, following the distribution of the lexicon.

[7]The surprising specificity of this word is due to the fact that Czech month names, like those of other Slavic languages like Polish and Ukrainian, are derived from native Slavic terms, not borrowed from Latin. The standard word for April is [dubɛn], named for the sprouting of the oak tree ([dub]).

[8]A term for early bicycles whose wooden wheels and heavy iron frames did little to cushion their riders from poorly maintained 19th century roadways; the Czech word is comprised of the same two roots as the English.

	noun	genitive		locative	
		tokens	*% -u*	*tokens*	*% -u*
rɪbɲiːk	'pond'	56 848	17.0%	40 015	74.6%
jazɪk	'tongue, language'	20 722	1.4%	48 808	22.3%
potok	'creek'	25 306	2.2%	18 839	14.2%
mliːn	'mill'	15 770	2.5%	14 852	7.4%
kout	'corner'	13 992	18.2%	14 280	11.1%
kliːn	'wedge, lap'	10 023	2.8%	9407	7.4%
komiːn	'chimney'	10 122	17.0%	8965	13.1%
zaːxod	'toilet'	2597	91.4%	8154	1.3%
bɛtlɛːm	'nativity scene'	4150	51.0%	2783	74.7%
naːrod	'nation'	1900	2.3%	3614	31.7%
tiːl	'back of the head, rear of troops'	2832	31.2%	2407	74.2%
naːɦon	'mill race, motor drive'	2223	96.3%	1409	97.2%
koʒɪx	'fur, fur coat'	914	97.4%	2379	95.2%
otsɛt	'vinegar'	1881	3.1%	1255	42.1%
vɛliːn	'control room'	796	33.5%	887	30.9%
dvuːr	'court'	36 442	1.6%	70 299	0.5%
sɛn	'dream'	8090	70.9%	23 551	100.0%
siːr	'cheese'	2365	21.6%	1027	100.0%
kalɪx	'chalice'	1944	30.8%	1327	100.0%
ɦr̝ɪb	'bolete'	226	85.8%	72	100.0%

Table 5.26: Variable nouns used in the variable word study, grouped by locative variability and ordered by combined frequency

Unlike the nonce stimuli, three of the stimuli used in this study undergo a fully regular stem alternation in suffixed forms: [otsɛt] 'vinegar' and [sɛn] 'dream' show a vowel–zero alternation (e.g. [sn-u]), while [dvuːr] 'court' shows vowel shortening with a change in quality: [dvor-u].

5.5.3 Procedure

The trials in this study are identical to those of the stimulus norming study in Section 5.4.3.2. Each participant saw 40 trials: each of the 20 stimuli was shown twice, once with genitive *-u* and once with genitive *-a*. Each half of the experiment contained the 20 stimuli; stimuli were randomly selected to appear with *-a* in either the first half or the second half. Thus, this experiment had a much higher proportion of stimuli shown with genitive *-a* (20 out of 40 trials) than the nonce word study (12 out of 50 trials).

The stimuli shown in Table 5.26 cover a wide semantic range, so many of the frame and target sentences used in the nonce word study would have produced discordant readings. Of these, I selected 7 frame sentences (presenting bare noun and genitive) and 12 target sentences (with blanks for genitive and locative) that yielded more or less compatible readings with all nonce words. One consequence of the semantic restrictions is that all of the target sentences presented the locative with the non-locational preposition [o] 'about'. Thus, I cannot test for the effect of syntactic context for the locative (which was significant in the nonce word study), though I do for the genitive.

5.5.4 Analysis

In this study, participants' choice of genitive for a word may be fully independent of the genitive with which it is presented, so all trials are kept, whether the genitives matched or not. In total, there were 3,560 trials. I fitted two logistic regressions testing the factors influencing speakers' choice of genitive and locative, respectively. The dependent variable of the first regression is the *genitive* suffix selected by the participant (*-a* or *-u*) with a random intercept for participant and item. I built up the regression one factor at a time using a forward stepwise algorithm using the `buildmer` function in R from the package of the same name (R Core Team, 2022; Voeten, 2022), which adds factors to the model one at a time such that each additional factor improves the model's Akaike Information Criterion (AIC), which measures how well the model fits the data while penalizing model complexity (that is, number of factors). Candidate factors included the experimental factors of genitive suffix presented in the frame, selected locative suffix, genitive preposition in the target sentence, and nuisance variables of trial number and order in which the genitive candidates were listed, and two lexical factors: preposition-adjusted genitive coefficient (a measure of how often a word appears with genitive in the lexicon, see Section 5.3.3.1) and preposition-adjusted locative coefficient (likewise for the locative). Given that genitive and locative are themselves correlated in the lexicon, I residualized locative coefficient on genitive coefficient to isolate the influence of locative coefficient not conflated with genitive coefficient. I also considered by-participant random slopes for each fixed effect added to the model.

The second regression, whose dependent variable is the *locative* suffix selected by the participant (-*ɛ* or -*u*), was built up similarly to the first, with analogous candidate factors: genitive presented in the frame sentence, genitive selected in the target sentence, trial number, order in which the locative candidates were listed, preposition-adjusted locative coefficient, and preposition-adjusted genitive coefficient residualized on the locative coefficient. Since all of the target sentences had the same locative preposition, I did not consider this as a factor.

5.5.5 Results

5.5.5.1 Descriptive summary

In Table 5.17 above, we saw the relationship between genitive and locative in the nonce word study: presented genitive is correlated with selected genitive, but this relation is mediated by selected genitive. In trials where speakers selected the opposite genitive from what was presented, the locative was correlated with selected genitive, but not presented genitive. In this real word study, the patterns are different. There is still a correlation between presented and selected genitive, suggesting that speakers' choice of genitive is influenced by the genitive they are shown (a priming effect of sorts). In addition, there is a correlation between selected genitive and locative. However, there is no correlation between *presented* genitive and locative. This is expected if there is no priming effect between genitive and locative, and the correlation between selected genitive and locative is instead due to the correlation between the two for individual nouns in the lexicon: in general, nouns with higher rates of genitive -*a* also have higher rates of locative -*ɛ*. This interpretation accords with the result of the statistical analysis in Section 5.5.5.3: the selected locative has a statistically significant effect of locative coefficient in the lexicon (which is itself correlated with the genitive coefficient), but not of selected genitive.

Table 5.27 shows how selected genitive varies with presented genitive. Participants assigned -*ɛ* in the locative substantially more often to words that were shown with -*a* in the genitive.

		selected genitive		
		-*u*	-*ɛ*	% -*u*
presented genitive	-*u*	935	845	52.5%
	-*a*	648	1132	36.4%

Table 5.27: Experimental frequency of selected genitive allomorphs -*u* and -*a* for real stimuli, by presented genitive

In Table 5.28, we see the effect of presented and selected genitive on locative. As with the nonce word study, shown in Table 5.17, selected genitive makes a large difference. However, unlike in the nonce word study, this difference does not show up as a substantial difference in the locative based on *presented* genitive: participants selected locative -*u* in 52.6% of trials where they were shown genitive -*u* and 51.8% of trials where they were shown genitive -*a*.

genitive		*locative*		
presented	*selected*	*-u*	*-ɛ*	*% -u*
-u	*-u*	581	354	62.1%
-u	*-a*	356	489	42.1%
-u	total	937	843	52.6%
-a	*-u*	403	245	62.2%
-a	*-a*	519	613	45.8%
-a	total	922	858	51.8%

Table 5.28: Experimental frequency of selected locative allomorphs *-u* and *-ɛ* for real stimuli, by presented and selected genitive

The variable word study also shows a difference in selected genitive according to syntactic context, shown in Table 5.29. Although the majority of trials asked for a genitive with the preposition [z] 'out of', two others were used in one target sentence each: [do] 'into' and [ofilɛdɲɛ] 'regarding'. Both of these prepositions saw higher use of genitive *-u* than [z], especially the latter. Bermel and Knittl (2012) predict higher use of [ɛ] with prepositions of motion like [z] and [do] than with other prepositions like [ofilɛdɲɛ] (see Section 5.1.2 for discussion), and I find this in the lexicon as well, as shown in Section 5.3.3.1. However, I did not find any substantive effect of genitive preposition in the target sentence for nonce words in Section 5.4.5.4.

		selected genitive		
preposition		*-u*	*-a*	*% -u*
z	'out of'	1247	1720	42.0%
do	'into'	153	144	51.5%
ofilɛdɲɛ	'regarding'	183	113	61.8%

Table 5.29: Experimental frequency of selected genitive allomorphs *-u* and *-a* for real stimuli, by preposition in frame sentence

5.5.5.2 Genitive

Table 5.30 shows the effects of the mixed logistic regression predicting selected genitive given random intercepts for participant and stimulus; fixed effects for presented genitive, preposition, preposition-adjusted genitive coefficient, and residualized preposition-adjusted locative coefficient (the latter two measures of rate of genitive and locative in the lexicon); and by-participant random slopes for all but locative coefficient. The strongest effect is that of presented genitive: speakers assigned *-a* to words substantially more often when they were also presented with *-a* in the genitive. The moderately high variance of the random slope for this effect indicates that speakers varied in their susceptibility to the frame sentence. This seems to be the case: while most participants chose a matching genitive about half the time (none had fewer than 15 matching trials out of 40, and 71 of 89 speakers matched between 15 and 25 trials), some matched substantially more often: two participants chose a matching genitive on *every* trial, and one more matched the genitive on all but one trial. This likely reflects a different approach to the experimental task: some speakers ignored the presented genitive and relied on their underlying lexicon, while others tended to copy the form shown. The similar variance of the random intercept for participant likewise shows that different people have different baseline rates of genitive *-a*, at least in this task. Participants also showed sensitivity to preposition: they assigned genitive *-u* significantly more often to nouns following the prepositions [do] 'into' and especially [ofilɛdɲɛ] 'regarding' than to those following [z] 'out of' (which was by far the most common). The stronger effect of [ofilɛdɲɛ] matches the lexicon and the prediction from Bermel and Knittl (2012) (see Section 5.1.2), while the weaker effect of [do] does not: Bermel and Knittl (2012) suggest that [z] and [do] should show similar rates of genitive *-a*, while Table 5.8 shows that [do] appears with genitive *-a* significanly *more* than [z]. This factor should be taken with a grain of salt: as described in Section 5.5.5.1, most of the target sentences have [z], and only one each have the other two prepositions. This factor could thus be conflating other effects of particular sentences used in the study. Adding a random slope for preposition did not significantly improve the model ($\chi^2 = 15.02$, $p = .090$) and the improvement did not make up for the additional model complexity according to the Akaike Information Criterion (AIC), so I omit it.

Random effects	*variance*	*SD*		
Participant				
Intercept	1.17	1.08		
Presented genitive (default: -u)				
-a	1.15	1.07		
Genitive coefficient	0.01	.11		
Item	0.21	.46		
Fixed effects	*β coef*	*SE*	*Wald z*	*p*
Intercept	–0.11	.17	–0.64	.5171
Presented genitive (default: -u)				
-a	**0.94**	**.14**	**6.56**	**<.0001**
Preposition (default: z)				
do	**–0.45**	**.15**	**–2.93**	**.0033**
ofilɛdɲɛ	**–1.09**	**.15**	**–7.13**	**<.0001**
Genitive coefficient	**0.32**	**.04**	**7.84**	**<.0001**
Residualized locative coefficient	**0.07**	**.02**	**4.52**	**<.0001**

Table 5.30: Effects of mixed logistic model predicting experimental use of genitive -*a*, with significant effects bolded

Both of the tested lexical effects made it into the model. As expected, preposition-adjusted genitive coefficient has a significant positive effect, meaning that words that take genitive -*a* in the lexicon were assigned genitive -*a* more often experimentally. Genitive coefficient is on the same scale as the model effect size, so the fact that its effect size is much less than 1 indicates that there is less differentiation among the experimental results than in the lexicon—as is usual in the studies in this dissertation, the experimental distribution of stimuli is much less extreme than the distribution of words in the lexicon. We can see this example at the extremes: [koʒɪx] 'fur, fur coat' has the lowest genitive coefficient of all the words used in this study at -5.90, corresponding to an expected rate of -*a* of $\frac{e^{-5.90}}{1+e^{-5.90}} = 0.3\%$ (see Section 4.3.3.1 for an explanation of the math); however, in the experiment, participants selected the genitive [koʒɪxa] in 30 of 178 trials, or 16.9%. The noun with the highest genitive coefficient is [jazɪk] ‘tongue, language’ at 4.03, corresponding to an expected rate of -*a* of $\frac{e^{4.03}}{1+e^{4.03}} = 98.2\%$. In the experiment, participants selected the genitive with -*a* in 143 of 178 trials, or 80.3%. The by-participant random slope for this model has a very low variance, indicating that speakers apply the distribution of genitive allomorphs from the lexicon to roughly the same degree. More surprisingly, the preposition-adjusted *locative* coefficient was also added to the model with a small but significant effect, even though this factor was residualized on the genitive coefficient in order to account for the fact that rates of genitive and locative -*u* for a given noun are correlated in the lexicon. I am not sure why this is the case, though I discuss possibilities in Section 5.5.6.

The effect of genitive coefficient and presented genitive can be seen in Figure 5.11 and Figure 5.12. These graphs plot the predicted likelihood of each stimulus taking genitive -*a*, calculated primarily from the genitive coefficient, on the x-axis and the experimental rate of selected genitive -*a* on the y-axis. The two graphs show the same data plotted on different scales: Figure 5.11 shows things in terms of log odds, which makes the relationship between genitive coefficient and predicted rate linear; Figure 5.12 shows the actual likelihood with untransformed scales. Each stimulus appears twice, split according to whether the word was presented with genitive -*u* (black) or -*a* (gray). Since the effect of genitive -*a* is positive in Table 5.30, the gray words are always to the right of the black words. The two tokens of each word are connected by a line to show the comparison; in every case, the line slopes upward to the right, indicating that every word was assigned -*a* more often when also presented with -*a*. Figure 5.11 shows that the rate of genitive -*a* (that is, the adjusted rate represented by the genitive coefficient) is nonetheless a very good predictor of experimental usage of -*a*: the words lie more or less on a line. Moreover, the effect of genitive coefficient is much stronger than that of presented genitive: the black and gray nouns are interspersed with one another, showing that a presented -*a* cannot make up for a low lexical rate of -*a*. Figure 5.12, on the other hand, shows that the predicted rates are more extreme than the actual rates: nouns that are predicted to take -*u* almost all the time ([koʒɪx], [naːɦon], and [zaːxod]) were assigned -*a* roughly

20–50% of the time, and nouns predicted to almost always take -*a*, like [potok] and [naːrod], often had a good number of responses with -*u* as well. These graphs also show lines corresponding to the fit of the model in Table 5.19 for nonce words with selected genitive -*u* (black) and -*a* (gray).

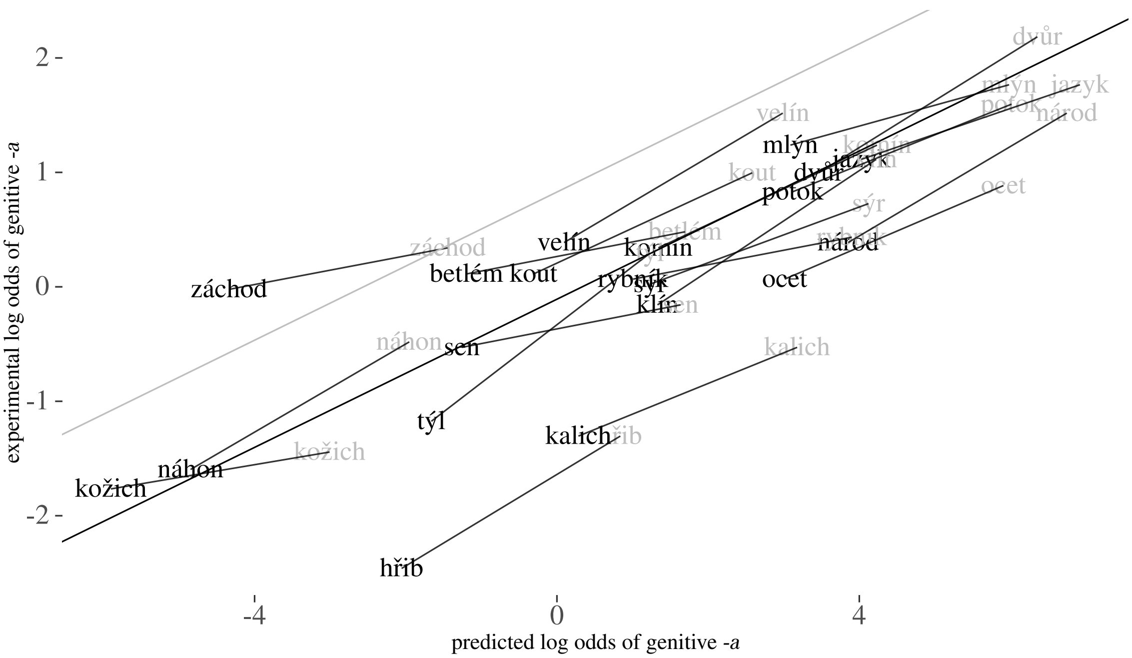

Figure 5.11: The relationship between predicted and experimental log odds of selected genitive -*a* for individual words with presented -*u* (black) and -*a* (gray) genitives, with a line showing the fit of the experimental model in Table 5.30

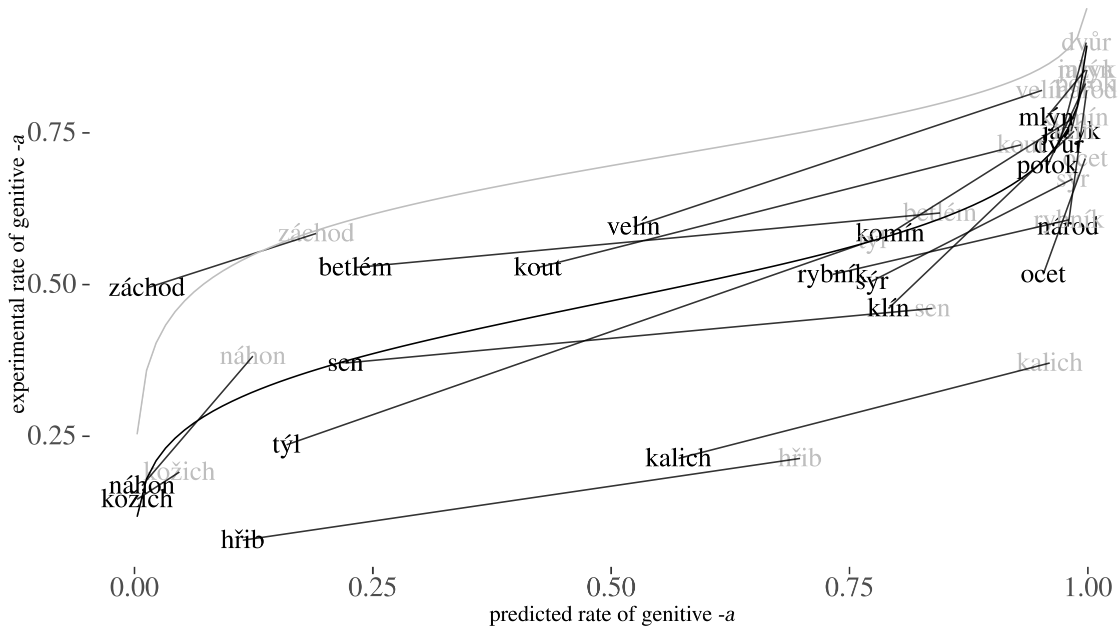

Figure 5.12: The relationship between predicted likelihood and experimental rate of selected genitive *-a* for individual words with presented *-u* (black) and *-a* (gray) genitives, with a line showing the fit of the experimental model in Table 5.30

5.5.5.3 Locative

The effects of the mixed logistic regression predicting selected locative are shown in Table 5.31. The factors added to this model, alongside random intercepts for participant and stimulus, include trial number, preposition-adjusted locative coefficient, and selected genitive, with both fixed effects and by-participant random slopes for all three factors. The fixed effects of this model are not very strong, but the strongest was locative coefficient. As was the case for selected genitive, a word's experimental rate of locative *-ɛ* is correlated with its rate in the lexicon, though here as well the distribution is less extreme than in the lexicon (the effect size is .23, which is much less than 1). Four of the stimuli ([sɛn] 'dream', [kalɪx] 'chalice', [siːr] 'cheese', and [ɦr̝ɪb] 'bolete') are attested exclusively with *-u* in the locative, giving them locative coefficients between −16.40 and −16.03 (predicting one *-a* every ten million tokens). In the experiment, however, these words received locative -*ɛ* 10.1–13.5% of the time—a relatively small percentage, but certainly not categorical. The random slope for locative coefficient has very low variance, indicating that participants applied patterns from the lexicon to roughly the same degree. The factor of selected genitive is added to the model, but as expected, the effect is not significant: although selected genitive and locative were correlated in Table 5.28, this is not significant when locative coefficient is taken into account. This is presumably because the correlation is driven by the fact that genitive and locative realization are correlated in the lexicon, so nouns with high experimental rates of selected genitive -*a* should also be more likely to have high rates of locative -*ɛ*. The random slope for selected genitive is quite high in variance—although the fixed effect was not significant, this suggests that some speakers did apply a correlation between genitive and locative more actively than others. Here too, though, the effects were relatively modest: participants matched genitive and locative (-*u* with -*u* or -*a* with -*ɛ*, respectively) in as many as 33 of 40 trials. By comparison, two participants matched genitives on *every* trial, as described above. Finally, the first effect added to the model (trial) shows a task effect: speakers assigned -*ɛ* in the locative more often in later trials than in earlier ones. This may represent some sort of familiarization effect: locative -*ɛ* is relatively uncommon in Czech, especially by type frequency (see Table 5.5), and the task actively suggested -*ɛ* as a possible locative suffix, so perhaps participants were more comfortable using it as the trials went on. This effect is fairly sizable (raising the baseline likelihood of -*ɛ* from 38.1% in the first trial to 65.0% by the last trial), but I have no other meaningful explanation for it, so I do not discuss it further. The variance of the random slope for trial number is very small, suggesting that most participants had a similar task effect.

Random effects	*variance*	*SD*		
Participant				
Intercept	0.82	.91		
Trial	0.00	.02		
Locative coefficient	0.01	.09		
Selected genitive (default: -u)				
-a	0.70	.09		
Item	0.25	.50		
Fixed effects	*β coef*	*SE*	*Wald z*	*p*
Intercept	**–0.51**	**.19**	**–2.73**	**.0064**
Trial	**0.03**	**.00**	**6.77**	**<.0001**
Locative coefficient	**0.23**	**.02**	**9.51**	**<.0001**
Selected genitive (default: -u)				
-a	0.15	.14	1.11	.2693

Table 5.31: Effects of mixed logistic model predicting experimental use of locative -*ɛ*, with significant effects bolded

Figure 5.13 and Figure 5.14 show the effect of preposition-adjusted locative coefficient and selected genitive from the model in Table 5.31. As with Figure 5.11 and Figure 5.12, these graphs plot the observed vs. predicted rates of locative -*ɛ* on the y- and x-axes, respectively; Figure 5.14 shows raw rates, while in Figure 5.13, the likelihoods are shown in terms of log odds. The predicted rates are derived primarily from the locative coefficients, which are a measure of the rate of locative -*ɛ* for a word in the lexicon adjusted for syntactic context (see Section 5.3.3.1). Figure 5.13 shows that these coefficients do a good job predicting the experimental results, except at the bottom end: words that are predicted to take -*u* (essentially) categorically, like [fiɽ̞ɪb], were actually assigned -*ɛ* occasionally (e.g. [fiɽ̞ɪbjɛ], with the added glide described in Table 5.3). Figure 5.14, in turn, shows that the experimental distribution is less extreme than that of the lexicon: no word received -*ɛ* in less than about 10% or more than about 80% of trials, although these words run the full range of locative allomorph distributions in the lexicon. To this we add the effect of selected genitive. Each word appears twice, once representing trials where speakers selected genitive -*u* (in black) and once with trials where speakers selected genitive -*a* (in gray). The size of each word corresponds to the number of trials—for example, the genitive selected for [jazɪk] was usually [jazɪka], so for this word, the gray is much larger than the black. The model in Table 5.31 has a positive coefficient for selected genitive, so the gray word is always to the right of the black word. The experimental comparison is much less clear than that between presented and selected genitive: for the majority of nouns, the gray word is above the black and the line connecting them slopes upward to the right, indicating that participants assigned this word locative -*ɛ* more often when they also assigned it genitive -*a*. However, the difference is generally quite small, and in some words, it is even reversed: for example, participants formed the locative of [otsɛt] as [otstu] *more* often when they selected the genitive form [otsta].

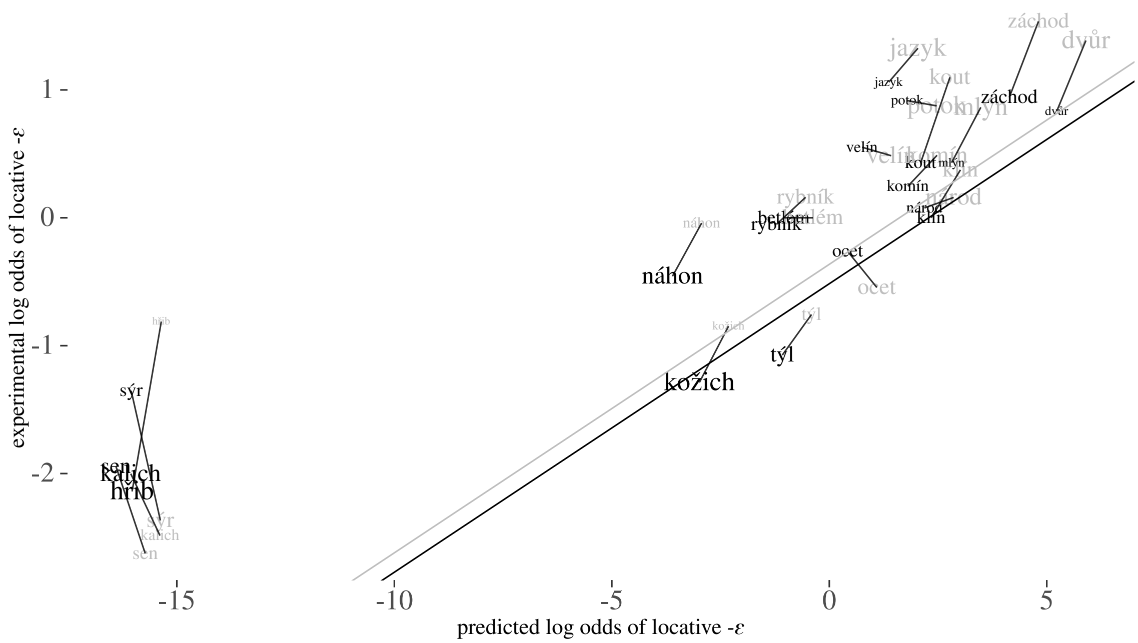

Figure 5.13: The relationship between predicted and experimental log odds of locative -*ɛ* for individual words with selected -*u* (black) and -*a* (gray) genitives, with a line showing the fit of the experimental model in Table 5.31

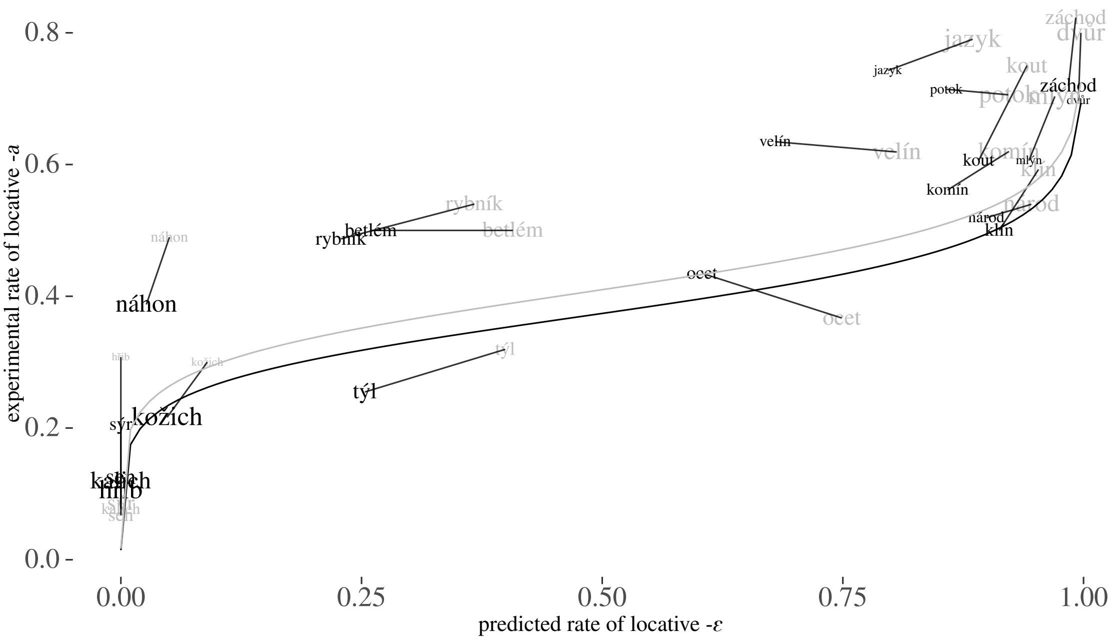

Figure 5.14: The relationship between predicted likelihood and experimental rate of locative -*ɛ* for individual words with selected -*u* (black) and -*a* (gray) genitives, with a line showing the fit of the experimental model in Table 5.31

5.5.6 Discussion

This study was intended as a companion to the nonce word study in Section 5.4 by distinguish between two possible interpretations of its results. In Section 5.4, I found that participants applied a morphological dependency between presented genitive and selected locative, mediated by selected genitive: they assigned locative -*ɛ* more often to words to which they also assigned genitive -*a*. In the intended interpretation, this dependency comes from a generalization over the lexicon that speakers are productively extending to new words: when a word's locative is unknown, its genitive can be used to make a better guess. However, another possibility is that the correlation represents a sort of priming effect: speakers observe the identity of -*u* at all times, not just in determining the behavior of nonce words. If the correlation between selected genitive and locative is just a priming effect, we would expect it to see it when participants are tested on *real* words as well; if it is a productive extension of a lexical pattern, we would only expect to see it in nonce words, not real words.

This experiment, which used the same trial design as the nonce word study in Section 5.4 but with real (variable) words as stimuli, produced a null result: the model in Table 5.31 shows no significant effect of selected genitive on selected locative. This result requires some interpretation. Participants did assign locative *-ɛ* more often to stimuli to which they assigned genitive *-a*, but this effect seems to be mostly attributable to the fact that words that have higher rates of genitive *-a* in the lexicon also tend to have higher rates of locative *-ɛ*, so some correlation would be expected if participants are simply applying the rates from the lexicon. Indeed, the lexical factor of preposition-adjusted locative correlation has a significant effect on locative realization, meaning that participants are selecting locative suffixes for words in accordance with their distribution in the lexicon—this seems to account for most of the correlation between genitive and locative in terms of raw counts. There does still seem to be some residual correlation: the effect of genitive on locative, while not significant, is slightly positive, and the random slope for selected genitive suggests that some participants are showing an active correlation between the two. However, even if this is the case, it is much less substantial than the same effect applied to nonce words. That is, the main result of this study is that there is *no significant evidence* for a priming effect between genitive and locative for real words. This contrasts with the very strong effect in the nonce word study. Taken together, these two results strongly indicate that speakers are productively extending a correlation between genitive *-a* and locative *-ɛ* to determine the behavior of unfamiliar words—in other words, Czech speakers have learned a morphological dependency between the genitive and locative.

We can contrast this null result in the locative with the effects of the model predicting selected genitive, shown in Table 5.30. As in the nonce word study, participants' choice of genitive suffix for a stimulus was significantly dependent on the genitive suffix presented for that stimulus (suggesting a sort of priming). As expected, a word's preposition-adjusted genitive coefficient was predictive of its experimental rate of genitive *-a* as well: the more often a word takes *-a* in the lexicon, the more often participants assigned it *-a* in the study. Participants seem to balance these factors in slightly different ways: a few always (or nearly always) copied the presented genitive and ignored the usual genitive suffix for the noun in question, while others were less swayed by the genitive suffix in the frame sentence. This suggests that speakers' choice of genitive for variable words can be subject to multiple factors, including a tendency to match previous genitive tokens of the same words in the discourse, and that different speakers approach this matching differently (in an experimental task, at least). Another contextual factor influencing speakers' choice of genitive is its syntactic context: in particular, speakers assigned genitive *-u* more often in the frame sentences where the target word was paired with prepositions [do] 'into' and [ofilɛdɲɛ] 'regarding' than when it appeared with the preposition [z] 'out of', which was used in the majority of the target sentences. In the lexicon, *-u* is more common with objects of [ofilɛdɲɛ] than with objects of [z]: for this preposition, speakers matched the lexical tendencies, as well as the description of Bermel and Knittl (2012) (see Section 5.1.2). In the nonce word study, I found an effect of syntactic context (preposition) for the locative but not the genitive; I do not have an explanation for why preposition was a significant factor in this study but not the nonce word study, but it is generally a welcome result that affrms the general influence of syntactic context on genitive and locative realization for variable Czech words.

Finally, I close this section by discussing one puzzling result: nouns with higher preposition-adjusted *locative* coefficients—that is, higher rates of locative *-ɛ* in the lexicon—were assigned genitive *-a* significantly more often. This is, in theory, expected given that rates of genitive and locative are positively correlated in the lexicon, so nouns with higher locative coefficients also have higher genitive coefficients; however, the effect held even when accounting for this by residualizing locative coefficient on genitive coefficient. I do not have a very satisfying explanation of this effect. One possibility is that the residualization did not account for all of the correlation between genitive and locative coefficient—that is, the locative coefficient effect is still just really a secondary expression of the much stronger correlation between genitive coefficient and selected genitive described above. However, this oddity is fairly minor and does not detract from the main results of this study: we see evidence of priming within a case (presented genitive predicts selected genitive), but not across cases (selected genitive does not significantly predict selected locative). This, in turn, strengthens the interpretation of the nonce word study that speakers are applying a morphological dependency between genitive -*a* and locative -*ɛ* to novel words.

5.6 Czech author corpus study

In Section 5.4.6.4, I argued that the model of learning variation I presented in Section 3.4 predicts that the correlation between Czech genitive and locative found in the nonce word study in Section 5.4 should also be detectable in the output of individual speakers. This argument went as follows: the rates of genitive and locative -*u* in the corpus are aggregates of different speakers who do different things with the word. For example, [rɪbɲiːk] 'pond' is variable in both genitive and locative, but not all speakers are variable for a given case: some always use -*u* for both cases, some never do, and some vacillate. However, each speaker's *input* may be more uniform, since a learner's input is an aggregation of the output of several different speakers. If a bias exists in the output of speakers that is plausibly not present in the input (an assumption which is to be tested at a later date), then that bias has to come from the influence of the grammar on the learning process itself. Moreover, the process of learning individual lexical items must allow for such bias to leave its mark on the adult forms of the items. By contrast, the simpler single grammar model claims that lexical entries directly track a speaker's input, so it predicts that the input should match the output, without any additional biases introduced.

I argued that one such learning bias is the correlation between genitive -*a* and locative -*ɛ*, and my multiple grammar model of learning variable lexical items predicts that a trace of this bias should be visible in the behavior of adult speakers, because it predicts noisy assignment of a baseline rate of locative -*ɛ* to new lexical items that is influenced by those words' genitive forms. Although speakers will adjust these rates as they hear a word more often and get a better sense of its "true" locative distribution, some trace of the noise should remain, and in this, we should see a correlation between genitive and locative endings *for a given speaker*.

