When Bases compete: a voting model of Lexical Conservatism*

Canaan Breiss, MIT†

April 19, 2023

Abstract

This paper examines Lexical Conservatism (Steriade, 1997), a phenomenon whereby the distribution of stem allomorphs in a morphological paradigm influences the way those paradigms accommodate derived members. Specifically, a phonological alternation only applies in a derived member if there is an existing form present elsewhere in the paradigm that offers the needed phonological material. Thus *compénsable* undergoes stress shift because the existing word *compénsatory* contains the *compéns-* allomorph. In contrast, **inúndable* is judged worse than *ínundable*, since there is no existing form with in *inúnd-*. In four experiments on English and Mexican Spanish, I demonstrate that the this dependency between paradigm structure and phonological process application generalizes to entirely novel words in a probabilistic manner. Further, I find that all stem allomorphs in a paradigm play a role in determining the form of the novel word, rather than only those that could reduce the markedness of the novel form as previous studies claimed. I propose a novel grammatical model where bases get to "vote" on the shape of the novel form: all stem allomorphs in a lexical entry stand in a correspondence relation to the novel form and exert their influence via multiple faithfulness constraints, which compete with standard markedness constraints in a probabilistic phonological grammar.

*Thanks to Bruce Hayes, Kie Zuraw, Claire Moore-Cantwell, Megha Sundara, Donca Steriade, Juliet Stanton, Dave Embick, audiences at the 95^{th} and 96^{th} LSA Annual Meetings, AMP 2020 and 2021, as well as UCLA, MIT, University of Massachusetts — Amherst, University of Southern California, for valuable feedback and suggestions. This work was supported in part by NSF Graduate Research Fellowship DGE-1650604 to CB, and by a UCLA Dissertation Year Fellowship. All errors are my own.

†Contact: canaanbreiss1@gmail.com

Keywords— Lexical Conservatism, word-formation, lexicon-phonology interaction

1 Introduction

This paper is about Lexical Conservatism (Steriade, 1997; Burzio, 1998), a correlation between the phonological properties of stem allomorphs in a paradigm, and the types of morphophonological alternations that word-formation processes induce in members of that paradigm. One well-known case from Steriade (1997) is that of English *-able*-derivatives, which fall into two classes with respect to stress placement: *íllustrate* yields an affixed coining *illústrable*, undergoing stress shift to relieve the lapse created by the affix, while *írrigate* yields *írrigable* with fixed stress. Steriade argued that this difference is not due to the forms *íllustrate* and *írrigate* themselves, which are termed *Local Bases* (Stanton and Steriade, 2014), but rather from the other members of the morphological paradigms they are embedded in: *íllustrate* has a morphologically related form *illústrative* with stress on the second syllable, while *írrigate* has no morphologically-related form in **irríg-*. Thus, speakers are *lexically conservative*: they shift stress rightward in *illústrable* because the presence of *illústrative* allows them to remain faithful to a morphological relation of *íllustrate*'s, while still repairing the *Extended Lapse violation. In *írrigable*, on the other hand, speakers have no such recourse, and so remain faithful to stress placement of *írrigate* at the expense of retaining the lapse.

In this paper, I follow Stanton and Steriade (2014) in using the term *Local Base* to refer to the paradigm member which is the semantically compositional source for the affixed form in question; I term the affixed form the *Derivative*. To refer to the form that is not the direct ancestor of the Derivative but rather a "sister" of the Local Base, the term *Remote Base* is used. Finally, I use the term *target syllable* to refer to the expected location of stress if it were shifted in the Derivative. Using these terms, we can summarise Steriade's insight in table 1, with the target syllable in the Local Base underlined.

Local Base	Remote Base	Derivative
ír<u>ri</u>gate	-	*írrigable* ∼ **irrígable*
íl<u>lu</u>strate	*illústrative*	**íllustrable* ∼ *illústrable*

Table 1: Distribution of paradigm members according to Steriade (1997), with target syllables underlined.

Since their publication, Steriade's findings have been widely recognized to be revealing about the interaction between grammar and lexicon and to have implications both for psycholinguistic models of grammar-lexicon interaction and phonological theories of word-formation (Burzio, 2002; Albright, 2002; Bermádez-Otero, 2011; Rolle, 2018).

Specifically, the existence of dependencies between non-cyclically-contained surface forms has been acknowledged to be particularly difficult for strictly derivational theories of the phonology-morphology interface to handle (Bermúdez-Otero, 2017). Lexical Conservatism has also accrued a growing number of empirical cases: a possibly exhaustive list at the time of writing is Burzio (1998) on Italian, Bat-El (2002); Asherov and Bat-El (2016) on Modern Hebrew, Pertsova (2005); Pertsova and Kuznetsova (2015) on Russian, O'Brien (2007) on Irish, Steriade (2008) on Romanian, Bonet and Torres-Tamarit (2010) on Catalan, Gunkel (2010, 2011) on Ancient Greek, Steriade (2012) on Latin, Simonović (2012); Simonović and Baroni (2014) on Serbo-Croatian, Steriade and Yanovich (2015) on Ukrainian, Steriade (1997); Steriade and Stanton (2020); Breiss (2021) on English, Steriade (1997) on French, Guekguezian and Jesney (2021) on Chukchansi Yokuts, and Breiss (2021) on Spanish.

1.1 Previous analyses, and their predictions

Two classes of analyses have been proposed for data on Lexical Conservatism. Here, I highlight a critical prediction of both theories, which will be central to this paper: the influence of a Remote Base on the Derivative should *only* be evident if the Remote Base has a form that could help relieve markedness of the Derivative relative to the Local Base.

Steriade (1997) proposes an account for the English *-able* data discussed above, as well as for similar facts about French liaison, that can be thought of as *quantifying* over Bases: as long as there is *some* Base in the lexicon that the Derivative can resemble, there is no penalty. An example of the Lex-X family of constraints employed is given here, where X is defined in reference to the primarily-stressed vowel of the Local Base.

- **Lex-Vowel**: Assign one violation if there is no Base in the input (here, the paradigm of the Local Base) that matches the vowel in the candidate in quality.

A different model for a broader range of English data comes from Steriade and Stanton (2020), who cast cyclic containment as a violable constraint.

- **C-Containment**: Assign one violation if the candidate does not stand in correspondence to the Local Base.

The model differs from Steriade (1997)'s quantification-based account in that C-Containment prefers Local Bases to Remote Bases, with correspondence between Remote Bases and candidate Derivatives only entertained when forced by an intervening Markedness constraint. However, as with the model of Steriade (1997), this model does not allow Remote Bases that are not markedness-improving to play a role in Derivative formation.

In section 5.4 I formalize the way that these models come up short, and provide an alternative model to account for the facts.

2 Contributions of this paper

This paper reports four experiments: an in-depth investigation into the Lexical Conservatism in English stress placement discussed first in Steriade (1997) and later in Steriade and Stanton (2020), and also a new case of Lexical Conservatism in Spanish mid-vowel diphthongization. The experiments serve in part to solidify the empirical ground on which theories of Lexical Conservatism rest by shoring up some methodological weaknesses of previous studies. Despite being typologically well-attested, it is still unclear what the synchronic status of Lexical Conservatism is. This is because all studies after Steriade (1997) have been based in existing lexical data (with the exception of Steriade and Stanton (2020)). Further, neither Steriade (1997) nor Steriade and Stanton (2020) verified that participants actually knew the Remote Bases whose knowledge was implied by the Lexical Conservatism effect they were studying.

To address this, in this paper I use possible-but-unattested combinations of existing affixes and stems. If speakers generalize the Lexical Conservatism dependency to such forms, we will have evidence that the grammar encodes the dependency in a more general fashion, beyond the specific lexical data instantiating the pattern. I further expand on Steriade's original insight by verifying the structure of each speaker's paradigm (that is, whether they know the relevant Local and Remote Bases), and basing conclusions about the dependency between Local and Remote Bases in Derivative formation only on data from speakers who know both forms.

2.1 Both paradigm shape and markedness influence Derivatives

As alluded to above, a major role of the experiments is to test the common prediction of the two existing models of Lexical Conservatism that the Derivative is sensitive only to Remote Bases whose form is markedness-reducing compared to the Local Base. Across experiments, I compare three types of Remote Base which differ in whether they provide a source for a less marked structure in the Derivative compared to the Local Base. The first familiar from the literature on Lexical Conservatism — the Remote Base exhibits a stem shape that resembles the result of applying a markedness-reducing phonological alternation to its Local Base. Returning again to English *-able*, I term Remote Bases like *labórious* as *helpful*, since its stem shape *labór-* is phonologically optimizing relative to the Local Base *lábor*, relieving the stress lapse. I address these types of data in Experiments 1, 2, and 3. In Experiment 4, I introduce data from Spanish that also contains helpful Remote Bases. In Spanish, unstressed diphthongs are marked, and many stressed diphthongs alternate with unstressed monophthongs (Harris, 1969). Thus for Local Bases like *muéble* "furniture" affixed with a stress-attracting affix like *-óso*, a marked structure is created when stress shifts off the diphthong in the base (**mueblóso* "full of furniture"). The presence of a Remote Base like *moblár* "to furnish", with an unstressed monophthong, is helpful because it contains a stem shape that repairs the diphthong. Thus the Spanish case is in line with the other cases in the literature where the presence of a helpful Remote Base facilitates repair in Derivatives, compared to Derivatives that don't have any Remote Base.

I contrast the presence of Derivatives with helpful Remote Bases in English and Spanish to Derivatives with different types of Remote Bases: *harmful*, in cases where the Remote Base has a structure that would be more marked than simply using the Local Base, and *unhelpful*, in cases where the Remote Base has the same structure as the Local Base.

In English, I term cases like Local Base *resíde* with Remote Base *résident harmful*, since if the Derivative reflected the stress placement of the Remote Base, it would actually result in a more marked form (here, more lapse-ful) than simply being faithful to the Local Base. In Experiment 3, participants produced almost zero forms like *résidable*, where the Derivative matched the harmful Remote Base *résident*, while there were numerous cases of Derivatives like *labórable*, where the Derivative matches the helpful Remote Base *labórious*.

Spanish presents a different view on non-markedness-reducing Remote Bases, as seen in cases like the Local Base *juérga* "spree" with the Remote Base *juerguísta* "reveller". Here, the Remote Base has a structure — the unstressed diphthong — which is the same, with respect to markedness-reduction, as simply failing to repair the marked structure in the Local Base created by affixation and its accompanying obligatory stress shift. I term these cases *unhelpful* since, unlike the English case, a Derivative that resembles them is as marked as one would be if it did not undergo repair, rather than more marked, as in the case of *resíde* and *résident*. Alongside these cases of unhelpful Remote Bases, Spanish also has cases like the Local Base *muéble* with a helpful Remote Base *moblár*, which follow the same pattern as in English. Thus, Experiments 3 and 4 can be seen acting together a pair to explore the limits of the extent to which Remote Bases can influence Derivatives even in the face of adverse markedness results.

Table 2 below summarises the types of Remote Base and their distribution throughout the experiments in this paper.

Local Base	Derivative	Remote Base	Remote Base type	Experiment
pláster	*plasterable*	–	None	1, 2, 3 (English),
noviémbre	*nov(i)embróso*	–		4 (Spanish)
lábor	*laborable*	*labórious*	Helpful	1, 2, 3 (English),
aciérto	*ac(i)ertóso*	*acertár*		4 (Spanish)
resíde	*residable*	*résident*	Harmful	3
diéta	*d(i)etóso*	*dietética/o*	Unhelpful	4

Table 2: Summary of the types of Remote Bases treated in this paper, with example Derivatives.

To account for these facts, I propose a model where Lexical Conservatism emerges from the interaction of grammar and lexicon where different listed allomorphs compete to have their structure reflected in the Derivative, with their influence scaled by their *resting activation*, alongside established principles of markedness and faithfulness. The behavior seen in English ("classical" Lexical Conservatism, with only helpful Remote Bases having an impact on Derivative formation) and the more unexpected behavior seen in Spanish (where both unhelpful and helpful Remote Bases play a role) both emerge from the proposed framework under differing strengths of markedness, which contrasts with other models of the phenomenon in the literature (Steriade, 1997; Steriade and Stanton, 2020).

2.2 Gradience is pervasive and cumulativity is key

Although discussed qualitatively in Steriade (1997), subsequent work has not probed variability in the impact of the Remote Base on Derivatives. Data from all experiments converge on a model of Lexical Conservatism where the influence of a Base is gradient, rather than absolute. That is, if a participant knew the Remote Base *illústrative*, they might still occasionally have produced Derivatives like *íllustrable* that resembled the Local Base *íllustrate*. Conversely, if a participant didn't know the Remote Base, sometimes the Derivative might have exhibited stress shift anyway (*illústrative*). Therefore, core to the voting theory of Bases I propose is that I propose is that Bases compete in model of the grammar that is inherently probabilistic, and that the observed influence of multiple Bases is only captured by a model that allows for constraint cumulativity. In section 5.2 I employ a Maximum Entropy (MaxEnt) grammar formalism (Smolensky, 1986; Goldwater and Johnson, 2003) to implement the Voting Bases model, but in principle other frameworks like Stochastic Harmonic Grammar (also known in the literature as Noisy Harmonic Grammar, (Boersma and Pater, 2016)) might well be equally suited to the task.

2.3 Non-phonological lexical characteristics influence Derivative formation

In Experiments 2 and 3, I delve more deeply into what makes a Remote Base influential. Is it simply the presence of the Base in the speaker's lexicon that facilitates an alternation in the Derivative, or is the Remote Base is actually accessed in real time during Derivative formation? I manipulated the Remote Base's accessibility explicitly by asking participants to declare their knowledge of the Remote Base before carrying out the Derivative-formation task, under the assumption that by raising the Base's salience (modeled formally with reference to *resting activation* (Morton, 1970), for details see section 6.2.1) makes it easier to access in memory, and thus more able to influence the Derivative. We find that Derivatives with primed Remote Bases more often exhibited stress shift, suggesting that not only the contents of the lexicon, but also the ease which which the speaker can access them in real time, influences outcomes of the phonological grammar. I propose an extension of the Voting Bases model that formalizes integration between psycholinguistic characteristics of lexical entries and the computation of the phonological grammar.

Finally, I follow up on Steriade's discussion of the semantic relationship between the Local and Remote Bases. Steriade suggests, based on her survey, that Remote Bases with transparent semantic roots in their Local Bases may be more "accessible" (Steriade, 1997:p. 1) than those whose relationship is opaque. I estimate the semantic similarity between Local and Remote Bases in Experiments 1-3 using a norming study, and find that semantic similarity does modulate the role of the Remote Base in Derivative formation, but in the opposite direction: Derivatives to Local Bases with a semantically-similar Remote Base are less likely than those with a semantically-dissimilar Remote Base to undergo stress shift. I discuss the way this unexpected finding nevertheless patterns with other studies of semantic interference in speech production, suggesting a tentative mechanism based in the way that semantic similarity modulates the spread of resting activation between lexical entries.

3 Lexical Conservatism in English stress

This section reports the results of three experiments on English that address questions of generalizability, the involvement of the lexicon, and the role of harmful Remote Bases. I discuss the phonological determinants of Derivative formation, as well as whether the *presence* of a Remote Base matters or not, and present qualitative summaries of the priming effects in Experiments 2 and 3. In order to increase statistical power and generalizability when examining lexical effects on Derivative formation I report a combined statistical analysis of all Derivatives with known helpful Remote Bases (collapsing across all three experiments) later in section 3.4.

3.1 Experiment 1

Experiment 1 replicates and extends Steriade (1997) on English stress shift.

3.1.1 Methods

3.1.1.1 Participants

Thirty-six UCLA undergraduate students participated and were compensated with course credit. Data from 5 were excluded because they did not self-report having spoken English since before the age of 7, leaving data from 31 subjects for analysis.

3.1.1.2 Materials

Fifty-seven Local Bases were drawn from Steriade (1997): 29 Local Bases with helpful Remote Bases (such as Local Base *íllustrate* with Remote Base *illústrative*) and 28 Local Bases without Remote Bases (such as *éducate* with no Remote Base in **edúc-*). I also included 62 new Local Bases: 30 Local Bases with helpful Remote Bases (such as Local Base *lábor* with Remote Base *labórious*), and 32 Local Bases without (such as *pláster*). To the extent possible, I chose Local Bases that had only one type of Remote Base stem allomorph (here, only helpful ones). The question of how multiple Remote Bases interact in Derivative formation is beyond the scope of the experiments presented here, but I take up the predictions of the Voting Bases model for such cases in section 7. Other than these restrictions, Local Bases were not chosen with any particular type of semantic or syntactic recombination restrictions in mind, other than what the author (a native speaker of mainstream North American English) deemed within the bounds of reason for a task involving creating novel forms.

Because the Local Bases contained a mixture of morphologically-complex (*-ate*-suffixed, as in *domesticate*) and morphologically-simple Local Bases (unsuffixed, as in *labor*), I used a reading-aloud task in Experiment 1 to limit participants' production choices to those of vowel quality and stress placement; stimuli from Experiment 1 are listed in table 24 in the Appendix.

Four affixes were chosen for testing: *-able, -ism, -ify,* and *-ity*. *-able* and *-ism* were combined with the 57 Local Bases drawn from Steriade (1997), and the other two were combined with the novel Local Bases, yielding 234 target Derivatives. The reason why the affixes were separated by Local Base source in this way (that is, not fully crossed) was to balance the desire to replicate Steriade's original findings using her own stimuli, with the interest in testing the generalizability of the *principle* of Lexical Conservatism to novel affixes, while keeping the total length of the experiment under an hour.

Some of the Local Base + affix combinations were existing words of English, and potentially known to participants (such as *illustrable*, largely drawn from Steriade's original stimuli) and some were designed to be novel (such as *plasterable*). For each Derivative, a carrier sentence was created which gave a periphrastic definition of the Derivative using the Local Base. For example, for the Local Base *illustrate* and the affix *-ism*, the carrier sentence was "An ideology which centers on illustrating could be called illustrism."

3.1.1.3 Procedure

Participants completed the experiment individually in a sound-attenuated room in the presence of a member of the study team. Participants were assigned to one of four randomization lists, and were told that they would be reading definitions of possible new English words. They were advised that some of the words might sound a little unusual, or might not be exactly how they'd choose to express a certain concept (for example, one might prefer to call an ideology centered around illustrating "illustrationism" or maybe just "an illustration cult"), but that they should pronounce the stimuli however felt most natural to them. Participants were instructed to read the sentence to themselves silently in their head, and then say the last word of the sentence — the Derivative — out loud. After the researcher guided participants through six practice trials, they completed the main task trials at their own pace. Each participant saw all Local Bases with all affixes.