In this study, I provide the basis for testnig this hypothesis using words that are variable in both the genitive and locative like [rɪbɲiːk] 'pond'. If a given speaker has -*a* as the genitive for this word, she should be more likely to also have -*ɛ* as the locative, and conversely: if she has -*u* in the genitive, she should be more likely to have -*u* in the locative as well. If it can then be shown that this bias does *not* appear in a speaker's input (which will require future work), the correlation must come from the productive application of existing lexical patterns linking the two cases in the grammar.

One result from this study is that variable words with multiple meanings often have very different rates of locative -*ɛ* for different meanings. As I discuss in Section 5.6.4.3, this sort of semantically conditioned variation is problematic for highly modular theories like Distributed Morphology.

5.6.1 Data

For this study, I used the same 15 doubly variable words (shown in Table 5.26) as in the variable word study in Section 5.5, excluding the five words used in that study that were variable in the genitive but not the locative ([dvuːr] 'court', [sɛn] 'dream', [siːr] 'cheese', [kalɪx] 'chalice', and [ɦr̝ɪb] 'bolete'). I looked at these doubly variable nouns in version 11 of the SYN corpus (Křen et al., 2022), a corpus of written texts provided by the Czech National Corpus. I looked at cases where the words are preceded by one of the "canonical" prepositions for each case: [do] 'into', [od] 'from', and [z] 'out of' for the genitive, and [na] 'on' and [v] 'in' for the locative. These prepositions should have the lowest rate of -*u*; this was done to keep the context as consistent as possible. For each lexical item, I included authors who used the word at least twice in both the genitive and locative. I excluded texts attributed to multiple authors or an editor, as well as texts attributed to a set of initials in parentheses (usually a byline for newspaper articles). This left a total of 6,626 locative tokens from 830 authors.

Some of the words chosen have multiple related meanings. This raises the possibility that polysemous variable words assign certain variants to particular meanings, as reported by Kiefer (1985) for Hungarian (see Section 3.4). To control for this, I coded individual tokens of the polysemous words for their meaning. The difference was not always clear-cut. For example, [jazık] means both 'tongue' and 'language', but the 'tongue' meaning often bleeds metaphorically into the 'language' meaning. In the case of [kliːn], it was sometimes difficult (especially in poetry) to distinguish between a literal use of 'lap' and a more metaphorical use of 'heart' (of a mountain range, forest, etc.). I defaulted to the 'lap' meaning, which was much more frequent, unless I had overt evidence to the contrary. The distirbution of affixes by meaning for polysemous words is found in Table 5.32.

word		genitive		locative	
number	*meaning*	*-a*	*-u*	*-ɛ*	*-u*
	jazɯk	**383**	**5**	**412**	**255**
1	language	321	1	407	59
2	tongue	62	4	5	196
	kliːn	**692**	**8**	**834**	**12**
1	lap/crotch	687	4	831	2
2	inner part	4	2	3	9
3	wedge	1	2	0	1
	koʒɪx	**1**	**7**	**2**	**9**
1	fur coat	0	4	0	8
2	fur	1	3	2	1
	tiːl	**89**	**71**	**64**	**126**
1	back of the head	80	15	57	12
2	rear of troops	9	52	7	103
3	back part of object	0	4	0	11

Table 5.32: Genitive and locative suffixes for polysemous doubly variable words

As can be seen in Table 5.32, words behave quite differently according to their meaning—for example, [tiːl] rarely takes *-u* in either case with the meaning of 'back of the head', but strongly prefers *-u* in both cases otherwise. However, in most cases there was still some variability within each meaning.

To avoid conflating this semantically conditioned variation with individual differences, I treated each meaning as a separate lexical item: [jazɯk_1] 'language', [jazɯk_2] 'tongue', etc. I removed [koʒɪx_1] 'fur coat' and [tiːl_3] 'back part of object' because these showed no variability among my tokens. I then recalculated my data set with the split lexical items. In the final data set, no tokens remain of [kliːn$_2$], [kliːn$_3$], or [koʒɪx$_2$]. This is because no author used these words twice in each case with that particular meaning. The updated data set includes 6,369 locative tokens from 811 authors.

5.6.2 Methods and analysis

For each author and lexical item, I classified its genitive suffix: *-u*, *-a*, and variable. Any word that appeared with both forms in a given case for a given author was classified as variable. This is a rough approximation, because most authors use each word in the genitive only a small number of times, and a larger corpus might reveal greater variability.

I fitted a logistic regression whose dependent variable was locative suffix. Its predictors included the noun (using the most frequent noun, [rɪbɲiːk] 'pond', as the baseline), the preposition, and the interaction between them, as well as the genitive suffix used by a given author for the noun. Although the interaction term made the model more complex, it did also improve it, and yielded interpretable results; thus, I include it. I did not include any additional terms for author: in theory, each author should have a random slope by noun, reflecting the hypothesized underlying grammar, where each speaker has an underlying locative distribution for each noun. However, the data are too sparse for such random effects to be calculated, so the combination of author and noun is represented by the genitive suffix predictor.

My hypothesis is that genitive suffix will have a significant effect in the second regression: authors that have genitive *-a* for a given noun are more likely to also have *-ɛ* for that same noun. If this is the case, it would suggest that one of the factors influencing the shape of individual speakers' grammars is a productive generalization between the choice of genitive and locative suffix. As a check, preposition should be significant as well, since [v] 'in' takes *-u* less than [na] 'on'.

5.6.3 Results

The results of the regression are shown below.

	β coef	*SE*	*Wald z*	*p*
Intercept	**–1.41**	**.37**	**–3.83**	**.0001**
Noun (default: rɪbɲiːk)				
bɛtlɛːm	0.06	.25	0.26	.7926
jazɪk$_1$	**2.75**	**.19**	**14.48**	**<.0001**
jazɪk$_2$	–17.35	751.71	–0.02	.9816
kliːn$_1$	**5.68**	**.71**	**7.96**	**<.0001**
komiːn	**3.08**	**.26**	**11.96**	**<.0001**
kout	**4.97**	**.28**	**17.82**	**<.0001**
mliːn	**5.01**	**.51**	**9.84**	**<.0001**
naːɦon	–17.16	1255.28	–0.01	.9891
naːrod	**2.53**	**.62**	**4.08**	**<.0001**
otsɛt	**1.72**	**.34**	**5.06**	**<.0001**
potok	**2.45**	**.16**	**15.77**	**<.0001**
tiːl$_1$	**2.23**	**.33**	**6.79**	**<.0001**
tiːl$_2$	**–1.54**	**.56**	**–2.76**	**.0057**
vɛliːn	1.44	.92	1.57	.1161
zaːxod	19.86	2656.71	0.01	.9940
Preposition (default: v)				
na	**–0.69**	**.12**	**–5.65**	**<.0001**
Genitive (default: -*u*)				
variable	0.37	.37	1.01	.3130
-a	**0.92**	**.37**	**2.50**	**.0125**
Noun * preposition (default: rɪbɲiːk * v)				
bɛtlɛːm * na	–1.22	1.07	–1.14	.2558
jazɪk$_1$ * na	**–2.27**	**.51**	**–4.46**	**<.0001**
jazɪk$_2$ * na	—	—	—	—
kliːn$_1$ * na	14.10	299.69	0.05	.9625
komiːn * na	0.25	.36	0.69	.4903
kout * na	–22.36	6522.64	0.00	.9973
mliːn * na	**–1.55**	**.60**	**–2.58**	**.0098**
naːɦon * na	0.69	4779.97	0.00	.9999
naːrod * na	–1.35	1.18	–1.15	.2524
otsɛt * na	—	—	—	—
potok * na	0.09	.26	0.33	.7382
tiːl$_1$ * na	0.02	1.28	0.01	.9891
tiːl$_2$ * na	—	—	—	—
vɛliːn * na	—	—	—	—
zaːxod * na	0.63	2680.14	0.00	.9998

Table 5.33: Regression model predicting locative suffix of the author study (with -*u* as default), with significant effects bolded

Most of the nouns have a baseline rate different from that of [rɪbniːk] ‘pond’, and the prefix [na] ‘on’ prefers *-u* relative to [v] ‘in’—the opposite result from my corpus and nonce word studies. We also see some significant interaction effects: for both [jazɪk$_1$] ‘language’ and [mliːn] ‘mill’, the preference of [na] to take *-u* is greater than expected.

We can see that the data in general are quite sparse: the standard error is very large for some terms, and blank interaction terms indicates that some nouns only appear with one preposition. For example, [vɛliːn] ‘control room’ only appears with one author, who uses the word five times, always with the preposition [v].

Despite this shortcoming, we see a significant effect of the morphological dependency: if a given author uses *-a* for a given word in the genitive, they are also more likely to use *-ɛ* for that word in the locative. Furthermore, adding this factor to the regression significantly improves its fit. Thus, the regression supports the hypothesis of a correlation between the two suffixes.

5.6.4 Discussion

5.6.4.1 Main results

The goal of this study was to show whether individual speakers show a correlation in the production of variable words in the genitive and locative. This serves as one piece of a study testing a bias in the mental grammar of Czech speakers, without directly testing them behaviorally on nonce words. Different speakers should have slightly different inputs, leading to different outputs for each word. However, I argued above that this difference *may not* lead to a correlation between a speaker’s genitive and locative for a *given* word (though showing this would require a later study on child-directed speech). Instead, if we do find that the correlation found in this study is not present in speakers’ input, my model of lexical productivity predicts that it should come from another source of individual difference: the grammar of lexical pattern learning, which includes a correlation between genitive and locative, influencing the initial assignment of baseline locative -ɛ rate to individual nouns under the multiple grammar hypothesis as described in Section 5.4.6.3. This study does indeed find the bias: although both genitive and locative can vary from speaker to speaker, there is a correlation between the two cases for a given speaker. This matches the finding of the nonce word study in Section 5.4 that speakers have a productive correlation between the two in their grammar, and—if it turns out that speakers are receiving unbiased input—lends novel support to my multiple grammar model of sublexicon learning in Section 3.4.5, in which such biases can influence the adult grammar.

5.6.4.2 Preposition sensitivity

The regression also included some unexpected effects. First of all, Table 5.33 shows that locative -ɛ is *more common* with the preposition [v] 'in' than with [na] 'on', which is the opposite effect from what I found in my study of the lexicon as a whole (Section 5.3) and for speaker behavior in the nonce word study (Section 5.4). I have no explanation for this discrepancy, which may be due to the small and specialized nature of this data set focused on particular words. However, it accords with the findings in the corpus study of Guzmán Naranjo and Bonami (2021).

The other unexpected effects are the interactions between certain nouns and their choice of preposition. In the model of variation presented in this work, the choice of locative suffix depends on the lexical item to which it attaches and its syntactic context (here operationalized as the preposition governing the noun). However, the two should not interact: the effect of syntactic context should apply uniformly across all lexical items. If our significant interaction terms are not explainable by quirks of the data, this would suggest the need for a more sophisticated model allowing syntactic context to affect *each* lexical item's choice of suffix at *its own rate*.

There are two significant interaction effects: both [jazɪk$_1$] 'language' and [mliːn] 'mill' take *-u* more often than expected in the context of the preposition [na] 'on'. In the case of 'language', this difference is interpretable: for semantic reasons, this word can only combine in a literal sense with [v] (that is, one says ‘in language’, not ‘on language’). Thus, instances of [na] with ‘language’ in the corpus are really verb–preposition collocations like [zaːlɛʒɛt na] ‘depend on’, and in general, locative prepositions selected by particular verbs take *-u* more than those in standard locational prepositional phrases (Bermel & Knittl, 2012). The inability of the corpus to distinguish between adverbial PPs and PPs selected by verbs is exacerbated in the case of ‘language’ due to its inability to combine with [na] for semantic reasons. (This may also contribute to the unexpected main effect of preposition discussed at the beginning of this section, in which [v] ‘in’ was more likely to appear with locative *-ɛ* than [na] ‘on’, contrary to the findings of the other Czech studies in this chapter.)

In the case of [mliːn] ‘mill’, the interaction also seems to be a factor of the data, though in a somewhat different way. The locative *-u* is quite uncommon for this word in general, appearing 16 times compared to *-ɛ*. Twelve of the 16 instances of locative [mliːn-u] appear in the writing of a single author, Petr Škotko, a local journalist who has written regularly about mills, usually in the context of former mills turned into agricultural museums. He also uses the preposition [na] (best translated in this case as ‘at the mill’) 28 times compared to only 4 for [v] (‘in the mill’). His distribution of prepositions is also quite different from the rest of that of the other authors in this study: altogether, [mliːn] follows [v] 362 times and [na] only 112 times. Thus, the significant interaction term is driven by the morphological and syntactic (and perhaps semantic) preferences of a single author. It is not indicative of a need for a more complicated model of syntactic and lexical variation in general.

5.6.4.3 Polysemy and semantically conditioned variation: a problem for Distributed Morphology

In Section 5.6.1, I showed that words with several closely related meanings often distribute genitive and locative variants unevenly among their meanings. For example, the locative suffix for [jazɪk] is usually *-ɛ* with the meaning ‘language’ and almost always *-u* with the meaning ‘tongue’. At first glance, this semantically conditioned variation resembles the contextual effect modelled in Section 3.4: the choice of locative suffix depends in part on the preposition assigning locative case. I accounted for this in the syntax: there are two locative cases that differ in some syntactic feature, and different prepositions assign the two locative cases to their objects at different rates. However, this approach will not work for variability conditioned on the semantics of the root, which is problematic for modern versions of Distributed Morphology, as I now show.

The core derivational mechanisms of Distributed Morphology, described by Harley and Noyer (1999), are as follows. First, morphosyntactic features are assembled into a syntactic structure that undergoes some syntactic computations. The output of the syntax is then sent off to two different modules. The first branch of the derivation undergoes post-syntactic morphological operations, after which the structure is converted to phonological material (PF, or Phonological Form). In the second branch of the derivation, the syntactic structure is interpreted semantically (LF, or Logical Form).

In the early version of Distributed Morphology described by Marantz, 1996 and Harley and Noyer (1999), there are three sites of language-specific knowledge. The Lexicon includes information that is relevant for the computation of syntax, so lexical items[9] comprise bundles of syntactically relevant features (category, person, number, gender, etc.). The Vocabulary comprises connections between lexical items and phonological realizations—so, for example, a vocabulary item for English would spell out plural number as *-z*. Finally, the Encyclopedia contains connections between syntactic nodes and semantic information that is not relevant for the syntax, like the difference in meaning between *cat* and *lizard*.

At some point in a derivation, the speaker has to decide whether she is talking about a tongue ([jazık]) or a back of a head ([tiːl]). However, the difference between these two is not relevant for the syntax. Accordingly, these two are not distinguished in the Lexicon, but rather in the

[9]Here, following the nomenclature of Distributed Morphology, "lexical item" specifically refers to a *syntactic* unit. Elsewhere I discuss the lexical entries of roots that are, strictly speaking, phonological "vocabulary items".

Vocabulary and the Encyclopedia. That is, there is a single, undifferentiated root element (or, at least, noun root) in the Lexicon (which I label $\sqrt{}$) that can be spelled out through different vocabulary items that compete with one another:

(44) *Vocabulary items with different phonology (early Distributed Morphology)*

a. $\sqrt{} \leftrightarrow$ jazık

b. $\sqrt{} \leftrightarrow$ tiːl

In this version of the theory, the speaker chooses what she is talking about by freely choosing one of the vocabulary items in (44). That is, the choice is made *in the PF branch*, after the syntactic computation has already been completed and shipped out. Accordingly, the difference in *meaning* between the two is contained in the Encyclopedia and cannot be read off of the syntactic structure alone. It must also depend on the choice of vocabulary item in the PF branch. The conclusion is that Encyclopedia entries must be able to look at the PF branch: the meaning of 'tongue' is associated in the Encyclopedia with (44a), while 'back of the head' is tied to (44b).

Marantz (1996) takes this conclusion even further: "We claim that any (unforced) choice made in the course of a derivation is input to semantic interpretation. Thus, for example, any choice of Lexical items for computation, any choice to combine X and Y rather than X and Z, any movement that is not forced by grammatical principle, any choice of Vocabulary item that is not forced by the 'Elsewhere' principle (choose the most highly specified item that fits) — any of these choices may be [. . .] subject to semantic interpretation with the help of the Encyclopedia."

As discussed in Section 2.2.2.3, diacritic features, like the weighted [+lvar] feature that controls the rate of locative *-u* and *-ɛ* for variable lexical items, are not relevant for the syntax and reside firmly in the PF branch. That is, the [+lvar] feature, with its lexically specified weight, is a feature of individual vocabulary items. Thus, we could have two vocabulary items that are identical phonologically but differ in their [+lvar] weight:

(45) *Vocabulary items with different* [+lvar] *weight (early Distributed Morphology)*

a. $\surd \leftrightarrow$ tiːl$_{[+\text{lvar}=1]}$

b. $\surd \leftrightarrow$ tiːl$_{[+\text{lvar}=3]}$

The [tiːl] inserted by (45b) will take locative -*ɛ* substantially more often than the [tiːl] inserted by (45a). Moreover, since this choice is made freely, the difference is interpretable: (45b) can be associated in the Encyclopedia with the meaning 'back of the head', and (45a), with 'rear of troops'. Thus, in this version of Distributed Morphology, the semantically conditioned variation described in Section 5.6.1 is unproblematic, and in fact predicted, since it is handled in the same way as lexically conditioned variation, as shown below:

(46) *Vocabulary items with different phonology and* [+lvar] *weight (early Distributed Morphology)*

a. $\surd \leftrightarrow$ jazık$_{[+\text{lvar}=1]}$

b. $\surd \leftrightarrow$ tiːl$_{[+\text{lvar}=3]}$

In both (45) and (46), the Encyclopedia ascribes a meaning difference to two vocabulary items that differ in the strength of their [+lvar] feature.

This view of root selection and unrestrained semantic interpretation has proven untenable (see Marantz, 2020). First of all, contextually dependent meanings do not seem to be completely unrestrained as predicted by the early model. Anagnostopoulou and Samioti (2013) argue that roots may mean different things in different contexts (a phenomenon known as contextual *allosemy*), but that these contexts seem to obey the same locality restrictions as contextual allomorphy (cf. Marantz, 2013). This suggests a model in which meaning (in non-idiomatic contexts, at least) is inserted on the LF branch by vocabulary items analogous to those proposed for the PF branch.

Another issue with the early version of Distributed Morphology is root suppletion. Marantz (1996, 1997) argues that this model does not allow roots to have contextually determined allomorphs, because this would be evidence for *competition* between vocabulary items spelling out the root, rather than *free choice*. Harley (2014) shows that root suppletion does in fact exist for Hiaki verbs, and thus, the choice between different verbs cannot be made freely by choice of vocabulary item in the PF branch. It must be made earlier, in the syntax, meaning that roots must be distinguished from each other in the syntax as well. However, this distinction cannot have any phonological or semantic content, because the same root can have two different phonological forms or two different meanings. It must be pure syntactic individuation: root lexical items differ from each other by having some sort of index.[10]

The resulting model (contemporary Distributed Morphology) is one in which vocabulary items connect phonological and semantic material, respectively, to individual roots. In (47), we see two such PF vocabulary items. Their phonological content on the right hand side is equivalent to those in (46), but the choice between them is no longer arbitrary, because they are spelling out different syntactic objects on the left: the individuated roots $\sqrt{\text{JAZYK}}$ and $\sqrt{\text{TÝL}}$.

(47) *PF vocabulary items with different phonology and* [+lvar] *weight (modern Distributed Morphology)*
 a. $\sqrt{\text{JAZYK}} \leftrightarrow$ jazık$_{[+\text{lvar}=1]}$
 b. $\sqrt{\text{TÝL}} \leftrightarrow$ tiːl$_{[+\text{lvar}=3]}$

These two vocabulary items also have LF counterparts that spell out their semantic content (presented informally here):

(48) *LF vocabulary items with different semantics (modern Distributed Morphology)*
 a. $\sqrt{\text{JAZYK}} \leftrightarrow$ 'tongue'
 b. $\sqrt{\text{TÝL}} \leftrightarrow$ 'back of the head'

In modern Distributed Morphology, the phonological form [jazık] is paired with the meaning 'tongue' because both spell out the same root, $\sqrt{\text{JAZYK}}$. No longer can the semantic derivation re-

[10]Roots may also need to contain some syntactic information as well, for example category or gender subcategorization.

fer directly to phonological form; the direct link between PF and LF has been severed and rerouted through the syntax.

Similarly, we can still have two different vocabulary items for [tiːl] with different [+lvar] weights, but now they both spell out the same abstract root $\sqrt{\text{TÝL}}$:

(49) *PF vocabulary items with different* [+lvar] *weight (modern Distributed Morphology)*

a. $\sqrt{\text{TÝL}} \leftrightarrow$ tiːl$_{[+\text{lvar}=1]}$

b. $\sqrt{\text{TÝL}} \leftrightarrow$ tiːl$_{[+\text{lvar}=3]}$

As before, we want these two PF vocabulary items to correspond to the meanings 'rear of troops' and 'back of the head', respectively. However, since they both spell out the same root, they will be *syntactically equivalent*: where meaning is determined, on the LF branch, the derivation will only see $\sqrt{\text{TÝL}}$, not the phonological material with the [+lvar] feature. As such, the choice of meaning for $\sqrt{\text{TÝL}}$ must be completely *independent* of its phonological material—that is, its [+lvar] value.

The one option available is to say that the meanings of [jazık], [tiːl], etc. are completely unrelated synchronically—that is, they spell out different roots that only happen to share a phonology and closely related meanings (this is known as accidental homophony—see Section 2.3.2.1). In this case, the variation would be lexically conditioned, equivalent to the paired PF and LF vocabulary items in (47) and (48):

(50) *PF vocabulary items with different* [+lvar] *weight (accidental homophony)*

a. $\sqrt{\text{TÝL}_1} \leftrightarrow$ tiːl$_{[+\text{lvar}=1]}$

b. $\sqrt{\text{TÝL}_2} \leftrightarrow$ tiːl$_{[+\text{lvar}=3]}$

(51) *LF vocabulary items with different semantics (accidental homophony)*

a. $\sqrt{\text{TÝL}_1} \leftrightarrow$ 'rear of troops'

b. $\sqrt{\text{TÝL}_2} \leftrightarrow$ 'back of the head'

This approach is counterintuitive: it is common for words to share a cluster of related meanings, and the rules in (51) assume a very regimented way of how these different meanings are determined in a derivation. Without a principled way to differentiate between polysemy and accidental homophony in synchronic grammars (see Marantz, 2020), we cannot evaluate whether this means of allowing semantically conditioned variation not mediated through syntax is plausible. However, the strict modularity of Distributed Morphology seems in this case to be unable to provide a satsifying account for this kind of semantically conditioned variation.

5.7 General discussion and summary

In this study, I looked at a correlation between the genitive and locative for a subset of nouns in Czech: the more common suffix for both cases is *-u*, while the suffixes *-a* in the genitive and *-ɛ* in the locative are used for a minority of nouns. The two forms are correlated in the lexicon: nouns that take *-a* in the genitive are more likely to also take *-ɛ* in the locative. Unlike the Hungarian study in Chapter 4, variable lexical items play a large role: most nouns that allow the suffix that is not *-u* do so variably, and nouns have different rates of *-u*. The likelihood of *-u* for a given token is also affected by its syntactic context: certain prepositions are more likely to trigger *-u* on their complements than others. Another difference between Czech and Hungarian is that phonology is much less predictive of locative case in Czech; the main effect is the morphological correlation between the cases.

These lexical trends were affirmed by the nonce word experiment: participants' choice of locative for nonce words was predicted in part by its phonology, but more strongly by the genitive that they chose to assign to that word, which mediates the relationship between choice of locative and the experimental condition, the genitive with which the word was presented. Syntactic context also had an effect: just as in the broader lexicon, speakers assigned locative *-ɛ* somewhat more often to words shown with the preposition [na] 'on' than [v] 'in', which in turn had a much higher rate of *-ɛ* than [o] 'about'. A variable word study confirmed that the dependency that speakers show between the genitive and locative of nonce words is a productive extension of lexical patterns: although speakers showed sensitivity to the genitive with which a real word was shown in choosing that word's *genitive*—understandable as a sort of priming effect—there was no such correlation between genitive and selected locative for real words. Thus, the nonce word results cannot be explained as a similar priming effect between the two cases: it must be a genuine morphological dependency.

Although the nonce word study confirmed the general hypothesis that speakers are mimicking the distribution of locative suffixes in the lexicon with respect to phonological, morphological, and syntactic factors, it cannot distinguish between various theories, described in Section 5.4.6.3, to account for how they do so. The multiple grammar model proposed in this work (see also Section 3.3) involves two layers of random selection for cases, like Czech, with variable lexical items: first, speakers stochastically assign a noun a "baseline" rate of locative -*ɛ* in its lexical entry, informed by its phonological and morphological properties. Then, they stochastically choose a locative suffix for a given token according to its lexical "baseline" rate.

In the alternative single grammar model, a locative suffix for a given token is chosen directly according to the phonological and morphological properties of its base. Real words each have their own default rate, so their lexical entry includes a factor that shifts the grammatically expected distribution of locative suffixes up or down to match the word's actual rate.

The two models differ in their predictions for both nonce word studies and naturalistic usage of real words, though in limited circumstances. In the single grammar model, the baseline rate of locative -*ɛ* for a given nonce word should directly correspond to its phonological and morphological properties, while the multiple grammar model predicts some random deviations from this. However, to test this using a nonce word study involves establishing the baseline rate for each participant and nonce word, which requires a huge amount of trials. For real words, the models differ in how they reflect the distribution of locative allomorphs in the input. The single grammar model predicts that a noun's baseline rate should exactly reflect its input, since it simply tracks the deviation between a noun's grammatically expected distribution and its actual distribution. The multiple grammar model, again, allows for some random, if biased, deviation, because lexical baseline rates are initially assigned without regard to the input distribution, although they are later adjusted to be closer towards the input distribution.

I argued that we can gather evidence for the multiple grammar model by looking at the distribution of genitive and locative allomorphs of words that are variable in both cases in the writing of individual authors. Although individual speakers receive somewhat different inputs, the genitive and locative suffixes for a given word may not be particularly correlated in the *input* for a given speaker. Accordingly, if the input can be shown to be unbiased, the single grammar model predicts that the two should not be correlated in individual speakers' *outputs*, either, because the output should be a faithful reflection of the input. In contrast, the multiple grammar model predicts that, if there is a grammatical bias reflecting a correlation between a word's genitive and locative, it should show up in the *output* of individual speakers even if it is not in their *input*. We do in fact see this in the output of individual speakers.

The variable author study thus does double duty. First, it supports the finding of the nonce word study that speakers have learned a correlation between the genitive and locative of individual nouns. Second, by showing that this pattern can be found in real words as well, it potentially provides support for my more complex multiple grammar model over a single grammar model in which lexical entries need not serve as an intermediary for productive generalizations of lexical patterns to influence the generation of inflected forms. This conclusion, however, is predicated on whether future work can show that the correlation between the two cases is not present for the input of individual speakers.

6 Russian plurals and diminutives

This case study looks at the morphological dependencies governing the choice of the Russian diminutive, a common derivational suffix: one common diminutive for masculine nouns, *-ók*, appears more often on words with suffix stress and words with irregular plural *-a*, and Russian speakers learn and apply this correlation to new nonce words. The morphological dependencies in this chapter are of interest for several reasons. First, the target suffix (the diminutive) has three allomorphs, unlike the others, which have two alternants for the relevant subset of the lexicon that I study. This means that my model from Chapter 3, which largely assumed features with at most two (binary) values, must be adapted to fit cases of allomorphy with a greater number of options. Second, this case involves a dependency between an inflectional suffix and a derivational suffix, whereas the previous cases only involved inflectional suffixes. In the theory of morphology I assume, there is no categorical distinction between inflection and derivation. However, theories that make a strong architectural distinction between the two may have issues with accounting for dependencies "across the border". Third, while previous studies looked at correlations between affixes, this study looks at a non-segmental inflectional property, namely stress pattern. As I discuss in Section 6.2.2, Russian stress is usually analyzed primarily through phonological means (differences in underlying stress marks on roots), so the correlation between stress and diminutive suffix may not be a morphological dependency in the strictest sense defined in the introduction to this dissertation. If stress patterns are analyzed with abstract stress marking, the correlation betwen stress pattern and diminutive must be captured as a source-oriented generalization over *underlying* *forms*, since different stress patterns are not always distinguished on the surface in unsuffixed base forms. This can be captured so long as sublexical grammars can operate on underlying rather than surface forms. The relationship between plural *-a* and diminutive *-ók*, on the other hand, is more clearly a morphological dependency. In fact, this study is unique in that the plural–diminutive dependency cannot be stated as a paradigm uniformity effect—that is, it is morphologically *ungrounded* as defined in Section 2.3.2.1. An account of morphological dependencies that explains my experimental data must thus be flexible enough to learn arbitrary correlations between patterns of allomorphy that are not grounded in paradigm uniformity. The sublexicon model described in Chapter 3 provides this needed flexibility.

6.1 Background

6.1.1 Inflection and diminutives in Russian

To provide background before describing the allomorphy of diminutive suffixes in Russian, I begin with an outline of Russian inflectional patterns. Table 6.1 shows the four major inflection classes for nouns, as has been presented at several previous points in this dissertation.

class	I	II	III	IV
example	‘law’	‘school’	‘bone’	‘wine’
nominative	zakon	ʂkol-a	kost^{j}	vin-o
accusative	zakon	ʂkol-u	kost^{j}	vin-o
dative	zakon-u	ʂkol-e	kost^{j}-i	vin-u
genitive	zakon-a	ʂkol-i	kost^{j}-i	vin-a
instrumental	zakon-om	ʂkol-oj	kost^{j}-ju	vin-om
locative	zakon-e	ʂkol-e	kost^{j}-i	vin-e

Table 6.1: Singular paradigms for Russian inflection classes (from Corbett, 1982)

This chapter focuses on inanimate nouns in class I, which are all canonically masculine. Most nouns in class I form their plural with *-i* (often realized as [ɨ]). However, a subset instead take *-a* (see Table 6.3 below). This irregularity interacts with other nominal morphology as well.

One important factor in Russian inflection is stress. Russian nouns can exhibit a number of stress patterns (described in more detail in Section 6.2.2): most nouns have stress fixed on a lexically specified syllable of the *stem* throughout their inflectional paradigm. A significant minority have stress fixed on the inflectional *suffix*, or, when there is no suffix, on the last syllable of the stem.[1] Finally, a smaller number of nouns show one of a number of *mobile* stress patterns, in which some inflected forms have stress on the stem and others, on the suffix. Examples of these stress patterns are shown in Table 6.2. In the nominative singular, which is unsuffixed, all three words bear stress on the one-syllable stem. From there, the words diverge: for [srok] 'period', the stress remains fixed on the stem throughout. For [doz̡dʲ] 'rain', stress is assigned to the suffix, when there is one. Finally, [sloj] 'layer' has some suffixed forms with suffix stress (like the nominative plural) and some with stem stress (like the dative singular). This is an example of a word with mobile stress.

meaning	*nom. sg.*	*dat. sg.*	*nom. pl.*	*stress pattern*
'period'	srok	sróku	srókʲi	stem
'rain'	doz̡dʲ	doz̡dʲú	doz̡dʲí	suffix
'layer'	sloj	slóju	slojí	mobile

Table 6.2: Russian fixed and mobile stress patterns

Certain stress patterns appear more frequently with particular inflection classes—that is, sets of inflectional suffixes (Brown et al., 1996)—but one uniform pattern, with only one exception (described in the paragraph below) is that plural -*a* attracts stress: that is, class I nouns with plural -*a* uniformly stress the suffix (including the irregular nominative plural -*a*) throughout the plural. An

[1] A few exceptional nouns with vowel–zero alternations and fixed suffix stress instead stress the first non-fleeting vowel in suffixless forms, e.g. [úgol] 'corner', whose genitive is [uglá].

example of this correlation is shown in Table 6.3 below.

An even smaller number of nouns have an extended stem with [j] in the plural. These words have *-a* in the nominative plural, though unlike other *-a* nouns, they bear *stem stress* throughout the plural. For example, [lʲist] 'leaf, page' has the nominative and dative plurals [lʲístʲja] and [lʲístʲjam] (stem stress, with the stem extension) with the meaning 'leaf' and [lʲistɨ́] and [lʲistám] (regular, suffix stress, no stem extension) otherwise. I do code these nouns separately (marking them as plural *-ja*) and include them in my data set. However, I do not analyze them, since they are so few in number: there are 20 such nouns altogether, and all but 8 are filtered out of my data set, mostly because they have alternate paradigms without the stem extension, as is the case with [lʲist].

With this background, I can now explain the inflectional patterns of Russian diminutives. Russian has a rich array of suffixes that express a combination of size (diminutive and augmentative) and expressive content (affectionate and vulgar/pejorative). Steriopolo (2008) argues that each inflection class is associated with a particular set of diminutive suffixes—for example, the diminutives discussed in this chapter all attach to class I masculine nouns. If this correlation is truly tied to inflection class (rather than gender or phonology, both of which are themselves partially predictive of inflection class), then it itself constitutes a morphological dependency between inflectional patterning and a derivational suffix. In this study, I focus on morphological dependencies within a narrower class of inflection patterns and diminutives. Three of the class I masculine diminituve suffixes are productive: *-ók*, *-ʲik*, and *-tʃʲik* (Gouskova et al., 2015; Magomedova, 2017; Magomedova & Slioussar, 2017). Magomedova and Slioussar (2017) argue that *-ók* is less productive than the other two: it bears primary stress (triggering a stress shift) and often triggers consonant alternations in the stem as well, and speakers are rather reluctant to extend it to loanwords (see also Gouskova et al., 2015). However, my study results suggest that this claimed lack of productivity is due at least in part a consequence of previous accounts not accounting for some of the factors that prefer *-ók*.

Most nouns with suffix stress in some or all inflected forms, including those with plural *-a*, form their diminutive with *-ók*, which likewise attracts stress. Examples of the intersection of plural and diminutive for class I nouns are shown in Table 6.3:

noun class	*example*	*gloss*	*plural*	*diminutive*
regular	motór	'motor'	motór-ɨ	motór-tʃʲik
regular	moróz	'frost'	moróz-ɨ	moróz-ʲik
-a plural	górod	'city'	gorod-á	gorod-ók

Table 6.3: Russian masculine plural and diminutive forms

This study tests the relationship between inflectional stress pattern and plural suffix and diminutives for class I nouns in the lexicon and in speakers' behavior. Unlike the Hungarian study in Chapter 4, this Russian study explores a morphological dependency between an inflectional suffix and a derivational suffix. This is unproblematic in theories where inflectional and derivational morphology are handled in the same module, like Distributed Morphology, but may not be predicted by theories that make a modular divide between them. In addition, the correlation between plural *-a* and *-ók*, unlike that discussed in Chapter 4, cannot be captured as a paradigm uniformity constraint (see Section 2.3.2.1). This highlights the need for a theory of morphological dependencies, like the sublexicon model described in Chapter 3, that can capture arbitrary correlations between inflectional patterns not necessarily grounded in particular kinds of markedness.

One final difference between this study and the previous ones is that this study involves *three* possible suffixes rather than two, with some overlap of words that admit multiple suffixes. Thus, I analyze a word's diminutive class as a set of three binary choices indicating whether each diminutive suffix is accepted or not, and focus on the opposition of *-ók* vs. the other two options. This is of theoretical interest, because my models of productivity (in Section 3.1.2) and variation (in Section 3.4) assumed that allomorphy can be modelled as a single binary choice. As I show in

Section 6.2, these models can also handle choices between more than two options.

6.1.2 Phonological effects

Before describing my own study, I will summarize previous corpus and nonce word studies of the Russian diminutives. My main point of departure is Gouskova et al. (2015), whose studies are very similar to mine. They test basic phonological generalizations about the distribution of diminutives from Polivanova (2008 [1967]) using a web corpus, a nonce word study, and a study eliciting diminutives of real words that do not have well-established diminutive forms. The restrictions observed by Polivanova (2008 [1967]) are shown in Table 6.4; patterns corroborated by the elicitation study of Gouskova et al. (2015) are bolded.

	-tʃʲik	*-ʲik*	*ók*
final coda	**no clusters**	—	**no clusters**
final C place	no dorsals	**no dorsals**	—
final C manner	**prefers sonorants**	—	—
stress	**must be final**	**must be final**	**not medial**

Table 6.4: Phonological generalizations about Russian diminutives according to Polivanova (2008 [1967]); bolded observations supported by elicitation study in Gouskova et al. (2015, p. 62)

In most cases, Polivanova's generalizations were affirmed by the elicitation study and other studies. In one case, speakers were more liberal than the lexicon—for example, they sometimes attached *-tʃʲik* to both real and nonce words ending in dorsals, despite such forms not appearing in the lexicon.

Gouskova et al. (2015) also report novel phonological patterns in both the lexicon and speaker behavior with nonce words. Their analysis of the nonce word study focuses on several phonological factors: final place and consonant, stress pattern, presence of a cluster (word-final and word-internal) and presence of word-internal hiatus. In Table 6.5, I list dispreferences in these areas that Gouskova et al. (2015, p. 54) find in the lexicon; those that were also significant in their nonce word study are bolded. Some of these line up with those of Polivanova (2008 [1967]) (for example, the dorsal effect), while others are novel.

	-tʃʲik	*-ʲik*	*ók*
dispreferred clusters	**word-final**	**word-medial**	**word-final**, word-medial
dispreferred hiatus	—	yes	**yes**
dispreferred final C place	**dorsal**	**dorsal**	—
dispreferred final C manner	—	**sonorant**	nasal
dispreferred stress	—	—	final, medial

Table 6.5: Phonological dispreferences found by Gouskova et al. (2015, p. 54) in the Russian lexicon; bolded observations supported by their nonce word study

One novel finding in Table 6.5 is that diminutive suffixes can impose well-formedness constraints on their bases whose loci are not at the right edge of the stem: for example, words with hiatus disprefer *-ók*, even though hiatus is stem-internal and not adjacent to the attachment point of the suffix. In these and other cases, as Gouskova et al. (2015) point out, the preferences that individual diminutives impose do not seem to be grounded in decreasing markedness of the derived form (see Section 2.3.1.2). They find that speakers nonetheless learn some of these "ungrounded" generalizations, such as the hiatus one.

Markedness-based explanations have been offered for some additional restrictions, but as discussed in Section 2.3.1, it is difficult to use this as a common-sense explanation for all phonological patterns. For example, Gouskova et al. (2015) find that the vowel in the last syllable of a noun root affects its diminutive selection in the lexicon: words that have *rounded* vowels in the last syllable, like [boj] 'battle', are *more* likely to combine with *-ók*, while words that have *high* vowels in the last syllable, like [mir] 'world', are *less* likely to combine with *-ʲik* (both of these words form diminutives with *-ók*: [boj-ók], [mir-ók], ignoring vowel reduction—see below). The authors suggest that the former is "a kind of vowel harmony pattern" (Gouskova et al., 2015, p. 54) with reference to a similar pattern for Hebrew plurals described by Becker (2009) (see Section 2.3.1.2)—that is, there is pressure for vowels in adjacent syllables to agree in rounding. However, Magomedova (2017, p. 136) calls the latter "an OCP-like pattern", referencing a constraint (the Obligatory Contour Principle) that puts pressure on vowels in adjacent syllables to *disagree* in some feature, in this case height. Both types of patterns occur in languages, but without a principled explanation of why *-ók* should want its vowel to agree with its base while *-ʲik* wants it to disagree, there is no a priori reason to favor one over the other in the name of markedness. The rounding harmony explanation for *-ók* is particularly unlikely because unstressed /o/ actually reduces to *unrounded* [ə] or [ɐ] (that is, the actual pronunciation of the diminutive of [boj] is [bɐjók]), so stem /o/ and suffix /o/ do not actually agree in rounding on the surface. Indeed, my studies find the opposite vowel rounding pattern from what Gouskova et al. (2015) report: nouns with underlying rounded vowels actually *disprefer -ók*.

Some other lexical patterns for the diminutive (whether speakers learn them or not) are likely grounded in functional and historical factors. For example, Gouskova et al. (2015) find that monosyllables are more likely to combine with *-ók* and *-ʲik* and less so with *-tʃʲik*. While *-ók* is the oldest form, it is losing ground to the others, especially *-tʃʲik*, which is a newer suffix often used with loanwords (Magomedova, 2017; Magomedova & Slioussar, 2017). Magomedova (2017) also claims that children prefer *-ʲik*—which she attributes to its affectionate nature (see Section 6.1.3). The distribution of monosyllables, then, may reflect the fact that older and monomorphemic native words and words used heavily by children tend to be shorter than more recent loanwords. This distributional fact may thus reflect the accretion of the lexicon over time and is not necessarily driven by grammatical pressures like phonological markedness.

The most studied phonological factor in the distribution of diminutives is the effect of word-final velars [k ɡ x]. Originally, nouns ending in velars were diminutivized with (usually) stressed *-ék*, which triggered mutation of those velars to [tʃʲ ʐ ʂ], while other nouns received *-ók*. However, a phonological change conflated the two vowels in many cases (not just in the diminutive), and the modern descendants of these suffixes are unstressed *-ək*, which only attaches to nouns ending in velars, and stressed *-ók*, which attaches to nouns ending in velars and other consonants (Kuznecov, 1953; Magomedova & Slioussar, 2017). The two suffixes are usually orthographically identical, and both generally trigger mutation of velars, examples of which are shown for *-ók* in Table 6.6.

word	*base*	*diminutive*
onion	luk	lutʃʲ-ók
meadow	luɡ	luʐ-ók
poem	stʲix	stʲiʂ-ók

Table 6.6: Alternation of velars with diminutive *ók* (Magomedova & Slioussar, 2017, p. 248)

The other diminutives, *-tʃʲik* and *-ʲik*, are traditionally described as rejecting velar-final nouns, but the latter does occasionally attach to nouns ending in velars. Kapatsinski (2010) finds one such diminutive, [bantʃʲik], from [bank] (with velar mutation), in a Russian dictionary (Ševeleva, 1974), although Zaliznjak (1977) does not include this word (very productive diminutives are often omitted from dictionaries). Magomedova (2017) finds some attestations of other velar-final nouns attaching to *-ik* and mutating the velar, such as [sapoʐɨk] as the diminutive of [sapoɡ] 'boot'.[2] (Both [bantʃʲik] and [sapoʐɨk] are attested in my corpus.) Kapatsinski (2010) and Magomedova and Slioussar (2017) find several examples of velar-final nouns taking *-ʲik* in web searches.

The examples I gave above show stem-final velars undergoing the same alternations before *-ʲik* as *-ók* (shown in Table 6.6). However, sometimes *-ʲik* attaches to velar-final stems without triggering a change in the stem consonant. Kapatsinski (2010) finds that the *-ók* diminutive consistently trig-

[2]The realization of the suffix as [ɨk] is due to a regular allophonic process following [ʐ].

gers velar stem alternations (as shown in Table 6.6), while *-ʲik* often fails to mutate velars. In particular, [g] is more likely to remain unmutated before *-ʲik* (e.g. [sapogʲik]) than [k] (e.g. [bankʲik]). Magomedova and Slioussar (2017) find more moderate results: *-ók* occasionally attaches to unmutated velars, while *-ʲik* almost always attaches to mutated stems of recent loanwords (however, older velar-final words attaching to [-ʲik] mutate much less often). In the nonce word study of Gouskova et al. (2015), speakers were asked to select the best of three possible diminutives of stimuli with the suffixes *-ók*, *-ʲik*, and *-tʃʲik*. Stimuli ending in velars were presented with mutations before *-ók* but not *-ʲik*: that is, the nonce word [xurʲak] was given the diminutives [xurʲatʃʲok], [xurʲakʲik], and [xurʲaktʃʲik]. In my nonce word study (Section 6.4), I also mutate stem-final velars before *-ʲik*: [xurʲatʃʲik].

6.1.3 Semantic effects

6.1.3.1 Expressive -ók and -ʲik: claimed problematic for Distributed Morphology

The previous discussion assumes that the three diminutive suffixes are lexically conditioned allomorphs without a distinction in meaning. Magomedova (2017) challenges this assumption, arguing that there is a meaning difference as well: in a nonce word study, she finds that *-ók* tends to be used in a pejorative sense, while *-ʲik* has more of an affectionate meaning (*-tʃʲik* remains neutral). These tendencies are not absolute, and interact with phonological factors: in a nonce word study, stimuli that were phonologically well-suited for *-ók*, in particular [snʲik] (which ends in a velar), were often assigned *-ók* even in affectionate contexts, and vice versa: prototypical *-ʲik* nonce words like [vʲips] (ending in a fricative-final cluster) were often given *-ʲik* diminutives even in pejorative contexts. She finds that *-ók* is, in general, used less productively than expected, and words that traditionally form their diminutives with *-ók* (like [sapog] 'boot', usual diminutive [sapoʐók]) are sometimes attested in corpora with *-ʲik* in affectionate contexts ([sapoʐik], as mentioned above). Thus, she calls them "pseudo-allomorphs": their distributions are not purely based on lexical phonological tendencies, but bear with them slight differences in meaning.

Magomedova (2017) argues that serialist theories like Distributed Morphology cannot account for the interaction of phonological and semantic conditioning in Russian diminutive assignment. However, her argument is based on a very specific analysis of Russian diminutives. Before discussing her analysis, I must clarify one point: if the conditioning were *exclusively* semantic, this would be problematic for Distributed Morphology, in which semantics and phonology are separate strands of the derivation downwind of syntax. Indeed, I discuss one such problematic case in Section 5.6.4.3. This is not the argument that Magomedova (2017) makes; rather, she assumes that various features of diminutives are encoded in the syntax (on a separate head, which she adjoins to the noun and labels DIM) using features like [+expressive], [+pejorative], and [+affectionate]. These features are then accessible to be spelled out phonologically by various diminutives.

Her analysis assumes a feature geometry where [+pejorative] and [+affectionate] are both dependent on [+expressive]—so *-ók* is marked as [+pejorative], *-^{j}ik* is marked as [+affectionate], and both are [+expressive]. However, spellout cannot involve a simple pairing of these feature sets, since (for example) *-ók* can sometimes be used in affectionate contexts. Morphologically, then, "the more specific affectionate feature may be traded for phonological well-formedness in a way to keep more general {+expressive}, despite the fact that pejorative nuance [associated with *-ók*] is actually right opposite to the speaker's original intention" (Magomedova, 2017, p. 138).

If *-ók* and *-^{j}ik* are the only diminutive suffixes marked for [+expressive], this is not a problem: whatever mechanism you have for handling variation in Distributed Morphology can weigh the factors correctly to match the distribution of the two suffixes. (For example, an analysis very similar to the one I present for syntactically conditioned variation in Czech in Section 3.4.) However, Magomedova (2017) notes that there are other diminutive suffixes than just the three discussed here, like the affectionate suffix *-ul^{j}a*, as in [d^{j}edul^{j}a], an affectionate form of [d^{j}ed] 'grandfather'. However, Vinogradov (1972) claims that *-ók*, *-^{j}ik*, and *-tʃ^{j}ik* are *non-expressive*, while other diminutive suffixes, like *-ulʲa*, are expressive. If *-ók* and *-ʲik* have taken on expressive content, they should now be syntactically and semantically similar to the other diminutives, in particular by bearing [+expressive]. Thus, if *-ók* can theoretically spell out [+expressive] alone, it should be in competition with all of the other diminutive suffixes in Russian, not just *-ʲik*.

Magomedova (2017) describes a prediction made by her straw Distributed Morphology analysis. Her nonce word study was a forced choice task, so speakers could only choose between *-ók*, *-ʲik*, and *-tʃʲik*. However, the test could be repeated with an additional suffix that is purely affectionate in nature, such as *-ulʲa*. In her straw analysis, *-ók* can spell out a syntactic node with [+affectionate] when it has enough of a phonological advantage over the diminutive that directly spells out [+affectionate], *-ʲik*. However, if *-ulʲa* is a competitor and the nonce word is a good phonological fit for both *-ulʲa* and *-ók*, *-ulʲa* should win every time because it is a better match syntactically (bearing [+affectionate]) without a phonological counterweight. If *-ók* "still appears considerably often" (Magomedova, 2017, p. 138), that would suggest that her Distributed Morphology analysis is wrong, because a suffix is occasionally winning competitions it should not.

6.1.3.2 -ók and -ʲik do not compete with other expressives

I agree with Magomedova (2017) that we would still expect *-ók* to be chosen even if *-ulʲa* is an option. However, contrary to her assertion, this is not a problem for theories like Distributed Morphology. The key question is whether diminutives like *-ók* and *-ʲik* are in direct competition with expressive suffixes like *-ulʲa*. In Magomedova's analysis, the *-ók* and *-ʲik* share a syntactic feature, [+expressive], which puts them into competition. (Presumably the expressive meanings bring them in line with the others semantically as well.) Steriopolo (2008), however, argues that *-ók* and *-ʲik* have always been expressives. In her account, they differ from other expressives like *-ulʲa* both syntactically and semantically.[3] Semantically, suffixes like *-ulʲa*, which she calls attitude

[3] She also groups another less productive diminutive, *-ʲets*, together with *-ók*, *-ʲik*, and the aforementioned *-ək*. She does not discuss *-tʃʲik*, which is etymologically derived from a sequence of *-ʲets* (with a vowel–zero alternation and a consonant mutation) and *-ʲik*.

suffixes (which I gloss as ATT), bear no descriptive content, only expressive content. We see this in (52): the contribution of *-ulʲa* is strictly about the speaker's attitude towards the grandmother:

(52) bab-ulʲa prʲi̯şla
grandmother-ATT came
Descriptive: 'Grandmother came.'
Expressive: speaker views grandmother positively (Steriopolo, 2008, pp. 16–17)

This can be contrasted with (53): suffixes like *-ók* and *-ʲik* contribute both a true diminutive meaning (hence she calls them size suffixes, or SIZE below) and an affective meaning:

(53) domʲ-ik stojit na gorʲe
house-SIZE stands on mountain
Descriptive: 'A small house stands on a mountain.'
Expressive: speaker views house positively (Steriopolo, 2008, p. 18)

In Steriopolo's judgement, neither component of the meaning of the suffix in (53) can be explicitly rejected: it is infelicitous to refer to a big [domʲik] or say "I hate my [domʲik]."[4] Thus, Steriopolo (2008) argues, it is not the presence of expressive content of attitude suffixes like *-ulʲa* that distinguish them from diminutives like *-ʲik*, it is the lack of descriptive content. Under a Distributed Morphology analysis like that of Magomedova (2017), we would thus expect *-ók* and *-ʲik* to agree in a feature like [+size] and/or [+diminutive] that *-ulʲa* lacks. Accordingly, if vocabulary items are specified for [+size] in addition to [+expressive], we would not expect *-ók* to compete with *-ulʲa* at vocabulary insertion, even if *-ók* takes on some new expressive meaning.

Even if the semantic distinction between attitude and size suffixes discussed by Steriopolo (2008) is not encoded featurally in the syntax, they are still different syntactically. Specifically, attitude suffixes are *heads*, while size suffixes are *modifiers*. She uses a number of criteria to argue for this distinction: for example, heads can change the category of their sisters, while modifiers cannot.

[4]Her judgements are not universal: for some speakers, the affective meaning of both *-ʲik* and *-ulʲa* are cancellable.

Accordingly, an attitude suffix like *-ulʲa* or the vulgar *-ʲuga* can attach to adjectives to form nouns, while *-ók* cannot:

(54) a. ʐ̺ad-n-ɨj
stingy-ADJ-M.NOM.SG
'stingy'

b. ʐ̺ad-nʲ-uga
stingy-ADJ-ATT
'stingy one (vulgar)' (Steriopolo, 2008, p. 65)

c. * ʐ̺ad-n-ok
stingy-ADJ-SIZE
'stingy one (diminutive)' (Steriopolo, 2008, p. 67)

Thus, even if size suffixes like *-ók* and *-ʲik* are not distinguished *featurally* from attitude suffixes like *-ulʲa* and *-ʲuga*, they would still be distinguished *structurally*, because they are in a different structural relation with the stem to which they are attaching. Thus, again, we would have a way to distinguish the two types of suffixes at the point of vocabulary insertion, so they do not compete directly for insertion in the same tree under a Distributed Morphology analysis. In Section 6.2, I only discuss the diminutives *-ók*, *-ʲik*, and *-tʃʲik*, and assume that attitude suffixes do not play a role in the derivation.

6.1.3.3 Another interpretation of Magomedova (2017)

Before discussing my own formal analysis, I argue that there is an alternate interpretation of the results of Magomedova (2017). Her claims that *-ók* is inherently pejorative and *-ʲik* is inherently affectionate rest on two pieces of evidence: a web search and a nonce word study. I address each of these in turn, concluding that both are compatible with an interpretation in which the diminutive suffixes are not obligatorily associated with pejorative and affectionate features in the syntax. If this is the case, then the Russian diminutive is not a case of lexically and syntactically conditioned variation in the sense described for Czech in Section 3.4. Avoiding syntactically conditioned variation is theoretically desirable, as it makes the derivational analysis simpler: a diminutive allomorph of a known noun depends only on the noun to which it attaches, not the syntactic (and semantic) context in which it appears. The purpose of this section is to justify the simpler formal analysis that does not encode pejorative and affectionate meaning directly into the syntax.

In the web study, Magomedova (2017) searched Google for *-ʲik* diminutives of four words that typically take *-ók*: [snʲeg] ‘snow’, [sapóg] ‘boot’, [xomʲák] ‘hamster’, and [sɨr] ‘cheese’ (standard diminutives [snʲez̨ók], [sapoz̨ók], [xomʲatʃʲók], [sɨrók]). In each case, she found *-ʲik* diminutives for these words as well: [snʲéz̨ɨk], [sapóz̨ɨk], [xomʲátʃʲik], [sɨ́rʲik]. She then manually inspected results for the first three words on a single site, *woman.ru*. The few tokens of these words with *-ʲik* appeared in affectionate contexts, as shown in Table 6.7.

word		*-ʲik hits*	*-ʲik context*	*-ók hits*	*-ók context*
[sapóg]	‘boot’	12	affectionate	470	affectionate or neutral
[snʲeg]	‘snow’	2	affectionate	3428	(not coded)
		34	nicknames/pet names		
[xomʲák]	‘hamster’	2	affectionate	10079	(not coded)

Table 6.7: Base alternations triggered by Russian diminutive suffixes

Magomedova (2017) takes this as evidence that *-ʲik* can be used in affectionate contexts, even for words with established *-ók* diminutives. This is true, but it does not necessarily follow that the use of *-ʲik* is specifically due to the affectionate context. For example, if a sufficiently large majority of diminutives are affectionate, we might expect *any* 12 randomly selected diminutives to all be affectionate, regardless of the suffix. That is, we must rule out the null hypothesis that there is no difference in the underlying proportions of affectionate diminutives for *-ók* and *-ʲik*. Of course, this depends on how many of the 482 total diminutives of [sapóg] (with either suffix) are affectionate. Let us call this number, which Magomedova (2017) does not report, x. We can then calculate the probability of randomly selecting 12 out of x affectionate diminutives given the null hypothesis as $\frac{x}{482} \cdot \frac{x-1}{481} \cdot \frac{x-2}{480} \cdots \frac{x-11}{471}$. Each of the terms in this product corresponds to a selection: first, we get to choose between one of the 482 diminutives, and the odds that this chosen diminutive is affectionate is $\frac{x}{482}$. Next, we must choose one of the remaining $x-1$ affectionate diminutives from the remaining total 481 diminutives, and so on until we have selected 12 diminutives. Clearly, then, the likelihood of the null hypothesis depends entirely on x. When $x = 400$, the likelihood is around 10%; when $x = 455$, it is approximately 50%. If affectionate diminutives of [sapóg] are very common in general, then, it is possible that a small number of speakers simply allow *-ʲik* diminutives for these words regardless of context.

In her nonce word study, Magomedova (2017) found that Russian speakers were more likely to assign *-ók* to stimuli presented in pejorative contexts, and *-ʲik* to stimuli presented in affectionate contexts (and vice versa); context had no effect on assignment of *-tʃʲik*. She interprets this result as evidence that *-ók* has an inherent pejorative meaning, while *-ʲik* has an inherent affectionate meaning. However, this does not necessarily follow. As in my nonce word studies, each speaker only saw each word once. In this dissertation, I have assumed that in a nonce word task, the frame provides information that can influence how speakers assign a word to a class when they are informed to choose between different allomorphs of (e.g.) the diminutive. In this case, the emotional valence of the frame sentence could be used to assign a word to its diminutive. That is, words that tend to appear in affectionate contexts would be more likely to be assigned *-ʲik*, while those that tend to appear in pejorative contexts would be more likely to take *-ók*—a semantic generalization similar to the phonological and morphological generalizations studied in this dissertation. Under this hypothesis, we would expect words that appear more often in affectionate contexts to be more likely to take *-ʲik*, *even in other contexts*. This is less direct than Magomedova's analysis, in which the context adds formal syntactic or semantic features to the derivation that then match up with the inherent features of the respective diminutives through a formal process. I call her interpretation of the results the *direct meaning* hypothesis and mine, the *pattern matching* hypothesis.

Her nonce word study results are compatible with both hypotheses. How do we distinguish them? One possibility is through repeated trials using the same nonce word. Suppose participants are asked to form diminutives first in an affectionate context, then in a neutral context. The direct meaning hypothesis predicts that participants should not show the same bias towards a nonce word taking *-ʲik* when it is presented for the second time in the neutral context, because the affectionate meaning is no longer there. The pattern matching hypothesis predicts that the *-ʲik* bias should persist even in the neutral context, because speakers form a bias from *initially seeing* the word in an affectionate context, and once a word has been assigned to the *-ʲik* class, it is more or less set.

Another test to distinguish the two hypotheses would be to ask speakers to form the diminutive of real words that can take both *-ʲik* and *-ók* diminutives, like [most] 'bridge', presented in affectionate, pejorative, and neutral contexts. The direct meaning hypothesis predicts that such words should be assigned *-ʲik* more in affectionate contexts and *-ók* more in pejorative contexts; the pattern matching hypothesis predicts that there should be no effect of context, because real words already have established lexical entries, and thus speakers should not be relying on semantic generalizations to fill in the words' lexical entries. In fact, these two studies pair together very well: the direct meaning hypothesis predicts a null result in the repeated nonce word task and a positive result in the variable word test, while the pattern matching hypothesis predicts the opposite.

In this dissertation, I assume the pattern matching hypothesis. However, the proposed semantic generalization for the Russian diminutive does not fall directly into my model: I do not consider the effects of semantics (distributional or otherwise) on assignment of morphological behavior, both because it is generally difficult to assign meaning to nonce words and because my model is situated in the PF branch of a model where phonology and semantics are spelled out separately from syntax (see Section 5.6.4.3). However, if a semantic generalization were to be shown, this would be in line with previous research, discussed in Section 2.1.4, showing that distributional semantics is predictive of a word's inflectional behavior (e.g. Guzmán Naranjo, 2020; Williams et al., 2020). In Section 6.2, I do not consider possible semantic differences between the diminutives in my formal analysis.

6.2 Formal analysis

In this section, I lay out a formal analysis of Russian diminutive suffixes and the morphological factors influencing choice of diminutive, namely inflectional stress pattern and nominative plural suffix. Russian stress is quite complicated, so I will not provide a full analysis. The relevant issue is: how are different stress patterns represented underlyingly? If stress patterns are represented through abstract underlying phonology, then the correlation between stress and diminutive is not a morphological dependency, but rather one between phonology and morphology. However, if these phonological properties are only visible in the underlying form, then the sublexical phonotactic grammars storing generalizations, described in Chapter 3, must be able to operate on underlying rather than surface representations. On the other hand, if inflectional stress is represented through diacritic features, then the correlation between stress pattern and diminutive is a morphological dependency in the sense used in this dissertation: a correlation between two diacritic features appearing on the same underlying form. In most analyses, stress is analyzed phonologically, sometimes with some diacritic features, so the influence of stress on diminutive suffix is not necessarily a morphological dependency. On the other hand, the choice of plural is clearly a selectional property indexed by a diacritic feature, so the dependency between plural *-a* and diminutive *-ók* is a morphological dependency like those in Hungarian and Czech.

However inflectional stress is encoded underlyingly, the relationship between stress pattern and diminutive can also be expressed as a *paradigm uniformity* constraint among surface forms: the diminutive *-ók*, which bears stress (that is, attaches to an unstressed stem), prefers nouns that have suffix or mobile stress paradigms (that is, have unstressed stems in at least some inflected forms). As discussed in Section 2.3.2.1, this dependency is thus morphologically grounded. Plural *-a* is also obligatorily stressed, so at first glance, it seems as if the correlation between this plural and diminutive *-ók* is also a paradigm uniformity effect. However, we are interested in this correlation to the extent that it is *even stronger* than what would be expected *given* the stress pattern: we wish to account for the fact that nouns with stressed plural *-a* take diminutive *-ók more often* than nouns with stressed plural *-i*. The residual effect of the segment *-a* is thus not a paradigm uniformity effect: it is the only morphological dependency discussed in this dissertation that cannot plausibly be expressed as morphologically grounded in a pressure for related forms to agree in some property. The existence of this effect is thus evidence for a less restricted pattern-matching model, like my sublexicon model, which can learn both grounded and arbitrary morphological dependencies.

The choice between diminutives is one of simple allomorph selection, and is described briefly in Section 6.2.1 before I go into the stress system. In this case, I assume that a given noun can have at most one diminutive. This is an oversimplification: nouns may be able to combine with two or even all three diminutive suffixes. I provide an account for nouns with variability in the diminutive in Section 6.2.3. Although not especially relevant for the empirical studies in this section, this variability is theoretically important because there are *three* outcomes, meaning that I cannot model the choice of diminutive as a binary outcome as I had in Section 3.4. The discussion in Section 6.2.3 shows that the assumption of binarity is not a necessary one, and that my model of variation can handle lexical itmes that may vary among more than two allomorphs.

6.2.1 The basic case

Here I describe the formal analysis of the Russian diminutive suffixes. This analysis is very simple. Following Steriopolo (2008) and Gouskova et al. (2015), I assume that these diminutives (which they call size diminutives, see Section 6.1.3) are modifiers to a nominal projection *n*. This structure, shown for [stol^{j}ik], the diminutive of [stol] ‘table’, is shown in (55).

(55) *Morphosyntactic structure of Russian diminutives*

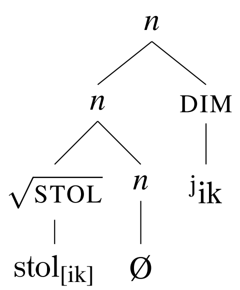

This derivation proceeds as follows: first, the syntactic structure is built up, with an uncategorized root $\sqrt{\text{STOL}}$, a categorizing head *n*, and the diminutive suffix. These are then spelled out via vocabulary insertion, starting with the root and moving outwards (Bobaljik, 2000) as discussed in Section 2.2.2.3. The underlying form of the root includes the sublexical diacritic feature [ik]. Since the *n* head is null, the phonological context available at the point when the diminutive is to be spelled out is just the underlying form of the root: /$\text{stol}_{[\text{ik}]}$/. The diminutive is then spelled out as ^{j}*ik* using rule (56b), which matches the [ik] feature on the noun. The representations in (56) ignore the stress properties of *-ók*, which will be explained in Section 6.2.2.4 below. Thus, the underlying form in (56a) is incomplete, but is sufficient for the purposes of showing how each diminutive is indexed to a lexical diacritic feature.

(56) *Rules of realization for the Russian diminutive (representation of stress omitted)*

a. DIM ↔ ok / [ok] ___

b. DIM ↔ ^{j}ik / [ik] ___

c. DIM ↔ tʃ^{j}ik / [chik] ___

This simple case holds for nouns[5] that take at most one of the three possible diminutives in (56). The vocabulary item spelling out a root is decorated with a diacritic feature indexing its diminutive: nouns that take *-ók* have an [ok] feature, nouns like [stol] that take *-^{j}ik* have an [ik] feature, and

[5]Under this theory, diminutive selection is controlled by the outermost non-zero affix, so "noun" here means the root or whatever derivational affix is closest to the diminutive. This predicts that nouns ending in the same derivational affix should take the same diminutive, since the selection is contained in the underlying representation of the affix itself.

nouns that take *-tʃʲik* have a [chik] feature. There is no default diminutive rule in (56): nouns lacking a lexical diminutive feature are unable to form diminutives; this can account for the fact that some nouns in Russian resist diminutive formation (cf. Gouskova et al., 2015, p. 50). For variable nouns that allow more than one diminutive, see the analysis in Section 6.2.3.

When a speaker attempts to form a novel diminutive, she must assign it to a sublexicon (that is, assign a feature to it). In this, I follow Gouskova et al. (2015): each of the three diacritic features in (56) is associated with a sublexical grammar (in my case, one that contains both phonological and morphological constraints, as in my other studies) that evaluates the novel form. The speaker then places the word into one of these sublexicons (that is, places the feature on the lexical item) according to its score: the better a word's score on a sublexical grammar, the more likely the word is to be placed into that lexicon. (See Chapter 3 for a full explanation of the sublexical grammar theory.) In this case, unlike in the previous studies, there are three possible options instead of two. However, this is not a problem for the theory: using a maximum entropy model (Goldwater & Johnson, 2003; Hayes & Wilson, 2008), a word's likelihood of being assigned to a sublexicon is still proportional to the exponential of its score evaluated by that sublexicon's grammar; this time there are three scores in the denominator instead of two. If $s(f,w)$ is the score assigned to word w by the sublexical grammar for feature f, then the probability $P(f,w)$ of the speaker placing f onto the lexical entry of w is equal to $\frac{e^{s(f,w)}}{\sum\limits_{f_i \in F} e^{s(f_i,w)}}$, where the denominator is the sum of the exponential of the scores assigned to w by all of the relevant features f_i. In the case of the Russian diminutive there are three such features, so the probability of a word w being assigned [ok] (for example) is $P([\text{ok}],w) = \frac{e^{s([\text{ok}],w)}}{e^{s([\text{ok}],w)} + e^{s([\text{ik}],w)} + e^{s([\text{chik}],w)}}$.

6.2.2 Encoding inflectional stress patterns

In this section, I give a descriptive and theoretical overview of Russian stress, focusing on inflectional stress patterns but also touching on the stress properties of the obligatorily stressed diminutive -*ók*. In particular, I will show that its obligatory stress is *different* in nature from that of class I plural -*a*, which also always bears stress.

The nouns in my corpus (inanimate nouns in class I) follow four stress patterns, shown in Table 6.8 for the nominative, dative, and genitive singular and plural: fixed stress, fixed suffix, and two mobile stress patterns. In both mobile patterns, stress is fixed on the stem in the singular. In one pattern, exemplified by [kolokol] 'bell', suffix stress appears throughout the plural; most nouns with plural -*a*, including [kolokol], fall into this class (though a couple have fixed suffix stress). However, this class also contains many words with plural -*i*, such as [nos] 'nose'. In the other, shown for [volos] 'hair', the nominative also has stem stress (as does the accusative, which is not shown but is identical to the nominative for inanimate nouns, as is typical for Slavic languages—see Section 2.1.5.2), but the rest of the plural cases have suffix stress. Nouns with fixed stem stress are by far the most common, while mobile stress is the least common pattern by type frequency. However, very frequent words are substantially overrepresented among words with mobile stress (Mołczanow et al., 2013, p. 166). In all cases, when a stressed suffix is longer than one syllable (the instrumental plural -$am^j i$), the first syllable of the suffix is stressed.

stress pattern	stem	suffix	mobile	
example	‘character’	‘pencil’	‘bell’	‘hair’
nominative singular	xaráktʲer	karandáʂ	kólokol	vólos
dative singular	xaráktʲeru	karandaʂú	kólokolu	vólosu
instrumental singular	xaráktʲerom	karandaʂóm	kólokolom	vólosom
nominative plural	xaráktʲerɨ	karandaʂɨ́	kolokolá	vólosɨ
dative plural	xaráktʲeram	karandaʂám	kolokolám	volosám
instrumental plural	xaráktʲeramʲi	karandaʂámʲi	kolokolámʲi	volosámʲi
number of noun types	7346	655	164	13
usual location of stress on stem	anywhere	final	initial	

Table 6.8: Stress patterns for Russian class I nouns (cf. Zaliznjak, 1977)

The abridged paradigms in Table 6.8 also show how the location of stem stress varies by stress type. For nouns with fixed suffix stress like [karandáʂ] ‘pencil’, the last syllable is stressed in the nominative singular, which has no suffix to bear stress. The only exception is a small number of stems with vowel–zero alternations like [úgol] (genitive [uglá]), where stress is on the last non-alternating vowel. However, other alternating nouns with suffix stress do bear stress on the alternating final vowel, as expected, like [konʲéts] ‘end’ (genitive [kontsá]). By contrast, nouns with mobile stress typically have stress on the *first* syllable of the stem when stressed. Exceptions include compounds or prefixed nouns, which bear stress on the initial syllable of the head (e.g. [elʲektroplúg] ‘electric plow’, derived from the mobile-stress noun [plúg] ‘plow’) and a scant few others, like the technical borrowing [kokʲílʲ], referring to a type of metal cast. Finally, nouns with fixed stem stress can bear stress on any syllable: compare penult-stressed [xaráktʲer] with initial-stressed [mʲínʲimum] ‘minimum’ and final-stressed [instʲitút] ‘institute’ (here I use loanwords for purposes of practical demonstration: most native three-syllable nouns are not monomorphemic).

6.2.2.1 Stem stress

As discussed previously, we are interested in how the different stress patterns are represented in the lexicon, because this tells us what speakers are generalizing over if they have learned a correlation between stress pattern and diminutive. Mołczanow et al. (2013, p. 167) provide a concise overview of the representation of Russian stress patterns in generative analyses (Alderete, 1999; Gouskova, 2010; Halle, 1973; Idsardi, 1992; Melvold, 1989; Revithiadou, 1999); the following discussion largely borrows from theirs. Given that nouns with fixed stem stress can bear stress on any syllable, these nouns must have stress marked underlyingly (here, as in most of the places in this chapter, I ignore vowel reduction in unstressed syllables):

(57) *Nouns with fixed stem stress have stress marked underlyingly*

a. /m^{j}ín^{j}imum/ [m^{j}ín^{j}imum] ‘minimum’

b. /xarákt^{j}er/ [xarákt^{j}er] ‘character’

c. /inst^{j}itút/ [inst^{j}itút] ‘institute’

When root and suffix are both underlyingly stressed, the root wins out—in constraint-based approaches, this is implemented with a positional constraint like MAX-ROOT-stress that penalizes deletion of stress marks on roots, specifically, outranking the general constraint MAX-stress (Alderete, 1999; Gouskova, 2010; Revithiadou, 1999). Thus, nouns with fixed stem stress can combine with either stressed or unstressed suffixes without ceding their stress.