After the Lexical Conservatism task, participants were asked to read each Local Base out loud and indicate whether they knew the word or not. They were instructed that "knowing the word" meant that they wouldn't need to stop and ask what the word meant if they heard it in conversation. Participants were also instructed that if they felt sure hadn't heard the word before but could guess its likely meaning from its parts (i.e., its morphemes), they should still indicate that they did not have prior knowledge of the word. After the list of Local Bases, participants were asked to read aloud and indicate whether they knew each Remote Base, for the half of Local Bases which had them. The experiment concluded with a short language-background questionnaire.

3.1.2 Data processing and analysis

Each Derivative was coded for whether it underwent stress shift or its stress matched the Local Base. Each Derivative was also coded for whether the participant indicated knowing the Local Base and (if there was one) the Remote Base. Trials on which the participant did not know the Local Base were excluded from analysis. I also annotated the list of Local and Remote Bases for various phonological qualities, including target syllable weight and secondary stress placement, to include in the statistical analysis. This annotation was performed on the list of stimuli submitted to the speaker, rather than their productions, and so is uniform across speakers.[1]

Statistical analysis was carried out in R (R Core Team, 2021) using Bayesian hierarchical logistic regression implemented using the *brms* package (Bürkner, 2017). Bayesian models estimate the posterior distribution of credible values for the statistical parameters of interest by integrating prior information (if any) about the likely values of the parameters with information in the data being analyzed. For a linguistically-oriented introduction to Bayesian methods for both theory-building and data analysis, see Nicenboim and Vasishth (2016); for tutorial materials on the `brms` package in a linguistic context, see Vasishth et al. (2018); Nalborczyk et al. (2019); for a more general primer in Bayesian modeling, see Kruschke (2014).

Common summary statistics for the posterior are the median value and the range of values contained in the central 95% the distribution known as a 95% Credible Interval (abbreviated "95% CI"). Another way of assessing the evidence for an effect is to calculate the proportion of posterior credible values which lie to one side of zero, denoted $P(|\hat{\beta}| > 0)$. This measure ranges from 0.5 (equal evidence for an effect in the opposite direction of the coefficient as one in the direction of direction of the coefficient) to 1 (extremely strong evidence for a nonzero effect in the direction of the parameter's coefficient). Both methods are reported in this paper. Full details of the models fit, including posterior samples, are provided in the supplementary materials.

3.1.3 Results

First, we find that the experimental data replicate the asymmetry between Derivatives with helpful Remote Bases and those with no Remote Bases described by Steriade (1997). Figure 1 plots stress shift in Derivatives by whether or not the participant knew the Remote Base. The facets of the plot divide the data by source, novel to this experiment vs. taken from Steriade (1997).

Local Base	Derivative	Remote Base	Source
organ	*organify, organity*	*organic*	Novel
tartan	*tartanify, tartanity*	–	

[1] A reviewer notes some deviation in their judgment from the coding of secondary stress in certain stimuli. Because this is not unexpected given the literature on the topic (for an overview, see Tokar, 2018), the raw data underlying all analyses in the paper is available in the supplementary materials, and so readers can experiment with the degree that alternative placements of secondary stress impacts the conclusions drawn here.

demonstrate	*demonstrable, demonstrism*	*demonstrative*	Steriade (1997)
venerate	*venerable, venerism*	–	

Table 3: Examples of Derivatives for categories displayed in figure 1.

Although the effect of the Remote Base is robust in both groups, it is also clear that the effect is not categorical. This implies that participants who know the relevant Remote Base still fail to produce stress-shifted Derivatives from time to time, and that sometimes participants who don't know the Remote Base still sometimes violate faithfulness to Local Base to repair the long lapse. I find that two types of factors condition the variation above and beyond the presence of a Remote Base: phonological, which I discuss here in the context of the same regression model used to verify the effect of the Remote Base, and lexical which I discuss in section 3.4. Because the characteristics of the stimuli differed substantially between those from Steriade and those which are novel — including in morphological complexity of the Local Base — it is not clear to which factor we should attribute the difference in overall propensity to undergo stress shift, which is higher in novel items. Nevertheless, the effect of the Remote Base remains, cross-cutting this distinction. We will see that in Experiments 2 and 3, which use all novel items, the stress shift rate is in line with novel items from Experiment 1.

It is well known that syllable weight and secondary stress placement are both related to primary stress placement in English (cf. Chomsky and Halle, 1968; Hayes, 1982; Burzio, 1994; Pater, 2000). Affixes also influence stress placement with some theories proposing a binary distinction between stress-neutral and stress-affecting affixes (Siegel, 1974), or proposing by-affix propensities to trigger stress shift (cf. Zuraw and Hayes, 2017; Zymet, 2018; Shih, 2018). Figure 2 plots Derivative stress shift against these two predictors, split by the source of the Local Base. Examples for target syllable weight are provided in table 4, and for target syllable secondary stress in table 5.

Local Base	Derivative	Remote Base	Target Syllable Weight	Source
tártan	*tartanify, tartanity*	–	Light	Novel
métal	*metalify, metality*	*metállic*	Light	
spándex	*spandexify, spandexity*	–	Heavy	

[2]Throughout the paper, I calculate these measures using the *binconf()* function with defaults settings, from the *Hmisc* package in R (Harrell Jr, 2023).

ágent	*agentify, agentity*	*agéntive*	Heavy	
tólerate	*tolerable, tolerism*	–	Light	
equílibrate	*equilibrable, equilibrism*	*equilíbrium*	Light	Steriade (1997)
obfuscate	*obfuscible, obfuscism*	–[3]	Heavy	
cómpensate	*compensable, compensism*	*compénsatory*	Heavy	

Table 4: Stimuli examples for the top row of figure 2

Local Base	Derivative	Remote Base	Target syllable stress	Source
spándèx	*spandexify, spandexity*	–	Secondary	Novel
ágènt	*agentify, agentity*	*agéntive*	Secondary	
tártan	*tartanify, tartanity*	–	None	
métal	*metalify, metality*	*metállic*	None	
állòcate	*allocable, allocism*	–	Secondary	
ímprègnate	*impregnable, impregnism*	*prégnant*	Secondary	Steriade (1997)
rélegate	*relegable, relegism*	–	None	
cústody	*custodiable, custodism*	*custódian*	None	

Table 5: Examples for Derivatives shown in the bottom row of figure 2.

The top row of figure 2 demonstrates the expected effect of heavy syllables increasing the probability of a stress-shifted Derivative to satisfy the Stress-to-Weight principle

[3]When Remote Base *obfúscatory* is not known to the participant. Steriade's stimuli did not include any Local Bases with heavy target syllables that categorically lacked Remote Bases.

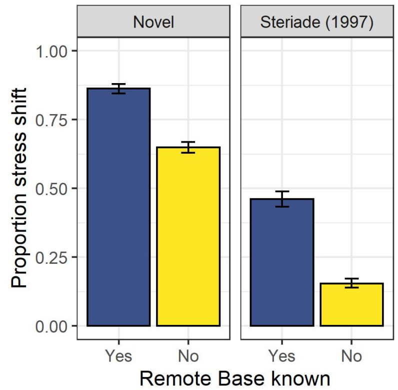

Figure 1: Means and binomial confidence interval[2] of the proportion of stress shift in Derivatives in Experiment 1, split by knowledge of the Remote Base (horizontal axis) and source of the Local Base (facets).

(Pater, 2000). The bottom row demonstrates the effect of secondary stress on Derivative stress placement. This influence, I argue, comes from the diminished faithfulness penalty for promoting to full stress a vowel which already has secondary stress in the Local Base. We also find that affixes differ in their propensity to trigger stress shift (not shown).

To examine the contribution of these factors to Derivative stress placement, I fit a model which predicted whether the Derivative underwent stress shift relative to the Local Base stress (levels 1 = undergoing stress shift, 0 = matching Local Base stress) on the basis of whether the Remote Base was known (levels 1 = *yes*, 0 = *no*), and Affix (levels *-able, -ism, -ity, -ify*). I also used two metrical well-formedness constraints, both referring to the status the target syllable, where stress would fall if it were shifted: the weight of the target syllable (levels 0 = *light* as in *indicate*, 1 = *heavy* as in *inundate*), and whether the target syllable hosted secondary stress in Local Base (levels 0 = *no* as in *cústŏdy*, 1 = *yes* as in *ícòn*). The model also included random intercepts for subject, Local Base, and unique Derivative combination of affix and Local Base, with random slopes of all fixed effects by subject, and of Remote Base Known and Affix by Local Base. Table 6 contains a summary of the fixed effects of the model.

The statistical model indicates that relative to the phonologically-bland intercept, having a heavy target syllable, or one with secondary stress, increases the log-odds of shifting stress rightward. We judge this based on the fact that the upper and lower bounds of the 95% Credible Interval (CI) exclude zero, and the fact that almost all credible values for the parameter lie to one side of zero ($P\left(|\hat{\beta}| > 0\right)$); we also find affix-specific effects. Critically, we can conclude that there is compelling evidence for the role of knowledge of the Remote Base: we conclude this based on the observation that the posterior distribution of likely values for the parameter falls quite far from zero, with

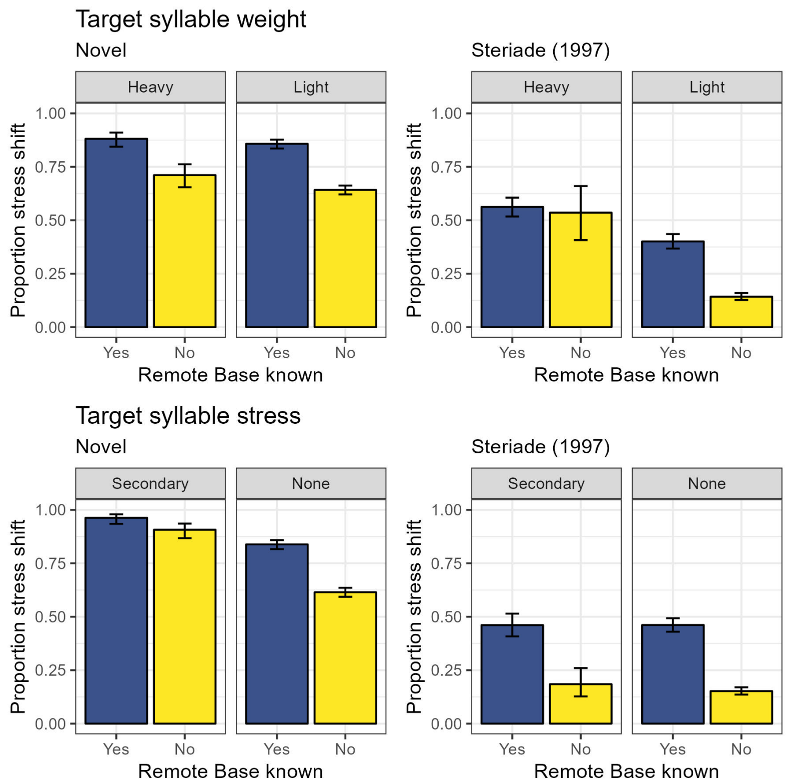

Figure 2: Means and binomial confidence interval from Experiment 1 of the proportion of Derivative stress placement (vertical axis), split by source of the Local Base (between plots), weight of the target syllable (facets within plots), and knowledge of the Remote Base (horizontal axis within plots).

a positive coefficient indicating an increase in stress shift for Derivatives with a known Remote Base.

In sum, Experiment 1 confirms that the effect of the Remote Base replicates in both Steriade's original stimuli, and that it extends to new stimuli where the Derivatives are entirely novel. The effect is also probabilistic, and that Remote Base makes itself known as one effect among many other well-known phonological ones that jointly influence stress placement in the Derivative.

3.2 Experiment 2

Experiment 2 adds a manipulation to see if priming can moderate the influence of the Remote Base. This acts as a more stringent test for the role of the Remote Base in Deriva-

Parameter	*Median*	*95% CI*	*P* $(\lvert\hat{\beta}\rvert > 0)$
Intercept: Affix = *-able*, Know Remote Base = *no* Target Syllable Heavy = *no*, Target Syllable Secondary Stress = *no*	-2.21	[-2.93, -1.50]	
Target Syllable Heavy = *yes*	0.86	[0.11, 1.66]	0.99
Target Syllable Secondary Stress = *yes*	1.27	[0.46, 2.15]	≈ 1
Affix = *-ify*	2.37	[1.53, 3.21]	≈ 1
Affix = *-ism*	-0.58	[-0.95, 0.22]	≈ 1
Affix = *-ity*	4.17	[3.34, 5.05]	≈ 1
Know Remote Base = *yes*	1.30	[0.74, 1.84]	≈ 1

Table 6: Model of Experiment 1, all Local Bases. Coefficients are in log-odds, with positive values indicating an increase in stress shift relative to the intercept.

tive formation, and can tell us whether our model of the mechanisms underlying Lexical Conservatism needs to be sensitive to only static characteristics of the Remote Base (existence, as well as possibly long-run frequency and semantic similarity to the Local Base), or both static and dynamic factors, such as the resting activation of the Remote Base in the lexicon in the moment the Derivative is formed. This can allow us to distinguish between two different theoretical mechanisms for implementing the relationship between Local and Remote Bases.

Cases where only static factors of the Remote Base matters, and Derivatives with primed Remote Bases shift stress at a rate similar to Derivatives with unprimed Remote Bases, are compatible with a representational account. Local Bases that have Remote Bases might be represented differently than those with Remote Bases, perhaps with a specific diacritic attached during acquisition when morphological relation is established, similar to ideas by (Bermúdez-Otero, 2017). This would enable them to be sensitive to the presence of the Remote Base, but not its real-time status in the lexicon.

If, on the other hand, both static and dynamic factors influence Derivative formation, a representational account where the Remote Base is not actively co-present in real time with the Local Base during the Derivative formation process is ruled out, and we must consider the specific mechanisms that allow multiple Bases in the lexicon to jointly influence the phonological grammar during Derivative formation.

Experiment 2 also more systematically controls the characteristics of the Local Bases, and fully crosses Local Bases with affixes so as to examine the effect of Steriade's original affix *-able* in entirely novel Derivatives.

3.2.1 Methods

3.2.1.1 Participants

Participant population, recruitment, screening, and compensation was the same as in Experiment 1. 34 participants were recruited and 4 excluded, leaving data from 30 participants to be analyzed.

3.2.1.2 Materials

Local Bases were 40 disyllabic nouns of English, balanced for target syllable weight, secondary stress, and presence of a Remote Base. Some of the Local Bases were also used in Experiment 1. All of the Local Bases were free-standing stems; there was no need for participants to strip an affix such as *-ate* from the Local Base *illustrate* to access the appropriate morphological stem for the intended Derivative. As in Experiment 1, Local Bases were chosen to have only one stem allomorph among Remote Bases. Stimuli for Experiment 2 are listed in table 25 in the Appendix. Two affixes were selected — *-able* and *-ic* — and were fully crossed with Local Bases so that each participant saw each affix attached to each Local Base. This yielded 80 unique Local Base + affix pairs in the Lexical Conservatism task. To accommodate the priming intervention in the experiment, Remote Bases were divided into two groups at random for separate vocabulary checks.

3.2.1.3 Procedure

Experiment 2 was conducted over the internet using the Experigen experimental platform (Becker and Levine, 2020). Participants were encouraged to sit in a quiet room and use headphones for the duration of the experiment, and their spoken responses were recorded. Instructions and practice trials were similar to Experiment 1.

Unlike Experiment 1, Experiment 2 used auditory presentation of Local Bases to avoid possible orthographic influences on responses. Affixes were displayed orthographically on screen, and participants were introduced in the instruction period to the task of hearing a Local Base and combining it with a written affix to form a Derivative, which they spoke aloud. Before the Lexical Conservatism task, participants completed a pre-task vocabulary check for all Local Bases, and the half of the Remote Bases to prime them. They then completed the Lexical Conservatism task, and afterwards a post-task vocabulary check for the other half of the non-primed Remote Bases, and a short language background survey. The experiment was entirely self-paced and took approximately

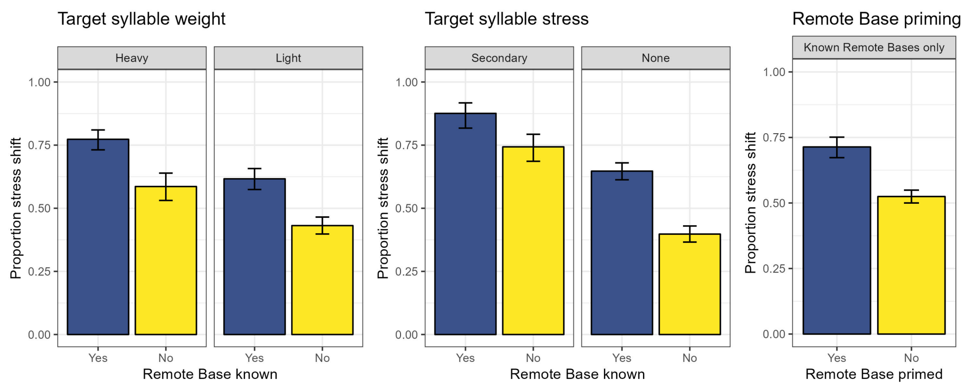

Figure 3: **Left and center:** Mean and binomial confidence intervals from Experiment 2 of the proportion stress shifted Derivatives (vertical axis) by target syllable weight (left), and target syllable stress (center), divided according to whether the participant knew the Remote Base (horizontal axis within each plot facet). Examples are listed in table 7. **Right:** Mean and binomial confidence intervals from Experiment 2 of the proportion stress shifted Derivatives (vertical axis) by whether the Remote Base was primed (horizontal axis). An example of this type are the Derivatives *laborable* and *laboric*, with the Local Base *lábor*, and the Remote Base *labórious* either known and also primed (left bar) or known but not primed (right bar).

40 minutes. As in Experiment 1, each participant saw all Local Bases with all affixes, regardless of their counterbalancing group in the pre-task vocabulary check.

Data processing and analysis followed Experiment 1, with the with the addition that the combinations of Local Bases *person* and *habit* with the affix *-able* were excluded because they were real words. All details of the model fit, including posterior samples, can be found in the supplementary materials.

3.2.2 Results

3.2.2.1 Confirming Lexical Conservatism and phonological determinants of Derivative stress

As in Experiment 1, we find that both lexical and phonological factors influence Derivative stress placement. These are plotted in the left two graphs in figure 3, and statistical analysis is reported in table 8. In contrast to the items from Steriade's original paper used in Experiment 1, these stimuli undergo stress shift at a relatively higher rate, more similar to the novel stimuli in Experiment 1.

Local Base	Derivative	Remote Base	Target syllable
scáffold	*scaffoldable, scaffoldic*	–	Heavy
túmult	*tumultable, tumultic*	*tumúltuous*	Heavy
pláster	*plasterable, plasteric*	–	Light
hábit	*habitable, habitic*	*habítual*	Light
nýlòn	*nylonable, nylonic*	–	Secondary stress
ínsèct	*insectable, insectic*	*insécticide*	Secondary stress
vélvet	*velvetable, velvitic*	–	No stress
coúrage	*courageable, courageic*	*courágeous*	No stress

Table 7: Examples for stimuli in the left two graphs of figure 3.

Also replicating Experiment 1, the presence of the Remote Base in an individual participant's lexicon leads to a higher rate of stress shift. This basic finding of Lexical Conservatism is again gradient, and sits alongside familiar phonological markedness avoidance effects and affix-conditioned behavior.

These findings were confirmed using a regression model fit to all Derivatives with known Local Bases. The dependent variable was whether the stress placement in the Derivative matched that of the Local Base (= 1) or the Remote Base (= 0), and the model included as fixed effects the weight of the target syllable (*light* = 0 vs. *heavy* = 1), whether the target syllable bore secondary stress (*no* = 0 vs. *yes* = 1), Affix (*-able* = 0 vs. *-ic* = 1), and whether the subject knew the Remote Base (*no / none exists* = 0 vs. *yes* = 1). The model contained random intercepts for subject, Local Base, and Local Base + Affix combination, with random slopes of all fixed effects by subject, a random slope of affix and whether the Remote Base was known by Local Base, and a random slope of whether the Remote Base was known by Local Base + Affix combination.

We see expected effects of target syllable weight and secondary stress on the propensity of speakers to produce Derivatives with right-shifted stress. We also observe, as expected, that *-ic* is much more strongly stress-attracting than *-able* (Marchand, 1960). Importantly, we replicate the findings of Experiment 1 in that knowledge of the Remote Base is associated with a higher probability of rightward stress shift, seen by the positive coefficient in the last line of table , the distribution of which excludes zero.

Parameter	*Median*	*95% CI*	$P(|\hat{\beta}| > 0)$
Intercept:			
Affix = *-able*			
Target Syllable Heavy = *no*			
Target Syllable Secondary Stress = *no*			
Know Remote Base = *no*	-1.45	[-2.00, -0.90]	
Affix = *-ic*	1.77	[1.18, 2.34]	≈ 1
Target Syllable Heavy = *yes*	0.48	[-0.19, 1.12]	0.92
Target Syllable Secondary Stress = *yes*	1.82	[1.07, 2.61]	≈ 1
Know Remote Base = *yes*	1.21	[0.66, 1.74]	≈ 1

Table 8: Model of Experiment 2, all Local Bases. Coefficients are in log-odds, with positive values indicating an increase in probability of stress shift relative to the intercept.

3.2.2.2 Priming the Remote Base

Turning to the point of interest for Experiment 2, we also find that a primed Remote Base exerts a greater impact on Derivative formation than an unprimed one; this is plotted in the right graph of figure 3, but statistical analysis is reserved for section 3.4.

In summary, Experiment 2 replicated the core findings of Experiment 1, and extended it by demonstrating that the Remote Base was accessed in the process of forming the Derivative, indicated by its influence being able to be manipulated by priming.

3.3 Experiment 3

As noted in section 2.1, Experiment 3 and Experiment 4 form a minimal pair testing the limits of the influence of the Remote Base. Experiment 3 probes whether a Remote Base that, if used as a model for the Derivative, would *increase* markedness relative to the more straightforward option of remaining faithful to the Local Base. An example of this type of harmful Remote Base is *résident* to a Local Base *resíde*: here, the Remote Base has a stress pattern that, if adopted in a form like *résidable*, would be even more marked than simply ignoring the Remote Base (yielding *resídable*).

3.3.1 Methods

3.3.1.1 Participants

Participant population, recruitment, screening, and compensation was the same as in Experiment 2. 54 participants were recruited, and 23 were excluded (15 for not having spoken English since before the age of seven consistently in some context, and 8 for technical problems relating to the sound quality), leaving 31 participants with data included in this study.

3.3.1.2 Materials

I report here results from a set of 50 disyllabic Local Bases, 20 with initial stress (ex., *cárrot, coúrage, hábit*) and 30 with final stress (ex., *presérve, propóse, províde*). Within each group, some Local Bases had Remote Bases with stress placed on the other syllable of the Base: for example, 10 of the initially-stressed disyllabic Local Bases had Remote Bases with final stress (ex. *coúrage* ∼ *courágeous, hábit* ∼ *habítual*), and 15 of the finally-stressed disyllabic Local Bases had Remote Bases with initial stress (ex. *províde* ∼ *próvidence, resíde* ∼ *résident*). As in the previous experiments, each Local Base was chosen with the goal of having only one type stem allomorph as a Remote Base.

I selected two affixes, *-able* and *-ist*, based on the description in Marchand (1960) that the suffixes were not stressed nor obligatorily stress-attracting (cf. also Aronoff, 1976). To avoid possible confounds associated with non-standard selection frames in Experiments 1 and 2, a mixture of the two affixes was paired with disyllabic Local Bases, depending on the lexical category (nouns or adjectives). Stimuli for Experiment 3 are listed in table 26 in the Appendix.

3.3.1.3 Procedure

Experiment 3 was conducted over the internet using the Labvanced experimental platform (Finger et al., 2017). Instructions, procedure, structure, and data annotation were identical to that of Experiment 2. Due to an error in the configuration of the randomization structure for the trials, each participant only saw a randomly-selected subset of 80 out of the 95 unique Lexical Conservatism trials; vocabulary-check trials were not affected. Since the missing trials were distributed randomly among item types and subjects, I do not judge this to be a reason to believe the results of Experiment 3 should be biased in any particular direction. Data exclusion criteria and statistical analysis followed Experiments 1 and 2, with the addition that trials of the Derivative *opposable* were excluded for being dictionary-listed, and therefore likely known to participants.

3.3.2 Results

In this set of Local Bases, we're interested in whether speakers treat harmful Remote Bases (ex., *résident*, a phonologically non-optimizing Remote Base for Local Base *resíde*) in the same way that they do helpful ones (ex., *habítual*, a phonologically-optimizing Remote Base for Local Base *hábit*). These results are displayed in the left pair of graphs in figure 4.

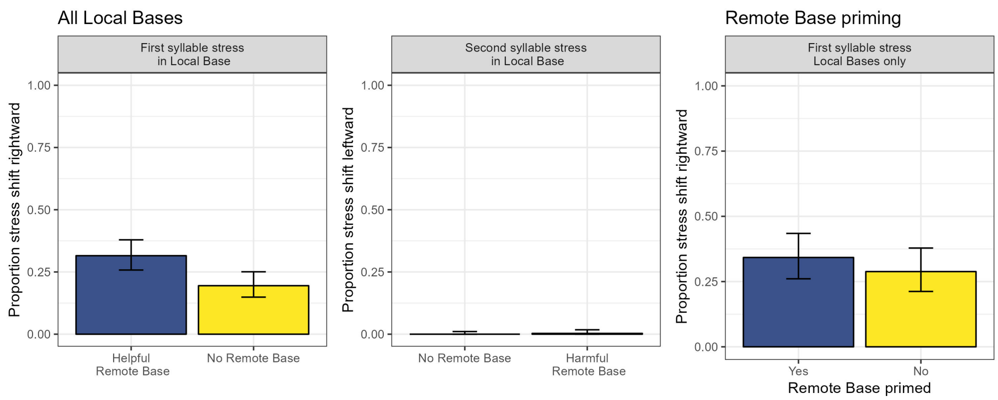

Figure 4: **Left and center:** Results of the Lexical Conservatism task with Local Bases from Experiment 3. The vertical axis plots the mean and binomial confidence interval of the proportion of stress shift in Derivatives, split by the stress of the Local Base (facets), and the type of Remote Base (horizontal axis within facets). Examples are as listed in table 9. **Right:** Results of priming the Remote Base in Experiment 3, only initially-stressed Local Bases. Mean and binomial confidence intervals from Experiment 3 of the proportion stress shifted Derivatives (vertical axis) by whether the Remote Base was primed (horizontal axis).

Local Base	Derivative	Remote Base	Remote Base type
pláster	*plasterable*	–	None
lábor	*laborable*	*labórious*	Helpful
finésse	*finessable*	–	None
resíde	*residable*	*résident*	Harmful

Table 9: Example stimuli for Derivatives plotted in figure 4.

Local Bases with helpful Remote Bases underwent stress shift at a rate higher than the baseline rate of unfaithfulness in Local Bases without any Remote Bases, while Local Bases with harmful Remote Bases did not form Derivatives that were meaningfully different from Local Bases without any Remote Bases. This is plotted in the left and center of figure 4. Strikingly, out of 340 trials where Derivatives were formed to final-stressed disyllabic Local Bases — those with harmful Remote Bases — only two were attested, *ópposist* and *ímposist*, and these occurred on trials when the relevant Remote Bases (*ópposite* and *ìmposítion*) were primed. Turning to the effect of priming in initial-stress Local Bases (right of 4), we see effects largely in line with the results of Experiment 2: some evidence in favor of primed Remote Bases yielding more Derivatives with stress mismatching the Local Base.

In the statistical model, the stress of Derivatives formed to disyllabic Local Bases was predicted by fixed effects of whether the target syllable was heavy, secondarily stressed, and whether the Remote Base was known to the participant. Because speakers treated Local Bases with helpful and harmful Remote Bases very differently, I fit the model only to those Derivatives formed to Local Bases with initial stress; a model with all Derivatives and Local Base stress placement as an interacting factor is available in the supplementary materials. As with Experiment 2, I reserve statistical analysis of the priming effect for section 3.4, immediately below.

Parameter	*Median*	*95% CI*	*P* $(\lvert\hat{\beta}\rvert > 0)$
Intercept: Target Syllable Heavy = *no* Target Syllable Secondary Stress = *no* Know Remote Base = *no*	-3.99	[-5.63, -2.65]	
Target Syllable Heavy = *yes*	2.37	[0.47, 4.61]	0.99
Target Syllable Secondary Stress = *yes*	2.43	[0.05, 4.92]	0.98
Know Remote Base = *yes*	1.28	[-0.16, 2.81]	0.96

Table 10: Model of Experiment 3, all disyllabic Local Bases with initial stress. Coefficients are in log-odds, with positive values indicating an increase in stress shift relative to the intercept.

The statistical model confirmed the observations made above: the range of credible values associated with knowing the Remote Base fell almost entirely above zero, indicating that we have strong evidence that the Remote Base played a meaningful role in stress shift. This finding sits alongside the consistent effects of syllable weight attracting stress, and secondary stress in Local Bases making better targets for primary stress in the Derivative.

In summary, Experiment 3 showed that if a Remote Base would result in *greater* markedness in the Derivative than remaining faithful to the Local Base — that is, if it is harmful — it seems to play very little role in Derivative formation. This is compatible with the traditional assumption of Lexical Conservatism, that only Bases which are markedness-reducing helpful should play a role in Derivative formation. What is unexplored, however, is whether Remote Bases that do not change the markedness of the Derivative relative to the Local Base nevertheless play a role in determining the rate of repair — that is, if they are merely *unhelpful*. This question is addressed in Experiment 4 in section 4.

3.4 Combined analysis of lexical characteristics of English

We now turn to the combined analysis of the influence of lexical characteristics of Remote Bases on Derivatives using combined data from Experiments 1, 2, and 3. The aim of the group analysis is to aggregate evidence from across the three English experiments, yielding a larger sample size (3,929 data points) and a wider range of semantic similarities and frequencies.

I fit a mixed-effects Bayesian logistic regression model to the data from Local Bases with helpful, known Remote Bases from Experiment 1, 2, and 3. The model had a dependent variable of Stress Shift (*yes* = 1, *no* = 0), and fixed effects of whether the target syllable was heavy (*yes* = 1, *no* = 0), whether the target syllable was secondarily-stressed (*yes* = 1, *no* = 0), whether the Remote Base was primed (*yes* = 1, *no* = 0), the source Experiment (a three-level unordered factor), and affix (a five-level unordered factor). The model also contained centered and scaled coefficients for semantic similarity and their interaction with centered and scaled Remote Base log-frequency taken from the English Lexicon Project (Balota et al., 2007).

Semantic similarity was estimated using data from a norming experiment where participants were instructed to use a 7-point Likert scale to rate how similar pairs of Local and Remote Bases were in meaning, where 7 was *extremely similar*. I collected data for all Local-Remote Base pairs in Experiments 1-3, and then used a mixed-effects model to derive item-level estimates of similarity for each pair, which were centered and scaled and used as predictor values in the model reported here. The full experiment is detailed in Breiss (2021:Appendix B).

The model fit to the combined set of Derivatives from Experiments 1-3 had random intercepts for Local Base, participant, and Derivative, with random slopes of all fixed effects except Experiment by participant, random slopes of priming and Experiment by Local Base, and Experiment by Derivative.

3.4.1 Results

For the sake of brevity, the full results of the model are presented in table 28 in the Appendix. Here, I focus on the effects of critical interest: how priming and semantic similarity of the Remote Base influenced the Derivative. I used the *emmeans()* function from the *emmeans* package in R (Lenth et al., 2019) to calculate the marginal effect of priming the Remote Base on Derivative stress shift, averaging over all levels of all other factors in the model. I found that priming increases the odds of stress shift robustly: $\hat{\beta} = 0.46$ [-0.04, 0.95], $|\hat{\beta}| > 0 = 0.97$. The median and 95% CI are plotted in the left panel of figure 5, using the *conditional_effects()* function in the `brms` package. I used the *emtrends()* function from the same package to calculate the marginal slope of the linear predictor of semantic similarity, and also found a reliable effect, but with the opposite sign: $\hat{\beta} =$ -0.30 [-0.65, 0.06], $|\hat{\beta}| > 0 = 0.95$. This indicates that semantic similarity decreases the chance of stress shift in the Derivative, even though that stress shift would led to the Derivative more closely resembling the Remote Base (median and 95% CI smooths in the right of figure 5). The direction of this effect is unexpected, and we take it up again in detail in section 6.2.1. There was no evidence suggesting any role for a main effect of Remote Base log-frequency, nor for an interaction with semantic similarity; the range of posterior credible values for these parameters straddle zero in both cases.

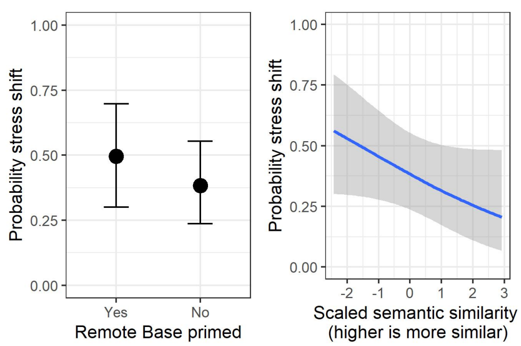

Figure 5: Plots of the marginal effects of priming (left) and of scaled semantic similarity (right). Note that since these are model-derived estimates, the point estimate represents the median of the posterior distribution, and the uncertainty the 95% Credible Interval for the posterior — thus, they naturally appear quite a lot wider than the binomial confidence intervals of the data plotted elsewhere in the paper.

3.5 Summary of English experiments

Experiments 1, 2, and 3 on English stress replicated and extended the findings of Steriade (1997). I found that the dependency between the presence a helpful Remote Base in a paradigm and the way it incorporates novel coinages replicates robustly in existing and novel Derivatives, but is probabilistic with counterexamples in both directions. Conditions on this variation include an independent role of phonological markedness avoidance, as well as static (semantic similarity) and dynamic (resting activation, as manipulated by priming) characteristics of the contents of the lexicon. Finally, Experiment

3 began a two-part exploration of non-optimizing relations between Remote Base and Derivative: in the case of Local Bases like *resíde* (harmful Remote Base *résident*), adopting the shape of the Remote Base *increases* the markedness of the Derivative relative to the Local Base. In the next section I consider the case of Spanish diphthongization which exemplifies the third possible relation of a Remote Base and Derivative, where a Remote Base has the same shape as the Local Base, and thus is neither markedness-increasing nor markedness-reducing for the Derivative.

4 Lexical Conservatism in Spanish vowels

Lexical Conservatism arises in Spanish from the distribution of mid-vowels *e* [e] and *o* [o], and their diphthongal counterparts *ie* [je] and *ue* [we]. In Experiment 4, I examine the contexts under which Spanish speakers monophthongize a diphthong when the licensing stress is moved off of it to probe whether phonologically non-optimizing Remote Bases, termed *unhelpful*, influence Derivative formation. This experiment, paired with Experiment 3 on English harmful Remote Bases, completes the exploration of possible ways non-optimizing Remote Bases could influence Derivative formation.

4.1 Background on Spanish diphthongization

All dialects of Spanish exhibit an alternation which affects some stressed diphthongs <ié>[ˈje] and <ué>[ˈwe], yielding alternation with corresponding unstressed monophthongs <e>[e] and <o>[o]. This can be seen in forms such as *truéno* ∼ *tronámos* "I thunder ∼ we thunder" and *siénto* ∼ *sentámos* "I sit ∼ we sit". The alternation is unpredictable, however, since some diphthongs do not alternate with monophthongs when unstressed, as in *miédo* ∼ *miedóso* "fear ∼ afraid" and *puéblo* ∼ *pueblíto* "(a) town ∼ (a) small town", and further not all unstressed monophthongs alternate with diphthongs under stress *podámos* ∼ *pódo* "We prune ∼ I prune" and *montámos* ∼ *mónto* "We mount ∼ I mount". The phenomenon has long been studied as an "old chestnut" of exceptionful phonology, with numerous analyses proposed focusing on different ways of encoding the distinction between alternating and non-alternating roots (Harris, 1969; Hooper, 1976; Carlson and Gerfen, 2011:among many others).