6.2.2.2 Mobile stress

Most analyses, with the exception of Alderete (1999), agree that nouns showing mobile stress patterns lack lexical stress. For these words, stress location depends on the stress properties of the suffix: if the suffix is underlyingly stressed, it will bear stress; if not, there is no underlying stress anywhere in the word and stress defaults to the first syllable. Nouns like [kok^{j}íl^{j}], which have mobile stress but have stress on the last syllable of the stem, require additional mechanisms (e.g.

Melvold, 1989, pp. 23–26). In constraint-based analyses, this is a straightforward consequence of an alignment constraint, ALIGN-L(PWd, Head), which requires the head (stressed) syllable to align with the left edge of the prosodic word; however, this constraint is outranked by the various MAX-stress constraints and only applies when there is no underlying stress.

Both of the mobile stress patterns feature stem stress in the singular, meaning that the singular case suffixes are unstressed (here I am not worried about the particulars of the case and number features):

(58) *Rules of realization for singular class I case suffixes*

a. [NOM, SG] ↔ Ø

b. [DAT, SG] ↔ u

c. [INS, SG] ↔ om

In both mobile stress patterns, the oblique plural cases (i.e. not nominative or accusative; here, dative and instrumental) are always stressed. Thus, these suffixes are underlyingly stressed. In the nominative plural, we have three different suffixes. Nouns like [kolokol] 'bell' and [nos] 'nose', which have suffix stress in the nominative plural, select for stressed plural suffixes: *-á* and *-í* (often realized as [ɨ́]), respectively. On the other hand, the small group of nouns like [vólos], which have stem stress in the nominative plural, select for an unstressed *-i*. As elsewhere in this dissertation, I assume that a noun's selectional properties are indexed by a diacritic feature referenced in the context of a rule of realization. Because *-í* is by far the most common plural suffix for class I mobile stress nouns, I assume that it serves as the unmarked default (for class I nouns); alternately, all nouns could be marked, and the majority of nouns would carry an [i] feature:

(59) *Rules of realization for plural class I case suffixes*

a. [NOM, PL] ↔ á / [a] ___

b. [NOM, PL] ↔ i / [i_unstr] ___

c. [NOM, PL] ↔ í

d. [DAT, PL] ↔ ám

e. [INS, PL] ↔ ám^{j}i

The rules in (59) pair up with the underlying representations in (60):

(60) *Nouns with mobile stress have no underlying stress*

a. /kolokol$_{[a]}$/ [kólokol] 'bell'

b. /volos$_{[i_unstr]}$/ [vólos] 'hair'

c. /nos/ [nós] 'nose'

For nouns with fixed stem (or suffix) stress, whether or not suffixes bear lexical stress is irrelevant. Thus, we can assume that the nouns with fixed stem stress in (57) also select for stressed plural suffixes by the rules in (59c), (59d), and (59e).

6.2.2.3 Suffix stress

For nouns with fixed suffix stress, like those with fixed stem stress, the stress properties of the suffix are irrelevant: stress is fully determined by the properties of the root. The literature is divided on exactly what the special root properties are. For Idsardi (1992) and Revithiadou (1999), for example, noun stems with fixed suffix stress are equipped with an underlying floating accent that docks, when possible, to the first syllable after the stem (that is, the first syllable of the suffix). Revithiadou (1999) attributes this stress placement to a constraint, *DOMAIN, which pushes floating accents to be realized on morphemes other than the ones with which they are associated. For these analyses, suffix stress is marked *phonologically*, with an accent mark that is not underlyingly associated with a foot or syllable in the root. I indicate this with a stress mark following the root, as in /karandaș´/ 'pencil'.

Alternately, nouns with suffix stress may not have phonological stress marking, but instead be indexed to a rule or constraint that determines their stress position. For example, Halle (1973) posits an Oxytone rule that inserts stress on suffixes following lexcially marked words, while Gouskova (2010) argues that suffix stress nouns are lexically indexed to a higher-ranked copy of a constraint ALIGN-R(PWd, Head), which prefers stress to be on the rightmost syllable (extra mechanisms are required to place stress on the first syllable of the instrumental plural suffix *-am^{j}i*). Under this analysis, suffix stress is marked purely morphologically, with a lexical diacritic feature, which I call [suff]: /karandaʂ$_{[suff]}$/.

Other analyses posit that nouns with suffix stress are marked both phonologically and morphologically, or neither. For Melvold (1989, p. 22), noun stems with suffix stress have underlying final stress; what distinguishes them from nouns with fixed stress on the last syllable of them is that they are lexically indexed to a Rule of Post-Accentuation that shifts stress one syllable to the right. This account works well in a rule-based framework, but is difficult to transfer into a modern constraint-based framework like Optimality Theory: while one could imagine an indexed constraint that could conceivably shift an underlying stress mark one syllable to the right, such a constraint would likely also map other underlying stress patterns to the same location. In this case, the underlying stress would be superfluous and nouns with suffix stress could be unmarked, as in the analysis of Gouskova (2010). On the other hand, for Alderete (1999), suffix stress is the unmarked default. However, his analysis makes use of output–output correspondence constraints (see Section 2.3.2.1) to account for mobile stress; such constraints are incompatible with my general program in this dissertation of keeping correlations between two related forms in the sublexical grammars and out of hard-coded derivational grammars. Thus, I do not consider either of these accounts further.

6.2.2.4 Dominant stress

The stress patterns of inflectional paradigms are quite complicated, but there is a further complication: some *derivational* suffixes (known as dominant suffixes) overwrite the stress pattern of the stems to which they attach. One such suffix is diminutive *-ók*. Like the plural suffixes in (59), this suffix is stressed in the nominative singular. However, unlike the inflectional suffixes, *-ók* imposes the same stress pattern on nouns to which it attaches regardless of their stress pattern, as shown in (61). In fact, *-ók* shows a vowel–zero alternation and fixed suffix stress—the plural of [mʲesʲats-ók] 'month (dim.)' is [mʲesʲats-kʲ-í]. For the purposes of illustration, I assume the floating stress analysis of fixed suffix stress. Melvold (1989) shows that suffixes must be lexically marked as dominant or not, for which she uses a feature [±dom(inant)]. Putting these together, the underlying form of *-ók* is /ok´$_{[+\text{dom}]}$/.

(61) *Dominant suffix -ók overwrites lexical stress*

a.	*stem stress:*	/mʲésʲats-ok´$_{[+\text{dom}]}$/	[mʲesʲatsók]	'month (dim.)'
b.	*suffix stress:*	/jazɨk´-ok´$_{[+\text{dom}]}$/	[jazɨtʃʲók]	'tongue (dim.), uvula'
c.	*mobile stress:*	/volos-ok´$_{[+\text{dom}]}$/	[volosók]	'hair (dim.)'

I do not attempt an analysis of stress dominance, the details of which would require complexity not necessary for this presentation. The main points are that dominance must be lexically marked, and that dominant suffixes behave differently from those that are simply underlyingly stressed, in that the former retain their stress even when attaching to an underlyingly stressed root.

6.2.2.5 The stress patterns of plural -a

As described previously, all class I plural suffix *-a* is obligatorily stressed—in fact, such nouns are obligatorily stressed throughout the plural. The majority have mobile stress (stem stress in the singular), although Zaliznjak (1977) lists two nouns with plural *-a* and fixed suffix stress. Curiously, the two nouns have closely related meanings: [obʂlág] 'cuff' and [rukáv] 'sleeve'.

This section addresses the question of *why* plural -*a* is always stressed. One possibility, assumed by Alderete (1999), is that -*a* is dominant like diminutive -*ók*. If this were the case, however, we would expect to see nouns that *only* stress -*a* and no other suffix: this would be the expected result of dominant -*a* attaching to an underlyingly stressed noun. Since we do not see such cases, it must instead be the case that plural -*a* never appears with underlyingly stressed nouns. There is nothing in the analysis of stress presented in this section that captures this: it must be a generalization that is not hard-coded into the grammar in the sense discussed in Section 2.3. That is, Russian speakers have learned a generalization that noun roots with the [a] feature, like /$\text{kolokol}_{[a]}$/ 'bell' ((60a)), do not have non-floating stress marks. Without committing to an analysis of how stress and prosodic information are stored, I write this as a constraint penalizing stressed syllable nuclei, presented in (62).

(62) *Constraint against underlying non-floating stress marks in the [a] sublexical grammar*
*[+stress, +syll]

Like the constraints described in Section 3.1.2, this is a phonological constraint in a sublexical grammar—in this case, the sublexical grammar for the [a] feature (which is borne by nouns with plural -*a*). That is, this constraint is not active in the course of regular phonological derivations, and individual forms may obey or violate it. It becomes active when a speaker encounters a novel word and uses the sublexical grammars to determine its behavior: since the constraint in (62) is in the [a] sublexical grammar (and likely heavily weighted), the speaker is (much) less likely to assign the [a] feature to words with underlying stress marks, and is thus less likely to form the plural of this word with -*a*. The other constraints described in this section are likewise constraints in sublexical grammars that are only invoked when speakers need to productively extend lexical patterns to determine a word's behavior when it is not already listed. In addition, the constraint in (62) crucially refers to *underlying* stress marks. This means that the sublexical grammar containing it must have access to underlying representations; that is, it is a *morpheme structure constraint* (see

Booij, 2011).

There is similarly a strong—but, in this case, not categorical—generalization against words with suffix stress. If suffix stress is analyzed with a diacritic feature, this is a morphological dependency: a constraint *[suff] in the [a] sublexical grammar. If suffix stress is analyzed as floating stress, then the [a] sublexical grammar can have a constraint against *all* stress markers, floating or not: *[+stress]. This will compound with the constraint against non-floating stress markers that explains the non-occurrence of *-a* with nouns with stem stress.

In the nonce word study in Section 6.4, all stimuli are disyllabic with stress on the last syllable in the base form. As explained above, final stress in the base form is typical of suffix stress, while nouns with mobile stress generally have initial stress in unsuffixed forms. Thus, suffix stress should lead speakers to assume that the stimuli have fixed suffix stress rather than mobile stress. This is uncommon for nouns with plural *-a*, which I include in my nonce word study, but not unattested: in all conditions, the stimuli follow attested, if not particularly common, patterns.

6.2.2.6 The -ók sublexical grammar

The main goal of this chapter is to study the influence of stress pattern and plural suffix on diminutive choice—in particular, the fact that *-ók* is more common with nouns that have suffix and mobile stress and preferred by nouns that take plural *-a* even beyond what would be expected given that *-a* itself is always stressed. The correlation between *-a* and *-ók* can be encoded as a constraint in the sublexical grammars for the other two diminutive features, [ik] and [chik], penalizing nouns with the [a] feature: *[a]. As I show in Table 6.11, nouns with mobile stress show an even stronger preference for *-ók* than nouns with suffix stress. This is neatly captured under the floating accent analysis of suffix stress: nouns with mobile stress are unpenalized in the [ok] sublexical grammar; nouns with suffix stress are penalized by a moderate constraint against lexical stress, *[+stress]; and nouns with stem stress are penalized by *both* *[+stress] and the more specific constraint against non-floating stress, *[+stress, +syll]. In this system, the correlation between *-a* and *-ók* is a morphological dependency, because plural *-a* is indexed with a diacritic feature, [a]. However, strictly speaking, the correlation between stress pattern and *-ók* is not: even though a noun's stress pattern is a fact about its inflectional paradigm, it is encoded through abstract *phonological* structure—see Section 2.2.2.4.

Before moving on from this section, I wish to emphasize that the *[a] constraints really are needed. At first, it may seem that we can do without them. After all, the sublexical grammars for the [a] and [ok] features are quite similar: both have a strong constraint *[+stress, +syll] and a weaker constraint *[+stress]. While these similarities may reinforce the relationship between [a] and [ok], they cannot explain the correlation without constraints directly referencing the [a] feature: words with plural *-a* take diminutive *-ók* more than would be expected from the stress constraints alone. Moreover, in my nonce word study, participants assign *-ók* more often to stimuli shown with plural *-a*, even though the setup leads them towards a conclusion that these nonce words have suffix stress (dispreferred by *-a*) rather than mobile stress (preferred by *-a*). This result requires an active correlation between *-a* and *-ók* that cannot be reduced to factors of stress.

This concludes the basic analysis of the inflectional factors influencing diminutive *-ók*, which lays the necessary groundwork for explaining the results of the nonce word study in Section 6.4: speakers have learned the morphological dependency between plural *-a* and diminutive *-ók*, as well as the dependency between suffix stress and diminutive *-ók*, which is arguably phonologically encoded. In Section 6.2.3, I address an issue of theoretical importance: words that can take more than one diminutive. However, this discussion has no direct relevance for the empirical portion of this chapter.

6.2.3 Words that allow multiple diminutives

6.2.3.1 Should we model variability with hybrid classes?

In the Hungarian study (Chapter 4), I ignored variable lexical items entirely, whereas in Czech (Chapter 5), variable items play a substantial theoretical role in my theoretical analysis. Variability is not central to the Russian case, and is not included in my study design: I model the nonce word task as a ternary choice assigning a word one of three categorical features. However, the model of variability I described in Section 3.4 handles variation under the assumption of a single binary feature. In this section, I describe how the model handles variation among multiple privative features. This accounts for variability in the diminutive for existing words, assuming that individual speakers sometimes do produce two diminutives for the same word.

Guzmán Naranjo (2019) argues that words with two acceptable diminutives fall into hybrid classes that inherit properties from their parent class: so, for example, nouns that take both -^{j}*ik* and *-ók* (by far the most common pairing of diminutives) belong to a mixed *ik∼ok* class that has all the properties of the *ik* class and the *ok* class, and can also have unique properties of its own. Thus, when speakers are forming generalizations over classes, nouns in the hybrid *ik∼ok* class should count towards generalizations over three categories: those derived over words that take *-ók*, those derived over words that take -^{j}*ik*, and those derived over the words that take both, which may show additional phonological (or other) tendencies. He shows that a classifier model is fairly successful at distinguishing *chik* words from *ik* and *ok* words, and *ik∼ok* words from the latter two, although *ik∼ok* words are often instead placed in the non-hybrid *ik* or *ok* classes. This, in turn, explains the relative frequency of this variable class: "The fact that we have more *ik∼ok* nouns than *chik∼ik* or *chik∼ok* nouns is due to the constraints for *IK* and *OK* being more compatible with each other, and producing a more relaxed set of constraints that *CHIK∼IK* or *CHIK∼OK*" (Guzmán Naranjo, 2019, p. 115). That is, this hybrid class is larger than the others because its two parent classes follow more similar generalizations.

Guzmán Naranjo (2019) argues that his results favor a hierarchical structure of hybrid classes and disfavor the model assumed by Gouskova et al. (2015) and by me, in which words can be placed into multiple classes, but no distinct hybrid class is formed. However, the results described above are compatible with both models: if nouns allowing multiple diminutives were simply in *both* sublexicons, containing both the [ik] and [ok] features, we would expect a similar conflation of predictors in the two classes. Instead, as Guzmán Naranjo and Bonami (2021) explain, a positive argument for hybrid classes requires that such classes hold some feature or generalization that is *not* inherited from the two parent classes.

6.2.3.2 Modelling variability with weighted features

Another problem—shared by both the hybrid class account and the multiple sublexicons account—is that they are imprecise. Neither of these accounts of variation allows individual words to have different distributions, but there are clear differences: for example, my data set includes 14 diminutive tokens for [borʃʲː] 'borscht', 9 with *-ók* and 5 with *-ʲik*—a relatively even balance. However, for [dom] 'house', the diminutive with *-ʲik* is much more common (9436 tokens), while *-ók* is used only occasionally (37 tokens). Thus, we need some additional complexity in our model of storage to allow speakers to represent the true distribution of diminutives.

The model of lexical variation I described in Section 3.4 accounts for arbitrary variable distributions through variable features that have a lexical weight *b* corresponding to their likelihood of apppearing with that feature. Thus, if [borʃʲː] appears about twice as often with *-ók* than with *-ʲik*, we could say that its underlying representation has both [ok] and [ik], but the [ok] feature with a higher weight than the [ik] feature: for the sake of demonstration, we can assign these weights as 2 and 1, yielding an underlying form which we can represent as /borʃʲː$_{b_{[\text{ok}]}=2, b_{[\text{ik}]}=1}$/. Suppose we wish to form the diminutive of this word. We spell it out using one of the rules in (56); which rule we choose (and thus which suffix) depends on the feature set of the noun in its context. Each variable feature causes a split in the derivation into two candidate derivations, one in which the feature is inserted and the other in which it is not; the feature's weight is assigned to the derivation in which the feature is inserted, while the derivation without the weight does not change its score. After the derivations are complete, one is chosen stochastically, such that higher-scoring derivations are more likely to be chosen.

In the case of [borʃʲː], we have two variable features, [ik] and [ok]. This yields a four-way split in the derivation: one in which [ik] is inserted, one with [ok], one with both, and one with neither. This derivation is shown in Table 6.9 (this example is slightly simplified compared to the model in Section 3.4: I omit the features' *a* parameters, since they are not relevant to the current example). Derivations 2 and 3 are well-behaved: in each, only one feature is present, so the competition between the rules in (56) is resolved unambiguously. Derivations 1 and 4, run into problems: when both features are present (Derivation 1), rules (56a) and (56b) are equally applicable according to standard assumptions of Distributed Morphology, and there is no way to choose without either an extrinsic ordering or random choice—which, by hypothesis, I do not allow in rule selection, only in feature insertion. On the other hand, when neither feature is present (Derivation 4), none of the rules in (56) have their conditions satisfied, and with no default, nothing can be inserted. In both cases, the derivation *crashes*; I assume that such derivations are eliminated from contention. The remaining competition is between the valid derivations, and of those, the *-ók* suffix gets chosen about 73% of the time.

output of syntax:	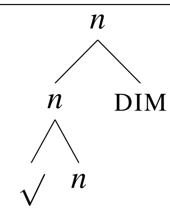			
noun:	borʃʲː			
[ok] parameter:	$b=2$			
[ik] parameter:	$b=1$			
	Derivation 1 [ok] present [ik] present	Derivation 2 [ok] present [ik] absent	Derivation 3 [ok] absent [ik] present	Derivation 4 [ok] absent [ik] absent
rule(s) of realization:	(56a), (56b)	(56a)	(56b)	—
score s_1:	3	2	1	0
output:	—	borʃʲː-ók	bórʃʲː-ik	—
final score s:	—	2	1	—
probability:	0	$\frac{e^2}{e^2+e^1} \approx .731$	$\frac{e^1}{e^2+e^1} \approx .269$	0

Table 6.9: Derivation process for the diminutive of [borʃʲː] ‘borscht’

If a noun does not have any diminutive feature in its lexical entry, none of the rules in (56) will apply and the derivation will crash. This is plausibly the situation for words that have no diminutive, like [son] ‘sleep, dream’ (Gouskova et al., 2015; Polivanova, 2008 [1967]).

As mentioned above, the model described in Section 3.4 and exemplified in Table 6.9 was designed for binary variable oppositions, and when used on privative features, it regularly generates derivational variants that crash and are removed from contention. One alternative within the sublexicon model is to allow lexical entries to belong fractionally to different sublexicons: for example, allowing [borʃʲː] to have 73% of the [ok] feature and 27% of the [ik] feature. Then, each derivation inserting the word requires a weighted random choice to attach either [ok] or [ik]. This approach to variation, which I call the coin flip model, is essentially that of Allen and Becker (2015) and Becker and Gouskova (2016), although these works do not include an articulated theory of morphosyntax; instead, sublexicons are directly associated with morphological processes. My model is more explicit about its derivational processes: I associate sublexicons with features in underlying representations; the morphological processes in this case (which are concatenative) are triggered by rules of realization that have the sublexicon features in their context.

For the case at hand, the two analyses make the same prediction: weighted lexical features on a word's lexical entry control the surface distribution of its diminutive variants. The main architectural difference is that the one shown in Table 6.9 involves *splitting the derivation*, while that described by Allen and Becker (2015) and Becker and Gouskova (2016) involves a weighted coin flip within *a single derivation*. For convenience, we can call these the single derivation and split derivation models. The single derivation model is simpler, in that it does not require the grammar to consider multiple derivational options in parallel. As explained in Section 3.4, though, the extra complexity is necessary to account for cases like Czech where lexical and syntactic variation interact. I refer the reader to that section for the full argument, but the gist is as follows: if syntactic variation is handled strictly before lexical variation (as required by theoretical concerns of modularity), then lexical variation can only operate within the bounds set by the coin flip from syntactic variation. For example, if a given preposition takes -ε 30% of the time on aggregate, even a very high lexical rate of -ε can only reach this 30%. However, a split derivation model provides an end run around modularity in this specific case: lexical variation can override the effects of syntactic variation and push it to the extreme end (see Section 3.4.3.3 for discussion). This is in fact the pattern we see in Czech. If further research provides evidence for the semantic (and syntactic) distinction argued by Magomedova (2017), described in Section 6.1.3, then we would need a split derivation model to account for semantic conditioning of allomorphy in Russian as well.

While variable items are important for any full account of Russian diminutive allomorphy, they are a relatively minor corner of nouns. In the following corpus and nonce word studies, I set aside the issue of item variability and focus on the phonological and morphological factors predicting whether a given word is attested with each diminutive.

6.3 Russian synchronic corpus study

As with the Hungarian corpus study in Section 4.3, this corpus study has two goals. The primary goal is to determine the phonological and morphological patterns fo diminutive allomorphy that speakers are expected to have learned from their lexicon. In particular, the model in Table 6.17 is used as my phonological representation of the lexicon, approximating speakers' sublexical grammar, in the nonce word study in Section 6.4. This model compares words that take diminutive *-ók* to words that take one of the other two diminutives, matching my nonce word study, which is a forced choice task between the three diminutives. Similarly, I test the strength of the predictive effect of stress pattern and plural *-a* on diminutive *-ók*. Although there is a clear overlap between plural *-a* and diminutive *-ók* in the lexicon, much of this is explained by a general preference for *-ók* to take nouns that have stressed suffixes, since plural *-a* always bears stress. Once lexical stress pattern is taken into account, the correlation between *-a* and *-ók* is positive but not significantly so. I suggest that Russian learners should respond to this pattern in different ways: some speakers should learn and encode a correlation between *-ók* and plural *-a* alongside the correlation with stress pattern, while others should only learn the stress correlation.

A secondary goal of this study is to confirm the results of previous descriptions and corpus studies of phonological factors influencing the Russian diminutive, in particular Gouskova et al. (2015). Most of the key results are the same, although there are some differences, which I attribute to a difference in corpora and the more complete view of the lexicon gained by my inclusion of morphological factors like plural suffix.

6.3.1 Data

6.3.1.1 Corpus construction

My corpus comprises class I inanimate nouns according to the inflectional data from Zaliznjak (1977). I searched the Russian National Corpus for possible diminutives for each word in texts dating from 1950 or later (the cutoff in the corpus for modern written texts), as well as the base (non-diminutive) lemma. Diminutives were constructed by attaching each of the three suffixes to base forms, either unmodified or with morphophonological alternations. Different suffixes trigger different alternations in their final consonant, as shown in Table 6.10 (Gouskova et al., 2015, p. 49). In addition, *-^{j}ik* generally triggers vowel–zero alternations for words that undergo them, except for monosyllables: the diminutive of [kov^{j}ór] 'carpet' is [kóvr^{j}ik], but the diminutive of [rot] 'mouth' (genitive [rta]) is [rót^{j}ik], not *[rt^{j}ik]. I searched for forms both with and without the expected alternations: diminutives that were expected to alter their bases usually did. A few forms showed variation—for example, [xr^{j}eb^{j}et] 'spine, ridge', which has an alternating vowel in its inflectional paradigm, appeared 35 times with diminutive *-^{j}ik*, 7 with the vowel ([xr^{j}eb^{j}et^{j}ik]) and 28 without ([xr^{j}ebt^{j}ik]).

diminutive suffix	*C alternation*	example	*triggers V–Ø alternation?*
tʃ^{j}ik	l∼l^j	kolokol ∼ kolokol^{j}tʃ^{j}ik 'bell'	no
^{j}ik	ts∼tʃ^j, C∼C^j	pal^{j}ets ∼ pal^{j}tʃ^{j}ik 'finger'	yes
ók	k∼tʃ^j, g∼ʐ, x∼ʂ	kr^{j}uk ∼ kr^{j}utʃ^{j}ók 'hook'	no

Table 6.10: Base alternations triggered by Russian diminutive suffixes

My corpus is similar to that of Gouskova et al. (2015), although they looked at diminutive forms in the Google Ngrams corpus. In comparison, the Russian National Corpus is much larger and cleaner, and is also lemmatized, allowing me to search for diminutive lemmas in all inflected forms without having to manually construct and search for multiple inflected forms as Gouskova et al. (2015) did. Thus, the results should be more accurate than those from Google Ngrams alone.

There are two main sources of bad data in my corpus. The first is that the Russian National Corpus sometimes mistags forms, so I manually filtered out tokens where the endings did not match the expected case endings. The second issue is the existence of "pseudo-diminutive" forms, which look like diminutives but are not. These are predominantly suffixes homophonous with the diminutives. For example, as Guzmán Naranjo (2019, p. 113) notes, the suffix *-ʲik* in [alkogolʲ-ik] 'alcoholic' is not a diminutive, and *-tʃʲik* can also form occupational nouns like [pulʲemʲot-tʃʲik] 'machine gun operator', from [pulʲemʲot] 'machine gun' (Gouskova et al., 2015, p. 50). Both of these suffixes typically form animate nouns, so they do not show up in my search, which is limited to nouns marked in the corpus as inanimate. Accordingly, most words with these suffixes are legitimate inanimate diminutives.[6]

Another set of "pseudo-diminutives" source of false positives is *-ok*, which has many meanings, some of which are stressed and some not (though the variants are orthographically identical, since Russian typically does not mark stress). Stressed *-ók*, is usually the desired diminutive. Unstressed *-ok* (reduced to [ək]) can *also* be a diminutive form (as described in Section 6.1.2), though it has different properties than the stressed *-ók* that is the subject of this study. These are mostly filtered out indirectly, as I describe below. Another common use of the suffix does require direct intervention: unstressed *-ok* can attach to a verb stem to form a noun representing the result of that verb's action (Švedova et al., 1980, p. 148). For example, the verb [obrʲez-at] 'trim' has a derived noun [obrʲézok] 'trimming' (usually plural, as in English). However, this verb stem also has a zero-derived noun [obrʲez] 'edge'. To ensure that [obrʲezok] does not get classified as an *-ók* diminutive of [obrʲez], I removed 54 potential *-ók* diminutives that are listed in Zaliznjak (1977)

[6]Gouskova et al. (2015) also mention pairs like [lʲingv^{i}st] 'linguist' and [lʲingv^{i}stʲikə]. These "pseudo-diminutives" are unproblematic for me for two reasons: first of all, the base of these forms is usually animate. Second, because the Russian National Corpus is lemmatized, I can limit my search to nouns that actually end in *-ʲik*, not those that have an additional feminine nominative suffix (*-a* [ə]) in the citation form.

with non-final stress. This is not foolproof, since there may be nouns in the corpus ending in unstressed *-ok* that are not listed in this dictionary, but it seems to be fairly effective.

Once these pseudo-diminutives are removed, there are still some forms that look like diminutives but are not. I did not want to remove false positives manually, as the data set was too large (and some edge cases too unclear) to comb by hand, and removing individual words that I happened to catch would introduce bias into the data set. Instead, I aimed to minimize false positives by programmatically removing cases that contained false diminutives but few if any true ones. The one such filter I implemented is the suffix *-enʲ*, as in the metallurgical term [plavʲenʲ] 'flux'. This suffix undergoes a vowel–zero alternation (genitive [plavnʲa]), so if the diminutive *-ʲik* is attached to the stem with a zero, the result would be [plavnʲik]. This, however, is homophonous with a noun-forming suffix *-nʲik*—in this case, the word means 'fin'. Although [plavʲenʲ] and [plavnʲik] are both derived from a verb root [plav], the latter is not the diminutive of the former. Accordingly, I removed forms created by adding *-ʲik* to the zero alternants of nouns ending in [enʲ].

Another issue is that diminutives of two words can be identical. For example, [palʲtʃʲik] is the diminutive formed by attaching *-ʲik* to [palʲets] 'finger' with a vowel–zero alternation and a consonant alternation (see Table 6.10), but it could also be the result of attaching *-tʃʲik* to [pal] 'burning' (with palatalization of [l], see Table 6.10 above). There is no way to easily identify which base a given diminutive "belongs" to without manual inspection, so I assign diminutives to the possible base noun that appears most frequently. Upon inspection, this seems to generally work well: for example, [pal] appears 64 times in the corpus, while [palʲets] has 40,899 tokens, so this heuristic correctly assigns [palʲtʃʲik] as the diminutive of [palʲets].

Similarly, homophones in Russian can have different declension patterns. For example, [ton] has two meanings corresponding to the English 'tone' ('tone of color' and 'musical tone') and are listed by Zaliznjak (1977) with three different patterns: the nominative plural is [toná] in the color meaning and [tónɨ] in the musical meaning; the latter meaning, in addition, has variable stress patterns in oblique plural cases, so the genitive plural can be [tónov] (showing fixed stem stress throughout the paradigm) or [tonóv] (showing mobile stress). Since stress pattern and plural suffix are important predictors of diminutive realization in my corpus study, I removed words with multiple possible inflection patterns and lexical stresses. Some of the words removed were quite frequent, including [rod] 'clan, generation' (plural [rodɨ́]), 'grammatical gender' (plural [ródɨ]), 'type' (plural [rodá]).

Next, I removed very infrequent bases (whose non-diminutive forms had less than five tokens in the corpus) and bases that were themselves diminutives, even though some double diminutives are attested. For example, the dictionary includes an entry for [lʲesók] 'grove', itself a diminutive of [lʲes] 'forest'. This, in turn, has the diminutive [lʲesótʃʲək], which is found by my search. One reason to discard such forms is that my theory of lexical exceptionality assumes that a stem's selectional properties are derived from its rightmost morpheme (see Section 4.3.1 for discussion). Accordingly, complex stems, including diminutives, should not be counted as separate types. My database does not mark derivational suffixes or compounds, so I do not make any attempt to remove complex stems from my corpus; however, removing double diminutives is an easy way to eliminate at least some derived forms that, strictly speaking, should not be counted. Another reason to exclude double diminutives is that, in the double diminutive suffix *-ótʃʲək*, the second suffix is not the stress-bearing *-ók* but the unstressed *-ok* described earlier in this section. Thus, eliminating these double diminutives avoids further false positives.

Finally, I removed the few words whose stress pattern information is not listed in the digitized version of Zaliznjak (1977). After all of these filters have been applied, my corpus has 8,178 base lemmas, of which 1,250 (15.3%) are attested in at least one diminutive form.

6.3.1.2 Descriptive summary

The distribution of the 1,250 nouns attested with at least one diminutive suffix is shown in Figure 6.1.

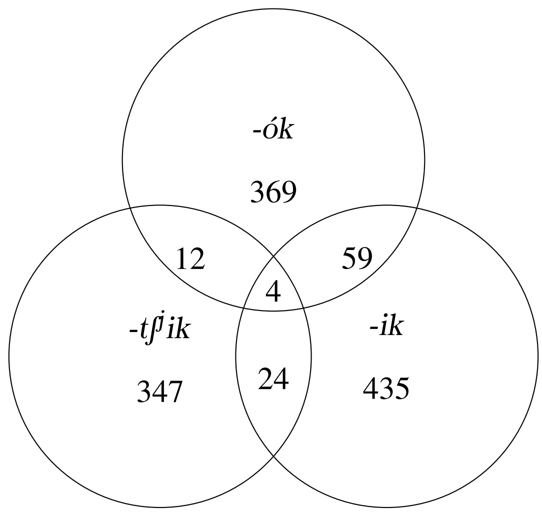

Figure 6.1: Distribution of Russian nouns among the three diminutives

The majority of nouns do not appear in any diminutive, and more frequent words are much more likely to have attested diminutives than less frequent words. Of the 1,515 nouns that have at least 1,000 tokens in my corpus, 666 (44.0%) have at least one diminutive.

Among nouns that have at least one diminutive, there is a slight frequency bias in choice of diminutive. This is shown in Table 6.11: while *-tʃʲik* is relatively steady across frequencies, more frequent words (that is, lemmas that are attested with more tokens, in any inflected form) show a slight preference for *-ʲik* and a slight dispreference for *-ók*. Given that *-ók* is older than *-ʲik* and both are fairly common with native words, I do not have a good explanation for this effect. The percentages in each row of Table 6.11 add up to more than 100% because nouns attested with multiple diminutives are double-counted; that is, nouns that take both *-ʲik* and *-ók* (by far the most common pairing) are included in the *multiple* column and both the % *-ʲik* and % *-ók* columns.

lemma frequency	*-t$\int^j$ik*	*-^{j}ik*	*-ók*	*multiple*	*% -t$\int^j$ik*	*% -^{j}ik*	*% -ók*
5–50	11	17	25	6	22.0%	37.3%	50.8%
50–500	102	103	125	12	31.3%	33.3%	38.9%
500–5,000	169	203	156	42	33.3%	41.4%	33.0%
5,000–50,000	63	101	58	32	29.1%	52.4%	31.9%
50,000–500,000	2	11	5	7	12.0%	68.0%	48.0%

Table 6.11: Distribution of diminutive allomorphs by token frequency (percentages double-count nouns that allow multiple diminutives)

The phonological generalizations over the Russian diminutives discussed in the literature can be found in Section 6.1.2, although my results do not entirely line up with previously reported ones (see the discussion in Section 6.3.4). In the remainder of this section, I present the generalizations over lexically specified effects that cannot be determined from the phonological form of the base noun alone: stress pattern, plural suffix, and vowel–zero alternations.

I code for three types of stress patterns: fixed *stem* stress; fixed *suffix* stress, in which the stress is on the suffix if there is one and reverts to the stem in unsuffixed forms; and one of a number of *mobile* stress patterns, in which some inflected forms have stress on the stem and some on the suffix. (See Section 6.2.2 for a more detailed account of Russian inflectional stress patterns.) As shown in Table 6.12, the majority of words have fixed stem stress, but of those with suffix or mobile stress, a disproportionate number are attested with diminutives. This is somewhat of a frequency effect: more frequent words with stem stress are more likely to be attested with diminutives. However, this is not the whole explanation: even very infrequent words with suffix or mobile stress are almost always attested with diminutives, and even at the highest frequency, nouns with stem stress are less likely to have diminutives. One possibility is that stem stress, being by far the most frequent pattern, attracts loan words, which are newer and thus less likely to have established diminutives. This would also explain the fact that non-stem stress is quite rare with *-t$\int^j$ik*, a diminutive used more often for newer words and loans (which are also more likely to have default stem stress). However, this explanation requires further study to confirm. Nouns with non-stem (suffix or mobile) stress have an especially strong preference for *-ók*: of the 832 nouns with non-stem stress, 381 of which take at least one diminutive, 281 are attested with *-ók* (including 53 in the *multiple* column).

stress pattern	*none*	*-tʃʲik*	*-ʲik*	*-ók*	*multiple*
stem	6477	339	345	141	44
suffix	384	8	78	153	32
mobile	67	0	12	75	23

Table 6.12: Distribution of diminutive allomorphs by stress pattern

The strong preference for *-ók* among nouns with at least some stressed suffixes is likely related to the stress patterns of the diminutives themselves: *-ók* attracts stress, while the other two diminutives do not. That is, the stressed diminutive suffix attaches more readily to words that appear with other stressed suffixes—a paradigm uniformity effect (see Section 2.3.2.1).

As shown in Section 6.2.2.5, lexically determined stress patterns interact with a segmental component of inflection class: a small number of class I nouns, like [górod] 'city' (the inflection class studied in this chapter) take *-a* as the nominative and accusative plural suffix (plural [gorod-á]), where most nouns take *-i* (sometimes realized as [ɨ])—for example, the plural of [motór] 'motor' is [motór-ɨ]. While the plural suffix *-i* can be stressed or unstressed, *-a* is always stressed when it occurs with these nouns. (This suffix appears much more commonly with class IV neuter nouns, in which case it can be stressed or unstressed.) The distribution of diminutives and plural suffixes, shown in Table 6.13, is even more skewed than that of stress patterns: of the 86 nouns with plural *-a*, nearly half take *-ók*, and very few appear with any other diminutives. This table omits eight nouns that take *-ja* as a plural, described in Section 6.1; five of these nouns appear with *-ók*, one with *-ʲik*, and the remaining two with both *-ók* and *-ʲik*.

plural	*none*	*-tʃʲik*	*-ʲik*	*-ók*	*multiple*
-i	6892	346	431	325	90
-a	36	1	3	39	7

Table 6.13: Distribution of diminutive allomorphs by plural suffix

The preference for plural *-a* nouns to take *-ók* may be rooted in the requirement for stress shared by the two suffixes. However, the preference for *-ók* among nouns that take plural *-a* is even stronger than that of nouns with suffix and mobile stress in general: including multiples, 46 nouns with plural *-a* are attested with *-ók*, which is 53.5% of all plural *-a* nouns and 92% of plural *-a* nouns that appear with at least one diminutive. By contrast, 281 nouns with at least some stressed suffixes take *-ók*, which is 33.7% of all nouns with non-stem stress and 73.8% of nouns with non-stem stress that take at least one diminutive. Thus, plural *-a* nouns represent the extreme of diminutive selection for nouns with non–fixed stem stress patterns.

I include one more morphological factor marking whether the last vowel of a noun's stem disappears when attached to inflectional suffixes. One stem with a vowel–zero alternation is [úgol] 'corner', whose genitive is [uglá]. This vowel–zero alternation must be represented abstractly, either through defective or otherwise marked phonemes or with a lexical diacritic feature (see Section 2.2.1 for empirical and theoretical discussion). Gouskova (2012) concludes that alternating morphemes should be marked with a diacritic feature and that non-alternation is the default. In addition, as described earlier in Section 6.3.1.1, *-ʲik* usually triggers vowel–zero alternations for polysyllabic stems that alternate in their inflectional paradigm, while the other two suffixes do not. From this, we might expect longer alternating nouns to appear more often with *-ʲik* (so they can express their alternation—a paradigm uniformity effect) and less often with *-tʃʲik*, which tends to appear with newer words that would not be grammaticalized with an alternation. The bias is quite slight, but alternating stems do appear quite rarely with *-tʃʲik* (only 5 of 530 alternating nouns). In fact, there is a slight preference for alternating nouns to take *-ók*: there are 43 such nouns (8.1%), compared to 29 that allow *-ʲik* (5.5%). However, compared to the other morphological effects, these effects are really quite small: altogether, 87.2% of alternating nouns are unattested in any diminutive, slightly more than the 84.5% of non-alternating nouns.

6.3.2 Methods and analysis

Given that each noun may take any or all of three possible diminutive suffixes, I fitted three series of regressions, one for each suffix (*-ók*, *-ʲik*, and *-tʃʲik*). Each has a binary dependent variable representing whether a given noun is attested with that suffix.

Similar to previous corpus studies in Section 4.3 and Section 5.3, I fitted three regressions in each series. The first includes a collection of phonological predictors: place and manner of the word's final consonant, height and rounding of the vowel in its last syllable, complexity of its final coda, whether it is monosyllabic, and whether its stress in the nominative singular is final. (Monosyllabic words are marked as having final stress.) These predictors are *source-oriented* (see Section 2.2.2.1 and Section 2.3.1.2): they refer to properties of the unaffixed base form, not the stem to which the diminutive actually attaches. Most saliently, nouns ending in dorsals, like [krʲuk] 'hook', are counted as having dorsals, even though *-ók* usually causes stem-final dorsals to mutate to coronals ([krʲutʃʲók]). Unlike in Section 4.3, before adding the morphological factor of plural suffix, I included an intermediate regression adding two other paradigmatic factors: whether the noun undergoes a vowel–zero alternation when suffixed, and the noun's stress pattern, divided into three categories: fixed stem, fixed suffix, and mobile. Finally, the third regression for each suffix adds the morphological factor of plural suffix (regular *-i* vs. irregular *-a*, also including the handful of words with irregular plural *-ja* as a third factor). As described in Section 6.1, the irregular plural *-a* always attracts stress onto itself, so I test whether correlations between plural form and diminutive can be attributed to the segmental plural suffix *-a* itself, or whether the unbalanced diminutive distribution of nouns that take plural *-a* can be reduced to the factor of stress pattern. All three regressions include a factor of the logarithm of base frequency to account for the fact that more frequent words are more likely to have attested diminutives. This factor is normalized by subtracting 4.25, which is very close to the mean and median log base frequency of all words that appear at least once in the corpus and to the peak of the normal distribution of base frequencies (truncated at the cutoff of 5 tokens), corresponding to a frequency of $e^{4.25} \approx 70$ tokens.

Each regression was assembled by forward stepwise comparison using the `buildmer` function in R from the package of the same name (R Core Team, 2022; Voeten, 2022). This function adds factors to the model one at a time such that each additional factor improves the model's Akaike Information Criterion (AIC), which measures how well the model fits the data while penalizing model complexity (that is, number of factors). If a given model does not include one of the factors listed above, it is because adding this factor to the model did not substantially improve it.

This study is intended in part as a replication of Gouskova et al. (2015) with a more modern data source, so I will now discuss differences between my analysis and theirs. In my data set, as in theirs, the number of nouns without attested diminutives far exceeds those with at least one diminutive attested. They handled this by randomly selecting a small number of nouns without diminutives to include in their analysis and omitting the rest. (However, they did note that a model fitted on all nouns yielded a qualitatively similar result.) I include all nouns in my data, but add a factor of base lemma frequency to account for the fact that more frequent nouns, regardless of their other properties, are more likely to have diminutives in the corpus.

Gouskova et al. (2015) present their results as a polytomous regression, which is a series of logistic regressions focusing on each outcome, e.g. words that take *-tʃʲik* vs. words that take a different diminutive or none. My analysis differs from theirs in two ways. First, they partition the set into four distinct groups: *-ók*, *-ʲik*, *-tʃʲik*, and none. It is unclear how they handle words that allow for multiple diminutives, like [blʲin] 'pancake', which is attested with both *-tʃʲik* (342 tokens) and

-ók (23 tokens). In my regressions, such words are placed into multiple categories. Relatedly, Gouskova et al. (2015) run a regression predicting each of the four groups, including those that take no diminutive. I only run three regressions predicting whether a word is attested with each of the three diminutives; I do not separately predict whether a word is attested with *no* diminutives.

Gouskova et al. (2015) only look at phonological factors, so their regression corresponds to my first regression, without the morphological predictors of stress pattern, plural, and vowel–zero alternation. However, their phonological factors are slightly different than mine. They also include some word-internal properties as predictors: whether a word has word-medial clusters or vowel hiatus. I focus only on the last syllable of the stem. Some of their factors have fewer levels than mine: they only single out certain values of the final consonant place (dorsal) and manner (glide, nasal, approximant) and last syllable vowel height (high), whereas I code for all possible values (although I count glides as approximants except for [v], which they count as a glide and I count as a fricative). Other than these small differences, our factors are largely the same.

The analysis of Gouskova et al. (2015) found that a number of phonological factors predicted a noun's choice of diminutive—following observations dating back to Polivanova (2008 [1967]). For example, nouns ending in velars are more likely to take *-ók* and less likely to take the others, while words ending in a consonant cluster are more likely to take -^{j}*ik* and less likely to take the others. I compare their results with mine in more detail in Section 6.3.4.

Finally, I also include a regression predicting a word's ability to take *-ók* from phonological characteristics, limiting my data set to nouns that appear with at least one diminutive (and excluding all those that lack diminutives entirely). This model best corresponds to my nonce word study, in which participants are forced to choose one of the three possible diminutive forms, and I use its predictions as the *phon_odds* factor representing the phonological patterns in the lexicon in my analysis of the nonce word study.

6.3.3 Results

I present the results of the Russian corpus study in a slightly different way from the others. Instead of showing every model, I only include the most complex one for each diminutive. I mention notable deviations from the simpler phonological models, but do not present the latter in full.

6.3.3.1 -tʃʲik

Table 6.14 shows the model with phonological and morphological factors predicting whether a word takes *-tʃʲik* as a diminutive, with factors listed in the order in which they were added to the model (roughly corresponding to their importance). Of the morphological factors, only stress pattern was added to the model; alternation pattern and plural suffix were not. Adding stress pattern to the phonological factors significantly improves the model ($\chi^2 = 15.96$, $p = .0003$) and does not substantially change the phonological effect sizes—the largest change in significant effect size is for polysyllabic words, which goes from 1.77 to 1.97 when stress pattern is added to the model. Words that always stress the stem in their inflectional paradigm are more likely to take *-tʃʲik* than those that sometimes or always stress the suffix. The intercept of this model is highly negative, reflecting the fact that only a small number of words are attested with *-tʃʲik*. That being said, base frequency is a very good predictor of attestation: the effect size of .67 corresponds to an increase of 1 in the log frequency, equivalent to multiplying the frequency by $e \approx 2.7$. Of the phonological effects, the strongest are affricates and dorsals: words ending in these consonants, like [krʲuk] 'hook' and [pálʲets] 'finger', *never* take *-tʃʲik*—the model has trouble handling truly categorical effects, hence the very high standard error and p value for these factors. Other than that, nouns ending in plosives and alveolars take *-tʃʲik* less than nouns ending in other consonants. Nouns ending in clusters, like [smɨsl] 'meaning', strongly disprefer *-tʃʲik*, as do nouns with non-final stress like [xaráktʲer] 'character' (which is not attested with any diminutive). Monosyllables take *-tʃʲik* less often than longer words, and the vowel of a noun's last syllable has a significant effect as well: nouns with high and low vowels take -*tʃʲik* more often than nouns with mid vowels (e.g. [plán-tʃʲik] 'plan (dim.)' and [kostʲúm-tʃʲik] 'suit (dim.)'), and nouns with rounded vowels take it more often than nouns with unrounded vowels. This model does a middling job at predicting whether a noun is attested with -*tʃʲik* ($R^2 = .29$).

	β coef	*SE*	*Wald z*	*p*
Intercept	**–7.26**	**.39**	**–18.45**	**<.0001**
Log frequency	**0.67**	**.04**	**19.16**	**<.0001**
C Manner (default: plosive)				
affricate	–15.01	1128.71	–0.01	.9894
fricative	**0.98**	**.27**	**3.70**	**.0002**
nasal	**2.84**	**.23**	**12.70**	**<.0001**
approximant	**2.12**	**.23**	**9.36**	**<.0001**
Coda (default: singleton)				
cluster	**–4.20**	**.53**	**–7.95**	**<.0001**
Stress (default: final)				
pre-final	**–3.22**	**.36**	**–8.93**	**<.0001**
Syllables (default: monosyllabic)				
polysyllabic	**1.77**	**.27**	**6.49**	**<.0001**
C Place (default: alveolar)				
labial	**0.48**	**.18**	**2.64**	**.0084**
dorsal	–15.37	371.98	–0.04	.9670
Stress pattern (default: stem)				
suffix	**–1.06**	**.30**	**–3.51**	**.0004**
mobile	–0.72	.52	–1.38	.1671
V Height (default: mid)				
high	**0.55**	**.17**	**3.09**	**.0020**
low	**0.78**	**.19**	**4.08**	**<.0001**
V Round (default: unrounded)				
rounded	**0.40**	**.17**	**2.38**	**.0171**

Table 6.14: Regression model with phonological and morphological predictors of diminutive -*tʃʲik*, with significant effects bolded

The effects of this regression are quite similar to those of Gouskova et al. (2015). In their Table 2, nouns ending in dorsals are much less likely to be placed in the -*tʃʲik* class, while those ending in nasals and approximants are much more likely. Nouns ending in clusters strongly disprefer -*tʃʲik*, while nouns with final stress take *-tʃʲik* more often than words with initial stress. Similarly, both of us find that monosyllables take *-tʃʲik* less often than longer words. The only difference is in the effects of the last vowel: in their model, nouns with high vowels in the last syllable are slightly less likely to take *-tʃʲik* than nouns with other vowels, though the effect is not significant. In my model, nouns with mid vowels disprefer *-tʃʲik* relative to nouns with high or low vowels. The discrepancy is likely due to how I coded for vowels (splitting mid and low vowels out), but given that their effect was not significant, the difference is not problematic. In addition, I found a vowel rounding effect: in my analysis, nouns with underlying rounded vowels [o u] (the former of which reduces to unrounded [ə] when unstressed) are slightly but significantly more likely to take *-tʃʲik* than those with unrounded vowels; for Gouskova et al. (2015), this effect was not significant.

As Gouskova et al. (2015) point out, certain phonological features tend to be correlated with one another—for example, consonant place and manner are not fully independent. To confirm that the factors in my model are sufficiently independent, I tested their variance inflation factor (VIF) using the `check_collinearity` function from R's `performance` package (Lüdecke et al., 2021). This measures whether different factors are describing the same effect. All of the effects had a VIF less than 2, far below the problematic range of 5 or 10 (see James et al., 2013), meaning that each effect in the lexicon model is predicting diminutive choice independently.

6.3.3.2 -ʲik

The model predicting whether a word takes *-ʲik* as a diminutive is shown in Table 6.15, whose factors are listed in the order in which they were added to the model. As with *-tʃʲik*, more frequent words are more likely to appear with *-ʲik* and stress pattern was the only morphological factor added to the model. Adding stress pattern to the phonological factors significantly improves the model ($\chi^2 = 46.79$, $p < .0001$), and only has a substantial effect on one factor: in the phonological model, nouns ending in affricates [ts tʃʲ] are significantly more likely to take *-ʲik* than nouns ending in plosives. However, once stress pattern is added to the model, showing that nouns with fixed-suffix stress prefer *-ʲik*, the affricate effect disappears (in fact, the effect size is now slightly *negative*, though no longer significant). This suggests that nouns with fixed suffix stress are over-represented among nouns ending in affricates, which is the case: 66 of 168 affricate-final nouns in my corpus (39.3%) have suffix stress; in the whole lexicon, the proportion of nouns with suffix stress is 8.0% (655 out of 8,178). Other phonological effects are more robust: for example, nouns ending in alveolars are more likely to take *-ʲik* (e.g. [momʲ'entʲik] 'moment, detail (dim.)') than those ending in labials ([gróbʲik] 'coffin (dim.)') and dorsals ([tántʃʲik] 'tank (dim.)', with a consonant alternation), and those ending in plosives prefer *-ʲik* relative to those ending in nasals and approximants. In addition, monosyllables and other words with final stress take *-ʲik* more often than longer words with non-final stress. This model performs similarly well to the model predicting *-tʃʲik* in Table 6.14, with $R^2 = .28$. Despite the masking of the affricate effect by stress pattern, the factors show very low collinearity, suggesting that all of the resulting effects are independent.

	β coef	*SE*	*Wald z*	*p*
Intercept	**–1.90**	**.17**	**–11.07**	**<.0001**
Log frequency	**0.56**	**.03**	**20.16**	**<.0001**
C Place (default: alveolar)				
labial	**–0.48**	**.18**	**–2.65**	**.0081**
dorsal	**–3.91**	**.40**	**–9.79**	**<.0001**
C Manner (default: plosive)				
affricate	–0.32	.29	1.10	.2701
fricative	–0.18	.14	–1.32	.1857
nasal	**–1.83**	**.21**	**–8.61**	**<.0001**
approximant	**–1.41**	**.16**	**–8.71**	**<.0001**
Syllables (default: monosyllabic)				
polysyllabic	**–1.14**	**.13**	**–9.07**	**<.0001**
Stress (default: final)				
pre-final	**–1.48**	**.24**	**–6.08**	**<.0001**
Stress pattern (default: stem)				
suffix	**1.19**	**.17**	**6.96**	**<.0001**
mobile	–0.05	.27	–0.19	.8535
V Height (default: mid)				
high	**–0.70**	**.15**	**–4.55**	**<.0001**
low	0.21	.14	1.51	.1318
V Round (default: unrounded)				
rounded	**0.29**	**.13**	**2.17**	**.0299**

Table 6.15: Regression model with phonological and morphological predictors of diminutive *-ʲik*, with significant effects bolded

Similar to *-tʃʲik*, nouns ending in dorsals take *-ʲik* very rarely; my data set contains seven such words. Some of these are false positives (e.g. [logʲik] ‘logic’ is recorded as a diminutive of [log] ‘ravine’), but some of are real diminutives (like [bantʃʲik] as a diminutive of [bank] ‘bank’), which is listed in the dictionary of Ševeleva (1974) as discussed in Section 6.1.2. These words have little effect on the model except to make the dorsal effect size somewhat smaller than it otherwise would be.

The effects in Table 6.15 are again quite similar to those reported in Table 2 of Gouskova et al. (2015): their reported effects for dorsals, nasals, approximants, high vowels, stress location, and word length go in the same direction as mine and are similar in relative effect size. They report a moderate effect of coda complexity: nouns ending in clusters are moderately more likely to take -ʲ*ik* than nouns ending in single consonants. However, this factor was not added to my model. Finally, I found that words with rounded vowels in the last syllable were slightly but significantly more likely to take -ʲ*ik*, whereas Gouskova et al. (2015) show a slight non-significant effect in the *opposite* direction.

6.3.3.3 -ók

Table 6.16 presents the model predicting whether a word takes *-ók*, with factors listed in the order in which they were added to the model. This model is very different from the previous ones: it contains fewer phonological factors but every morphological factor, including plural marker. Once stress pattern is added, it is the most important factor, even outpacing base frequency (which goes in the expected direction, though with a smaller effect size): nouns that stress the suffix some or all of the time are much more likely to take stressed *-ók* as a dimunitive (for example, [jazɨ́k] 'tongue' has plural [jazɨkʲí] and diminutive [jazɨtʃʲók]). Adding this and vowel–zero alternation *substantially* improves the model ($\chi^2 = 571.47$, $p < .0001$). Nouns that have vowel–zero alternations in their inflectional paradigm (like [kovʲór] 'carpet', genitive [kovr-á]) take *-ók* significantly *less* than those that do not—for example, the diminutive of 'carpet' is [kóvr-ʲik]. The third model, which also includes plural suffix as a predictor, improves the model slightly ($\chi^2 = 4.06$, $p < .0001$). Nouns with plural *-a* (which is always stressed) are more likely to take *-ók even when stress pattern is factored in*, but this effect does not reach the level of significance at $p < .05$. Of the eight nouns that have the plural *-ja* pattern, seven allow *-ók*, so this factor is significant as well. The addition of stress pattern allows one phonological factor to emerge: stress position was not included in the model with only phonological factors, but once stress pattern is added, we get a strong significant effect: nouns with non-final stress, like [mólot] 'sledgehammer' in the base form are much more likely to take *-ók* (in this case, [molot-ók] 'hammer'). This makes sense given that almost all nouns with suffix stress have stress on the last syllable of the stem in the unsuffixed base form (see Section 6.2.2). The only other significant phonological effects are that nouns ending in dorsals, like [grʲex] 'sin' are more likely to take *-ók* ([grʲes̩ók]) than those ending in alveolars, monosyllables take *-ók* much more often than longer words, and nouns that end in clusters substantially disprefer *-ók*. This model is a better fit than the other two ($R^2 = .34$). As expected given its overlap with plural and stressed syllable, the factor of stress pattern has a somewhat higher VIF (3.12) than any in the other two models, suggesting that this factor overlaps somewhat with others. In addition to being interpretable, this number is still well in the range of being unproblematic, so I do not address it further.

	β coef	*SE*	*Wald z*	*p*
Intercept	**–3.25**	**.17**	**–18.78**	**<.0001**
Stress pattern (default: stem)				
suffix	**3.05**	**.17**	**18.42**	**<.0001**
mobile	**2.85**	**.27**	**10.57**	**<.0001**
Log frequency	**0.38**	**.03**	**12.93**	**<.0001**
C Place (default: alveolar)				
labial	–0.40	.25	–1.58	.1148
dorsal	**1.13**	**.13**	**8.50**	**<.0001**
Syllables (default: monosyllabic)				
polysyllabic	**–1.89**	**.17**	**–10.87**	**<.0001**
Coda (default: singleton)				
cluster	**–1.98**	**.26**	**–7.63**	**<.0001**
Stress (default: final)				
pre-final	**0.93**	**.19**	**4.97**	**<.0001**
Vowel–zero alternation (default: no alternation)				
alternation	**–0.90**	**.20**	**–4.42**	**<.0001**
Plural (default: -ɨ)				
-a	0.63	.38	1.69	.0912
-ja	**5.23**	**1.10**	**4.77**	**<.0001**

Table 6.16: Regression model with phonological and morphological predictors of diminutive *-ók*, with significant effects bolded

Phonological effects play only a secondary role in Table 6.16, and some significant effects from

Gouskova et al. (2015) do not make it into my model at all. For example, they found (in Table 2) that nouns ending in nasals slightly but significantly disprefer *-ók*, while nouns with rounded vowels in the last syllable slightly but significantly prefer *-ók*. Otherwise, their phonological effects are similar to mine: nouns ending in dorsal prefer *-ók* relative to alveolar-final nouns, nouns ending in clusters disprefer *-ók*, monosyllabic words take it more often than longer words. In their model, words with final stress (including monosyllables) take *-ók* less often than words with initial stress. This matches my model, which is somewhat surprising given that this effect only emerged once the factor of stress pattern was added.

In anticipation of the nonce word study, I present one more regression predicting whether a given noun takes *-ók* as a diminutive. In this case, I limit the data to nouns that appear with at least one diminutive (1,250 of the 8,178 total nouns) and only use phonological predictors. The effects of this regression will be used as the representation of the phonological distribution of the lexicon in the nonce word study, Section 6.4.

The effect sizes of this regression are shown in Table 6.17. This model contains the same four phonological factors as Table 6.16, with all individual effects pointing in the same direction; the dorsal preference for *-ók*, in particular, becomes much stronger: if a dorsal-final word takes any diminutive at all, it likely takes *-ók*. This model also adds three more phonological factors: noun stems with low and rounded vowels in the last syllable accept *-ók* significantly less than those with mid and rounded vowels, respectively. In addition, final consonant manner was added to the model, though none of the effects are significant. This model does a fairly good job of predicting whether a noun will take *-ók* ($R^2 = .55$), and there is very little collinearity among the factors (the VIF for the factors are all 1.44 or smaller).

	β *coef*	*SE*	*Wald z*	*p*
Intercept	**0.70**	**.27**	**2.64**	**.0082**
C Place (default: alveolar)				
labial	**–0.61**	**.27**	**–2.27**	**.0234**
dorsal	**6.39**	**.75**	**8.57**	**<.0001**
Stress (default: final)				
pre-final	**3.63**	**.31**	**11.67**	**<.0001**
Syllables (default: monosyllabic)				
polysyllabic	**–2.55**	**.21**	**–12.04**	**<.0001**
Coda (default: singleton)				
cluster	**–1.87**	**.31**	**–6.07**	**<.0001**
V Height (default: mid)				
high	0.35	.22	1.58	.1144
low	**–0.87**	**.26**	**–3.37**	**.0008**
V Round (default: unrounded)				
rounded	**–0.59**	**.21**	**–2.86**	**.0043**
C Manner (default: plosive)				
affricate	0.97	.53	–1.82	.0684
fricative	–0.25	.26	–0.95	.3422
nasal	–0.27	.27	–1.02	.3095
approximant	0.46	.25	1.89	.0588

Table 6.17: Regression model with phonological predictors of diminutive *-ók* over nouns that take at least one diminutive suffix, with significant effects bolded

A more complex model that includes the morphological effects of stress pattern and plural suffix (as before, plural *-a* is a substantial positive predictor of taking *-ók* but not a significant one) is a subtantially better fit than the model in Table 6.17 ($\chi^2 = 159.34$, $p < .0001$), bringing the R^2 up to .64. This more complex model retains all of the phonological factors—adding these morphological factors barely disturbs most of the factors in Table 6.17, with the exception of affricates, which become significantly *less* likely to take *-ók* once phonological factors are taken into account. Thus, the inclusion of additional phonological factors in Table 6.17 relative to Table 6.16 is not due to the absence of morphological factors, but rather the increased clarity gained from removing from consideration nouns that appear with no diminutive (previously accounted for somewhat with the factor of frequency). Of the newly added phonological factors (vowel height and roundedness, consonant manner) two are significant in the lexicon study of Gouskova et al. (2015): as mentioned above, they found that nouns ending in nasals take *-ók* significantly less often, while nouns with rounded vowels take *-ók* significantly more often. When looking at this slice of the lexicon, my model yields the same nasal effect (though not significant) and a significant effect of rounding in the *opposite* direction: in Table 6.17, nouns with rounded vowels in the last syllable are *less* likely to take *-ók*. I compare my results with those of Gouskova et al. (2015) more extensively in the following discussion.

6.3.4 Discussion

6.3.4.1 Phonological effects

Previous work, surveyed in Section 6.1.2, studied the phonological factors that influence the choice of Russian diminutive. In Table 6.18, I describe the significant phonological effects found in the corpus study of Gouskova et al. (2015). The more salient effects have already been discussed, like the tendency of dorsals to go with *-ók*. Others have clear markedness-based explanations: for example, nouns ending in clusters disprefer *-tʃʲik*, the only diminutive suffix that begins with a consonant. The length effect, as described previously, may be a historical residue: *-tʃʲik* is a more recently innovated suffix, and often goes with loanwords. If loanwords and other neologisms (which can often be derived or complex) have a tendency to be longer than older native words, we would expect the effect that we see. However, this hypothesis requires further study.

My corpus study, whose results are shown in Table 6.19, was intended in part as a replication of Gouskova et al. (2015), and despite some slightly different coding of the variables, most of the major effects were replicated. This is shown by the bolded effects in the two tables. Indeed, the main effects that differed between the two studies are those involving properties of the last-syllable vowel (height and rounding). Given that these are the primary effects involving units not directly adjacent to the suffixes, we would expect them to be less robust. The additional effect of rounding came out as significant in the second model predicting *-ók* in Table 6.17, which was limited to nouns that took at least one diminutive. However, the effect was *opposite* that of Gouskova et al. (2015): nouns with rounded vowels in the last syllable took *-ók less* than those with unrounded vowels. They offered a markedness-based explanation of this effect as rounding harmony. Given that unstressed /o/ reduces to an unrounded vowel like [ə] or [ɐ], this explanation does not hold on the surface anyway: unstressed and stressed /o/ do not agree in rounding.

		likelihood of accepting …		
phonological factor		*-tʃʲik*	*-ʲik*	*-ók*
final C place	dorsal	≪ **other**	≪ **other**	≫ **other**
final C manner	glide		< obstruent	
	approximant	≫ **obstruent**	< **obstruent**	
	nasal	≫ **obstruent**	< **obstruent**	< obstruent
final coda	cluster	≪ **singleton**	> singleton	< **singleton**
stress	final	≫ **initial**	> **initial**	< **initial**
	medial			≪ initial
length	polysyllabic	> **monosyllabic**	< **monosyllabic**	< **monosyllabic**
final V height	high		< **other**	
final V round	rounded			> unrounded

Table 6.18: Summary of significant phonological effects on possessive allomorphy from Gouskova et al. (2015, p. 54) (doubled ≪ and ≫ represent effect sizes greater than 2, **bolded** effects match Table 6.19)

		likelihood of accepting …		
phonological factor		*-tʃʲik*	*ʲik*	*-ók*
final C place	labial	> alveolar	< alveolar	
	dorsal	≪ **alveolar**	≪ **alveolar**	> **alveolar**
final C manner	approximant	≫ **plosive**	< **plosive**	
	nasal	≫ **plosive**	< **plosive**	
	fricative	> plosive		
	affricate	≪ plosive		
final coda	cluster	≪ **singleton**		< **singleton**
stress	final	≫ **other**	> **other**	< **other**
length	polysyllabic	> **monosyllabic**	< **monosyllabic**	< **monosyllabic**
final V height	high	> mid	< **mid**	
	low	> mid		
final V round	rounded	> unrounded	> unrounded	

Table 6.19: Summary of significant phonological effects on possessive allomorphy from Section 6.3.3 (doubled ≪ and ≫ represent effect sizes greater than 2, **bolded** effects match Table 6.18)

The present corpus study relies on the automatic tagging of the Russian National Corpus (with some cleaning), while that of Gouskova et al. (2015) used searches in a Google Ngram corpus. The latter mostly comprises scanned books, while the former contains a more diverse set of source materials: books, journalism, internet texts, and so on. Given that diminutive usage is somewhat associated with informality, the source materials of the two corpora might be expected to affect these results. However, the close alignment between their results and mine suggests that they are robust, despite the possibility of false positives and differences in corpus makeup, corpus construction (for example, I excluded animates, while they did not), and variable coding.

6.3.4.2 The effect of stress

The primary innovation of this study is its inclusion of factors for inflectional patterns, most importantly a noun's stress pattern across its inflectional paradigm and its segmental plural suffix. Most Russian words have stress fixed on the stem throughout their paradigm, though a minority exhibit one of a number of patterns of stress on the suffix in some or all cases (not counting unsuffixed forms, in which stress necessarily reverts to the stem). Sometimes inflectional stress patterns and segmental content interact. The irregular nominative plural *-a* studied in this section is always stressed, and nouns with variable plural suffixes, or with different plural suffixes for closely related meanings, often have concomitant variance in their stress pattern: for example, [ton] 'tone' has plural [toná] or [tónɨ] depending on the particular meaning of the word (Zaliznjak, 1977), as mentioned in Section 6.3.1.1. As discussed in Section 6.2.2, stress patterns are typically analyzed using underlying stress marks that have different realizations on the surface. To capture the relationship between stress pattern and diminutive, then, sublexical grammars must be able to make generalizations over underlying forms of roots, not just surface forms of base (unsuffixed) nouns. On the other hand, plural *-a* must be marked through a diacritic feature, so its correlation with diminutive *-ók* can be analyzed with a morphological dependency between diacritic features as in the other cases in this dissertation.

Brown et al. (1996) showed that certain stress patterns are typical of certain inflection classes. However, the relationship between stress pattern and diminutive suffix has not previously been explored quantitatively. There are tendencies: for example, nouns with fixed suffix stress prefer *-^{j}ik* and disprefer *-tʃ^{j}ik*. As mentioned previously, *-tʃ^{j}ik* is a newer diminutive often used on loanwords, and fixed stem stress is the default majority pattern. The slight preference of *-^{j}ik*, which is at least in part a frequency effect (words with suffix stress are more frequent and thus more likely to appear with diminutives), pales in comparison with the preferences of *-ók*. This diminutive suffix strongly prefers nouns that have a stressed suffix in all (coded as suffix stress) or some (coded as mobile stress) paradigm cells. Of the 444 nouns that appear with *-ók*, 281 have a stressed suffix in at least some inflected forms (63.3%). For all nouns, the percentage is only 10.2% (832 of 8,178); of nouns that appear with at least one diminutive, the percentage is 30.5% (381 of 1,250).

This heavily skewed distribution is a major determining factor of *-ók*, and previous studies that focused on the phonological generalizations behind diminutive choice omitted an important part of the picture. This incomplete description may have also affected the results of nonce word studies: Gouskova et al. (2015) and Magomedova (2017) both found that *-ók* was less productive than the other two suffixes (that is, selected less often in a forced choice task). However, these studies only presented words in their base form, without any information about stress patterns. If positive evidence of suffix stress is a main factor required to push a word towards *-ók* for speakers in this task, then we would expect apparent underproductivity if stress pattern is not included. In Section 6.4, I present a nonce word task in which participants are forced to choose between the three possible diminutive suffixes. Each nonce word in this task is also shown in its plural form; 25 of 40 trials for each speaker have stress on the plural suffix. If *-ók* is still fully productive, we would expect it to be chosen frequently in this experiment, and this is indeed what I find. This suggests that the previous reported claim that *-ók* is losing productivity in Russian is due to an incomplete picture of its distribution.

The tendency for *-ók* to attach to nouns with some inflected forms with stressed suffixes is a sort of paradigm uniformity effect (see Section 2.3.2.1), given that *-ók* itself, unlike *-ʲik* and *-tʃʲik*, always attracts stress. That is, a noun is more likely to have suffix stress in its diminutive if it has suffix stress in at least one of its inflected forms. As with the phonological patterns studied, this is a variable tendency which would be strange to hard-code into the grammar as an inviolable principle—see Section 2.3 for discussion. Like the other morphological dependencies in this dissertation, the correlation between diminutive *-ók* and inflectional suffix stress should be considered a correlation between two discrete lexically specific patterns.

6.3.4.3 Stress and segmental inflection patterns

A small number of nouns (86 in my corpus) take stressed *-a* in the plural instead of the usual *-i* (my representation of allophonic [i∼ɨ]), which can be stressed or unstressed. Of these, a majority (46) take diminutives with *-ók*, a remarkable number given that the vast majority of nouns do not combine with any diminutives at all. In fact, there are only four nouns with plural *-a* that combine with *-ʲik* or *-tʃʲik* but not *-ók*. This is quite a strong tendency. However, since plural *-a* is always stressed, we must ask: can the proclivity of plural *-a* nouns to take diminutive *-ók* be reduced to the general tendency of nouns with suffix stress to take diminutive *-ók*, or is it even stronger? Stress pattern and plural suffix are included as separate factors in the regression in Table 6.16. Plural *-a* nouns are indeed moderately more likely to take *-ók*, even on top of the much stronger effect for nouns with fixed suffix or mobile stress, but this effect does not reach significance ($p = .09$).

All of the phonological and morphological effects studied in this dissertation are somewhat conflated with each other: although the collinearity does not reach problematic levels, properties of the final consonant (place, manner, presence in clusters, etc.) are not fully independent. The conflation between suffix stress and plural *-a* raises an issue: do we expect speakers to learn a correlation between *-ók* and plural *-a* on top of the correlation between *-ók* and suffix stress? As Divjak et al. (2016, p. 27) note, there are often many possible grammars that can be learned from a distribution of input forms, because many of the various factors predicting this distribution also predict one another. Thus, individual speakers seem to have slightly different grammars, because they have learned different combinations of relevant factors. If this is the case, we would expect some speakers to have learned a correlation between *-a* and *-ók* and others to have only learned the more general correlation with suffix stress. Thus, in Section 6.4, I look not just for overall effects, but also for evidence of individual differences. There will be substantial variance given that each subject sees only ten stimuli with plural *-a*, but if some speakers have a different underlying grammar, we would expect the distribution of the plural suffix effect for individual speakers to follow a bimodal distribution.

6.4 Russian nonce word study

This study tests whether speakers show sensitivity to phonological properties and plural form in selecting the diminutive of nonce words. In order to test these correlations, I presented nonce words with one of three plural suffixes: unstressed *-i*, stressed *-i*, and stressed *-a*. Speakers were expected to uniformly show sensitivity to stress pattern (unstressed *-i* vs. stressed *-i*) and variably show sensitivity to plural suffix (stressed *-a* vs. stressed *-i*). Unlike the previous nonce word studies, this one only includes a stimulus testing study, where participants had to select diminutive forms for the stimuli; since I borrowed stimuli from Gouskova et al. (2015), I did not conduct a stimulus norming study.