This unpredictable alternation is the result of a historical merger between low-mid *[ɛ, ɔ], which exhibited exceptionless alternation between stressed *[je, we] and unstressed *[ɛ, ɔ], and high-mid vowels *[e, o] which did not alternate with stress (Penny, 2002). Because of the markedness-reducing neutralization of diphthongs and mid-vowels in unstressed positions, it was not obvious to speakers which mid-vowels alternated in this way, and which did not. This in turn led to further remodelling and analogical change in the paradigms, and gave rise to the present state of etymologically-informed but synchronically-arbitrary alternation, which exhibits type-level variation of individual roots with different affixes (ex., for the base *puébl-o* "town", both *pueblíto* "small town" and *población* "population"), as well as token-level variation within roots (ex., *cientóso* ∼ *centóso* "muddy"). Thus, the synchronic situation is one where both roots and affixes exhibit lexical propensities to alternate when the appropriate phonological conditions are met, giving rise to a complex landscape of cross-cutting conditioning factors.

A small body of experimental work has examined how speakers extend these lexical generalizations to novel words. Eddington (1996, 1998) conducted two experiments in which speakers of Iberian Spanish were asked to attach 10 stress-attracting affixes to novel bases with a stressed diphthong. He recorded the rate at which the affixes induced monophthongization, and found rates of monophthongization varying from 4.7% for the diminutive *-(c)íllo* to 86.2% for adjectivizing *-óso*, indicating that speakers have internalized affix-specific propensity information about the contents of their lexicon. Carlson and Gerfen (2011) examines similar data on lexically-specific affix behavior, and finds evidence for a relation between the productivity of the affix and its propensity to trigger alternation. Bybee and Pardo (1981) also adduced evidence of speakers being willing to generalize the alternation to novel forms, and Albright et al. (2001) found that speakers used the distributional phonotactic information in the environment of unstressed mid-vowels to predict whether they will alternate with diphthongs under stress.

4.2 Contexts for Lexical Conservatism

Unlike English, the suffixes of Spanish are almost always stress-bearing, and so we cannot examine contexts where affixation moves stress onto an unstressed monophthong that is part of the root. However, there are many cases where stressed diphthongs in roots have their stress removed under affixation. Since unstressed diphthongs are assumed to be phonotactically marked in the literature cited above, and were also found to be dispreferred to corresponding monophthongs in a phonotactic well-formedness experiment reported in Breiss (2021:ch. 4), we can ask whether the phonological grammar shows sensitivity to other paradigm members when determining if a newly-unstressed diphthong should be monophthongized. Words with a stressed diphthong that don't have any morphological relatives with differing stress, like *siniéstro* "sinister" or *ungüénto* "ointment", constitute a control case where the behavior of the phonological grammar can be observed in isolation: any repair in these environments is due to the conflict of markedness and faithfulness, without the interference of paradigm structure.

Local Bases that have morphological relatives with differing stress placement can be divided into those where the vowel in the Remote Base is monophthongized, which I term *helpful* Remote Bases as in *niébla* ∼ *neblína* "fog ∼ mist" or *muéble* ∼ *moblár* "furniture ∼ to furnish", which admit classical Lexical Conservatism, or is left unrepaired as in *unhelpful* Remote Bases *ambiénte* ∼ *ambientál* "environment ∼ environmental" or *juérga* ∼ *juerguísta* "spree ∼ reveler". Lexical Conservatism predicts that speakers have the option of relying on the stem allomorph of Remote Bases that have an unstressed mononphthong to ease the penalty for monophthongizing the Local Base. The alternation also provides the context for an unhelpful Remote Base — a paradigm-member having an unstressed diphthong — to influence the Derivative even though it would not decrease markedness. This outcome is not predicted by classical theories of Lexical Conservatism, as discussed above in section 1.1. I summarise the types of Local Bases and the relevant aspects of their paradigm structure in table 11 below.

Local Base	Helpful Remote Base	Unhelpful Remote Base	Derivative
ungüénto, siniéstro	-	-	*ungüentóso ∼ungontóso, siniestróso ∼sinestróso*
muéble, niébla	*moblár, neblína*	-	*mueblóso ∼moblóso, nieblóso ∼neblóso*
juérga, ambiénte	-	*juerguísta, ambientál*	*juergoso ∼jorgoso, ambientóso ∼ambentóso*

Table 11: Demonstration of the paradigmatic structure and relations relevant to the current study of Spanish monophthongization.

4.3 Experiment 4

Experiment 4 asked speakers of Mexican Spanish to create novel morphologically-complex words by affixing the adjectivizing suffix *-óso* to existing nouns.

4.3.1 Methods

4.3.1.1 Participants

Thirty native speakers of Mexican Spanish were recruited using the Prolific online subject pool. Recruitment was subject to the restrictions that participants had no self-reported reading difficulties, were born and resided in Mexico at the time of the study, and identified Spanish as their first language. Participants were paid approximately $9 for their time.

4.3.1.2 Materials

Ninety Local Bases were selected for the study through the use of the dictionary *Diccionario de la Lengua Española (DLE)* and the assistance of a linguistically-trained native speaker. 45 Local Bases contained a stressed front-diphthong *ié*, and 45 contained a stressed back-diphthong *ué*. Within each diphthong set of 45, 15 Local Bases had no Remote Bases, 15 had helpful Remote Bases, and 15 had unhelpful Remote Bases; stimuli are listed in the Appendix. A phonetically-trained female native speaker of Mexican Spanish recorded each of the Local and Remote Bases for use in the experiment.

In line with Experiments 1-3, Local Bases were chosen to have Remote Bases with only one stem allomorph. In the few cases where a Local Base had two types of stem allomorphs among Remote Bases, including Local Base *huévo* "egg" with unhelpful Remote Base *huevón* "a lazy person" (slang, literally "a big egg") and helpful Remote Base *ovíparo* "egg-laying, oviparous", the unintended Remote Base (here *ovíparo*) was of a type that, if it had exerted a strong influence on Derivative formation, would have resulted in a weakened distinction between the groups of Local Bases with unhelpful Remote Bases and those with no Remote Bases. Since in aggregate the two groups undergo repair at meaningfully different rates, I conclude that individual items of this type do not undermine the conclusions we can draw from Experiment 4, since the direction of the confound (here, the helpful Remote Base *ovíparo*) cut against the direction of the effect we observe. Future work is needed to probe the behavior of Local Bases with multiple types of Remote Base; this topic is taken up again in section 7.

A single derivational affix *-óso* was chosen for the study, and Local Bases were selected such that none of the Derivatives formed through their combination with *-óso* were listed in the DLE, in an effort to ensure that as many as possible of the Derivatives in the study would be nonce-forms to the participants. I chose this affix because it caused the most stem monophthongization among affixes were tested in studies by Eddington (1996), in the hopes that it might balance out the low rates of alternation in observed in the *wug*-tests of Bybee and Pardo (1981); Albright et al. (2001).

One small further nuance comes from the fact that Spanish is a language that has thematic vowels that generally mark grammatical gender, as in *entierr<u>o</u>* or *tiend<u>a</u>*. These vowels regularly and obligatorily delete before vowel-initial suffixes, so throughout I assume that viable Derivative candidates are, for example, the forms *tiendóso* or *tendóso*, but never **tiendaóso* or **tendaóso*, since such forms are never attested in the literature, and were not observed in the experimental responses reported here.

4.3.1.3 Procedure

The experiment was conducted over the internet using the Experigen in-browser platform (Becker and Levine, 2020) with instructions, structure, and procedure identical to Experiment 2 and 3, except translated into Spanish. Each participant saw all Local Bases.

4.3.2 Data processing and analysis

Derivatives were annotated for whether the vowel they contained was an unstressed diphthong or monophthong; there were no cases where participants did not shift stress to the penult of the Derivative, nor any instances where the theme vowel was not truncated. In cases where a participant gave more than one response, I considered the last one produced as their response for that trial. Each Derivative was annotated for whether the speaker knew the Local Base and, if extant, the Remote Base. Phonetically ambiguous tokens were extremely uncommon (at most one or two per participant), and were checked with a native speaker. Data exclusion criteria followed Experiments 1-3, with the additional exclusion of responses to the Local Base *priesa* "(a) rush" due to experimenter error ($n = 30$). The statistical analysis was similar to those for Experiments 1-3, but omitted the fixed effect of priming and corresponding random slopes.

4.3.3 Results

The primary question of interest in this experiment is whether participants were sensitive to the presence of both the helpful and unhelpful Remote Bases in forming Derivatives. Figure 6 plots the proportion of Derivatives containing a monophthong according to the type of Local Base, and examples of each type are given in table 12.

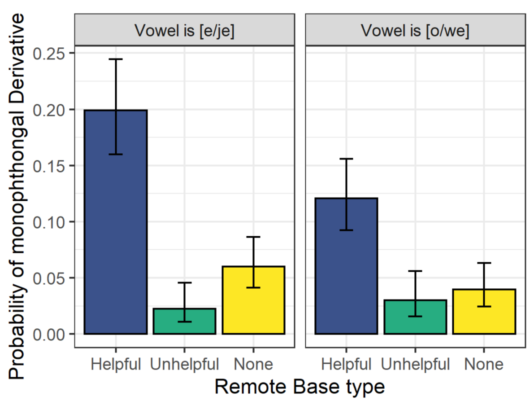

Figure 6: Results of the lexical conservatism task in Experiment 4. Vertical axis plots the proportion (mean and binomial confidence interval) of monophthongal Derivatives formed for each type of Local Base (horizontal axis); panels plot type of diphthong. Examples are in table 12.

Local Base	Derivative	Remote Base	Remote Base type	Vowel
viérne	*v(i)ernóso*	–	None	
ciérro	*c(i)erróso*	*cerrár*	Helpful	[e/je]

higiéne	*hig(i)enóso*	*higienísta*	Unhelpful	
suélo	*sue/olóso*	–	None	
vuélo	*vue/olóso*	*volár*	Helpful	[o/we]
huévo	*hue/ovóso*	*huevón*	Unhelpful	

Table 12: Examples of stimuli plotted in figure 6.

Knowing a helpful Remote Base clearly influenced participants' responses, eliciting much more monophthongization in such Derivatives compared to those without any Remote Base. This is the canonical Lexical Conservatism effect described by Steriade and others in the literature. Turning to Derivatives with unhelpful Remote Bases, there is less monophthongization than in the Local Bases without Remote Bases. The same pattern obtains across both types of diphthong, with slightly more repair in the front diphthong cases.

One quirk of the results is that overall rate of monophthong production in the study is quite low; for all categories of Local Base below twenty percent, and in most well below ten percent. This finding is puzzling given how extensive the stress-conditioned alternation of diphthongs is throughout the language, but is in line with previous experimental work on the topic, which found that the alternation was often difficult to elicit in experimental or quasi-experimental contexts (Bybee and Pardo, 1981; Albright et al., 2001).

These conclusions were verified using a Bayesian mixed-effects logistic regression model fit to the results of Experiment 4. The Derivative vowel type was the dependent variable ($diphthong = 0$, $monophthong = 1$). The model contained a three-level categorical variable of Remote Base type (*none, unhelpful, helpful*), a binary categorical variable of diphthong type ([e/je] or [o/we]), and random intercepts for subject and Local Base, with a random slope of the fixed effects by subject. The results of the model are displayed in table 13.

Experiment 4 demonstrated that the opportunity for Lexically-Conservative behavior provided by the lexicon is indeed represented in the grammars of individual speakers. Speakers showed an increased willingness to monophthongize a newly-unstressed diphthong if, for that Local Base, there exists a Remote Base in which the corresponding unstressed vowel is a monophthong. Further, we find evidence for unhelpful Remote Base activity: Local Bases with an unstressed diphthong in the corresponding Remote Base vowel are *even less* likely to form monophthongal Derivatives compared to Local Bases with no Remote Bases. The statistical model confirms this association between helpful Remote Bases and an increase in monophthongization (the posterior values of the *helpful* level of Remote Base all lie above zero), and between unhelpful Remote Bases and a decrease in monophthongization (posterior values lying below zero) relative to Local Bases with no Remote Base (the intercept). We can also conclude that the marked-

Parameter	*Mean*	*95% CI*	*P* $(\|\hat{\beta}\| > 0)$
Intercept: Diphthong type = [e/je] Remote Base = *none*	-3.80	[-4.82, -2.91]	
Diphthong type = [o/we]	-1.28	[-2.46, -0.26]	0.99
Remote Base = *unhelpful*	-1.15	[-2.66, -0.01]	0.96
Remote Base = *helpful*	1.17	[0.24, 2.07]	0.99

Table 13: Model of Experiment 4. Coefficients are in log-odds, with positive values indicating an increase in probability of diphthongization relative to the intercept.

ness of the unstressed diphthong is quite weak, since Derivatives that have no Remote Bases only monophthongize a small minority of the time This corresponds to the distinctly negative value of the intercept in the regression model in table 13, and is in line with previous experimental work on Spanish diphthongization.

The force exerted by unhelpful Remote Bases on Derivatives is unexpected from the traditional point of view of Lexical Conservatism, since it flies in the face of the markedness-reducing goal of the phenomenon as originally formulated, demonstrated in Experiments 1-3. Experiment 4 demonstrates that a theory that only pays attention to *some* Remote Bases is not sufficient to model the data; at the same time, though, the theory must be sufficiently restrictive to not allow *any* Remote Base to influence the Derivative, since we observed in Experiment 3 that harmful Remote Bases exerted next-to-no effect on Derivatives. In the next section, I propose a phonological model that is able to accommodate the coexistence of weak markedness with an attractive effect of the unhelpful Remote Base (the case in Spanish), and forbids it in cases where the Remote Base is unhelpful and markedness is stronger (as in English).

5 The Voting Bases model

In this section, I lay out a new theory of Lexical Conservatism that is based on the principle that each Base gets a "vote" in how the Derivative is realized. These competing demands are modeled using multiple faithfulness constraints, and I demonstrate that both the English and Spanish data follow from the same principles of the theory under different strengths of markedness and faithfulness. In order to simplify the exposition, I first outline the principles of the theory using schematic examples, then scale up to a full model of first the Spanish, then the English data using Maximum Entropy Harmonic Grammar (Goldwater and Johnson, 2003). I also briefly return to a formal model of how the Spanish data in particular allow us to adjudicate between the proposed voting theory and the two other theories of Lexical Conservatism in the literature discussed in section 5.4. Note that in this paper, the role of formal modeling is to illustrate the tenets of the theory and how it derives data patterns from Experiments 1-4; I leave the role of learning in setting constraint weights in response to corpus data of attested paradigms for future work (see also section 7.

5.1 Stage-setting assumptions

Before turning to the core theoretical proposal, I lay out and motivate some assumptions about the nature of the grammar and the lexicon that underpin it.

With regard to phonological formalism, the non-categorical nature of the phenomenon indicates we must employ a probabilistic framework: Derivative repair is not obligatory, and is attested even in the absence of a Remote Base. Further, the fact that the form of the Derivative is jointly conditioned by multiple phonological factors suggests that a weighted-constraint model which derives constraint cumulativity by default is most appropriate. I use the Maximum Entropy Harmonic Grammar framework (Smolensky, 1986; Goldwater and Johnson, 2003) to implement my analysis, although in principle a model using the Stochastic Harmonic Grammar framework (also called Noisy Harmonic Grammar, (Boersma and Pater, 2016)) might also be possible.

With regard to the contents and structure of the lexicon, I follow Steriade's position: "[a]ny non-nonce word, any non-hapax form is, I assume, accessible as a base of affixation for the creation of a novel form. In other terms, I assume that any non-nonce form is lexically recorded..." (Steriade, 1997:p. 2). In the decades since her paper, this position has been largely vindicated by psycholinguistic research. Evidence for whole-sale listing comes from the work of Bybee and Pardo (1981); Hay (2003); Hay and Baayen (2005) and others. The lexicon redundantly lists regular stem allomorphs, evidenced by frequency effects in processing that are sensitive to sub-lexical morphological structures (Hay, 2001). Finally, there is also evidence that the lexicon contains robust amounts of word-specific phonetic detail (see evidence summarized in Pierrehumbert, 2016).

In sum, ample evidence has accumulated that the lexicon redundantly lists both regular and irregular, whole and decomposed forms. I rely on this literature to support the existence of listed stem allomorphs as available "inputs" to the phonological grammar, which I make use of in the model directly below.[4]

[4]A reviewer notes that, in contrast to the Voting model, one might see the data presented in this paper as an instance of Output-Output faithfulness, following Benua (2000). To my knowl-

In modeling the effects of priming on Derivative formation, it will also be important to make reference to the concept of *resting activation*, a scalar quantity that reflects how "active" or "accessible" a specific lexical entry is. This construct is core to almost all psycholinguistic theories of language production, comprehension, and acquisition, and is grounded in a broad and deep literature; for overviews see Levelt (1999); Storkel and Morrisette (2002); Weber and Scharenborg (2012). I discuss the integration of this psycholinguistic quantity into the phonological grammar in more detail in section 6.2.1.

5.2 The core proposal

The model has two basic principles. First, each listed allomorph gets a say (that is, a "vote") in how the Derivative is realized, regardless of its status with respect to markedness (optimizing (helpful), non-optimizing (unhelpful), or anti-optimizing (harmful)). This influence is operationalized via multiple faithfulness constraints enforcing identity between the Derivative and different listed allomorphs, each scaled by the resting activation of that lexical item. Second, markedness constraints evaluating candidate Derivatives cross-cut this network of faithfulness constraints.

I use an extended schematic example to illustrate the basic workings of the model. First, let us consider a single Local Base with no Remote Bases, with two candidates, one that undergoes a markedness-improving alternation (the *changed candidate*) and the other which does not (the *faithful candidate*).[5] There are two constraints, Faith-Local, that is violated by the changed candidate, and Markedness, which is violated by the faithful candidate. This scenario is demonstrated in table 14 below.

/Local Base/ Weight:	Faith-Local 1	Markedness 1
a. *faithful*		1
b. *changed*	1	

Table 14: A tableau illustrating the schematic violation profile of a Derivative to a Local Base with no Remote Bases.

edge, such theories are not explicit with regard to whether they assume the output form being referenced is derived on-line in parallel with the form it influences, or whether it is stored, so I do not lean on this angle here. In either case, the novel contribution of the Voting theory is to tie the degree of influence the Remote Base has to the psycholinguistic quantity *resting activation*, which is not addressed in traditional Output-Output faithfulness analyses.

[5] Note that in this terminology, "faithful" and "changed" are defined with reference to the Local Base, and are meant purely as descriptive terms for what happened to the Derivative relative to the phonologically-uninteresting "default" outcome of being faithful to the Local Base, when there are no strong markedness constraints or Remote Bases involved.

In this scenario, we can note that although there are two constraint weights to set, it is only the difference between the two that is critical. The graph below in figure 7 demonstrates that as the weight of MARKEDNESS changes with FAITH-LOCAL held constant, the probability of the changed candidate likewise differs, from low when MARKEDNESS is below FAITH-LOCAL, to middling when the two weights are equal, to high when the weight of MARKEDNESS exceeds that of FAITH-LOCAL. The shape that the relationship traces is the well-known *sigmoid* of the logistic curve relating Harmony to probability in a MaxEnt grammar; for an extensive review of the relevance of this shape to phonology specifically and the places it appears in linguistic phenomena more broadly, see Hayes (2022).

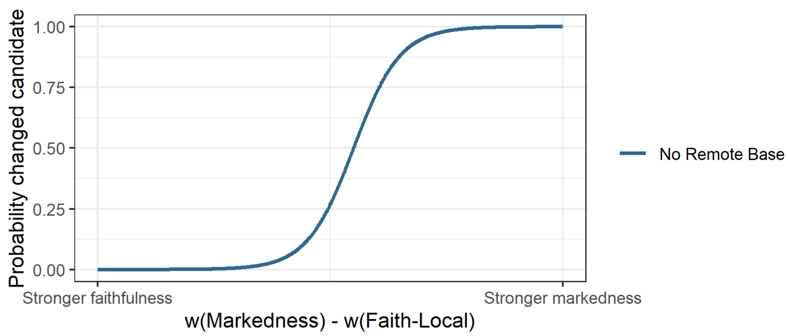

Figure 7: Schematic change in probability of the changed candidate based on the difference between the weight of MARKEDNESS and FAITH-LOCAL.

Consider next a Local Base with a helpful Remote Base as in table 15. We see that the presence of a helpful Remote Base increases the odds of Derivatives resembling it: classical Lexical Conservatism. Here, the faithful candidate violates not only MARKEDNESS, but also FAITH-REMOTE, a constraint enforcing faithfulness to the Remote Base. The degree to which FAITH-REMOTE is less than FAITH-LOCAL governs the strength of attraction of the Remote Base — that is to say, the horizontal displacement of the sigmoids from one another (for the sake of this schematic example, I hold this value constant). This is implemented in the tableau in figure 8.

Adding the new tableau to the visual typology in 8, we see that its changed candidate is consistently higher (greater probability) than that of one with no Remote Base. This follows from the lower harmony penalty because the changed candidate with a helpful Remote Base satisfies both MARKEDNESS and and FAITH-REMOTE.

We can also add lines for the different types of non-optimizing Remote Base. First let us consider the unhelpful Remote Base, where the Remote Base has the same markedness as the faithful candidate. This is seen in the tableau below.

For any vertical position on the horizontal axis, the line for the changed candidate is consistently below (less probable than) the changed candidate for the Local Base with

/Local Base/ Weight:	Faith-Local 1	Markedness 1
a. *faithful*		1
b. *changed*	1	

/Local/, /Helpful Remote/ Weight:	Faith-Local 1	Faith-Remote 1	Markedness 1
a. *faithful*		1	1
b. *changed*	1		

Table 15: Tableaux illustrating the schematic violation profile of a Derivative to a Local Base with no Remote Base (top) and a helpful Remote Base (bottom), for a fixed weight of Faith-Remote.

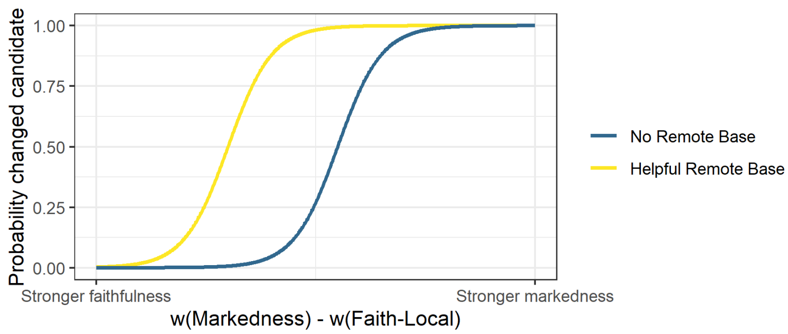

Figure 8: Schematic change in probability of the changed candidate based on the difference between the weight of Markedness and Faith-Local, for a Local Base with a helpful Remote Base and for a Local Base with no Remote Base.

no Remote Base. This again follows from its increased harmony penalty relative to the changed candidate with no Remote Base: the changed candidate with an unhelpful Remote Base is more penalized because it violates both Faith constraints.

Finally, we add the schematic line for the harmful Remote Base: these are cases where, if the Remote Base were adopted in the Derivative, the changed candidate would actually be *more* marked than the faithful candidate.

Plotting the probability of the changed candidate in figure 10 reveals that the changed candidate influenced by the harmful Remote Base is only likely when the weight of Faith-Local is high relative to Markedness, and even then is somewhat marginal.

We are now in a position to apply the Voting Bases model to the experimental data. I first fit the Spanish data because it illustrates the symmetry between helpful and unhelpful Remote Bases, and because it is a case where the Voting Bases theory of Lexical Conservatism makes different formal predictions than those of Steriade (1997, 2000).

/Local Base/ Weight:	FAITH-LOCAL 1	FAITH-REMOTE	MARKEDNESS 1
a. *faithful*			1
b. *changed*	1		
/Local/, /Helpful Remote/ Weight:	**FAITH-LOCAL 1**	**FAITH-REMOTE 1**	**MARKEDNESS 1**
a. *faithful*		1	1
b. *changed*	1		
/Local/, /Unhelpful Remote/ Weight:	**FAITH-LOCAL 1**	**FAITH-REMOTE 1**	**MARKEDNESS 1**
a. *faithful*			1
b. *changed*	1	1	

Table 16: Tableaux illustrating the schematic violation profile of a Derivative to a Local Base with no Remote Base (top), a helpful Remote Base (center), and an unhelpful Remote Base (bottom).

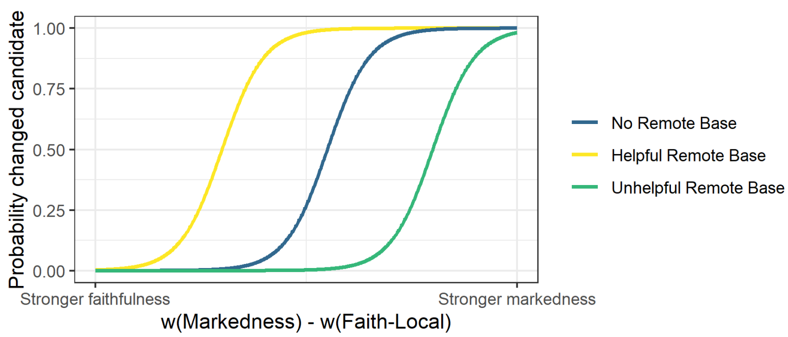

Figure 9: Schematic change in probability of the changed candidate based on the difference between the weight of MARKEDNESS and FAITH-LOCAL at a fixed weight of FAITH-REMOTE, for a Local Base with a helpful Remote Base, a Local Base with no Remote Base, and a Local Base with an unhelpful Remote Base.

5.3 Evaluation on Spanish data

The data to be modeled are the responses from Experiment 4, graphed in figure 6. The constraint set follows the principles of the Voting Bases theory outlined above: markedness constraints penalize unstressed diphthongs in candidates, and each Base exerts its influence via a separate faithfulness constraint. I use two distinct markedness constraints because front and back diphthongs show different rates of repair in the experiment. The markedness constraints are the following:

/Local Base/ Weight:	Faith-Local 1	Markedness 1
a. *faithful*		1
b. *changed*	1	

/Local/, /Helpful Remote/ Weight:	Faith-Local 1	Faith-Remote 1	Markedness 1
a. *faithful*		1	1
b. *changed*	1		

/Local/, /Harmful Remote/ Weight:	Faith-Local 1	Faith-Remote 1	Markedness 1
a. *faithful*		1	
b. *changed*	1		1

Table 17: Tableaux illustrating the schematic violation profile of a Derivative to a Local Base with no Remote Base (top), a helpful Remote Base (center), and a harmful Remote Base (bottom).

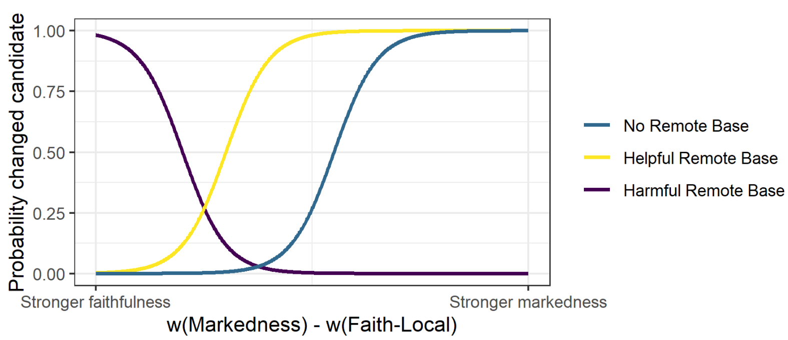

Figure 10: Schematic change in probability of the changed candidate based on the difference between the weight of Markedness and Faith-Local, for a Local Base with a helpful Remote Base, a Local Base with no Remote Base, and a Local Base with a harmful Remote Base.

- ***Unstressed -*ie*-**: Assign one violation for each unstressed front diphthong -*ie*- in a candidate.

This constraint is violated in forms *ambientál* "environmental" and *dietético* "dietitian". This constraint is motivated by the literature documenting avoidance of unstressed diphthongs in static lexical patterns and pathways of diachronic change, and also corroborated by the phonotactic markedness task described in Breiss (2021:ch. 3).

- ***Unstressed -*ue*-**: Assign one violation for each unstressed front diphthong -*ue*- in a candidate.

This constraint is the back counterpart of *Unstressed -*ie*-, and is violated in forms such as *crueldád* "cruelty" and *suegrástra* "stepmother-in-law".

As discussed in section 5.2, faithfulness constraints play a special role in the Voting Bases model: rather than a single constraint enforcing identity between a single UR and corresponding segments in a range of candidates, multiple constraints serve this same function, embodying the pull of each paradigm member on each candidate. Faithfulness constraints used in this analysis are the following:

- **Id-V-Local:** Assign one violation for each vowel in the Local Base that is not identical to its corresponding vowel in the candidate

Violations of this constraint are found in the UR-SR pairs /mwebloso/ → [moblóso] meaning "full of furniture", where the Local Base is *muéble* "furniture", and /djetoso/ → [detóso] "pertaining to a diet" where the Local Base is *diéta* "diet".

- **Id-V-Remote:** Assign one violation for each vowel in a Remote Base that is not identical to its corresponding vowel in the candidate.

Violations of this constraint are found in the UR-SR pairs /apwestoso/ → [apwestóso] "dashing, daring" with the Remote Base *apostár* "to bet", and /wevoso/ → [ovóso] "eggy" with the Remote Base *huevón* "a lazy person" (slang, literally "a big egg").

5.3.1 Model fitting and evaluation

I fit a MaxEnt model to counts of Derivative productions with monophthongal vs. diphthongal unstressed vowels using Excel's *Solver* utility (Fylstra et al., 1998), with a Gaussian prior on weights with a mean of 0 and standard deviation of 100. The goal of using a prior was to endow the model with a mild preference for lower constraint weights, which can be overcome with a sufficient amount of data; for precedents see Wilson (2006); White (2017). For each Local Base there are two candidates, one having an unstressed diphthong and the other an unstressed monophthong. The fitted weights are displayed in table 18.

Constraint	Weight
*Unstressed -*ie*-	0.45
*Unstressed -*ue*-	0.00
Id-V-Local	3.04
Id-V-Remote	1.09

Table 18: Constraint weights of model fit to data from Experiment 4, on Spanish.

I test the model against one that does not involve Remote Bases (forcing weights for faithfulness constraints referring to them to be zero), and one that does not distinguish between Local and Remote Bases (forcing weights to be equal for all faithfulness constraints). Both significantly underperform the Voting Bases model on a likelihood ratio test (ΔLoglikelihood $= 46$, $p < 0.001$, with two degrees of freedom; and ΔLoglikelihood $= 211.5$, $p < 0.001$, with two degrees of freedom, respectively).

Turning to the weights themselves, we can see that Local Bases have a stronger role in influencing the Derivative than Remote ones, and that the weight of markedness is nonzero but mild. Quantitatively, the R^2 was 0.25, and the fit between predicted and observed data is displayed in figure 11 below. Since the phonological model is not missing obvious grammar-wide constraints that could explain the low R^2 value, I speculate that this is due to the high amount of within-word variability or idiosyncrasy, a property of the alternation also observed by Eddington (1996); Albright et al. (2001); Carlson and Gerfen (2011).

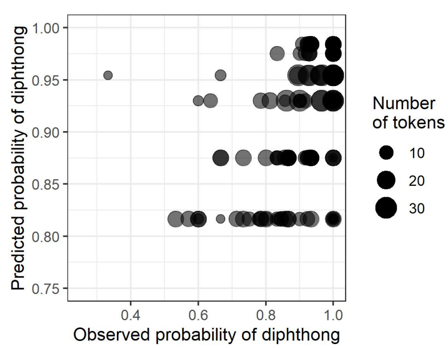

Figure 11: Predicted vs. observed fit to data from Experiment 4 on Spanish.

We can see where the Spanish data fit into the predicted typology of the model, in figure 12. Note that the finding that the helpful Remote Base exerts a stronger effect in probability space than the unhelpful Remote Base does falls out automatically as a consequence of the model; the difference between the blue line and the yellow line is greater than between the blue line and the green line, even though the difference in weight between Faith-Local and Faith-Remote is identical in the model.

5.4 Comparison to alternative models

A mentioned in section 1.1, the Spanish data allow us to distinguish the voting model of Lexical Conservatism from the two other formal models in the literature. Here I briefly demonstrate formally that neither Steriade's original method of analysis from her 1997

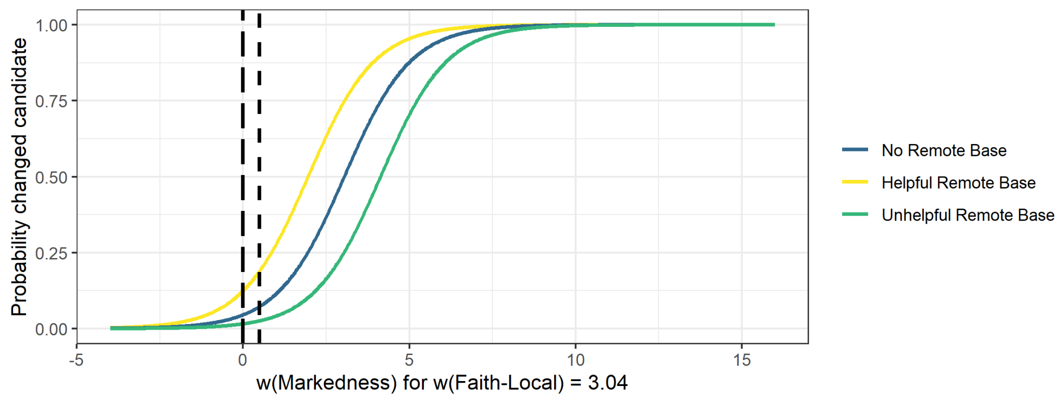

Figure 12: Schematic place of the Spanish data in the predicted typology of my model. The left, long-dashed vertical line is the typological place of Derivatives to Local Bases with a back diphthong, and the right short-dashed line represents the place of Derivatives to Local Bases with a front diphthong. These correspond to the weights of the markedness constraints in table 18.

work, nor the more recent proposal by Steriade and Stanton (2020), is able to capture the activity of the unhelpful Remote Base in Experiment 4.

5.4.1 Steriade (1997)'s analysis: quantifying over bases

Recall that Steriade (1997)'s model quantified over Bases: any Base in the Derivative's paradigm served to license its structures in the Derivative. The Lex-X family of constraints she used is re-supplied here, and makes reference reference to the primarily-stressed vowel of the Local Base.

- **Lex-Vowel**: Assign one violation if there is no Base in the input that matches the vowel in the candidate in quality.

This model is unable to distinguish rates of monophthongization in Derivatives that show an influence of an unhelpful Remote Base from those that have no Remote Base, as demonstrated in table 19 below.

Since the Lex-Vowel constraint does not actually encourage identity between Bases and candidates but instead simply licenses the possible existence of candidates that resemble any Base in the lexicon, the fact that the Remote Base resembles the Local Base makes no difference in how probability is allotted to forms. This stands in contrast to the voting model's violation profile for these cases in table 16, which does distinguish these cases.

/wéko/ Weight:	Lex-Vowel 1	*Unstressed -*ue*- 1
a. [wekóso]		1
b. [okóso]	1	

/mwéble/, /mobl-/ Weight:	Lex-Vowel 1	*Unstressed -*ue*- 1
a. [mweblóso]		
b. [moblóso]		

/xwérga/, /xwerg-/ Weight:	Lex-Vowel 1	*Unstressed -*ue*- 1
a. [xwergóso]		1
b. [xorgóso]	1	

Table 19: A schematic example demonstrating that an analysis in the style of Steriade (1997) fails to capture the influence of the unhelpful Remote Base.

5.4.2 Steriade and Stanton (2020)'s analysis: one Base per candidate

The model of Lexical Conservatism proposed in Steriade and Stanton (2020); Steriade (2018) assumes that each candidate Derivative stands in correspondence to a single Base — Local or Remote — that is in the input to the tableau. Their constraint C-Containment, is defined below; it intuitively penalizes candidates that do not have the Local Base cyclically contained within within the Derivative.

- **C-Containment**: Assign one violation if the candidate does not stand in correspondence to the Local Base.

The constraint BD-IdStress — referencing the Base-Dependent relationship defined by the indexation of the candidates to the Bases in the input (L(ocal) or R(emote)) — is violated if the stressed syllable in the candidate does not correspond to the stressed syllable in the Base it depends on; the markedness constraint is as in table 19.