In choosing *-ók* vs. the other two suffixes, participants showed significant but not very strong sensitivity to phonology, even though the stimuli were chosen with a limited phonological range. Results also show substantial morphological effects: nouns with suffix stress on the plural are assigned *-ók* much more often than those with fixed stem stress, and nouns with stressed plural *-a* are *even more* likely to take *-ók*. This shows that speakers have learned both morphological tendencies separately. In Section 6.3.4.3, I predicted that some speakers would assign *-ók* to nouns with stressed plural *-a* than those with stressed plural *-i*. However, this prediction is not borne out: instead, we see that speakers have all learned the plural suffix effect equally, without differences in underlying grammars: *all* speakers grammatically associate plural *-a* with higher likelihood of *-ók*.

6.4.1 Participants

I recruited 120 participants through Prolific. Because there were not enough Prolific users in Russia proper, participants could be born and located anywhere, so long as they identified as Russian speakers. Four participants were eliminated due to technical issues, while another two were removed for responding that they had not spoken Russian their whole lives, leaving a total of 114 speakers.

6.4.2 Stimuli

I selected 87 of the 300 nonce words used as stimuli by Gouskova et al. (2015), which they generated to test the influence of a number of phonological factors (such as hiatus, final coda properties, stress, etc.) on diminutive choice. In their study, stress was indicated in two ways: first, as in my study, participants were given audio recordings for all forms; second, stress was marked orthographically using an acute accent on the stressed vowel, which is common for didactic purposes but not normal in everyday texts. I wished to avoid any potential artificiality involved with orthographic stress marking, so in my study stimuli were presented in normal orthography (without stress marks) with accompanying audio recordings in all forms, so listening to the recordings was the only way to determine stress and distinguish between orthographically identical forms. The stimuli chosen comprised a subset of the range of phonological forms: they were all disyllabic and had final stress, no word-internal hiatus, and no word-final clusters (one stimulus with a cluster was accidentally included in the set; trials with this word were removed). Each had three possible diminutives (with *-ók*, *-ʲik*, and *-tʃʲik*) and three possible plurals (stem stress with *-i*, suffix stress with *-i*, suffix stress with *-a*).

As I described in Section 6.2.2, nouns with suffix stress usually have stress on the last syllable in unsuffixed forms, while nouns with mobile stress usually have stress on the first syllable in forms where the stem is stressed. My stimuli all have final stress when unsuffixed, so when the stress is on the plural suffix, this suggests that the nonce word has suffix stress rather than mobile stress. In Section 6.2.2.5, I mentioned that most nouns with plural *-a* have mobile stress, and only a couple have suffix stress. Thus, the stimuli with plural *-a* in my study can be inferred to have suffix stress (the rare pattern) rather than mobile stress (the more common pattern). As far as I know, this design choice did not affect the results.

6.4.3 Procedure

Participants were shown 40 stimuli in frame sentences showing the stimulus twice: once in the nominative or accusative singular, and once in the nominative or accusative plural. An example of a trial is shown in Figure 6.2. Both instances of the stimulus was accompanied by a button that played an audio recording; participants had to listen to both recordings before continuing. (In Figure 6.2, each button is accompanied by a phonetic representation of the word that played when it was clicked.) As an attention check, participants had to correctly select the plural form appearing in the first sentence from the three possibilities, in order to ensure that they were really registering the plural form. Since two of the plural forms differed only in stress, participants had to listen to at least one of the possible plural forms (if the plural ended in *-i*, they had to listen to at least one of the two possible plurals in *-i*) before selecting the plural. Once the correct plural form was selected, a second frame sentence appeared, in which participants had to select from the range of possible diminutive forms (again with corresponding audio). Participants had to listen to all three possible diminutives before choosing, to ensure that they registered that *-ók* was stressed. Upon choosing a diminutive, participants had to indicate whether the choice was a difficult one. The example in Figure 6.2 shows an example where the nonce word is shown with stressed *-i* in the plural; in other trials, the frame sentence used one of the other two options listed.

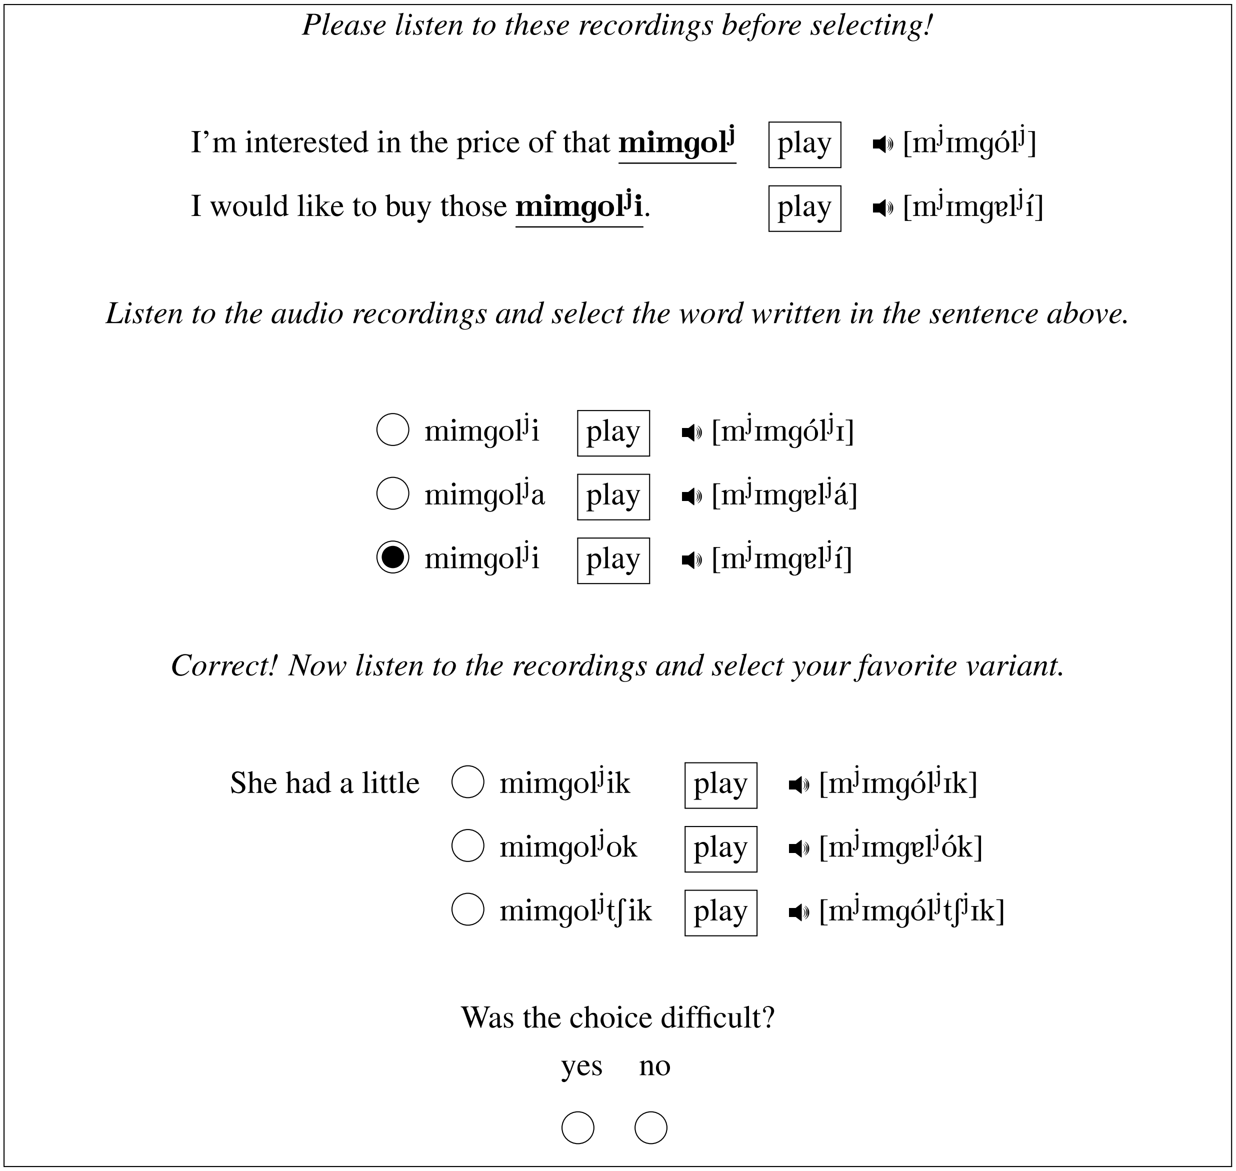

Figure 6.2: Trial for Russian stimulus testing study (with phonetic transcriptions of recordings next to the buttons that played them)

6.4.4 Analysis

The previous nonce word studies included a check where I prompted speakers for the target morphological form (in this case, the plural) a second time. Due to the added complexity of the task in this study (the audio files and the choice between three options), I did not include a second prompt. All trials were kept except for those with the one nonce word ending in a cluster, which was thrown out. This left a total of 4,588 trials. I fitted three series of mixed logistic regressions, one for each of the three diminutive suffixes, whose dependent variable was whether that diminutive was selected for a given nonce word. My main focus is the pair of regressions predicting *-ók* vs. the other two options, and I only mention the results for *-ʲik* and *-tʃʲik* briefly in the discussion. The first regression of each pair includes a fixed effect of phonology. If speakers are matching the lexicon, then the experimental results should correlate with a noun's likelihood of taking *-ók* in the lexicon. As with the previous studies, I represent the phonological model of the lexicon with a single *phon_odds* coefficient, which predict whether a word takes *-ók* using the model in Table 6.17. This model is limited to words that take at least one diminutive, and thus best represents the forced choice task. See Section 4.3.3 for an explanation of how *phon_odds* are calculated. I also include random intercepts for participant and item and consider a by-participant random slope for *phon_odds* that reflects whether individual participants show different sensitivity towards the phonological effects. I consider a binary variable indexing whether speakers found the choice difficult and nuisance variables marking trial number (1–40) and presentation order (whether the selected diminutive was listed first, second, or third, coded as -1, 0, and 1); only presentation order improved the model according to the Akaike Information Criterion (AIC), which rewards model fit and penalizes model complexity (number of factors), and the others were not added to the model.

The second regression in each series adds the morphological factors: stress pattern (whether the plural was presented with stress on the stem or the suffix) and plural suffix (*-i*, which may be stressed or unstressed, vs. *-a*, which must be stressed). I also considered by-participant random slopes for these two effects. Finally, the third regression in each series replaces *phon_odds* with individual phonological factors to see whether some factors were more heavily applied than others.

I also fitted a regression whose dependent variable was *difficulty* of choosing the diminutive.

Because particular phonological characteristics might make the choice difficult, I included each phonological variable that differed among the stimuli (final C place and manner and final V height and rounding) separately, rather than including a single *phon_odds* factor. I also included random intercepts for participant and item, the diminutive chosen on that trial, and all of the fixed effects mentioned above: stress pattern, plural suffix, and the nuisance variables of trial number and presentation order. Because I am interested in *-ók* compared to the other two diminutives, I set *-ók* as the base level for chosen diminutive. I built this model by stepwise comparison using the `buildmer` function from the R package of the same name (Voeten, 2022): factors were added to the model if they improved the model's AIC.

6.4.5 Results

In this section, I focus on the factors driving participants for or away from choosing *-ók*, specifically, over the other two options. This was the most popular diminutive, with 1,664 choices compared to 1,185 for *-ʲik* and 1,186 for *-tʃʲik*. These proportions were due in part to the high prevalence of plurals with suffix stress (25 out of 40 trials), which moderately prefer *-ók*, as shown in Table 6.20. The preference for *-ók* is even stronger among nouns with plural *-a*, which were assigned *-ók* in an absolute (though slight) majority of trials.

plural	*-ók*	*other*	% *-ók*
stem stress *-i*	585	1155	33.6%
suffix stress *-i*	797	943	45.8%
suffix stress *-a*	583	577	50.3%

Table 6.20: Experimental frequency of *-ók* and other diminutives, by plural condition

6.4.5.1 Phonology

Table 6.21 shows the effects of the mixed logistic regression using phonological factors predicting whether participants chose a diminutive with *-ók*. The model has two fixed effects: the *phon_odds* coefficient derived from the phonological model of the lexicon in Table 6.17, and presentation order: the higher up *-ók* was presented in the list of the three diminutive options, the more likely it was to be chosen, though the difference is not quite significant. The other candidate factors (trial number and whether the choice was hard) did not improve the model and were thus not added to it. The model also includes random intercepts for participant and item and a by-participant random slope for *phon_odds*. This model is not a very good fit: it has an R^2 of .231 and a Somers' D of .536 (Lüdecke et al., 2021; Somers, 1962), meaning that the model assigns a higher probability to the correct outcome 53.6% of the time.

Random effects	*variance*	*SD*		
Participant				
Intercept	0.35	.59		
Phon_odds	0.00	.06		
Item	0.25	.50		
Fixed effects	*β coef*	*SE*	*Wald z*	*p*
Intercept	**−0.27**	**.08**	**−3.17**	**.0015**
Phon_odds	**0.20**	**.02**	**8.98**	**<.0001**
Presented order	−0.08	.04	−1.88	.0599

Table 6.21: Effects of mixed logistic model with predictions of the phonological model of the lexicon (Table 6.17) for experimental selection of diminutive *-ók*, with significant effects bolded

The effect size for the *phon_odds* coefficient is quite low, .20. This means that the effect of phonology on the experimental results is only 20% as large as on the lexicon model. In fact, this coefficient is overwhelmingly driven by dorsals, which yielded far and away the strongest effect in Table 6.17—other moderately large effects from the lexicon (stress location, word length, final coda complexity) are uniform for the stimuli in this experiment. The random intercepts have moderate variance, suggesting that some items were judged slightly better or worse for *-ók* even beyond their phonology, and that participants have different baseline preferences for or against *-ók*—the number of diminutives with *-ók* selected by each participant ranged from 4 to 29. The random slope has almost no variance, indicating that participants were roughly equally sensitive to the effect of phonology.

We can see the strong effect of dorsals in Figure 6.3 and Figure 6.4, which show the predicted and actual rates of *-ók* for individual nonce words plotted on a log odds scale and an untransformed scale, respectively. The words are grouped into two tight clusters: on the bottom left are words ending in alveolars and labials, which have a predicted likelihood of *-ók* between around 0% and 30%, though the actual experimental rate goes as high as about 70% for [xakóts]. On the right is the cluster of words ending in dorsals, which are predicted to take *-ók* nearly categorically, although their actual rate ranges as low as about 30%. While the cluster of dorsals on the right is situated higher than the cluster on the left, the vertical difference between them (reflecting their experimental rates) is much smaller than in the lexicon, and there is a lot of overlap between the experimental rates of dorsal-final and other words. The graphs show each nonce word twice: in black, the word's position on the x-axis assumes a fixed intercept, so each word's position is solely a function of its phonology; the lighter gray shows the adjustment of the random intercepts. The purpose of the random intercepts is to bring each word closer to the line of best fit. Thus, words below the line (which received *-ók* less often than expected given their *phon_odds*) have a negative random intercept, so the word in gray is to the left of its black counterpart; words above the line, which get *-ók* more than expected by the phonological model, have positive random intercepts, so the gray word is to the right of the black. In general, the greater the distance between a word and the line of best fit, the greater the magnitude of its random intercept.

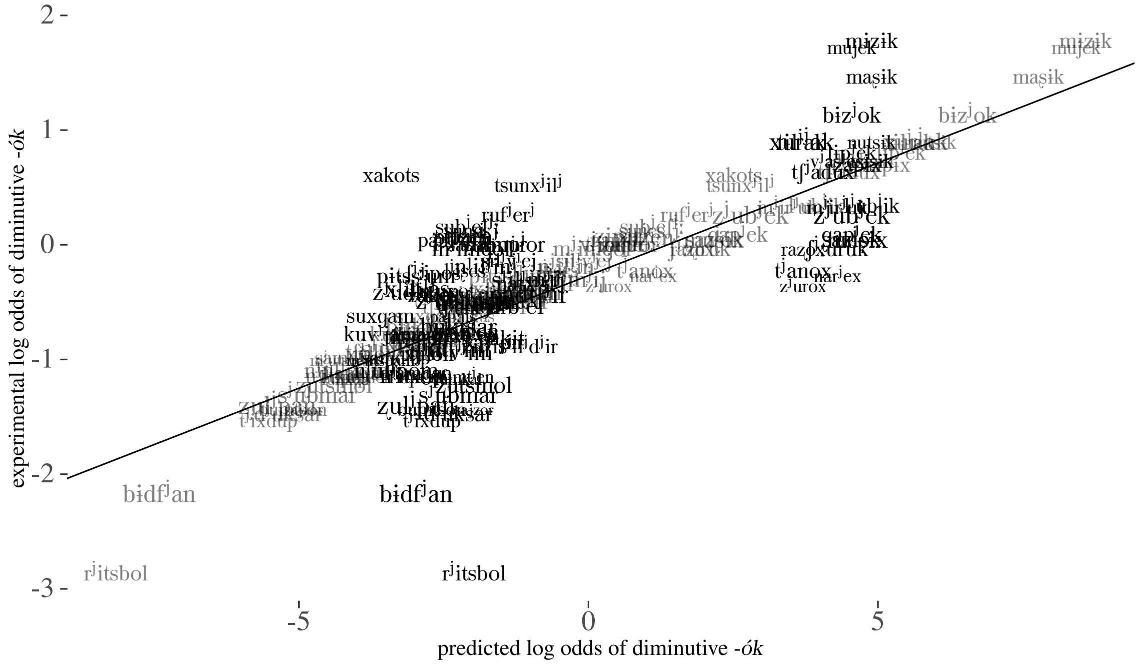

Figure 6.3: The relationship between predicted and experimental log odds of diminutive -*ók* for individual nonce words with (gray) and without (black) the random intercept, sized according to number of trials, with a line showing the fit of the experimental model in Table 6.17

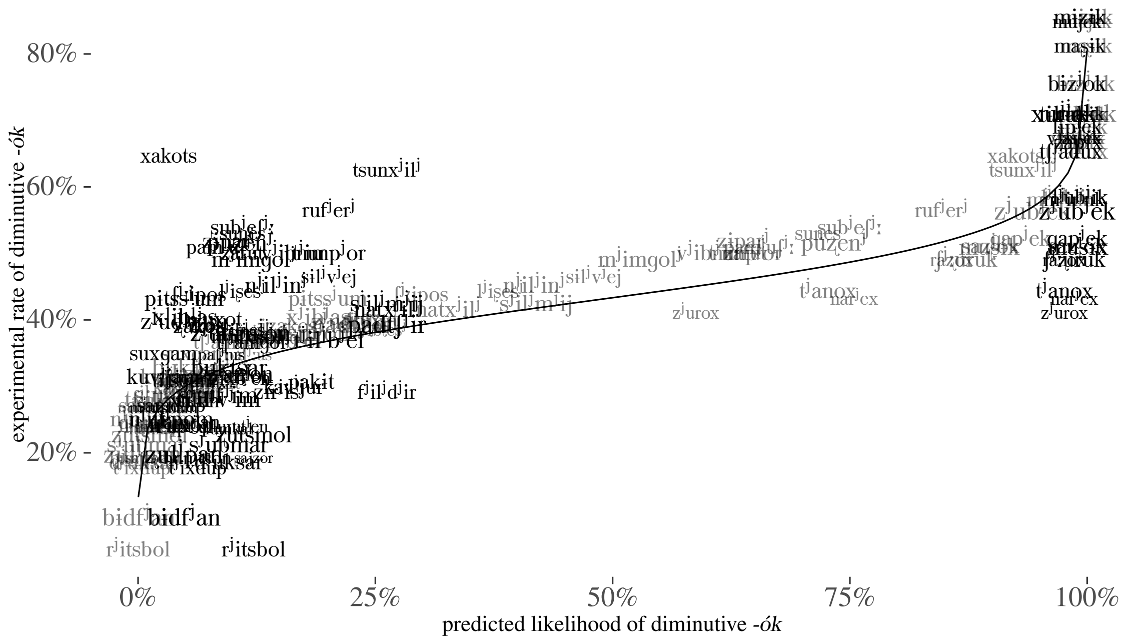

Figure 6.4: The relationship between predicted likelihood and experimental rate of diminutive *-ók* for individual nonce words with (gray) and without (black) the random intercept, sized according to number of trials, with a line showing the fit of the experimental model in Table 6.17

6.4.5.2 Phonology and morphology

The model in Table 6.22 includes two new factors: plural stress (stem vs. suffix) and plural suffix (*-i* vs. *-a*). These two factors covary: stimuli with plural *-i* can have stem stress or suffix stress (15 trials each), while stimuli with plural *-a* can only have suffix stress (10 trials). I present the two factors separately below, but because they are mutually exclusive, it is equivalent to represent the plural as a single factor with a three-way distinction comparing a baseline of suffix stress *-i* to stem stress *-i* and plural *-a*. This model is a much better fit than the pure phonological model ($\chi^2 = 114.7$, $p < .0001$), but still not very good: $R^2 = .267$ (.112 if only the fixed effects are taken into account), with a Somers' D of .566, indicating that the model correctly guesses the outcome in 56.6% of trials. The effects shared between this model and the phonological model in Table 6.21 have very similar sizes, with one exception: adding the morphological factors makes the intercept much lower, as much of the work done by the intercept in the phonological model is now shifted onto the positive morphological effects. However, the fixed effect of phonology, presented order, and the random effects are much the same. As before, the effect of presented order does not quite meet the significance threshold of $p < .05$. Adding by-participant random slopes for plural stress and suffix to this model did not significantly improve its fit, and they showed modest variance, especially plural suffix (.28 for stress, .02 for suffix). These random slopes also did not substantially affect the size or significance of the fixed effects, so I do not include them in the presented model.

Random effects	*variance*	*SD*		
Participant				
Intercept	0.37	.61		
Phon_odds	0.00	.07		
Item	0.27	.52		
Fixed effects	*β coef*	*SE*	*Wald z*	*p*
Intercept	**–0.73**	**.10**	**–7.39**	**<.0001**
Plural stress (default: stem)				
suffix	**0.64**	**.08**	**8.12**	**<.0001**
Phon_odds	**0.20**	**.02**	**8.86**	**<.0001**
Plural suffix (default: -i)				
-a	**0.23**	**.09**	**2.65**	**.0080**
Presented order	–0.07	.04	–1.78	.0744

Table 6.22: Effects of mixed logistic model with predictions of the phonological model of the lexicon (Table 6.17), plural stress, and plural suffix for experimental selection of diminutive *-ók*, with significant effects bolded

Before looking at the visual representation of individual nonce words, let us look at two example stimuli and compare them to the predicted morphological effect, which is that rate of *-ók* selection should increase going from stem stress *-i* to suffix stress *-i* to suffix stress *-a* (the three plural conditions). Table 6.23 shows the effect of plural condition on diminutive selection for [t^janóx] and [z^juróx]. Both of these stimuli end in dorsals, so they are expected to have fairly high baseline rates of *-ók*—however, these words show a less extreme distribution than many of the stimuli ending in dorsals, especially [g k]. Each of these stimuli shows one of the expected effects. Participants assigned *-ók* more often to [tʲanóx] when it was presented with suffix stressed plural *-i* than when it was shown with stem stress, but they assigned *-ók less* often to [tʲanóx] with plural *-a*, contrary to the prediction. On the other hand, participants assigned *-ók* to [zʲuróx] at similar rates across the two plural *-i* conditions, but assigned it much more often when the word was shown with plural *-a*, as predicted. This example shows that the results were somewhat messy for individual stimuli, even though the aggregate effects were as expected.

	tʲanóx			zʲuróx		
plural	*-ók*	*other*	*% -ók*	*-ók*	*other*	*% -ók*
stem stress *-i*	9	13	40.9%	4	10	28.6%
suffix stress *-i*	14	7	66.7%	9	19	32.1%
suffix stress *-a*	4	14	22.2%	12	7	63.2%

Table 6.23: Effect of plural condition on selected diminutive for two Russian nonce stimuli, [tʲanóx] and [zʲuróx]

The effect of plural stress for individual nonce words can be seen in Figure 6.5 and Figure 6.6. Unlike the previous graphs, these do not indicate the random intercepts for item (which serve, predictably, to bring words closer to the line of best fit). Instead, each nonce word appears twice, once for its trials with stem stressed plural with *-i* (black) and once with suffixed stressed plural (gray), where the suffix can be either *-i* or *-a*. Since the majority of trials have suffix stress, the gray words are generally slightly larger than the black ones. Although the size of the stress effect in Table 6.22 is substantially higher than that of phonology, we can see that the effect of dorsal consonant is considerably greater than that of stress. This is because the effect size of dorsals in the phonological model of the lexicon, in Table 6.17, is 6.39. Even when this is compressed considerably by the effect size of *phon_odds* in Table 6.22, the resulting effect of dorsals in this model is $6.39 \cdot .20 = 1.28$, which is twice as large as the effect of suffix stress (.64). The graphical interpretation in Figure 6.5 is that the horizontal distance between a word with stem stress (black) and suffix stress (gray) is about half as large as the distance between the cluster of dorsal-final nouns and the cluster of nouns ending in other consonants. For most of the words, the gray is higher than the black, and the lines connecting them slope upward and to the right, indicating that most nouns, as predicted, were assigned *-ók* more often when presented with suffix stress. The lines connecting these graphs are light for the purpose of legibility.

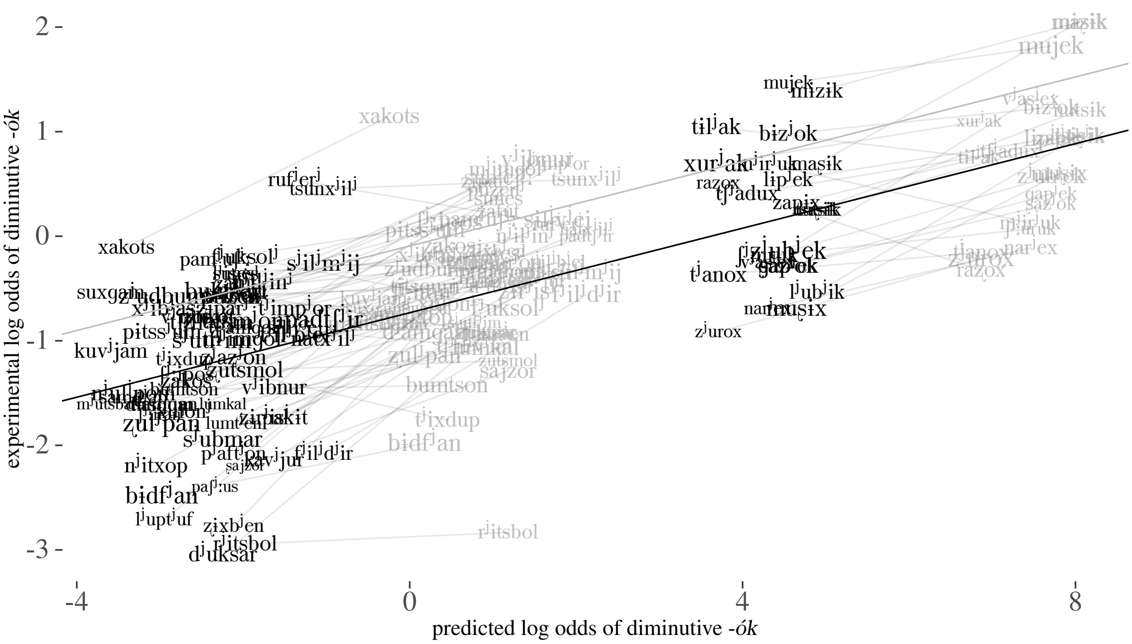

Figure 6.5: The relationship between predicted and experimental log odds of diminutive *-ók* for individual nonce words with stem stress (black) and suffix stress (gray) in the plural, sized according to number of trials, with a line showing the fit of the experimental model in Table 6.17

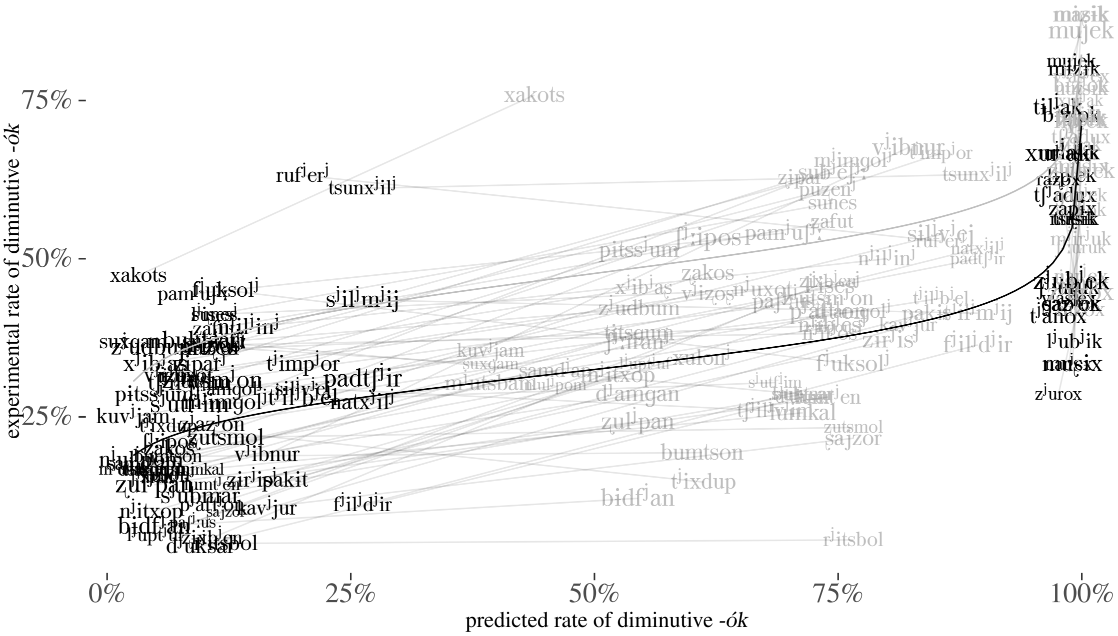

Figure 6.6: The relationship between predicted likelihood and experimental rate of diminutive *-ók* for individual nonce words with stem stress (black) and suffix stress (gray) in the plural, sized according to number of trials, with a line showing the fit of the experimental model in Table 6.17

Let us now look at the effect of plural suffix. Stimuli presented with plural *-a* were given *-ók* at a higher rate than average, but this was due in large part to the fact that this plural suffix is obligatorily stressed. According to Table 6.22, there is a significant difference: stimuli presented with plural *-a* were given *-ók* more often than those shown with plural *-i*, *even once stress is taken into account*. To isolate the effects of the suffix itself, we must compare these trials to trials where stimuli were presented with stressed suffix *-i* in the plural. We see this in Figure 6.7 and Figure 6.8. These graphs only include trials where the plural was presented with suffix stress. In black are trials with stressed plural *-i* (15 per participant), while trials with plural *-a* (10 per participant) are in gray. This comparison is much less clear than the previous one: there is not much separation between the two colors, and individual words can go in either direction. There is nonetheless a visible pattern: many of the words (likely the majority) have lines pointing upward and to the right, suggesting that, on the whole, stimuli tended to take *-ók* more often when presented with plural *-a*. It is unclear why the results are so noisy: perhaps the underlying effect is small enough that a clear item-by-item result requires a larger number of participants.

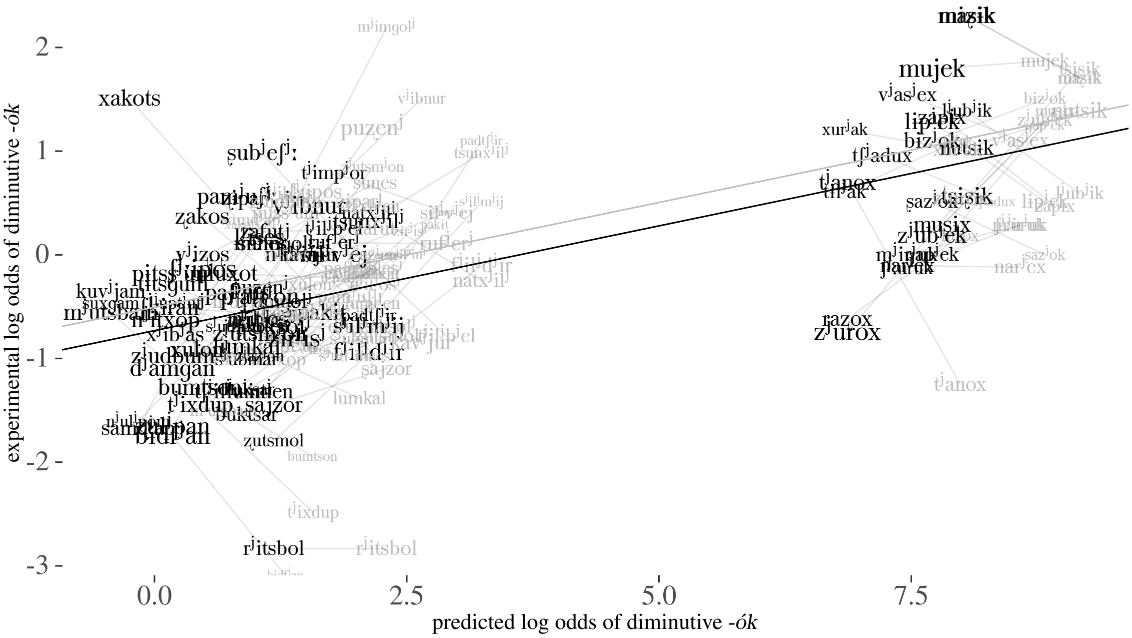

Figure 6.7: The relationship between predicted and experimental log odds of diminutive -*ók* for individual nonce words with stressed plural -*i* (black) and stressed plural -*a* (gray), sized according to number of trials, with a line showing the fit of the experimental model in Table 6.17 (trials with plural stem stress omitted)

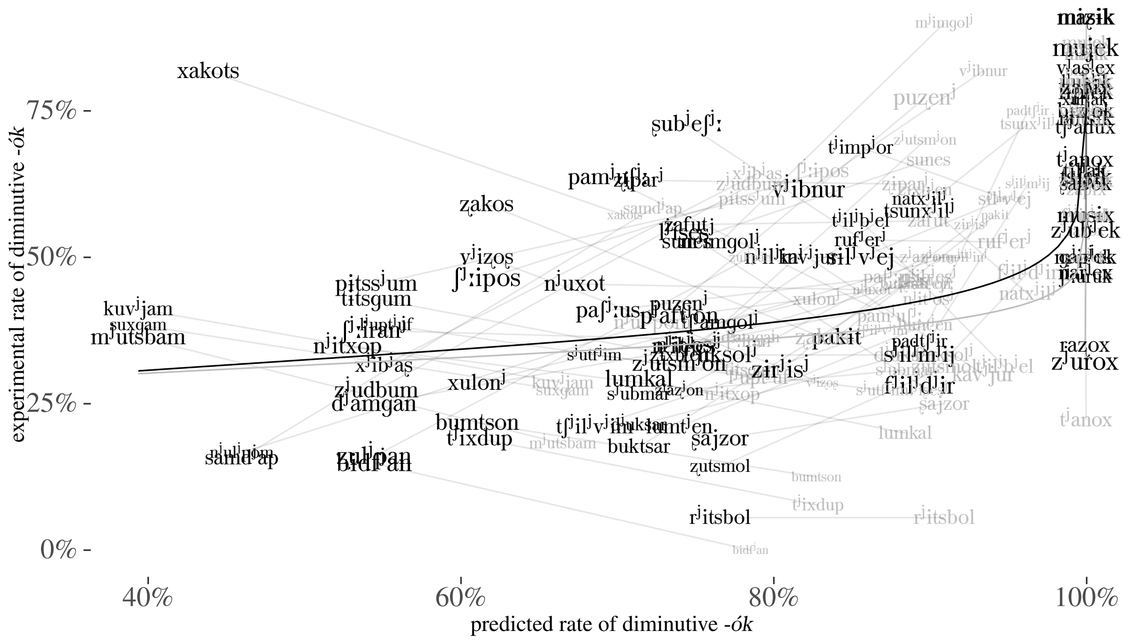

Figure 6.8: The relationship between predicted likelihood and experimental rate of diminutive *-ók* for individual nonce words with stressed plural *-i* (black) and stressed plural *-a* (gray), sized according to number of trials, with a line showing the fit of the experimental model in Table 6.17 (trials with plural stem stress omitted)

6.4.5.3 Specific phonological effects

The figures above show that the *phon_odds* factor is predominated by the effect of word-final dorsals: in both the lexicon and the experimental results, words ending in dorsals, like the nonce word [tʲanóx] receive *-ók* ([tʲanoʂók]) much more often than others. The model of the lexicon in Table 6.17 has other significant effects as well: nouns ending in labials take *-ók* less than those ending in coronals, while nouns with low and rounded vowels in the last syllable take *-ók* less often than those with mid and unrounded vowels, respectively. (Other significant effects are not tested in this study due to the choice of stimuli.) Do speakers in the experiment apply these individual phonological effects as well, or only the dorsal effect?