In this model, satisfaction of markedness is achieved by candidate (a), which is faithful to the Remote Base at the expense of faithfulness to the Local Base. Candidate (b) satisfies faithfulness to the Base it depends on (the Local one), but violates markedness in doing so. Candidate (c) is in correspondence with the Local Base but violates faithfulness to it.

However, like the model put forward by Steriade (1997), the model cannot account for cases where an unhelpful Remote Base exerts a role in shaping the Derivative: there is no weight of faithfulness and markedness that allows for the Local Base with an

/wéko/$_L$ Weight:	BD Id-V 1	*Unstressed -*ue*- 1	C-Containment 1	H	*p*
a. [wekóso]$_L$		1		1	.5
b. [okóso]$_L$	1			1	.5

/mwéble/$_L$, /mobl-/$_R$ Weight:	BD Id-V 1	*Unstressed -*ue*- 1	C-Containment 1	H	*p*	sum *p*
a. [mweblóso]$_L$		1		1	.31	.36
b. [mweblóso]$_R$	1	1	1	3	.04	
c. [moblóso]$_L$	1			1	.31	.63
d. [moblóso]$_R$			1	1	.31	

/xwérga/$_L$, /xwerg-/$_R$ Weight:	BD Id-V 1	*Unstressed -*ue*- 1	C-Containment 1	H	*p*	sum *p*
a. [xwergóso]$_L$		1		1	.36	.5
b. [xwergóso]$_R$		1	1	2	.13	
c. [xorgóso]$_L$	1			1	.36	.5
d. [xorgóso]$_R$	1		1	2	.13	

Table 20: A schematic example demonstrating that an analysis in the style of Steriade and Stanton (2020) fails to capture the influence of the unhelpful Remote Base. The probabilities are derived using a MaxEnt grammar from the constraints and weights shown, and probabilities over surface-identical candidates that differ only in their correspondence relations are summed to yield a predicted rate of monophthongization.

unhelpful Remote Base, here *juérga* "spree" ∼ *juerguísta* "reveller", to have a rate of monophthongization which differs from that of the Local Base without any Remote Base, like *huéco* "gap". This follows from Steriade's original assumption about the strictly markedness-improving role of the Remote Base, and is evident in the example in table 20 below.

To quantify the degree of misfit induced by the inability to capture the effect of unhelpful Remote Bases in Spanish, I fit MaxEnt grammars based on both Steriade (1997)'s and Steriade and Stanton (2020)'s models to the full data from Experiment 4 that was treated in section 5.3, and examined the evidence ratio (Burnham and Anderson, 2002; Anderson and Burnham, 2002; Burnham and Anderson, 2004) for each of the two alternative analyses to the one I proposed in section 5.2; this is displayed in table 21 below.

Model	*LL*	*ΔAICc*	*Evidence ratio favoring the first model*
Voting Bases model (this paper)	-562.24		
Steriade (1997)	-567.08	7.68	46 : 1
Steriade & Stanton (2020)	-567.08	9.68	126 : 1

Table 21: Statistical measures of model fit and comparison for the voting theory of Remote Bases proposed in section 5.2, the quantification-based theory of Steriade (1997), and the one-Base-at-a-time theory of Steriade and Stanton (2020).

We can see that the inability to distinguish between Derivatives with and without a non-optimizing Remote Base leads for the weight of evidence to favor the Voting Bases model.

5.5 Evaluation on English data

We now turn to the English data. To aid in a unified phonological analysis, I aggregate the data across experiments; however, I exclude Experiment 1 because it used written stimuli, and in Steriade's original stimuli the *-ate* suffix was pre-removed in the presentation of the Derivative, making direct between-experiment comparison impossible.

5.5.1 Constraints used in analyzing data from Experiments 2 and 3

The analysis uses traditional constraints drawn from the literature on English stress. The markedness constraints employed in the analysis are the following:

- ***ExtendedLapse**: Assign one violation for each string of three unstressed syllables in the Derivative.

This constraint is violated in forms such as *bállotable* and *láborable*, and obeyed in shifted candidates like *ballótable* and *labórable*. This constraint drives use of helpful Remote Bases, and its strength suppresses analogical influence of faithfulness to harmful Remote Bases in English. For precedents, see Gordon (2002); Stanton (2016); see also Steriade and Stanton (2020) for use in the analysis of cases of Lexical Conservatism in English stress.

- ***WeakFinalTernary:** Assign one violation for each sequence of two unstressed, word-final syllables directly preceded by a syllable bearing secondary stress.

This constraint is violated in forms such as *bánkrùptable* and *cúckòldable*. Speaking in terms of SPE stress numbering (Chomsky and Halle, 1968), it forbids the sequence *200#*, where *2* indicates secondary stress. This principle of English metrical structure can be found described in Pater (2000), where it was cast as a ban on non-right-aligned main stress (assuming final-syllable extrametricality), and much work before and since has found this principle emergent from comprehensive analyses of English stress.

- **Pre-stress *-ic*:** Assign one violation if syllable directly preceding the suffix *-ic* does not bear primary stress.

This constraint is violated in forms such as *lúmberic* and *résinic*. Although analyses of the stress preferences of English affixes broadly construed generally make use of a distinction between Level 1 and Level 2 affixes to regulate this behavior, for the simpler present case I use a "brute force" constraint like this one to model the degree to which *-ic* prefers to be pre-stressed (cf., for instance, Chomsky and Halle, 1968:who posit an affix-specific rule).

- ***WeakHeavy:** Assign one violation when a post-tonic heavy syllable in the Derivative does not also bear stress (primary or secondary).

This constraint is violated in forms such as *séquencable*. This constraint enforces one aspect of the stress-to-weight principle, a typological propensity for heavy syllables to attract stress (see Ryan, 2016:for an overview of the literature), and of the Latin Stress Rule of English (Chomsky and Halle, 1968; Liberman and Prince, 1977:et seq.).

The faithfulness constraints included in the analysis are also standard in format, except that they are differentiated based on the status of the Base that they refer to. Here, this includes Local, Remote, and within Remote, primed and non-primed versions. Following the general approach of (Boersma and Hayes, 2001; Linzen et al., 2013; Coetzee and Kawahara, 2013; Coetzee, 2016), we can model the influence of resting activation on the phonological grammar by treating it as a scaling factor that can change the impact of violating faithfulness to a Base. Here, I treat that scaling factor as binary — the Remote Base was either primed or not — and for ease of exposition modeled the two levels of the scaling factor as two different constraints. However, it would also be possible to have a continuous measure of resting activation, and multiply a single grammar-wide weight of Faith-Remote by that factor, similar to the model of Coetzee and Kawahara (2013). I do not model the role of semantic similarity between Local and Remote Bases in this paper, since it relies on a more nuanced model of resting activation in the lexicon, as discussed in section 6.2.1 below. For the purposes of the formal analysis, the faithfulness constraints that reference priming participate in the phonological grammar in exactly the same way as constraints that do not reference resting activation.

- **Id-[stress]-Local:** Assign one violation if the primary stressed syllable in the Local Base does not correspond to the primary stressed syllable in the candidate.

Violating UR-SR pairs include /sénator/ → [senátorist] and /sénator/ → [sènatórist].

- **Id-[stress]-Remote-primed:** Assign one violation if the primary stressed syllable in a primed Remote Base does not correspond to the primary stressed syllable in the candidate.

Violating UR-SR pairs include /túmult/→ [túmultist] if *tumúltuous* is primed, and /próverb/ → [próverbist] if *provérbial* is primed.

- **Id-[stress]-Remote-unprimed:** Assign one violation if the primary stressed syllable in an unprimed Remote Base does not correspond to the primary stressed syllable in the candidate.

Violating UR-SR pairs include /túmult/ → [túmultist] if *tumúltuous* is not primed, and /próverb/ → [próverbist] if *provérbial* is not primed.

5.5.2 Model fitting and evaluation

Model fitting followed the same procedure as with the Spanish model. An example of the Local Base *lábor* with Remote Base *labórious* is below in table 22.

The best-fitting weights for the constraints are listed in table 23.

Constraint	Weight
*ExtendedLapse	1.17
*WeakFinalTernary	2.65
Pre-stress *-ic*	3.13
*WeakHeavy	0.23
Id-[stress]-Local	2.48
Id-[stress]-Remote-primed	1.23
Id-[stress]-Remote-unprimed	0.75

Table 23: Constraint weights for model fit to English data from Experiments 2 and 3.

/ˈle͡ɪbɚ/$_L$, /ləˈbɔɹ-/$_R$	*ExtendedLapse	*WeakFinalTernary	Pre-stress -*ic*	*WeakHeavy	Id-[stress]-Local	Id-[stress]-Remote-primed	Id-[stress]-Remote-unprimed
a. [ˈle͡ɪbɚəbl̩]	1					1	
b. [ləˈbɔrəbl̩]					1		

Table 22: A sample candidate and violation set from a tableau where the Remote Base *laborious* is known and primed.

5.5.3 Discussion

Qualitatively examining the weights in table 23, the markedness constraints are in line with what one might expect from an experimental test of the principles of English metrical phonology: there is a strong effect of matching weight to stress, as well as the avoidance of long lapses. Turning to faithfulness constraints, we find that the status of the Base in the lexicon is reflected directly in the weights. Most prominent is the Local Base, which exerts a powerful influence on the Derivative; Remote Bases have lower weights of faithfulness, and reflect a distinction of priming such that primed Remote Bases are more influential on the Derivative than unprimed ones. In quantitative terms, the model achieves a reasonable fit to the data, with an R^2 of 0.67; the model predictions are plotted below in figure 13.

I do not carry out constraint-by-constraint significance testing because many of the narrower effects have been more rigorously assessed in section 3 using inferential statistics. However, I do test the fit of the full model against two alternative theoretical claims which do not correspond to previous statistical tests done on the data. First, I compare the full model to one where the lexicon plays no role in regulating the accessibility of Bases. This model does not allow faithfulness constraints to have different weights based on whether they refer to Local and Remote Bases; it significantly underperforms the full model (ΔLoglikelihood $= 6.4$, $p = 0.002$, with two degrees of freedom). The second model is one that denies a role for Remote bases to play in generating the data, so the weights

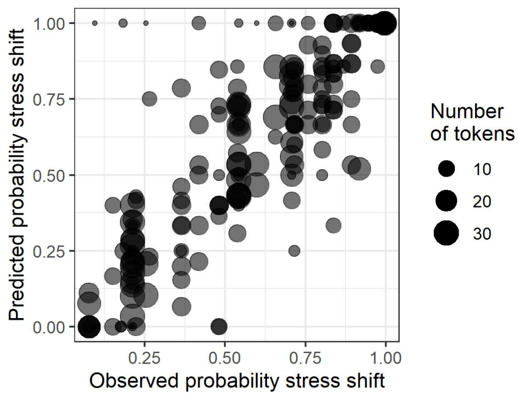

Figure 13: Predicted vs. observed fit to English data from Experiments 2 and 3.

of faithfulness to Remote Bases are forced to zero; this model also significantly underperforms (ΔLoglikelihood $= 63.1$, $p < 0.001$, with two degrees of freedom). I take these findings as points in favor of a general model of Lexical Conservatism that relies on Remote Bases, and further allows the lexicon to modulate access to them (cf. section 6.2.1 below).

Finally, we can approximately locate the effects observed in English on the typology graph, shown in figure 14. Note that this is not a quantitative estimate (this would involve converting the more articulated markedness constraint structure summarized in table 23 to the schematic three-parameter scenario), but simply a visual aid to the intuition about where the English data lie in the typology that the Voting Bases model can encompass. In English, we find that there is some rightward stress shift in Local Bases without Remote Bases, and more if there is a helpful Remote Base. However, we see hardly any leftward stress shift (two tokens total in all of Experiment 3), which exhibit an influence of the harmful Remote Base. This dramatic difference between helpful and harmful Remote Bases falls out from the differing harmony penalty felt by the changed candidate Derivative: in the helpful Remote Base case, Markedness and Faith-Remote are relatively well-balanced against Faith-Local, but with a harmful Remote Base, Markedness and Faith-Local together far outweigh Faith-Remote.

6 Broader considerations

6.1 Why is Lexical Conservatism not ubiquitous?

The proposed Voting Bases theory of Lexical Conservatism relies on a language-general proposal about the way Base competition in the lexicon interacts with the phonological

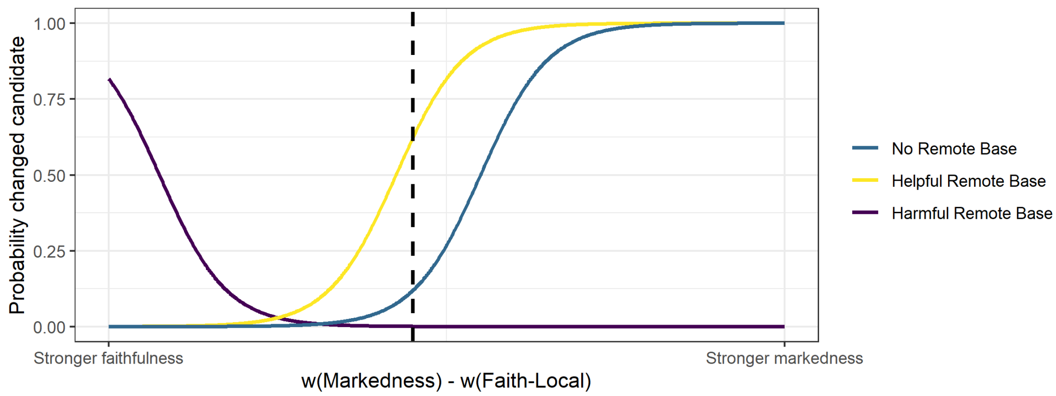

Figure 14: Schematic place of the English data in the predicted typology of the Voting Bases model.

grammar. Therefore it is important to ask why Lexical Conservatism and other Base effects are not attested in all languages. Aside from the reasonable possibility that such effects might be more widespread than generally assumed and phonological description has not caught up with this reality yet, I note that the voting mechanism only allow Remote Bases to yield noticeable effects on Derivative formation when language- and context-specific morphological and phonological conditions hold: there must be a sufficiently strong motivating markedness constraint to encourage deviation from the Local Base, there must exist a Remote Base in the paradigm of the Derivative, and there must be a sufficiently weak faithfulness constraint indexed to the Local Base. The regions of weight space for Markedness and Faith-Local where the voting model predicts vanishingly small effects of the Remote Base are marked by the vertical dashed lines on the graph in figure 15.

Finally, if the weight of Faith-Remote is near-zero, even given the appropriate balance between the weights of Markedness and Faith-Local the Remote Base will hardly have any detectable effect on the distribution of Derivatives. This dependence on both amenable paradigm structures and constraint weights is demonstrated in this paper by the non-effect of harmful Remote Bases on Derivative formation seen in Experiment 3.

6.2 Connecting processing to the phonological grammar

Turning now to the effects of processing factors on Derivative formation, I suggest that the influence of the lexical characteristics observed in Experiments 1-3 form part of a broader emerging body of literature documenting cases where non-phonological characteristics of a lexical item influence phonological outcomes in the grammar. In his paper, we saw in English that priming the Remote Base increased its influence on the

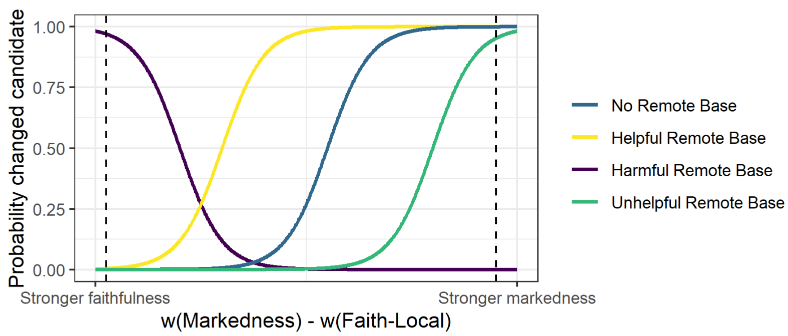

Figure 15: Schematic example where Remote Base activity is suppressed at extreme values of markedness or faithfulness; strong local faithfulness to the left, and strong markedness to the right.

Derivative. We also observed that semantic similarity of the Local and Remote Bases played a role in influencing whether the Remote Base had an effect on the Derivative in question, albeit one in a direction opposite to the one Steriade anticipated (Steriade, 1997:p. 19). Further afield, Eddington (2006); Kim (2021); Breiss et al. (2021) have observed that the token frequency of non-local surface forms of paradigm members conditions variability in Paradigm Uniformity, and Baroni (2001); Zuraw (2007); Zuraw and Peperkamp (2015); Zuraw et al. (2020); Wurm (1997); Caramazza et al. (1988) have demonstrated that token frequency can interact with lexical retrieval to bias phonological behavior. Work by Wagner (2012); Kilbourn-Ceron et al. (2016); Lamontagne and Torreira (2017); Kilbourn-Ceron and Sonderegger (2018); Tamminga (2018); Kilbourn-Ceron and Goldrick (2021:among others) under the banner of the Production Planning Hypothesis has found that lexical characteristics of morphemes can influence phonological behavior by facilitating or inhibiting planning in speech production. Thus in general there is robust evidence that the relative salience or prominence of a lexical item can influence how that item is treated by the phonological grammar. We take this to be a fact about the way domain-general cognitive resource allocation interacts with language-specific phenomena and therefore do not motivate it based on language-specific knowledge, learned or innate.

6.2.1 Resting activation

The question remains, however, of how our phonological theories should be set up to interact with these cognition-general factors. I suggest that the intuitive notion of salience motivated above maps well onto the psycholinguistic construct of *resting activation* (see, for example, Morton, 1970: and a multitude of others since). Activation is a concept that is motivated on the basis of independent psycholinguistic studies of the lexicon (see

Levelt, 1999; Storkel and Morrisette, 2002; Weber and Scharenborg, 2012:for overviews of word production, acquisition, and comprehension respectively). Detailed accounts of how it is computed differ between theories, but a common property is that a lexical entry is associated with a largely unconscious scalar property that is influenced by static and dynamic lexical factors, computed by the language processing system quasi-deterministically on the basis of the speakers' local and global language experience. The Voting Bases theory proposed here is not committed to a specific theory of the structure and dynamics of the lexicon, other than it makes reference to this property, however it might be computed or derived. For present purposes, I suggest that priming the Remote Base can raise the resting activation of the stem allomorph that it contains, and that high-frequency Remote Bases have long-run higher resting activations than low-frequency ones.

I now turn to the puzzling finding that increased semantic similarity between the Local and Remote Bases inhibits, rather than facilitates, the creation of Derivatives which resemble the Remote Base. For a tentative explanation, I rely on the findings of Wheeldon and Monsell (1994); Wheeldon (2003) that semantic similarity does not influence resting activation itself, but rather modulates connections between lexical representations (here, of allomorphs of a Base), and thus how much "spill-over" in resting activation there is from one high-activation form to another, lower-activation one. Harley (1993) demonstrates that in the on-line production of a given target word, highly activated related words undergo "suppression" of their resting activation to minimize competition and allow the correct word to be uttered (Rahman and Aristei, 2010). This provides a possible mechanism for the findings for semantic similarity discussed in section 3.4. For a given Local and Remote Base in the same paradigm, a Remote Base that is very semantically similar to the Local one would be a source of interference for a participant seeking to utter a Derivative built to the Local Base, and so their language processing system might require more suppression to access the Local Base in producing the Derivative. A less semantically-similar Remote Base would not be as strong a source of interference and so would not be suppressed as much, allowing it to have an influence on the Derivative-formation process. In this account, semantic similarity acts on resting activations rather than directly influencing the phonological grammar; thus I do not include it in the phonological model of the English experimental data given in section 5.2. Future work that jointly models the dynamics of the lexicon in theory own right alongside outcomes of the phonological grammar is required to further probe the role of semantic similarity between the Local and Remote Bases.

Connecting the phonological grammar to the dynamics of the lexicon comes with the added benefit that lexically-scaled faithfulness constraints can only take on values that are compatible with factors known to influence the resting activation of lexical items. This makes logically-possible constraint weightings where faithfulness to an unprimed Remote Base is much stronger than faithfulness to a primed Remote Base *a priori* unlikely on grounds unrelated to the phonological grammar.

6.3 Theoretical implications

This paper is hardly the first to suggest that non-phonological characteristics of stored lexical items influence outcomes of the phonological grammar. The proposed voting theory joins (at minimum) Output-Output faithfulness (Benua, 2000), Gradient Symbolic Representation Smolensky and Goldrick (2016), Representational Strength Theory (Moore-Cantwell, *in progress*), the UseListed-based theory of Zuraw (2000), and UR Constraints (Pater et al., 2012). The novel contributions that the Voting Bases theory makes are linking scaled faithfulness to resting activation, and allowing multiple faithfulness constraints to exert analogical pressure on the Derivative. However, future investigation is needed to find out whether both of these features are critical in modeling Lexical Conservatism, or whether some of the other theories just mentioned, enriched with one assumption or the other, might achieve broader empirical coverage and predictive power.

The data presented here also motivate some rethinking of other assumptions about the roots of alleged universals in word-formation and other domains of the grammar, as long acknowledged by Steriade (Steriade, 1997, 2018; Steriade and Stanton, 2020, 2021). Specifically, cases like Lexical Conservatism where surface forms make reference to non-local members of the morphological paradigm pose difficulties for theories that make cyclic inheritance a core, automatic, or "universal" aspect of their process of word-formation. These include most contemporary syntax-based theories of word-formation like Distributed Morphology (see the overview in Bobaljik, 2017) and Nanosyntax (see the overview in Baunaz et al., 2018), as well as theories like Stratal Optimality Theory which are also based on a notion of (phonological) cyclicity (Bermúdez-Otero, 2017). The facts about Lexical Conservatism, rather, support model of word-formation where the structure of the morphological paradigm, combined with language-specific strength of markedness and lexical characteristics, are the driving factors in novel word formation, such as the one advanced here, in which "cyclic" inheritance of features of the Local Base by Derivatives is the norm, but also allowing the emergence of non-cyclic behavior — typologically rare but not at all unattested — as an automatic outcome of specific relations between the structure of the lexicon and markedness.

Finally, the data presented here are not compatible with an account of Lexical Conservatism that relies on abstract or "cobbled" URs, where a single UR which is not present in any surface cell in a paradigm underlies multiple members of the paradigm derived from it by regular phonology (Bermúdez-Otero, 2017). One way to see this is by looking at the Spanish data: although Derivatives to Local Bases with no Remote Bases (such as *Viena*), and those to Local Bases with unhelpful Remote Bases (such as *ciéncia* $\sim$ *cienciólogo*) have the same distribution of vowel qualities in their allomorphs, nevertheless Derivatives that have the unhelpful Remote Base exhibit less diphthongization than those with no Remote Base. Such effects cannot be attributed to a more abstract UR that underlies both Local and Remote Bases, as Bermúdez-Otero (2017) suggests for English. Such effects also cannot be captured by simple item-indexation theories where each stem is indexed as alternating or non-alternating, since although the evidence for non-alternation is equal between Local Bases with no Remote Bases and those with unhelpful Remote Bases, the two are distinct in their rate of alternation in novel forms.

7 Summary and outlook

This paper has presented evidence for synchronic robustness of Lexical Conservatism, a dependency between paradigm shape and phonological process application in novel words laid out in Steriade (1997). It also demonstrated that the phenomenon is probabilistic, with Derivative formation influenced by phonological and lexical factors. It used a comparison of the English case from Steriade (1997) with novel data from Mexican Spanish to demonstrate that the dependencies characteristic of Lexical Conservatism arise from competitive voting between listed Base allomorphs, cross-cut by phonological markedness.

I modeled these data by proposing a framework of "voting" Bases that predicts the influence of Remote Bases for any language, dependent on its paradigm shape and the relative weight of markedness and faithfulness constraints. The model uses traditional markedness constraints, and proposed that faithfulness scaled based on resting activation allow the lexical characteristics of different Bases in the lexicon to exert an influence on the phonological computation of the Derivative. By virtue of its roots in psycholinguistic evidence for the contents and dynamics of the lexicon, the Voting Bases model is also consistent with consensus models of speech production and parsing, providing a link between phonological theory and psycholinguistic data.

Although this paper set out to resolve existing ambiguities about the status of Lexical Conservatism in the synchronic grammar and its relationship to the contents of the lexicon, in the process new questions have arisen that merit future work. I discuss two here: multiple Remote Bases for a single Local Base, and the role of learning in Lexical Conservatism.

This paper focused on Local Bases with only one type of Remote Base stem allomorph (helpful, harmful, or unhelpful). However, there is no reason to think that the voting mechanism holds only for these types of paradigms. In cases where there are multiple Remote Base stem allomorphs, some helpful and some not, the voting model predicts that each Base's stem allomorph will exert a pull on the Derivative via its own faithfulness constraints, and exhibiting cumulativity with other constraints in the grammar as outlined in section 5.2. Future experimental work that specifically looks at such forms is needed to test this prediction.

Turning now to learning, there is much that remains to be understood. Narrowly, we have not yet explored whether there is evidence in the learners' input to set the weight of FAITH-REMOTE on a language-specific basis for each case of Lexical Conservatism. More broadly, I have simply assumed that the architecture assumed by the voting Bases theory is "always on", a manifestation of independently-motivated lexical or cognitive structures. Future modeling studies may reveal that this is better viewed as a parameter set by the learner in response to language-specific data.

Appendix

Materials for Experiment 1

Local Base	Remote Base	Affixes	Source
carrot	—	*-ify, -ity*	Novel
cotton	—	*-ify, -ity*	Novel
denim	—	*-ify, -ity*	Novel
diamond	—	*-ify, -ity*	Novel
fennel	—	*-ify, -ity*	Novel
flannel	—	*-ify, -ity*	Novel
fungus	—	*-ify, -ity*	Novel
garlic	—	*-ify, -ity*	Novel
granite	—	*-ify, -ity*	Novel
leather	—	*-ify, -ity*	Novel
lettuce	—	*-ify, -ity*	Novel
lumber	—	*-ify, -ity*	Novel
marble	—	*-ify, -ity*	Novel
mushroom	—	*-ify, -ity*	Novel
muslin	—	*-ify, -ity*	Novel
nylon	—	*-ify, -ity*	Novel
onion	—	*-ify, -ity*	Novel
pepper	—	*-ify, -ity*	Novel
plaster	—	*-ify, -ity*	Novel
protein	—	*-ify, -ity*	Novel
pumice	—	*-ify, -ity*	Novel
resin	—	*-ify, -ity*	Novel
rubber	—	*-ify, -ity*	Novel
rubric	—	*-ify, -ity*	Novel
silver	—	*-ify, -ity*	Novel
spandex	—	*-ify, -ity*	Novel
spinach	—	*-ify, -ity*	Novel
tartan	—	*-ify, -ity*	Novel
turnip	—	*-ify, -ity*	Novel
velvet	—	*-ify, -ity*	Novel
acid	acidic	*-ify, -ity*	Novel
agent	agentive	*-ify, -ity*	Novel
angel	angelic	*-ify, -ity*	Novel

Table 24 continued from previous page

Local Base	Remote Base	Affixes	Source
artist	artistic	*-ify, -ity*	Novel
atom	atomic	*-ify, -ity*	Novel
autumn	autumnal	*-ify, -ity*	Novel
carbon	carbonic	*-ify, -ity*	Novel
cherub	cherubic	*-ify, -ity*	Novel
courage	courageous	*-ify, -ity*	Novel
demon	demonic	*-ify, -ity*	Novel
ether	ethereal	*-ify, -ity*	Novel
habit	habitual	*-ify, -ity*	Novel
icon	iconic	*-ify, -ity*	Novel
insect	insecticide	*-ify, -ity*	Novel
justice	justiciable	*-ify, -ity*	Novel
logic	logician	*-ify, -ity*	Novel
magic	magician	*-ify, -ity*	Novel
metal	metallic	*-ify, -ity*	Novel
moment	momentous	*-ify, -ity*	Novel
moron	moronic	*-ify, -ity*	Novel
music	musician	*-ify, -ity*	Novel
novice	novitiate	*-ify, -ity*	Novel
office	official	*-ify, -ity*	Novel
organ	organic	*-ify, -ity*	Novel
parent	parental	*-ify, -ity*	Novel
person	personification	*-ify, -ity*	Novel
pirate	piratical	*-ify, -ity*	Novel
program	programmable	*-ify, -ity*	Novel
sentence	sentential	*-ify, -ity*	Novel
super	superfluous	*-ify, -ity*	Novel
abdicate	—	*-able, -ism*	Steriade
accelerate	—	*-able, -ism*	Steriade
agitate	—	*-able, -ism*	Steriade
allocate	—	*-able, -ism*	Steriade
ameliorate	—	*-able, -ism*	Steriade
annihilate	—	*-able, -ism*	Steriade
communicate	—	*-able, -ism*	Steriade
dedicate	—	*-able, -ism*	Steriade
educate	—	*-able, -ism*	Steriade
eradicate	—	*-able, -ism*	Steriade

Table 24 continued from previous page

Local Base	Remote Base	Affixes	Source
examine	—	*-able, -ism*	Steriade
exterminate	—	*-able, -ism*	Steriade
generate	—	*-able, -ism*	Steriade
illuminate	—	*-able, -ism*	Steriade
investigate	—	*-able, -ism*	Steriade
irrigate	—	*-able, -ism*	Steriade
nominate	—	*-able, -ism*	Steriade
penetrate	—	*-able, -ism*	Steriade
pollinate	—	*-able, -ism*	Steriade
precipitate	—	*-able, -ism*	Steriade
predicate	—	*-able, -ism*	Steriade
procrastinate	—	*-able, -ism*	Steriade
prognosticate	—	*-able, -ism*	Steriade
propagate	—	*-able, -ism*	Steriade
relegate	—	*-able, -ism*	Steriade
remunerate	—	*-able, -ism*	Steriade
resuscitate	—	*-able, -ism*	Steriade
segregate	—	*-able, -ism*	Steriade
tolerate	—	*-able, -ism*	Steriade
venerate	—	*-able, -ism*	Steriade
attribute	attribution	*-able, -ism*	Steriade
compensate	compensatory	*-able, -ism*	Steriade
concentrate	concentric	*-able, -ism*	Steriade
confiscate	confiscatory	*-able, -ism*	Steriade
contemplate	contemplative	*-able, -ism*	Steriade
contribute	contribution	*-able, -ism*	Steriade
(take) custody	custodian	*-able, -ism*	Steriade
demonstrate	demonstrative	*-able, -ism*	Steriade
domesticate	domesticity	*-able, -ism*	Steriade
equilibrate	equilibrium	*-able, -ism*	Steriade
infiltrate	filter	*-able, -ism*	Steriade
illustrate	illustrative	*-able, -ism*	Steriade
impregnate	pregnant	*-able, -ism*	Steriade
incorporate	incorporeal	*-able, -ism*	Steriade
inculpate	inculpable	*-able, -ism*	Steriade
indicate	indicative	*-able, -ism*	Steriade
influence	influential	*-able, -ism*	Steriade

Table 24 continued from previous page

Local Base	Remote Base	Affixes	Source
integrate	integrative	*-able, -ism*	Steriade
interrogate	interrogative	*-able, -ism*	Steriade
intuit	intuition	*-able, -ism*	Steriade
obfuscate	obfuscatory	*-able, -ism*	Steriade
oblige	obligate	*-able, -ism*	Steriade
expurgate	purgatory	*-able, -ism*	Steriade
reciprocate	reciprocity	*-able, -ism*	Steriade
remediate	remedial	*-able, -ism*	Steriade
remonstrate	remonstrance	*-able, -ism*	Steriade
sequester	sequestrate	*-able, -ism*	Steriade
designate	signatory	*-able, -ism*	Steriade
assimilate	similitude	*-able, -ism*	Steriade

Table 24: Local Bases for the Lexical Conservatism task in Experiment 1, listed with Remote Base (if any), the affixes they were combined with, and the source (novel in this experiment, or from Steriade (1997)).

Materials for Experiment 2

Local Base	Remote Base
ballot	—
bankrupt	—
blizzard	—
buzzard	—
carrot	—
cuckold	—
denim	—
diamond	—
fungus	—
granite	—
lumber	—
nylon	—
orange	—
plaster	—
resin	—
scaffold	—
spandex	—
spinach	—
thermos	—
velvet	—
autumn	autumnal
commerce	commercial
courage	courageous
essence	essential
ether	ethereal
finance	financial
habit	habitual
insect	insecticide
labor	laborious
major	majority
mammal	mammalian
modern	modernity
moment	momentous
office	official
parent	parental
person	personify

Table 25 continued from previous page

Local Base	Remote Base
proverb	proverbial
sequence	sequential
substance	substantial
tumult	tumultuous

Table 25: Local Bases for the Lexical Conservatism task in Experiment 2, with the Remote Base if any.

Materials for Experiment 3

Local Base	Remote Base	Affix	Local Base stress	Remote Base stress	Remote Base Type
abuse	—	able	final	—	—
achieve	—	able	final	—	—
alert	—	able	final	—	—
appraise	—	able	final	—	—
approve	—	able	final	—	—
behave	—	able	final	—	—
bequeath	—	able	final	—	—
demote	—	able	final	—	—
diffuse	—	able	final	—	—
enclose	—	able	final	—	—
finesse	—	able	final	—	—
infuse	—	able	final	—	—
peruse	—	able	final	—	—
secede	—	able	final	—	—
traverse	—	able	final	—	—
accuse	accusation	able	final	initial	harmful
compose	composition	able	final	initial	harmful
confide	confidant	able	final	initial	harmful
conserve	conservation	able	final	initial	harmful
dispose	disposition	able	final	initial	harmful
divide	dividend	able	final	initial	harmful
expose	exposition	able	final	initial	harmful
impose	imposition	able	final	initial	harmful
oppose	opposition	able	final	initial	harmful
precede	precedence	able	final	initial	harmful
preserve	preservation	able	final	initial	harmful
propose	proposition	able	final	initial	harmful
provide	providence	able	final	initial	harmful
reserve	reservation	able	final	initial	harmful
reside	residence	able	final	initial	harmful
bankrupt	—	able	initial	—	—
cuckold	—	able	initial	—	—
plaster	—	able	initial	—	—
scaffold	—	able	initial	—	—
ballot	—	ist	initial	—	—

Table 26 continued from previous page

Local Base	Remote Base	Affix	Local Base stress	Remote Base stress	Remote Base Type
blizzard	—	ist	initial	—	—
buzzard	—	ist	initial	—	—
denim	—	ist	initial	—	—
granite	—	ist	initial	—	—
velvet	—	ist	initial	—	—
courage	courageous	ist	initial	final	helpful
habit	habitual	ist	initial	final	helpful
moment	momentous	ist	initial	final	helpful
proverb	proverbial	ist	initial	final	helpful
tumult	tumultuous	ist	initial	final	helpful
finance	financial	able	initial	final	helpful
labor	laborious	able	initial	final	helpful
major	majority	able	initial	final	helpful
parent	parental	able	initial	final	helpful
sequence	sequential	able	initial	final	helpful

Table 26: Stimuli for the Lexical Conservatism task in Experiment 3, listed with the affix they took in the experiment, their Remote Base (if any), as well as the stress placement in the Local and Remote bases, and the type of Remote Base.