To answer this question, I built a model with phonological and morphological factors similar to that in Table 6.22, but with individual phonological factors as candidates rather than the omnibus *phon_odds*. Only four of the phonological variables tested in the lexicon model varied among my stimuli: final C place and manner, and final V height and rounding. These were the candidate phonological variables used in the model built up by stepwise comparison. The resulting model, which also includes random intercepts for participant and item, is shown in Table 6.24; it contains the two morphological factors as above (stress pattern and plural suffix), presented order, and all four phonological factors. The non-phonological effects are very similar in size to those in Table 6.22, and the phonological effects are all in the same direction as those in the phonological model of the lexicon in Table 6.17 except for that of approximants, which is significant in neither model. The dorsal effect is much smaller here than in the lexicon model (.90 vs. 6.39, or .14 times as large, not far from the .20 effect size of *phon_odds* in Table 6.22). Most other effects are also smaller than in the lexicon, but the difference is more moderate: for example, the effects of low vowels and rounded vowels in the lexicon model in Table 6.17 are $-.87$ and $-.59$, respectively, compared to $-.36$ and $-.44$ in the experimental model. The effect of final consonant manner is about as strong, if not somewhat stronger, in the experimental results, Here the effect of nasals and affricates reach significance, whereas in the lexicon model none of the manner effects are significant.

Random effects	*variance*	*SD*		
Participant	0.38	.62		
Item	0.19	.44		
Fixed effects	*β coef*	*SE*	*Wald z*	*p*
Intercept	**−0.65**	**.22**	**−3.02**	**.0025**
Plural stress (default: stem)				
suffix	**0.63**	**.08**	**8.03**	**<.0001**
C Place (default: alveolar)				
labial	−0.25	.21	−1.21	.2279
dorsal	**0.90**	**.19**	**4.87**	**<.0001**
Plural suffix (default: -i)				
-a	**0.23**	**.08**	**2.73**	**.0063**
V Height (default: mid)				
high	0.22	.14	1.63	.1037
low	**−0.36**	**.18**	**−2.01**	**.0441**
V Round (default: unrounded)				
rounded	**−0.44**	**.14**	**−3.19**	**.0014**
Presented order	−0.08	.04	−1.84	.0657
C Manner (default: plosive)				
affricate	**1.39**	**.58**	**2.38**	**.0174**
fricative	−0.21	.18	−1.14	.2544
nasal	**−0.45**	**.21**	**−2.13**	**.0330**
approximant	−0.31	.22	−1.43	.1521

Table 6.24: Effects of mixed logistic model with phonological and morphological predictors for experimental selection of diminutive *-ók*, with significant effects bolded

This model shows that the *phon_odds* effect in Table 6.22 is not driven entirely by dorsals: speakers have also learned patterns from the lexicon about the final vowel and the manner of the final consonant. While these effects are still substantially smaller than the dorsal effect, their strength in the experimental results is closer to their strength in the lexicon.

6.4.5.4 Individual variation in the plural suffix effect

I predicted that individual speakers would vary in whether they treated plural *-a* as a signal for *-ók* beyond the effects of stress pattern. The lexicon is compatible with two interpretations, one in which plural *-a* is itself a predictor of *-ók* and one in which the level of *-ók* among nouns with plural *-a* is due entirely to its suffix stress. Accordingly, we should expect some speakers to show sensitivity to the plural suffix beyond the effect of stress, while others should treat all words with stressed plurals the same way. This cannot be detected by a high by-participant random slope for plural, since mixed-effects regressions assume that random effects are sampled from a single population (see Houghton & Kapatsinski, 2023). As discussed previously, this random slope did not improve the model and was omitted from it. Thus, we must use other means to investigate the prediction.

Let us look at the behavior of individual speakers more closely. Again, the prediction is that some speakers have a correlation between plural *-a* and diminutive *-ók*, while others do not. On the other hand, all speakers are predicted to have a correlation between plural suffix stress and diminutive *-ók*. Accordingly, we should see a relative difference between the two comparisons in the distribution of individual speakers.

In Figure 6.9, we see the distribution of how strong the effect of stress pattern was for individual participants.[7] This graph groups participants by the difference in rate of *-ók* assigned to stimuli presented with stem stressed plural *-i* and suffix stressed plural *-i*, respectively (trials with plural *-a* are omitted). For example, one participant selected *-ók* for 7/15 trials in the stem stressed plural *-i* condition (46.7%) and 9/15 trials in the suffix stressed plural *-i* condition (60%). Thus, this participant is placed into the bin for $60 - 46.7 = 13.3\%$, which is the slight trough between the two peaks. The distribution appears roughly unimodal, with a notional peak indicating that most participants assigned *-ók* 10–20% often to stimuli with suffix stress. This distribution is what we would expect from a constant effect with some amount of random noise (in this case, possibly quite a large amount given the relatively small number of trials for each participant).

[7] In these graphs, all trials are included, even those with the one cluster-final nonce word omitted elsewhere. This makes the numbers more even, as all denominators are the same for all participants.

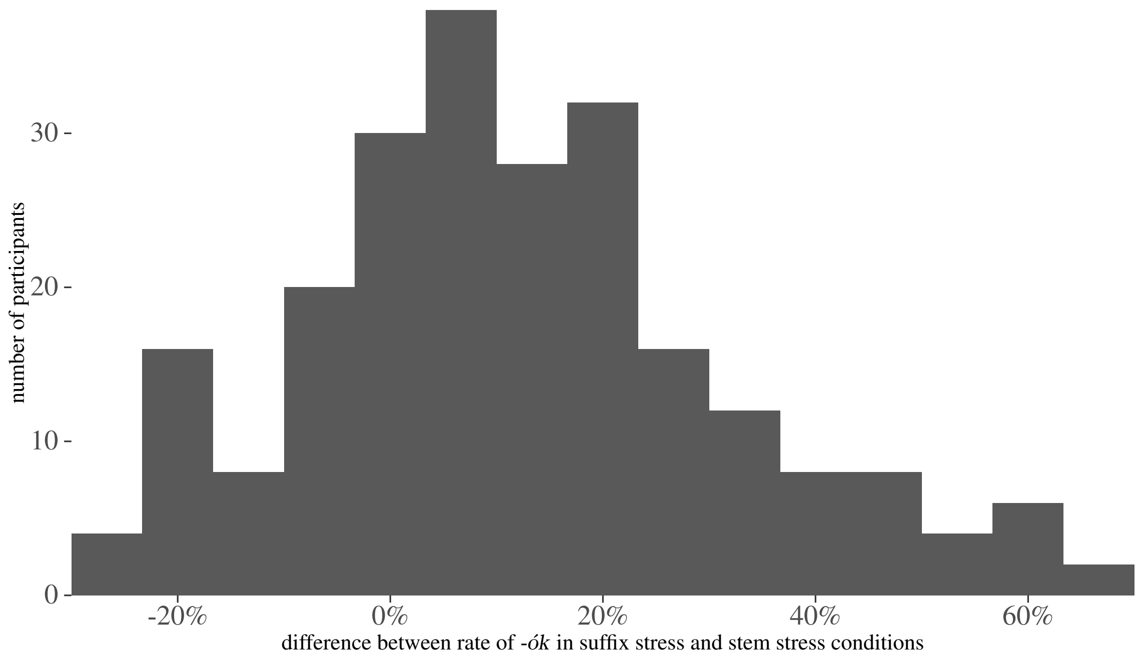

Figure 6.9: Distribution of participants by effect size of plural stress (difference between each participant's rates of -*ók* assigned to stimuli with stem stressed plural -*i* and suffix stressed plural -*i*)

I now investigate whether the distribution in Figure 6.9 is normal and unimodal using statistical tests. First, the Kolmogorov–Smirnov test (Massey, 1951), implemented as the `ks.test` function of the `stats` package in R (R Core Team, 2022), tests whether a given collection of data points is drawn from a given distribution—in this case, the normal distribution. According to this test, the stress differences shown in Figure 6.9 do not differ significantly from a normal distribution ($D = .11$, $p = .134$): we cannot reject the null hypothesis that this distribution is normal. (If we test the by-participant random slopes for stress pattern, which were omitted from the model in Table 6.22 because they did not improve it, we get a similar result.) To test whether the data are unimodal (have a single peak), I use the folding test (Siffer et al., 2018) implemented using the `folding.test` function from the R package `Rfolding` (Siffer, 2018). This test assesses whether a data distribution is compatible with the null hypothesis of *multimodality*—that is, the null hypothesis is that the data have several peaks. The folding test concludes that the distribution of differences is unimodal ($p = .0001$): that is, there is only one peak in the data, as expected. (The result for the random slopes for stress pattern is similar.)

Let us now compare this distribution with the graph showing the effect of plural suffix, shown in Figure 6.10. This graph shows the difference in how often individual speakers assigned *-ók* to stimuli presented with suffix stressed plural *-i* and suffix stressed plural *-a*, respectively (trials with stem stress are omitted). Here, the distribution is predicted to be *bimodal*, with one peak around 0% (for those speakers who have not learned a correlation between plural *-a* and diminutive *-ók*) and another peak somewhere above 0% (for speakers who have learned such a correlation). However, this is not what we see in Figure 6.10. Instead, we see a roughly unimodal distribution. The general shape of this distribution is roughly similar to that of Figure 6.9, although its peak is closer to 0%. This means that the effect size of plural suffix is smaller than that of stress pattern, but there is no evidence that the underlying distribution is bimodal.

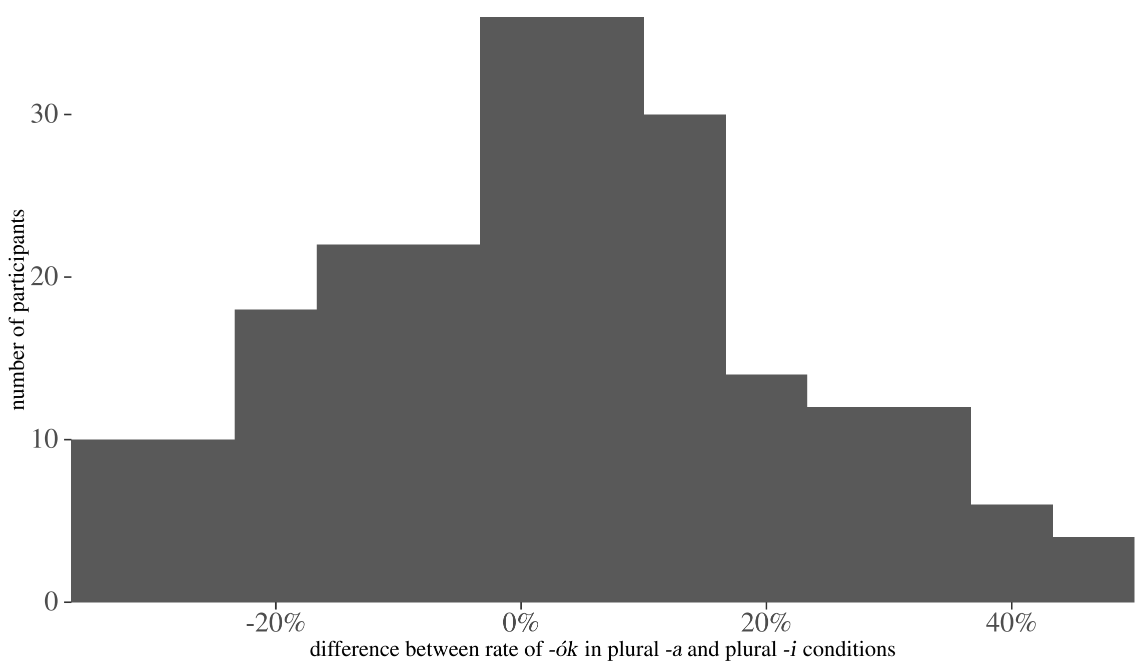

Figure 6.10: Distribution of participants by effect size of plural suffix (difference between each participant's rates of *-ók* assigned to stimuli with suffix stressed plural *-i* and suffix stressed plural *-a*)

The Kolmogorov–Smirnov test for normality suggests that the data in Figure 6.10 is compatible with a normal distribution ($D = .06$, $p = .796$): there is no evidence to reject the null hypothesis that the suffix differences are normally distributed (and likewise the random slopes for suffix). Similarly, the folding test concludes that the suffix differences (and the random slopes) follow a unimodal distribution ($p = .0001$). These tests suggest that the strength of the plural suffix effect does not follow a bimodal distribution. However, we should be cautious in rejecting our hypothesis. The effect in Figure 6.10 is, in general, smaller than the stress pattern effect. If the two peaks are very close together, with one peak at 0% (speakers with no difference) and another around 5–10% (speakers with a small difference), we would get something very similar to what we see in Figure 6.10: we do not have the resolution (in particular, the number of trials per participant) to distinguish between a unimodal distribution and a bimodal one with two peaks very close together. However, we certainly do not have any positive evidence to reject a normal unimodal distribution: it seems that all speakers have learned the correlation between plural *-a* and diminutive *-ók* equally.

6.4.5.5 Difficulty of choosing a diminutive

One of the nuisance variables tested in the models above was the answer to the question of whether the choice between diminutives was hard. This factor did not make it into any of the models, but here, I explore the question of what made the choice of diminutive difficult. As with the other models, the model in Table 6.25 was built using stepwise comparison with random intercepts for participant and item. In addition to *phon_odds* and morphological nuisance variables, I included selected diminutive as a candidate factor. I set *-ók* as the baseline level, so that I could directly compare *-ók* with each of the other two diminutives. The factors that make it into the model are those that have to do with the task rather than the stimuli themselves. The most salient factor is the random intercept for participant. Participants varied greatly in how often they marked the choice as hard: 24 participants said the choice was never hard, and 77 of 116 said that ten or fewer choices were difficult. However, one participant marked the choice as difficult in 30 trials (three quarters of trials). The other significant factors are task effects. Diminutive order marks whether the chosen diminutive was listed first, second, or third among the options: the lower it was listed, the more likely participants were to mark the choice as difficult. The factor of trial number suggests that participants got more accustomed to the task as it went on, though this effect is very small and not significant: they were more likely to mark earlier trials as difficult than later trials. Finally, the choice of diminutive had an effect: choosing *-ʲik* was deemed harder than choosing *-ók*. This may be due to the particular design of the study, which included more words expected to take *-ók* than *-ʲik*, while *-tʃʲik* is generally a safe choice for unfamiliar words.

Random effects	*variance*	*SD*		
Participant	2.57	1.60		
Item	0.09	.29		
Fixed effects	*β coef*	*SE*	*Wald z*	*p*
Intercept	**–1.93**	**.19**	**–10.28**	**<.0001**
Presented order	**0.15**	**.05**	**2.95**	**.0031**
Trial number	–0.01	.00	–1.84	.0652
Chosen diminutive (default: -ók)				
-ʲik	**0.23**	**.10**	**2.24**	**.0254**
-tʃʲik	–0.01	.11	–0.09	.9320

Table 6.25: Effects of mixed logistic model with predictions of whether the choice of diminutive was hard, with significant effects bolded

The variance of the random intercept for item is quite small, suggesting that individual stimuli were roughly equally difficult in their trials. To check for phonological factors, I built a second model without the random intercept for item and with individual phonological factors as candidates (final C place and manner, final V height and rounding). Only consonant manner was added to the model, with one significant effect: stimuli ending in nasals, like [xulónʲ], were deemed more difficult than those ending in plosives. This second model was a slightly worse fit than that of Table 6.25, though not significantly so. Thus, while there are some phonological tendencies towards predicting difficulty in choosing a diminutive, these effects are quite small.

6.4.6 Discussion

6.4.6.1 Inflectional patterns

The main goal of this experiment was to test whether Russian speakers learn and apply the inflectional dependencies for the diminutive suffix *-ók*: in the lexicon, *-ók* appears more often with words that have suffix stress in at least some forms, including one small subset whose nominative plural is stressed *-a*, not the usual plural marker *-i*. The effect of stress pattern is clear in the lexicon, so speakers are expected to apply it across the board. Indeed, this is what we find: overall, participants assigned *-ók* significantly more often to nonce words presented with suffix stress in the plural than to nonce words presented with fixed stem stress. If all speakers have a similar underlying effect size for stress pattern, we would expect the difference between the rate of *-ók* in the stem stressed plural *-i* condition and the suffix stressed plural *-i* condition for individual participants to follow a normal distribution, and this is what we see in Figure 6.9.

The predicted effect of plural suffix was different. While nouns with plural *-a* take *-ók* far more often than nouns with plural *-i*, a large part of this effect is due to the fact that plural *-a* obligatorily takes stress, meaning that all such nouns have at least some inflected forms with suffix stress. In the statistical model of the lexicon in Table 6.17, plural suffix had an effect over and beyond that of stress pattern, but this was not significant. Thus, there are two grammars that are roughly compatible with the input: one in which the effect of plural *-a* compounds that of suffix stress (that is, both morphological factors are encoded), and one which only encodes suffix stress, such that the plural *-a* effect is derived entirely from the asymmetry in stress patterns for such nouns, as discussed in Section 6.3.4.3. This means that some speakers should treat nonce words with suffix stressed plural *-a* the same as nonce words with suffix stressed plural *-i*, while others should assign *-ók* more often to nouns with plural *-a*. The distribution of individual speakers' difference in *-ók* rates for the two conditions is predicted to be *bimodal*, with one peak at 0 (no effect, the two conditions are equal) and one above 0 (higher rate of *-ók* for plural *-a*). However, the actual distribution, shown in Figure 6.10, seems to be normal and have only one peak. This peak is closer to 0 than that of the stress pattern effect shown in Figure 6.9, suggesting a smaller effect size. Contrary to the prediction, we do not find evidence that individual speakers are treating the influence of plural *-a* on diminutive *-ók* differentially. However, we do confirm the basic hypothesis that lexically specific inflectional factors—stress pattern and plural suffix—actively influence diminutive choice in the grammar of Russian speakers. In Russian, as in Hungarian and Czech, we have evidence that speakers learn and apply morphological dependencies.

6.4.6.2 Comparison with previous studies

The morphological effects described in the previous section have not been tested in earlier nonce word studies of the Russian diminutive (Gouskova et al., 2015; Magomedova, 2017). Those studies found that *-ók* was selected less often than expected, suggesting that it is losing productivity. In my study, however, speakers used *-ók* freely—in fact, it was the most common diminutive selected. This suggests that earlier studies may have undercounted the use of *-ók* by omitting some of the factors that are most closely associated it. Given evidence of suffix stress or plural *-a*, Russian speakers are happy to productively extend *-ók*. These morphological factors may themselves not be particularly productive, which would explain the fairly limited reach of *-ók* to new words—this continues a trend, from the other studies in this dissertation, that less productive minority patterns tend to cluster together.

While this study can only partially replicate the phonological effects of Gouskova et al. (2015) and Magomedova (2017), I briefly compare the results here. Magomedova (2017) tested speakers on ten monosyllabic nonce words presented in different semantic contexts (see Section 6.1.3). Four of these nouns ended in velars (favoring *-ók*), four ended in fricatives (favoring *-ʲik*), and two ended in nasals (favoring *-tʃʲik*). Although her participants selected *-ók* for velar-final nouns more often than for other nouns, *-ʲik* was still selected more often than *-ók*. My results were different: *-ók* was assigned over twice as often to velar-final nouns than -^{j}*ik* was (755 vs. 304), and even in trials where the nonce words had fixed stem stress (for which *-ók* is less likely), *-ók* was almost twice as common as -^{j}*ik* for velar-final nouns (248 vs. 132). The difference between her study and mine may be, at least in part, a task effect: my design included stress shifts and stimuli presented auditorily, so participants may have been nudged to focus more on stress, a factor saliently implicated in the use of *-ók*, which in turn may have pushed up rates of *-ók* higher than expected.

The stimuli used in my nonce word study comprised a subset of those used in Gouskova et al. (2015). In particular, all of my words are disyllabic with stress on the final syllable, end in singleton consonants, and have no word-internal hiatus. Of the hand-picked phonological factors studied in Gouskova et al. (2015, p. 67), I only test two: they found that nouns ending in dorsals prefer *-ók* and disprefer -^{j}*ik* and -*tʃ*^{j}*ik*, while nouns ending in sonorants prefer *-ók* and -*tʃ*^{j}*ik* and disprefer -^{j}*ik*. The model in Table 6.24 breaks the *phon_odds* coefficient into individual phonological factors. This model confirmed the preference of dorsals for *-ók*; the factor of final consonant manner was not added to the model, so the preferences of sonorants (or a subset) were not confirmed. A similar model for -*tʃ*^{j}*ik* confirms the effect of sonorants: nouns ending in nasals and approximants take -*tʃ*^{j}*ik* more than those ending in plosives. However, the dorsal effect was not replicated, as final consonant place was not added to the model. Likewise, a model for -^{j}*ik* includes three phonological factors: final consonant manner, vowel height, and vowel rounding. Nouns ending in fricatives prefer -^{j}*ik* relative to those ending in plosives—a preference also found by Magomedova (2017)—while nouns ending in nasals disprefer -^{j}*ik*. Previous studies found that nouns with [i] disprefer -^{j}*ik*; I found, somewhat similarly, that nouns with low vowels prefer -^{j}*ik* relative to those with mid vowels; the dispreference for -^{j}*ik* nouns with high vowels was not significant. Finally, I found a previously unreported result that matches the lexicon: nouns with rounded vowels in the last syllable significantly prefer -^{j}*ik*.

My study also yielded two other significant phonological effects for *-ók*: nonce words with low vowels in the last syllable took *-ók* more than those with mid vowels, and likewise, nonce words with unrounded vowels took *-ók* more than those with rounded vowels. Both of these effects match my model of the lexicon in Table 6.17. However, the corpus study of Gouskova et al. (2015) found that nouns with rounded vowels take *-ók more* often than those with unrounded vowels, the opposite of both my corpus and nonce word studies. This suggests that my corpus is a better representation of the lexicon than that of Gouskova et al. (2015), because its results more closely reflect the generalizations that Russian speakers have learned.

In general, this nonce word study yields similar phonological results to previous nonce word studies of the Russian diminutive. The primary difference is the *morphological* effects of suffix stress and plural *-a*, which are both correlated with *-ók*. The failure of previous studies to look at these factors may have contributed to the relative lack of productivity reported for *-ók*. A fuller accounting of the factors associated with the various diminutives suggests that *-ók* is still quite productive, but in contexts that were previously overlooked.

6.5 General discussion and summary

This study looked at two properties of Russian inflectional paradigms associated with the choice of Russian diminutive for masculine inanimate nouns: first, a minority of nouns with suffix stress in their inflectional paradigm show a strong preference for *-ók*, a stressed allomorph of the diminutive, over two other options, unstressed *-ʲik* and *-tʃʲik*. Inflectional stress is usually analyzed with underlying stress markers, so this can be captured as a correlation between words that have a diacritic feature governing diminutive *-ók* and words that have particular kinds of stress markers (or a lack thereof) in their underlying form, as described in Section 6.2.2.6. Speakers were expected to learn this pattern, and did: participants in a nonce word study assigned *-ók* more often to words presented with plurals bearing suffix stress.

One subset of nouns with suffix stress is particularly likely to take diminutive *-ók*: those with irregular plural *-a* instead of the regular *-i*. While *-i* is compatible with any stress pattern, *-a* only appears with nouns that stress the suffix throughout the plural. Following the observation of Divjak et al. (2016) that the input is often compatible with multiple grammars, I predicted variation in the effect of plural suffix: some speakers should have a correlation between plural *-a* and diminutive *-ók* in addition to the effect of stress pattern, while others should only encode the stress pattern effect and subsume the plural *-a* effect into the general effects of stress. In the nonce word study, speakers consistently assigned *-ók* more often to nouns presented with stressed plural *-a* than those with stressed plural *-i*, suggesting that they had directly learned a correlation between *-a* and *-ók*. I found evidence suggesting that speakers uniformly learned this effect. Since this pattern must be learned separate from stress, it cannot be encoded as a paradigm uniformity effect for stress: there must be a separate correlation between the two paradigm cells that cannot be reduced to shared properties.

The Russian diminutive study makes several contributions to the program pursued in this dissertation. First, it shows that speakers learn morphological dependencies between properties of a word's inflectional patterns and an affix that is traditionally seen as derivational. Second, my results show that speakers robustly learn salient correlations (in this case, between plural *-a* and diminutive *-ók*) even when they can arguably be reduced to other factors (the correlation between stressed suffixes and diminutive *-ók*). Third, if stress pattern is analyzed (as is typical) with underlying stress marking on roots, then this study shows that source-oriented generalizations can be made over underlying forms (with stress marks) rather than unsuffixed base surface forms (which are often ambiguous between multiple stress patterns). Finally, while the effect of stress pattern can be plausibly explained through paradigm uniformity, the plural effect cannot. Speakers nonetheless learn it, showing that morphological dependencies can be learned even if they are not grounded in uniformity constraints.

7 Conclusion

The main goal of this dissertation was to extend the systematic study of productivity to a new domain: morphological dependencies. That is, speakers use information within known suffixed forms of novel words to productively extend correlations between lexically specific patterns to unknown forms. These results show the general validity of the observation that speakers tend to extend patterns from their lexicon to nonce words, stochastically matching the distribution of alternants in the lexicon (e.g. Ernestus & Baayen, 2003; Gouskova et al., 2015; Hayes et al., 2009). This has previously been applied to productive extension of *phonological* patterns to nonce words; my studies show that it holds for morphological dependencies as well. These results also complement work studying the efficient organization of morphological paradigms (e.g. Ackerman & Malouf, 2013; Ackerman et al., 2009; Bonami & Beniamine, 2016). This line of research has studied morphological dependencies in the lexicon, but has not tested whether speakers use these dependencies to efficiently fill in gaps in their knowledge of the morphological behavior of individual words. This dissertation shows that they do.

7.1 Case studies

Each of the three case studies in this dissertation uses a novel nonce word experimental paradigm to show that speakers have learned a morphological dependency: a correlation in the lexicon between two lexically specific patterns. This productive extension of lexical patterns is robust across different languages and configurations. In Chapter 4, I showed that Hungarian speakers have learned the tendency for a small class of irregular nouns, known as “lowering stems” for their conditioning of low vowels in certain inflectional suffixes, to occur with the possessive suffix *-ɒ*, not the other variant, *-jɒ*. Hungarian possessive allomorphy also observes several phonological patterns, and speakers applied these patterns to nonce words as well. This study shows that speakers consider phonological and morphological patterns side-by-side in determining the behavior of nonce words.

Chapter 5 studied a morphological dependency in Czech. Most nouns in a particular inflection class (hard-stem inanimates) have *-u* in both the genitive and locative case in the singular, but both cases have an alternative suffix: *-a* in the genitive and *-ɛ* in the locative. In the lexicon, words that take *-a* in the genitive are also more likely to take *-ɛ* in the locative, and speakers productively applied this correlation to nonce words as well: their choice of locative for a nonce word was influenced by the genitive with which that word was shown, mediated by the genitive they chose for that word. When the experiment was repeated with *real* words that are *variable* in one or both cases, speakers’ choice of locative for a word was not influenced by the genitive with which that word was presented. These two results, when taken together, indicate that the experimental correlation between genitive and locative really is located in the productive extension of lexical patterns to nonce words, rather than at some other point (for example, priming between forms suffixed with *-u*). This chapter concluded with a corpus study showing that individual authors have a correlation in their use of genitive and locative suffixes for a given word. This shows that speakers have a bias in their production between genitive and locative; if this bias is not present in their input (to be tested at a later date), the bias must have come from a morphological dependency in their grammar.

Finally, in Chapter 6, I looked at inflectional factors influencing the choice of a Russian derivational suffix, the diminutive. One of the three common diminutive allomorphs for Russian class I nouns, *-ók*, is substantially preferred by nouns that have certain stress patterns (either fixed on the suffix or alternating between stem and suffix) and nouns that take the uncommon plural *-a* rather than the regular plural *-i*. Stress pattern is typically analyzed in Russian through abstract phonology (underlying stress marks), so the correlation between stress pattern and diminutive must make reference to underlying forms. In a nonce word study, speakers productively extended both the stress pattern correlation and the plural correlation, assigning *-ók* more often to words presented with a stressed plural suffix, especially when that plural was *-a*. This study shows that morphological dependencies are not limited to inflectional paradigms: speakers also extend correlations between inflectional and derivational suffixes. It also shows that speakers learn morphological dependencies that cannot be described in terms of paradigm uniformity, or identity between related forms: nouns with both plural *-a* and diminutive *-ók* do not necessarily have more similar plural and diminutive forms than nouns that have a different plural or a different diminutive.

7.2 Modelling morphological dependencies

The results of my three studies present a broad overall picture of how how speakers productively apply patterns to new words: they can learn arbitrary phonological and morphological factors influencing the distribution of allomorphs and weigh them against one another. These patterns can apply to any complex form built off a stem, inflectional or derivational, and need not be grounded in a particular notion of phonological or morphological "naturalness", as described in Section 2.3.

These results can be accounted for by the sublexicon model (cf. Allen & Becker, 2015; Becker & Gouskova, 2016; Gouskova et al., 2015), presented in Chapter 3, a flexible approach that allows speakers to learn a relatively unrestricted range of phonological and morphological dependencies and apply them stochastically to new words. In this model, lexically conditioned patterns of allomorphy are indexed with diacritic features, and speakers learn generalizations over lexical items that share a feature; in particular, morphological dependencies are correlations between two features that tend to cooccur on the same lexical items. These generalizations are stored as constraints in phonotactic grammars: each feature defines a sublexicon comprised of words that have that feature, and each sublexicon has its own constraint-based phonotactic grammar exemplifying what it means to be a typical word of that sublexicon. When a new lexical item does not have a listed feature determining its behavior, speakers evaluate it on the relevant sublexical grammars to decide which feature to assign to it.

My implementation of the sublexicon model combines with Distributed Morphology, a theory in which words are composed of smaller pieces arranged in a syntactic structure. As discussed in Section 2.1.5, this model shows that piece-based theories of morphology can also allow speakres to learn paradigmatic relations between forms built off of the same root, responding to criticism of piece-based models of morphology from Ackerman and Malouf (2013) and others. Beyond this general theoretical point, many of the particular architectural choices of the sublexicon model are not directly tested by the studies in this dissertation. Nonetheless, the empirical and theoretical work presented here should serve as a proof of concept for nonce word studies testing the psychological reality of morphological dependencies and inspiration for further studies exploring the kinds of morphological dependencies speakers learn and refining the theoretical framework used to account for them.

Bibliography

Ackerman, F., Blevins, J. P., & Malouf, R. (2009). Parts and wholes: Implicative patterns in inflectional paradigms. In J. P. Blevins & J. Blevins (Eds.), *Analogy in grammar: Form and acquisition* (pp. 54–82). Oxford University Press. https://doi.org/10.1093/acprof:oso/9780199547548.003.0003

Ackerman, F., & Malouf, R. (2013). Morphological organization: The low conditional entropy conjecture. *Language*, *89*(3), 429–464. https://doi.org/10.1353/lan.2013.0054

Adger, D., & Smith, J. (2010). Variation in agreement: A lexical feature-based approach. *Lingua*, *120*(5), 1109–1134. https://doi.org/10.1016/j.lingua.2008.05.007

Albright, A., & Hayes, B. (2003). Rules vs. analogy in English past tenses: A computational/experimental study. *Cognition*, *90*(2), 119–161. https://doi.org/10.1016/s0010-0277(03)00146-x

Alderete, J. D. (1999, May). *Morphologically governed accent in Optimality Theory* [Doctoral dissertation, University of Massachusetts Amherst].

Allen, B., & Becker, M. (2015). *Learning alternations from surface forms with sublexical phonology* [lingbuzz/002503].

Anagnostopoulou, E., & Samioti, Y. (2013). Allosemy, idioms, and their domains: Evidence from adjectival participles. In R. Folli, C. Sevdali, & R. Truswell (Eds.), *Syntax and its limits* (pp. 218–250). Oxford University Press. https://doi.org/10.1093/acprof:oso/9780199683239.003.0012

Aronoff, M. (1994). *Morphology by itself: Stems and inflectional classes*. MIT Press.

Arsenijević, B. (2021). No gender in 'gender agreement': On declension classes and gender in Serbo-Croatian. *Balcania et Slavia*, *1*(1), 11–46. https://doi.org/10.30687/bes/0/2021/01/001

Baerman, M., Brown, D., & Corbett, G. G. (2017). *Morphological complexity*. Cambridge University Press. https://doi.org/10.1017/9781316343074

Bartos, H. (1999). *Morfoszintaxis és interpretáció. a magyar inflexiós jelenségek szintaktikai háttere* [Doctoral dissertation, Eötvös Loránd University].

Baunaz, L., & Lander, E. (2018). Nanosyntax: The basics. In L. Baunaz, K. De Clercq, L. Haegeman, & E. Lander (Eds.), *Exploring nanosyntax* (pp. 3–56). Oxford University Press. https://doi.org/10.1093/oso/9780190876746.003.0001

Becker, M. (2009, February). *Phonological trends in the lexicon: The role of constraints* [Doctoral dissertation, University of Massachusetts Amherst].

Becker, M., & Gouskova, M. (2016). Source-oriented generalizations as grammar inference in Russian vowel deletion. *Linguistic Inquiry*, *47*(3), 391–425. https://doi.org/10.1162/ling_a_00217

Becker, M., Ketrez, N., & Nevins, A. (2011). The surfeit of the stimulus: Analytic biases filter lexical statistics in Turkish laryngeal alternations. *Language*, *87*(1), 84–125. https://doi.org/10.1353/lan.2011.0016

Benua, L. (1997). *Transderivational identity* [Doctoral dissertation, University of Massachusetts Amherst].

Benus, S., & Gafos, A. I. (2007). Articulatory characteristics of Hungarian 'transparent' vowels. *Journal of Phonetics*, *35*, 271–300.