Materials for Experiment 4

Local Base (ie)	Remote Base (ie)	Local Base (ue)	Remote Base (ue)	Class
audiencia	—	ungüento	—	*(none)*
sapiencia	—	güecho	—	*(none)*
pierna	—	güemul	—	*(none)*
biela	—	atuendo	—	*(none)*
adviento	—	duende	—	*(none)*
priesa	—	elocuencia	—	*(none)*
siesta	—	hueco	—	*(none)*
nieto	—	pirueta	—	*(none)*
mies	—	buega	—	*(none)*
siniestro	—	sabueso	—	*(none)*
aliciente	—	silueta	—	*(none)*
Viena	—	suela	—	*(none)*
noviembre	—	cruento	—	*(none)*
vieira	—	güeña	—	*(none)*
viernes	—	sueco	—	*(none)*
ariete	arietar	dueño	adueñarse	*unhelpful*
lienzo	liencillo	consecuente	consecuentemente	*unhelpful*
ambiente	ambiental	cruel	crueldad	*unhelpful*
cielo	cielito	deshueso	deshuesadao	*unhelpful*
ciencia	cienciólogo	juete	juetazo	*unhelpful*
experiencia	experiencial	encuesta	encuestada	*unhelpful*
cliente	clientela	juerga	juergista	*unhelpful*
conciencia	concienciar	huevo	huevón	*unhelpful*
oriente	oriental	secuencia	secuencial	*unhelpful*
paciencia	impacientar	secuestro	secuestrador	*unhelpful*
dieta	dietético	suegra	suegrastro	*unhelpful*
higiene	higienista	güero	güerito	*unhelpful*
riel	rielero	delincuente	delincuentemente	*unhelpful*
ciemo	aciemar	huebra	huebrero	*unhelpful*
rienda	riendazo	buey	bueyero	*unhelpful*
acierto	acertar	compuesta	compostura	*helpful*
fiebre	febril	mueble	moblar	*helpful*
asiento	asentar	escuela	escolar	*helpful*
niebla	neblar	grueso	grosor	*helpful*
cierna	cerner	muerte	mortero	*helpful*
obediencia	obedecer	rueda	rodar	*helpful*

Table 27 continued from previous page

Local Base (ie)	Remote Base (ie)	Local Base (ue)	Remote Base (ue)	Remote Base type
sosiego	sosegar	apuesta	apostar	*helpful*
tienda	tender	ruego	rogar	*helpful*
viejo	vejez	cuento	contar	*helpful*
incienso	incensar	almuerzo	almorzar	*helpful*
cierre	cerrar	encuentro	encontrar	*helpful*
ciego	cegar	acuerdo	accordar	*helpful*
aprieto	apretar	vuelo	volar	*helpful*
entierro	enterrar	trueno	tronar	*helpful*
friega	fregar	consuelo	consolar	*helpful*

Table 27: Local Bases for the Lexical Conservatism task in Experiment 4, listed with Remote Base (if any).

Model of combined English experiments

Parameter	*Median*	*95% CI*	*P* $(\lvert\hat{\beta}\rvert > 0)$
Intercept:			
Target syllable secondary stress = *no*			
Target syllable heavy = *no*			
Affix = *-able*			
Experiment = *1*			
Remote Base = *unprimed*			
Remote Base log-freq. = *average values*			
Semantic similarity = *average values*	-0.47	[-1.17, 0.22]	
Target syllable heavy = *yes*	0.79	[0.10, 1.47]	0.99
Target syllable secondary stress = *yes*	1.05	[0.21, 1.89]	0.99
Affix = *-ic*	1.80	[1.10, 2.54]	≈ 1
Affix = *-ify*	2.08	[1.24, 2.94]	≈ 1
Affix = *-ism*	-0.57	[-1.08, -0.06]	0.99
Affix = *-ist*	-0.89	[-3.19, 1.29]	0.81
Affix = *-ity*	4.01	[3.03, 5.06]	≈ 1
Experiment = *2*	0.14	[-0.76, 1.06]	0.62
Experiment = *3*	-0.89	[-2.63, 0.82]	0.86
Remote Base = *primed*	0.46	[-0.02, 0.96]	0.97
Semantic similarity (*scaled 1-unit increase*)	-0.30	[-0.66, 0.06]	0.95
Remote Base log-freq. (*scaled 1-unit increase*)	0.02	[-0.30, 0.34]	0.55
Freq. × sim. (*scaled 1-unit increase*)	0.03	[-0.29, 0.35]	0.58

Table 28: Model of combined data from the wug test in Experiments 1, 2, and 3. Coefficients are in log-odds, with positive values indicating an increase in chance of stress shift relative to the intercept.

References

Albright, A., Andrade, A., and Hayes, B. (2001). Segmental environments of spanish diphthongization. *UCLA Working Papers in Linguistics*, 7:117–151.

Albright, A. C. (2002). *The identification of bases in morphological paradigms*. PhD thesis, University of California, Los Angeles.

Anderson, D. R. and Burnham, K. P. (2002). Avoiding pitfalls when using information-theoretic methods. *The Journal of wildlife management*, pages 912–918.

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

Asherov, D. and Bat-El, O. (2016). Multiple defaults: feminine-et and-a in hebrew present tense. *Morphology*, 26(3-4):399–423.

Balota, D. A., Yap, M. J., Hutchison, K. A., Cortese, M. J., Kessler, B., Loftis, B., Neely, J. H., Nelson, D. L., Simpson, G. B., and Treiman, R. (2007). The English lexicon project. *Behavior research methods*, 39(3):445–459.

Baroni, M. (2001). The representation of prefixed forms in the italian lexicon: Evidence from the distribution of intervocalic [s] and [z] in northern italian. In *Yearbook of morphology 1999*, pages 121–152. Springer.

Bat-El, O. (2002). True truncation in colloquial hebrew imperatives. *Language*, pages 651–683.

Baunaz, L., Lander, E., De Clercq, K., and Haegeman, L. (2018). Nanosyntax: the basics. *Oxford Studies in Comparative Syntax*, pages 3–56.

Becker, M. and Levine, J. (2020). Experigen: an online experiment platform.

Benua, L. (2000). *Phonological relations between words*. Psychology Press.

Bermádez-Otero, R. (2011). Cyclicity. *The Blackwell companion to phonology*, pages 1–30.

Bermúdez-Otero, R. (2017). Stratal phonology. In *The Routledge handbook of phonological theory*, pages 100–134. Routledge.

Bobaljik, J. D. (2017). Distributed morphology. In *Oxford Research Encyclopedia of Linguistics*.

Boersma, P. and Hayes, B. (2001). Empirical tests of the gradual learning algorithm. *Linguistic inquiry*, 32(1):45–86.

Boersma, P. and Pater, J. (2016). Convergence properties of a gradual learning algorithm for harmonic grammar.

Bonet, E. and Torres-Tamarit, F. (2010). Allomorphy in pre-clitic imperatives in formenteran catalan: An output-based analysis. In *Romance Linguistics 2009*, pages 337–352. John Benjamins.

Breiss, C., Katsuda, H., and Kawahara, S. (2021). Paradigm uniformity is probabilistic: Evidence from velar nasalization in Japanese. In *Proceedings of WCCFL 39*. Cascadilla Press.

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

Bürkner, P.-C. (2017). brms: An R package for Bayesian multilevel models using Stan. *Journal of statistical software*, 80(1):1–28.

Burnham, K. P. and Anderson, D. R. (2002). *Model selection and multimodel inference: A practical information-theoretic approach.* Springer New York, 2 edition.

Burnham, K. P. and Anderson, D. R. (2004). Multimodel inference: understanding aic and bic in model selection. *Sociological methods & research*, 33(2):261–304.

Burzio, L. (1994). *Principles of English stress.* Number 72. Cambridge University Press.

Burzio, L. (1998). Multiple correspondence. *Lingua*, 104(1-2):79–109.

Burzio, L. (2002). Missing players: Phonology and the past-tense debate. *Lingua*, 112(3):157–199.

Bybee, J. L. and Pardo, E. (1981). On lexical and morphological conditioning of alternations: A nonce-probe experiment with spanish verbs. *Linguistics*, 19(9-10):937–968.

Caramazza, A., Laudanna, A., and Romani, C. (1988). Lexical access and inflectional morphology. *Cognition*, 28(3):297–332.

Carlson, M. T. and Gerfen, C. (2011). Productivity is the key: Morphophonology and the riddle of alternating diphthongs in spanish. *Language*, pages 510–538.

Chomsky, N. and Halle, M. (1968). The sound pattern of English.

Coetzee, A. W. (2016). A comprehensive model of phonological variation: Grammatical and non-grammatical factors in variable nasal place assimilation. *Phonology*, 33(2):211–246.

Coetzee, A. W. and Kawahara, S. (2013). Frequency biases in phonological variation. *Natural Language & Linguistic Theory*, 31(1):47–89.

Eddington, D. (1996). Diphthongization in spanish derivational morphology: an empirical investigation. *Hispanic Linguistics*, 8(1):1–13.

Eddington, D. (1998). Spanish diphthongization as a non-derivational phenomenon. *Rivista di Linguistica*, 10:335–354.

Eddington, D. (2006). Paradigm uniformity and analogy: The capitalistic versus militaristic debate. *International Journal of English Studies*, 6(2):1–18.

Finger, H., Goeke, C., Diekamp, D., Standvoß, K., and König, P. (2017). Labvanced: a unified javascript framework for online studies. In *International Conference on Computational Social Science (Cologne)*.

Fylstra, D., Lasdon, L., Watson, J., and Waren, A. (1998). Design and use of the microsoft excel solver. *Interfaces*, 28(5):29–55.

Goldwater, S. and Johnson, M. (2003). Learning OT constraint rankings using a maximum entropy model. In *Proceedings of the Stockholm workshop on variation within Optimality Theory*, volume 111120.

Gordon, M. (2002). A factorial typology of quantity-insensitive stress. *Natural Language & Linguistic Theory*, 20(3):491–552.

Guekguezian, P. and Jesney, K. (2021). "epenthetic" vowels are not all equal: Gradient representation in yokuts roots and suffixes. In *38th West Coast Conference on Formal Linguistics*, pages 221–230. Cascadilla Proceedings Project.

Gunkel, D. (2010). *Studies in Greek and Vedic Prosody, Morphology, and Meter.* PhD thesis, University of California, Los Angeles.

Gunkel, D. C. (2011). The emergence of foot structure as a factor in the formation of greek verbal nouns in-$\mu\alpha$ (τ)-. *Münchener Studien zur Sprachwissenschaft*, 65:77.

Harley, T. A. (1993). Phonological activation of semantic competitors during lexical access in speech production. *Language and Cognitive Processes*, 8(3):291–309.

Harrell Jr, F. E. (2023). *Hmisc: Harrell Miscellaneous.* R package version 5.0-1.

Harris, J. W. (1969). *Spanish Phonology.* MIT Press.

Hay, J. (2001). Lexical frequency in morphology: Is everything relative?

Hay, J. (2003). *Causes and consequences of word structure.* Psychology Press.

Hay, J. B. and Baayen, R. H. (2005). Shifting paradigms: Gradient structure in morphology. *Trends in cognitive sciences*, 9(7):342–348.

Hayes, B. (1982). Extrametricality and English stress. *Linguistic inquiry*, 13(2):227–276.

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

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

Kilbourn-Ceron, O. and Goldrick, M. (2021). Variable pronunciations reveal dynamic intra-speaker variation in speech planning. *Psychonomic Bulletin & Review*, pages 1–16.

Kilbourn-Ceron, O. and Sonderegger, M. (2018). Boundary phenomena and variability in Japanese high vowel devoicing. *Natural Language & Linguistic Theory*, 36(1):175–217.

Kilbourn-Ceron, O., Wagner, M., and Clayards, M. (2016). The effect of production planning locality on external sandhi: A study in/t. In *The Proceedings of the 52nd Meeting of the Chicago Linguistics Society, lingbuzz/003119.*

Kim, J. (2021). The capitalistic versus militaristic debate: Paradigm uniformity revisited. *Korean Journal of English Language and Linguistics*, 21:636–655.

Kruschke, J. (2014). *Doing Bayesian data analysis: A tutorial with R, JAGS, and Stan.* Academic Press.

Lamontagne, J. and Torreira, F. (2017). Directionality in production planning: Crossword hiatus resolution in spanish. In *Meeting of the Canadian Linguistics Association.*

Lenth, R., Singmann, H., Love, J., Buerkner, P., and Herve, M. (2019). Package 'emmeans'.

Levelt, W. J. (1999). Models of word production. *Trends in cognitive sciences*, 3(6):223–232.

Liberman and Prince (1977). On stress & linguistic rhythm. *LI*, 8(2):249–336.

Linzen, T., Kasyanenko, S., and Gouskova, M. (2013). Lexical and phonological variation in russian prepositions. *Phonology*, 30(3):453–515.

Marchand, H. (1960). *The categories and types of present-day word formation*. Harassowitz, Weissbaden.

Morton, J. (1970). A functional model for memory. *Models of human memory*, pages 203–254.

Nalborczyk, L., Batailler, C., Lœvenbruck, H., Vilain, A., and Bürkner, P.-C. (2019). An introduction to Bayesian multilevel models using brms: A case study of gender effects on vowel variability in standard Indonesian. *Journal of Speech, Language, and Hearing Research*, 62(5):1225–1242.

Nicenboim, B. and Vasishth, S. (2016). Statistical methods for linguistic research: Foundational Ideas — Part II. *Language and Linguistics Compass*, 10(11):591–613.

O'Brien, J. (2007). Lexical conservatism in irish prepositions. *Ms., University of California, Santa Cruz.[Available online at http://people. ucsc. edu/~ jpobrien/papers/obrien_irish_prep. pdf]*.

Pater, J. (2000). Non-uniformity in English secondary stress: the role of ranked and lexically specific constraints. *Phonology*, pages 237–274.

Pater, J., Jesney, K., Staubs, R., and Smith, B. (2012). Learning probabilities over underlying representations. In *Proceedings of the twelfth meeting of the Special Interest Group on Computational Morphology and Phonology*, pages 62–71. Association for Computational Linguistics.

Penny, R. J. (2002). *A history of the Spanish language*. Cambridge University Press.

Pertsova, K. (2005). How lexical conservatism can lead to paradigm gaps. *UCLA working papers in phonology*, 11.

Pertsova, K. and Kuznetsova, J. (2015). Experimental evidence for lexical conservatism in russian: Defective verbs revisited. In *Proceedings of FASL*, volume 24.

Pierrehumbert, J. B. (2016). Phonological representation: Beyond abstract versus episodic. *Annual Review of Linguistics*, 2:33–52.

R Core Team (2021). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria.

Rahman, R. A. and Aristei, S. (2010). Now you see it… and now again: Semantic interference reflects lexical competition in speech production with and without articulation. *Psychonomic Bulletin & Review*, 17(5):657–661.

Rolle, N. R. (2018). Transparadigmatic output-output correspondence: A case study from ese ejja. In *Proceedings of the Annual Meetings on Phonology*, volume 5.

Ryan, K. M. (2016). Phonological weight. *Language and Linguistics Compass*, 10(12):720–733.

Shih, S. (2018). Learning lexical classes from variable phonology. In *Selected Papers from Asian Junior Linguists Conference*, volume 2, pages 1–15. ICUWPL.

Siegel, D. C. (1974). *Topics in English morphology.* PhD thesis, Massachusetts Institute of Technology.

Simonović, M. (2012). The emergence of postcyclic prosody in loanword integration: Toneless latinate adjectives in serbo-croatian. *Acta Linguistica Hungarica*, 59(1-2):221–243.

Simonović, M. and Baroni, A. (2014). Lexicon, markedness and grammar in the serbo-croatian wobbly a. *U B. Arsenijević, A. Janić & B. Stanković (ur.), SinFoniJA*, 7.

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

Smolensky, P. and Goldrick, M. (2016). Gradient symbolic representations in grammar: The case of French liaison. *Rutgers Optimality Archive*, 1552.

Stanton, J. (2016). Learnability shapes typology: the case of the midpoint pathology. *Language*, 92(4):753–791.

Stanton, J. and Steriade, D. (2014). Stress windows and base faithfulness in english suffixal derivatives. Talk at 22mfm.

Steriade, D. (1997). Lexical conservatism. *Linguistics in the morning calm*, pages 157–179.

Steriade, D. (2000). Paradigm uniformity and the phonetics-phonology boundary. *Papers in laboratory phonology*, 5:313–334.

Steriade, D. (2008). A pseudo-cyclic effect in romanian morphophonology. *Inflectional identity*, 18:313–360.

Steriade, D. (2012). The cycle without containment: Latin perfect stems. *Unpublished ms.*

Steriade, D. (2018). Cyclic inheritance, generalized. Colloquium at NYU.

Steriade, D. and Stanton, J. (2020). Productive pseudo-cyclicity and its significance. Talk at LabPhon 17.

Steriade, D. and Stanton, J. (2021). Stress and the cycle, in English and elsewhere. in progress manuscript, MIT NYU.

Steriade, D. and Yanovich, I. (2015). Accentual allomorphs in east slavic: an argument for inflection dependence. *Understanding Allomorphy*, pages 254–314.

Storkel, H. L. and Morrisette, M. L. (2002). The lexicon and phonology.

Tamminga, M. (2018). Modulation of the following segment effect on English coronal stop deletion by syntactic boundaries. *Glossa: a journal of general linguistics*, 3(1).

Tokar, A. (2018). Secondary stress variation in contemporary english. *AAA: Arbeiten aus Anglistik und Amerikanistik*, 43(2):159–191.

Vasishth, S., Nicenboim, B., Beckman, M. E., Li, F., and Kong, E. J. (2018). Bayesian data analysis in the phonetic sciences: A tutorial introduction. *Journal of phonetics*, 71:147–161.

Wagner, M. (2012). Locality in phonology and production planning. *McGill working papers in linguistics*, 22(1):1–18.

Weber, A. and Scharenborg, O. (2012). Models of spoken-word recognition. *Wiley Interdisciplinary Reviews: Cognitive Science*, 3(3):387–401.

Wheeldon, L. (2003). Inhibitory form priming of spoken word production. *Language and cognitive processes*, 18(1):81–109.

Wheeldon, L. R. and Monsell, S. (1994). Inhibition of spoken word production by priming a semantic competitor. *Journal of memory and language*, 33(3):332–356.

White, J. (2017). Accounting for the learnability of saltation in phonological theory: A maximum entropy model with a P-Map bias. *Language*, 93(1):1–36.

Wilson, C. (2006). Learning phonology with substantive bias: An experimental and computational study of velar palatalization. *Cognitive science*, 30(5):945–982.

Wurm, L. H. (1997). Auditory processing of prefixed English words is both continuous and decompositional. *Journal of memory and language*, 37(3):438–461.

Zuraw, K. (2007). Frequency influences on rule application within and across words. In *Proceedings from the Annual Meeting of the Chicago Linguistic Society*, volume 43, pages 283–309. Chicago Linguistic Society.

Zuraw, K. and Hayes, B. (2017). Intersecting constraint families: an argument for harmonic grammar. *Language*, 93(3):497–548.

Zuraw, K., Lin, I., Yang, M., and Pepperkamp, S. (2020). Competition between whole-word and decomposed representations of English prefixed words. *Morphology*.

Zuraw, K. and Peperkamp, S. (2015). Aspiration and the gradient structure of English prefixed words. In *ICPhS*.

Zuraw, K. R. (2000). *Patterned exceptions in phonology*. PhD thesis, University of California, Los Angeles.

Zymet, J. (2018). *Lexical propensities in phonology: corpus and experimental evidence, grammar, and learning*. University of California, Los Angeles.