Berko, J. (1958). The child's learning of English morphology. *Word*, *14*(2–3), 150–177. https://doi.org/10.1080/00437956.1958.11659661

Bermel, N., & Knittl, L. (2012). Morphosyntactic variation and syntactic constructions in Czech nominal declension: Corpus frequency and native-speaker judgements. *Russian Linguistics*, *36*(1), 91–119. https://doi.org/10.1007/s11185-011-9083-x

Bermúdez-Otero, R. (2012). The architecture of grammar and the division of labor in exponence. In J. Trommer (Ed.), *The morphology and phonology of exponence* (pp. 8–83). Oxford University Press. https://doi.org/10.1093/acprof:oso/9780199573721.003.0002

Bermúdez-Otero, R. (2013). The Spanish lexicon stores stems with theme vowels, not roots with inflectional class features. *Probus: International Journal of Romance Linguistics*, *25*(1), 3–103. https://doi.org/10.1515/probus-2013-0009

Bobaljik, J. D. (2000). The ins and outs of contextual allomorphy. In K. Grohmann & C. Struijke (Eds.), *University of Maryland working papers in linguistics* (pp. 35–71).

Boersma, P., & Pater, J. (2016). Convergence properties of a gradual learning algorithm for Harmonic Grammar. In J. J. McCarthy & J. Pater (Eds.), *Harmonic Grammar and Harmonic Serialism* (pp. 389–434). Equinox.

Bolozky, S., & Becker, M. (2006). Living lexicon of Hebrew nouns [University of Massachusetts Amherst].

Bonami, O., & Beniamine, S. (2016). Joint predictiveness in inflectional paradigms. *Word Structure*, *9*(2), 156–182. https://doi.org/doi.org/10.3366/word.2016.0092

Bonet, E., Lloret, M.-R., & Mascaró, J. (2007). Allomorph selection and lexical preferences: Two case studies. *Lingua*, *117*, 903–927. https://doi.org/10.1016/j.lingua.2006.04.009

Booij, G. (2011). Morpheme structure constraints. In M. van Oostendorp, C. J. Ewen, E. Hume, & K. Rice (Eds.), *The Blackwell companion to phonology*. *Vol. 4*. *Phonological interfaces* (pp. 2049–2069). Wiley-Blackwell. https://doi.org/10.1002/9781444335262.wbctp0086

Brown, D., Corbett, G., Fraser, N., Hippisley, A., & Timberlake, A. (1996). Russian noun stress and network morphology. *Linguistics*, *34*, 53–107. https://doi.org/10.1515/ling.1996.34.1.53

Bybee, J. (1995). Regular morphology and the lexicon. *Language and Cognitive Processes*, *10*(5), 425–455. https://doi.org/10.1080/01690969508407111

Bybee, J. (2001). *Phonology and language use*. Cambridge University Press. https://doi.org/10.1017/cbo9780511612886

Bybee, J. L., & Moder, C. L. (1983). Morphological classes as natural categories. *Language*, *59*(2), 251–270. https://doi.org/10.2307/413574

Bybee, J. L., & Pardo, E. (1981). On lexical and morphological conditioning of alternations: A nonce-probe experiment with Spanish verbs. *Linguistics*, *19*, 937–968. https://doi.org/10.1515/ling.1981.19.9-10.937

Caha, P. (2021). Modeling declensions without declension features: The case of Russian. *Acta Linguistica Academica*, *68*(4), 385–425. https://doi.org/10.1556/2062.2021.00433

Cameron-Faulkner, T., & Carstairs-McCarthy, A. (2000). Stem alternants as morphological stigmata: Evidence from blur avoidance in Polish nouns. *Natural Language and Linguistic Theory*, *18*, 813–835. https://doi.org/10.1023/a:1006496821412

Carstairs-McCarthy, A. (2001). Grammatically conditioned allomorphy, paradigmatic structure, and the Ancestry Constraint. *Transactions of the Philological Society*, *99*(2), 223–245. https://doi.org/10.1111/1467-968x.00081

Chan, E. (2008). *Structures and distributions in morphology learning* [Doctoral dissertation, University of Pennsylvania].

Chomsky, N. (2004). Beyond explanatory adequacy. In A. Belletti (Ed.), *Structures and beyond* (pp. 104–131). Oxford University Press.

Chomsky, N., & Halle, M. (1968). *The sound pattern of English*. Harper & Row.

Chvany, C. V. (1986). Jakobson's fourth and fifth dimensions: On reconciling the cube model of case meanings with the two-dimensional matrices for case forms. In R. D. Brecht & J. S. Levine (Eds.), *Case in Slavic* (pp. 107–129). Slavica Publishers.

Cinque, G. (2013). Cognition, universal grammar, and typological generalizations. *Lingua*, *130*, 150–165. https://doi.org/10.1016/j.lingua.2012.10.007

Clahsen, H. (1999). Lexical entries and rules of language: A multidisciplinary study of German inflection. *Behavioral and Brain Sciences*, *22*(6), 991–1013. https://doi.org/10.1017/s0140525x99002228

Coetzee, A. W., & Pater, J. (2011). The place of variation in phonological theory. In J. Goldsmith, J. Riggle, & A. C. L. Yu (Eds.), *The handbook of phonological theory* (Second, pp. 401–434). Wiley-Blackwell. https://doi.org/10.1002/9781444343069.ch13

Corbett, G. G. (1982). Gender in Russian: An account of gender specification and its relationship to declension. *Russian Linguistics*, *6*, 197–232. https://doi.org/10.1007/bf03545848

Corbett, G. G. (1991). *Gender*. Cambridge University Press. https://doi.org/10.1017/cbo9781139166119

Corbett, G. G. (2009). Canonical inflectional classes. In F. Montermini, G. Boyé, & J. Tseng (Eds.), *Selected proceedings of the 6th décembrettes: Morphology in bordeaux* (pp. 1–11). Cascadilla Proceedings Project.

Corbett, G. G., & Baerman, M. (2006). Prolegomena to a typology of morphological features. *Morphology*, *16*, 231–246. https://doi.org/10.1007/s11525-006-9105-3

Dąbrowska, E. (2001). Learning a morphological system without a default: The Polish genitive. *Journal of Child Language*, *28*(3), 545–574. https://doi.org/10.1017/s0305000901004767

Dékány, É. (2018). The position of case markers relative to possessive agreement: Variation within Hungarian. *Natural Language & Linguistic Theory*, *36*, 365–400. https://doi.org/10.1007/s11049-017-9379-7

de Lacy, P. (2004). Markedness conflation in Optimality Theory. *Phonology*, *21*, 145–199. https://doi.org/10.1017/s0952675704000193

den Dikken, M. (2015). The morphosyntax of (in)alienably possessed noun phrases: The Hungarian contribution. In K. É. Kiss, B. Surányi, & É. Dékány (Eds.), *Approaches to Hungarian*.

Vol. 14. *Papers from the 2013 Piliscsaba conference* (pp. 121–145). John Benjamins. https://doi.org/10.1075/atoh.14.05dik

Dickins, T. (1998). Prepositional vocalization in contemporary Czech. *The Slavonic and East European Review*, *76*(2), 201–233.

Divjak, D., Dąbrowska, E., & Arppe, A. (2016). Machine meets man: Evaluating the psychological reality of corpus-based probabilistic models. *Cognitive Linguistics*, *27*(1), 1–33. https://doi.org/10.1515/cog-2015-0101

É. Kiss, K. (2002). *The syntax of Hungarian*. Cambridge University Press. https://doi.org/10.1017/cbo9780511755088

Elsen, H. (2002). The acquisition of German plurals. In S. Bendjaballah, W. U. Dressler, O. E. Pfeiffer, & M. D. Voeikova (Eds.), *Morphology 2000: Selected papers from the 9th Morphology Meeting, Vienna, 24–28 February 2000* (pp. 117–127). John Benjamins. https://doi.org/10.1075/cilt.218.10els

Embick, D. (2003). Locality, listedness, and morphological identity. *Studia Linguistica*, *57*(3), 143–169. https://doi.org/10.1111/j.0039-3193.2003.00102.x

Embick, D., & Halle, M. (2005). On the status of stems in morphological theory. In T. Geerts, I. van Ginneken, & H. Jacobs (Eds.), *Romance languages and linguistic theory 2003: Selected papers from 'Going Romance' 2003, Nijmegen, 20–22 November* (pp. 37–62). John Benjamins. https://doi.org/10.1075/cilt.270.03emb

Embick, D., & Marantz, A. (2008). Architecture and blocking. *Linguistic Inquiry*, *39*(1), 1–53. https://doi.org/10.1162/ling.2008.39.1.1

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

Flack, K. (2007). Templatic morphology and indexed markedness constraints. *Linguistic Inquiry*, *38*(4), 749–758. https://doi.org/10.1162/ling.2007.38.4.749

Gagliardi, A., & Lidz, J. (2014). Statistical insensitivity in the acquisition of Tsez noun classes. *Language*, *90*(1), 58–89. https://doi.org/10.1353/lan.2014.0013

Goldwater, S., & Johnson, M. (2003). Learning OT constraint rankings using a maximum entropy model. In J. Spenader, A. Eriksson, & Ö. Dahl (Eds.), *Variation within Optimality Theory: Proceedings of the Stockholm Workshop on 'Variation within Optimality Theory', April 26–27, 2003 at Department of Linguistics, Stockholm University, Sweden* (pp. 111–120).

Gouskova, M. (2010). The phonology of boundaries and secondary stress in Russian compounds. *Linguistic Review*, *27*(4), 387–448. https://doi.org/10.1515/tlir.2010.015

Gouskova, M. (2012). Unexceptional segments. *Natural Language and Linguistic Theory*, *30*(1), 79–133. https://doi.org/10.1007/s11049-011-9142-4

Gouskova, M., & Bobaljik, J. D. (2022). The lexical core of a complex functional affix: Russian baby diminutive *-onok*. *Natural Language & Linguistic Theory*, *40*, 1075–1115. https://doi.org/10.1007/s11049-021-09530-1

Gouskova, M., & Linzen, T. (2015). Morphological conditioning of phonological regularization. *The Linguistic Review*, *32*(3), 427–473. https://doi.org/10.1515/tlr-2014-0027

Gouskova, M., Newlin-Łukowicz, L., & Kasyanenko, S. (2015). Selectional restrictions as phonotactics over sublexicons. *Lingua*, *167*, 41–81. https://doi.org/10.1016/j.lingua.2015.08.014

Gribanova, V., & Harizanov, B. (2017). Locality and directionality in inward-sensitive allomorphy: Russian and Bulgarian. In V. Gribanova & S. S. Shih (Eds.), *The morphosyntax-phonology connection: Locality and directionality at the interface* (pp. 61–90). Oxford University Press. https://doi.org/10.1093/acprof:oso/9780190210304.003.0003

Guzmán Naranjo, M. (2019). *Analogical classification in formal grammar*. Language Science Press. https://doi.org/10.5281/zenodo.3191825

Guzmán Naranjo, M. (2020). Analogy, complexity and predictability in the Russian nominal inflection system. *Morphology*, *30*, 219–262. https://doi.org/10.1007/s11525-020-09367-1

Guzmán Naranjo, M., & Bonami, O. (2021). Overabundance and inflectional classification: Quantitative evidence from Czech. *Glossa: a journal of general linguistics*, *6*(1), 88. 1–31. https://doi.org/10.5334/gjgl.1626

Hajič, J., Hlaváčová, J., Mikulová, M., Straka, M., & Štěpánková, B. (2020). MorfFlex CZ 2.0 [LINDAT/CLARIAH-CZ digital library at the Institute of Formal and Applied Linguistics (ÚFAL), Faculty of Mathematics and Physics, Charles University]. http://hdl.handle.net/11234/1-3186

Hall Jr., R. A. (1953). *Haitian Creole: Grammar, texts, vocabulary* [With the collaboration of Suzanne Comhaire-Sylvain, H. Ormonde McConnell, and Alfred Métraux]. The American Anthropological Association.

Halle, M. (1973). Prolegomena to a theory of word formation. *Linguistic Inquiry*, *4*(1), 3–16.

Halle, M. (1997). Distributed morphology: Impoverishment and fission. In B. Bruening, Y. Kang, & M. McGinnis (Eds.), *PF: Papers at the interface* (pp. 425–449). MITWPL.

Halle, M., & Marantz, A. (1993). Distributed morphology and the pieces of inflection. In K. Hale & S. J. Keyser (Eds.), *The view from Building 20: Essays in linguistics in honor of Sylvain Bromberger* (pp. 111–176). MIT Press.

Halle, M., & Marantz, A. (2008). Clarifying "Blur": Paradigms, defaults, and inflectional classes. In A. Bachrach & A. Nevins (Eds.), *Inflectional identity* (pp. 55–72). Oxford University Press.

Harley, H. (2014). On the identity of roots. *Theoretical Linguistics*, *40*(3–4), 225–276. https://doi.org/10.1515/tl-2014-0010

Harley, H., & Noyer, R. (1999). Distributed Morphology. *Glot International*, *4*(4), 3–9.

Harrell, R. S. (1962). *A short reference grammar of Moroccan Arabic: With an appendix of texts in urban Moroccan Arabic by Louis Brunot*. Georgetown University Press.

Haugen, J. D. (2016). Readjustment: Rejected? In D. Siddiqi & H. Harley (Eds.), *Morphological metatheory* (pp. 303–340). John Benjamins. https://doi.org/10.1075/la.229.11hau

Hay, J. (2003). *Causes and consequences of word structure*. Routledge. https://doi.org/10.4324/9780203495131

Hayes, B., & Londe, Z. C. (2006). Stochastic phonological knowledge: The case of Hungarian vowel harmony. *Phonology*, *23*, 59–104. https://doi.org/10.1017/s0952675706000765

Hayes, B., & White, J. (2013). Phonological naturalness and phonotactic learning. *Linguistic Inquiry*, *44*(1), 45–75. https://doi.org/10.1162/ling_a_00119

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

Hayes, B., Zuraw, K., Siptár, P., & Londe, Z. (2009). Natural and unnatural constraints in Hungarian vowel harmony. *Language*, *85*(4), 822–863. https://doi.org/10.1353/lan.0.0169

Hockett, C. (1954). Two models of grammatical description. *Word*, *10*(2–3), 210–234. https://doi.org/10.1080/00437956.1954.11659524

Houghton, Z. N., & Kapatsinski, V. (2023). *Are your random effects normal? A simulation study of methods for estimating whether subjects or items come from more than one population by examining the distribution of random effects in mixed-effects logistic regression* [Ms.].

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/10.1207/s15473341lld0102_3

Idsardi, W. J. (1992, July). *The computation of prosody* [Doctoral dissertation, Massachusetts Institute of Technology].

Inkelas, S., Orgun, O., & Zoll, C. (1997). The implications of lexical exceptions for the nature of grammar. In I. Roca (Ed.), *Derivations and constraints in phonology* (pp. 393–418). Clarendon Press.

Ito, J., & Mester, A. (2008). Lexical classes in phonology. In S. Miyagawa & M. Saito (Eds.), *The Oxford handbook of Japanese linguistics* (pp. 84–106). Oxford University Press. https://doi.org/10.1093/oxfordhb/9780195307344.013.0004

Itô, J., & Mester, A. (1999). The phonological lexicon. In N. Tsujimura (Ed.), *The handbook of Japanese linguistics* (pp. 62–100). Blackwell. https://doi.org/10.1002/9781405166225.ch3

Jakobson, R. (1984). Morphological observations on Slavic declension (the structure of Russian case forms). In L. R. Waugh & M. Halle (Eds.), *Russian and Slavic grammar: Studies 1931–1981* (pp. 105–133). Mouton. https://doi.org/10.1515/9783110822885.105

James, G., Witten, D., Hastie, T., & Tibshirani, R. (2013). *An introduction to statistical learning with applications in R*. Springer. https://doi.org/10.1007/978-1-4614-7138-7

Janků, L. (2022). *The nanosyntax of Czech nominal declensions* [Doctoral dissertation, Masaryk University].

Jarosz, G. (2022, February). *Learning hidden generalizations* [Colloquium at New York University].

Kager, R. (1996). On affix allomorphy and syllable counting. In U. Kleinhenz (Ed.), *Interfaces in phonology* (pp. 155–171). Akademie Verlag.

Kalin, L. (2022). Infixes really are (underlyingly) prefixes/suffixes: Evidence from allomorphy on the fine timing of infixation. *Language*, *98*(4), 641–682. https://doi.org/10.1353/lan.2022.0017

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. M. (2005). Characteristics of a rule-based default are dissociable: Evidence against the Dual Mechanism Model. In S. Franks, F. Y. Gladney, & M. Tasseva-Kurktchieva (Eds.), *Annual workshop on Formal Approaches to Slavic Linguistics: The South Carolina meeting 2004* (pp. 136–146). Michigan Slavic Publications.

Kenstowicz, M. (1997). Quality-sensitive stress. *Rivista di Linguistica*, *9*(1), 157–187.

Kenstowicz, M., & Rubach, J. (1987). The phonology of syllabic nuclei in Slovak. *Language*, *63*(3), 463–497. https://doi.org/10.2307/415002

Kiefer, F. (1985). The possessive in Hungarian: A problem for natural morphology. *Acta Linguistica Academiae Scientiarum Hungaricae*, *35*(1–2), 85–116.

Kiparsky, P. (1982). Explanation in phonology [Written in 1968 and originally published in 1973]. Foris Publications.

Kiparsky, V. (1979). *Russian historical grammar*. *Vol. 1*. *The development of the sound system* (J. I. Press, Trans.). Ardis.

Kraska-Szlenk, I. (2007). *Analogy: The relation between lexicon and grammar*. LINCOM.

Křen, M., Cvrček, V., Hnátková, M., Jelínek, T., Kocek, J., Kovářı́ková, D., Křivan, J., Milička, J., Petkevič, V., Procházka, P., Skoumalová, H., Šindlerová, J., & Škrabal, M. (2022). Korpus SYN, verze 11 z 14.12.2022. *Ústav Českého národního korpusu FF UK*.

Kuznecov, P. S. (1953). *Istoričeskaja grammatika russkogo jazyka: Morfologija*. Izdatel'stvo Moskovskogo universiteta.

Labov, W. (1969). Contraction, deletion, and inherent variability of the English copula. *Language*, *45*(4), 715–762. https://doi.org/10.2307/412333

Legendre, G., Miyata, Y., & Smolensky, P. (1990). *Harmonic Grammar - a formal multi-level connectionist theory of linguistics well-formedness: Theoretical foundations* (ICS Technical Report No. 90-5) (CU-CS-465-90). University of Colorado at Boulder.

Legendre, G., Sorace, A., & Smolensky, P. (2006). The Optimality Theory–Harmonic Grammar connection. In P. Smolensky & G. Legendre (Eds.), *The harmonic mind: From neural computation to optimality-theoretic grammar*. *Vol. 2*. *Linguistic and philosophical implications* (pp. 339–402). MIT Press.

Lieber, R. (1980). *On the organization of the lexicon* [Doctoral dissertation, Massachusetts Institute of Technology].

Lightner, T. (1965). *Segmental phonology of Modern Standard Russian* [Doctoral dissertation, Massachusetts Institute of Technology].

Lignos, C., & Gorman, K. (2014). Revisiting frequency and storage in morphological processing. In A. Beltrama, T. Chatzikonstantinou, J. L. Lee, M. Pham, & D. Rak (Eds.), *Proceedings of the forty-eighth annual meeting of the Chicago Linguistic Society* (pp. 447–461).

Linzen, T., Kasyanenko, S., & Gouskova, M. (2013). Lexical and phonological variation in Russian prepositions. *Phonology*, *30*, 453–515. https://doi.org/10.1017/s0952675713000225

Lüdecke, D., Ben-Shachar, M. S., Patil, I., Waggoner, P., & Makowski, D. (2021). performance: An R package for assessment, comparison and testing of statistical models. *Journal of Open Source Software*, *6*(60), 3139. https://doi.org/10.21105/joss.03139

Lundová, I. (2018). *Vokalizace předložek* [Master's thesis, Masarykova univerzita].

Magomedova, V. (2017). Pseudo-allomorphs in Modern Russian. *University of Pennsylvania Working Papers in Linguistics*, *23*(1), 16.

Magomedova, V., & Slioussar, N. (2017). Stem-final consonant mutations in modern Russian. In Y. Oseki, M. Esipova, & S. Harves (Eds.), *Annual workshop on Formal Approaches to Slavic Linguistics: The NYU meeting 2015* (pp. 239–259). Michigan Slavic Publications.

Marantz, A. (1996, December). *"Cat" as a phrasal idiom: Consequences of late insertion for Distributed Morphology* [Ms.].

Marantz, A. (1997). No escape from syntax: Don't try morphological analysis in the privacy of your own lexicon. *University of Pennsylvania Working Papers in Linguistics*, *4*(2), 201–225.

Marantz, A. (2013). Locality domains for contextual allomorphy across the interfaces. In O. Matushansky & A. Marantz (Eds.), *Distributed morphology today: Morphemes for morris halle* (pp. 95–116). MIT Press. https://doi.org/10.7551/mitpress/9701.003.0008

Marantz, A. (2020, August). Contextual allosemy and idioms.

Marantz, A. (2023, January). *Understanding the role of a generative grammar in language use: The* ***slithy toves*** *slay* Words and Rules [Talk presented at the Linguistics Society of America 97th Annual Meeting, Denver].

Marcus, G. F., Brinkmann, U., Clahsen, H., Wiese, R., & Pinker, S. (1995). German inflection: The exception that proves the rule. *Cognitive Psychology*, *29*(3), 189–256. https://doi.org/10.1006/cogp.1995.1015

Massey, F. J., Jr. (1951). The Kolmogorov-Smirnov test for goodness of fit. *Journal of the American Statistical Association*, *46*(253), 68–78. https://doi.org/10.1080/01621459.1951.10500769

McCarthy, J. J., & Prince, A. (1993a). Generalized alignment. In G. Booij & J. Van Marle (Eds.), *Yearbook of morphology 1993* (pp. 79–153). Springer.

McCarthy, J. J., & Prince, A. (1993b). Prosodic morphology: Constraint interaction and satisfaction. *Linguistics Department Faculty Publication Series*, *14*.

McCurdy, K., Goldwater, S., & Lopez, A. (2020). Inflecting when there's no majority: Limitations of encoder-decoder neural networks as cognitive models for German plurals. In D. Jurafsky, J. Chai, N. Schluter, & J. Tetreault (Eds.), *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics* (pp. 1745–1756). Association for Computational Linguistics. https://doi.org/10.18653/v1/2020.acl-main.159

Melvold, J. L. (1989, February). *Structure and stress in the phonology of Russian* [Doctoral dissertation, Massachusetts Institute of Technology].

Mołczanow, J., Domahs, U., Knaus, J., & Wiese, R. (2013). The lexical representation of word stress in Russian: Evidence from event-related potentials. *The Mental Lexicon*, *8*(2), 164–194. https://doi.org/10.1075/ml.8.2.03mol

Müller, G. (2004). On decomposing inflection class features: Syncretism in Russian noun inflection. In G. Müller, L. Gunkel, & G. Zifonun (Eds.), *Explorations in nominal inflection* (pp. 189–227). Mouton de Gruyter. https://doi.org/10.1515/9783110197501.189

Nakisa, R. C., Plunkett, K., & Hahn, U. (2001). A cross-linguistic comparison of single and dual-route models of inflectional morphology. In P. Broeder & J. Murre (Eds.), *Models of language acquisition: Inductive and deductive approaches*. MIT Press.

Nevins, A., & Parrott, J. K. (2010). Variable rules meet Impoverishment theory: Patterns of agreement leveling in English varieties. *Lingua*, *120*(5), 1135–1159. https://doi.org/10.1016/j.lingua.2008.05.008

Nosofsky, R. N. (1990). Relations between exemplar-similarity and likelihood models of classification. *Journal of Mathematical Psychology*, *34*(4), 393–418. https://doi.org/10.1016/0022-2496(90)90020-a

Oravecz, C., Váradi, T., & Sass, B. (2014, May). The Hungarian gigaword corpus. In N. Calzolari, K. Choukri, T. Declerck, H. Loftsson, B. Maegaard, J. Mariani, A. Moreno, J. Odijk, & S. Piperidis (Eds.), *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC'14)* (pp. 1719–1723). European Language Resource Association.

Papp, F. (1969). *A magyar nyelv szóvégmutató szótára*. Akadémiai Kiadó.

Parker, J., & Sims, A. D. (2020). Irregularity, paradigmatic layers, and the complexity of inflection class systems: A study of Russian nouns. In P. Arkadiev & F. Gardani (Eds.), *The complexities of morphology* (pp. 23–51). Oxford University Press. https://doi.org/10.1093/oso/9780198861287.003.0002

Paster, M. (2015). Phonologically conditioned suppletive allomorphy: Cross-linguistic results and theoretical consequences. In E. Bonet, M.-R. Lloret, & J. Mascaró (Eds.), *Understanding allomorphy: Perspectives from Optimality Theory* (pp. 218–253). Equinox. https://doi.org/10.1558/equinox.25219

Paster, M. E. (2006). *Phonological conditions on affixation* [Doctoral dissertation, University of California, Berkeley].

Pater, J. (2006). The locus of exceptionality: Morpheme-specific phonology. In L. Bateman, M. O'Keefe, E. Reilly, & A. Werley (Eds.), *Papers in optimality theory III* (pp. 259–296).

Pater, J. (2010). Morpheme-specific phonology: Constraint indexation and inconsistency resolution. In S. Parker (Ed.), *Phonological argumentation: Essays on evidence and motivation* (pp. 123–154). Equinox.

Pater, J. (2016). Universal grammar with weighted constraints. In J. J. McCarthy & J. Pater (Eds.), *Harmonic Grammar and Harmonic Serialism* (pp. 1–46). Equinox.

Patz, E. (1991). Djabugay. In R. M. W. Dixon & B. J. Blake (Eds.), *The handbook of Australian languages. Vol. 4. The aboriginal language of Melbourne and other grammatical sketches* (pp. 245–347). Oxford University Press.

Pertsova, K. (2004). *Distribution of genitive plural allomorphs in the Russian lexicon and in the internal grammar of native speakers* [Master's thesis, University of California, Los Angeles].

Pierrehumbert, J. (2001). Stochastic phonology. *Glot International*, *5*(6), 195–207.

Pinker, S., & Prince, A. (1988). On language and connectionism: Analysis of a parallel distributed processing model of language acquisition. *Cognition*, *28*(1–2), 73–193. https://doi.org/10.1016/0010-0277(88)90032-7

Polivanova, A. K. (2008 [1967]). Obrazovanije umen'šitel'nyx suščestvitel'nyx mužskogo roda. In *Obščeje i russkoje jazykoznanije: Izbrannyje raboty* (pp. 8–22). Rossijskij gosudarstvennyj gumanitarnyj universitet.

Potts, C., Pater, J., Jesney, K., Bhatt, R., & Becker, M. (2010). Harmonic Grammar with linear programming: From linear systems to linguistic typology. *Phonology*, *27*, 77–117. https://doi.org/10.1017/s0952675710000047

Prasada, S., & Pinker, S. (1993). Generalisation of regular and irregular morphological patterns. *Language and Cognitive Processes*, *8*(1), 1–56. https://doi.org/10.1080/01690969308406948

Prince, A., & Smolensky, P. (2004). *Optimality theory: Constraint interaction in generative grammar* [Originally published in 1993]. Blackwell. https://doi.org/10.7282/T34M92MV

Privizentseva, M. (2022). *Mixed agreement in Russian: Gender, declension, and morphological ineffability* [Ms.].

R Core Team. (2022). *R: A language and environment for statistical computing*. R Foundation for Statistical Computing. Vienna. https://www.R-project.org/

Rácz, P., & Rebrus, P. (2012). Variation in the possessive allomorphy of Hungarian. In F. Kiefer, M. Ladányi, & P. Siptár (Eds.), *Current issues in morphological theory: (ir)regularity, analogy and frequency: Selected papers from the 14th International Morphology Meeting, budapest, 13–16 may 2010* (pp. 51–64). John Benjamins. https://doi.org/10.1075/cilt.322.03rac

Rebrus, P. (2013). Miért nincs *j*? In A. Benő, E. Fazakas, & E. Kádár (Eds.), *„…hogy legyen a víznek lefolyása…": Köszöntő kötet Szilágyi N. Sándor tiszteletére* (pp. 383–401). Erdélyi Múzeum-Egyesület.

Rebrus, P., & Szigetvári, P. (2018, November). *Morphosyntactically conditioned morphophonological variation* [Presentation at PhonVar 1, Barcelona].

Rebrus, P., Szigetvári, P., & Tŏrkenczy, M. (2012). Dark secrets of Hungarian vowel harmony. In E. Cyran, H. Kardela, & B. Szymanek (Eds.), *Sound, structure and sense: Studies in memory of Edmund Gussmann* (pp. 491–508). Wydawnictwo KUL.

Rebrus, P., Szigetvári, P., & Törkenczy, M. (2017). Asymmetric variation. In G. Lindsey & A. Nevins (Eds.), *Sonic signatures: Studies dedicated to John Harris* (pp. 164–187). John Benjamins. https://doi.org/10.1075/lfab.14.c10

Rebrus, P., Szigetvári, P., & Törkenczy, M. (2019, October). *Vacillation and lexical variation in Hungarian front/back harmony* [Poster presented at the 7th Annual Meeting on Phonology, Stony Brook University].

Revithiadou, A. (1999, January). *Headmost accent wins: Head dominance and ideal prosodic form in lexical accent systems* [Doctoral dissertation, Universiteit Leiden].

Ritter, N. A. (2023). Hungarian possessive allomorphy in the lexicon. In J. van de Weijer (Ed.), *Representing phonological detail: Part i: Segmental structure and representations* (pp. 33–70). De Gruyter Mouton. https://doi.org/10.1515/9783110730098-003

Rodina, Y., & Westergaard, M. (2012). A cue-based approach to the acquisition of grammatical gender in Russian. *Journal of Child Language*, *39*, 1077–1106. https://doi.org/10.1017/s0305000911000419

Rounds, C. (2008). *Hungarian: An essential grammar* (2nd Edition). Routledge. https://doi.org/10.4324/9780203886199

Rubach, J. (1984). *Cyclic and lexical phonology: The structure of Polish*. Foris Publications. https://doi.org/10.1515/9783111392837

Rubach, J. (2013). Exceptional segments in Polish. *Natural Language & Linguistic Theory*, *31*(4), 1139–1162. https://doi.org/10.1007/s11049-013-9204-x

Rumelhart, D. E., & McClelland, J. L. (1986). On learning the past tenses of English verbs. In J. L. McClelland, D. E. Rumelhart, & the PDP Research Group (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition. Vol. 2. Psychological and biological models* (pp. 216–271). The MIT Press. https://doi.org/10.7551/mitpress/5237.003.0008

Rysling, A. (2016). Polish yers revisited. *Catalan Journal of Linguistics*, *15*, 121–143. https://doi.org/10.5565/rev/catjl.187

Schreuder, R., & Baayen, R. H. (1995). Modeling morphological processing. In L. B. Feldman (Ed.), *Morphological aspects of language processing* (pp. 131–154). Lawrence Erlbaum Associates.

Schuler, K. D., Yang, C., & Newport, E. L. (2021, March). *Testing the Tolerance Principle: Children form productive rules when it is more computationally efficient* [10.31234/osf.io/utgds]. https://doi.org/10.31234/osf.io/utgds

Ševeleva, M. S. (Ed.). (1974). *Obratnyj slovar' russkogo jazyka*. Sovetskaja Enciklopedija.

Siffer, A. (2018). *Rfolding: The folding test of unimodality* [R package version 1.0]. https://CRAN.R-project.org/package=Rfolding

Siffer, A., Fouque, P.-A., Termier, A., & Largouët, C. (2018). Are your data gathered?: The folding test of unimodality. *KDD '18: Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2210–2218. https://doi.org/10.1145/3219819.3219994

Siptár, P., & Törkenczy, M. (2000). *The phonology of Hungarian*. Oxford University Press.

Skousen, R. (1989). *Analogical modeling of language*. Kluwer Academic Publishers. https://doi.org/10.1007/978-94-009-1906-8

Solomyak, O., & Marantz, A. (2010). Evidence for early morphological decomposition in visual word recognition. *Journal of Cognitive Neuroscience*, *22*(9), 2042–2057. https://doi.org/10.1162/jocn.2009.21296

Somers, R. H. (1962). A new asymmetric measure of association for ordinal variables. *American Sociological Review*, *27*(6). https://doi.org/10.2307/2090408

Steriopolo, O. (2008). *Form and function of expressive morphology: A case study of Russian* [Doctoral dissertation, University of British Columbia].

Švedova, N. J., Arutjunova, N. D., Bondarko, A. V., Ivanov, V. V., Lopatin, V. V., Uluxanov, I. S., & Filin, F. P. (Eds.). (1980). *Russkaja grammatika*. *Vol. 1. Fonetika, fonologija, udarenije, intonacija, slovoobrazovanije, morfologija*. Izdatel'stvo Nauka.

Szeredi, D. (2016, September). *Exceptionality in vowel harmony* [Doctoral dissertation, New York University].

Tessier, A.-M. (2011). USELISTEDERROR: A grammatical account of lexical exceptions in phonological acquisition. In S. Lima, K. Mullin, & B. Smith (Eds.), *NELS 39: Proceedings of the thirty-ninth annual meeting of the North East Linguistic Society* (pp. 813–827, Vol. 2). Graduate Linguistics Student Association, University of Massachusetts Amherst.

Trommer, J. (2021). The subsegmental structure of German plural allomorphy. *Natural Language & Linguistic Theory*, *39*, 601–656. https://doi.org/10.1007/s11049-020-09479-7

Vinogradov, V. V. (1972). *Russkij jazyk: Grammatičeskoje učenije o slove* (2nd). Vysšaja škola.

Voeten, C. C. (2022). *buildmer: Stepwise elimination and term reordering for mixed-effects regression* [R package version 2.6]. https://CRAN.R-project.org/package=buildmer

Williams, A., Pimentel, T., McCarthy, A. D., Blix, H., Chodroff, E., & Cotterell, R. (2020). Predicting declension class from form and meaning. In D. Jurafsky, J. Chai, N. Schluter, & J. Tetreault (Eds.), *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics* (pp. 6682–6695). Association for Computational Linguistics. https://doi.org/10.18653/v1/2020.acl-main.597

Wolf, M. (2008, September). *Optimal Interleaving: Serial phonology-morphology interaction in a constraint-based model* [Doctoral dissertation, University of Massachusetts Amherst].

Wurzel, W. U. (1989). *Inflectional morphology and naturalness*. Kluwer.

Yang, C. (2016). *The price of linguistic productivity: How children learn to break the rules of language*. MIT Press. https://doi.org/10.7551/mitpress/9780262035323.001.0001

Yearley, J. (1995). Jer vowels in Russian. In J. Beckman, L. W. Dickey, & S. Urbanczyk (Eds.), *Papers in optimality theory* (pp. 533–571, Vol. 18). Graduate Student Linguistics Association.

Zaliznjak, A. A. (1977). *Grammatičeskij slovar' russkogo jazyka*. Russkij Jazyk.

Zaretsky, E., & Lange, B. P. (2016). No matter how hard we try: Still no default plural marker in nonce nouns in Modern High German. In H. Christ, D. Klenovšak, L. Sönning, & V. Werner (Eds.), *A blend of MaLT: Selected contributions from the Methods and Linguistic Theories Symposium 2015* (pp. 153–178). University of Bamberg Press.