Modeling frequency-conditioned paradigm uniformity in Japanese voiced velar nasalization*

Canaan Breiss, University of Southern California†
Hironori Katsuda, University of Kansas
Shigeto Kawahara, Keio University/International Christian University

Under review at *Phonology*

Abstract

Recent quantitative work on the variable [g] ~[ŋ] alternation in compounds of certain dialects of Japanese has revealed *token frequency* of the compound as a whole, and of the compound's second-member (N2) in its freestanding form, to be important predictors of the alternation (Breiss et al., 2021a, Breiss et al. *to appear*). In this paper, we propose a formal phonological analysis of data presented in Breiss et al. (*to appear*) that integrates usage-based factors like frequency with the action of the phonological grammar, extending the mechanisms of lexicon-grammar interaction proposed in Breiss (2024). We demonstrate that the model fits experimental data better than—or at least comparably to—a theoretically-naïve statistical model proposed in the previous work. Based on the success of our modeling, we discuss the role of token frequency in phonological patterning more broadly, and how the mechanism that we propose in this paper might be extended to unify a range of contradictory frequency-dependent processes that have been observed in the literature.

1 Introduction

This paper is about how to integrate information about *usage frequency*—here, the token frequency of morphemes in the language experience of an individual speaker—into a constraint-

*This work was supported in part by JSPS grant #22K00559 to Shigeto Kawahara. Thanks to audiences at the University of Southern California and the University of Pennsylvania for valuable discussion and feedback. Supplementary material can be accessed at *https://osf.io/wcvyx/?view_only=ca3fb8d413d04a76b60844fe78376594*.

†Contact: cbreiss@usc.edu

based phonological grammar formalism that characterizes that speaker's generative linguistic knowledge.

We take as our empirical case the frequency-conditioned variability in optional paradigm uniformity in voiced velar nasalization (henceforth "nasalization") in phonologically-conservative Japanese dialects, recently studied using corpus data by Breiss et al. (2021b) and experimentally verified in Breiss, Katsuda and Kawahara (*to appear*) — henceforth BKK. These studies are the latest in a long research tradition centered on the allophonic distribution of /g/ in conservative Japanese dialects, where a static phonotactic restriction enforces /g/ to be realized as [g] prosodic-word-initially and [ŋ] elsewhere (e.g. Kindaichi 1942; Trubetskoy 1969; Labrune 2012). This correspondence is disrupted in compounds with /g/-initial second member (N2) that can occur as a free morpheme: in compounds with N2s that do not occur as free-standing words, the /g/ $\rightarrow$ [ŋ] alternation is exceptionless, but in compounds where N2 may additionally occur as a free-standing word (that is, with initial [g]) the nasalization process is optional (Ito and Mester, 1996, 2003).

The contribution of recent work by BKK (reviewed in detail in section 2) is to characterize this variation in quantitative detail, and in particular to highlight how the token frequency of both the compound and the free N2 impact the outcome of optional nasalization: higher frequency compounds encourage more nasalization of medial /g/ to [ŋ], while higher frequency free N2s encourage more retention of medial /g/ as [g], remaining uniform across the paradigm of their free-standing forms and compound forms (Steriade, 2000; Benua, 2000).

The novel contribution of this paper is to provide a formally explicit model of the experimental data. The model builds upon the Voting Bases model of lexicon-grammar interaction (Breiss, 2024), originally proposed to model Lexical Conservatism (Steriade, 1997). Lexical Conservatism is a type of paradigm uniformity where the distribution of stem allomorphs (referred to as "bases") in a paradigm influences the way that paradigm accommodates new members. The canonical example comes from Steriade (1997), who observed that the phonologically-similar forms *rémedy* and *párody* differ in their behavior when affixed with *-able*, yielding *remédiable* with shifted stress, but *párodiable*, with fixed stress. She argued that this difference stems not from the forms *remedy* and *parody* themselves, but from the fact that *remedy* has a stem allomorph *remédi-* in *remédial* that satisfies the marked lapse arising from affixation.

Breiss (2024) examined the same Lexical Conservatism dependency using novel derived forms (like *lábor* + *-able*, with related form *labórious*, and *pláster* + *-able* with no phonologically-advantageous related form), and found that in experimental settings, speakers are sensitive not only to the *presence* of the phonologically-beneficial stem allomorph (like *remédial* and *labórious*), but also to its salience in the lexicon as manipulated by priming. To account for these data, he proposed a formal phonological model that integrates the influence of the contents of the lexicon along with their resting activation, enabling the phonological grammar to be sensitive to the psycholinguistic properties of the morphemes which it manipulates. Breiss (2024) termed this formal model of lexicon-grammar interaction the *Voting Bases* model.

In this paper, we demonstrate that the Voting Bases model extends, without modification, to the separate case of lexicon-grammar interaction found in Japanese nasalization. The success of the model suggests that the foundational principles of the Voting Bases model may be a good candidate for a general theory of the way that the lexicon and grammar interact. This finding also underscores the explanatory value to be gained for phonological phenomena by adopting a more psycholinguistically-nuanced portrait of the lexicon as a dynamic substrate that can influence the computations of the grammar on the items which it contains.

The layout of the paper is as follows: the first two sections of the paper review in some depth basic facts about Japanese nasalization drawn from the literature (section 2), and then specifically reviews in detail Experiment 1 of BKK (section 3). The following section, 4, focuses on the Voting Bases model, and how we apply it to the context of optional paradigm uniformity. Section 5 then actually fits the model to the experimental results, and discusses relative and absolute model fit in comparison to minimally-different models that incorporate only some of the assumptions of the Voting Bases model. The paper closes in section 6 with a discussion of broader issues, touching on how such a system might come to be in the mind of the learner, on the merits of a joint model of psycholinguistic and grammatical influence on word formation, and on what a unified theory of token frequency effects on the phonological grammar might look like.

2 The traditional picture of Japanese nasalization

The data that we model in this paper comes from Experiment 1 of BKK, which investigated the variation between [g] and [ŋ] induced by the phonotactics of phonologically-conservative dialects of Japanese. The pattern, which has been well-studied in both in generative and non-generative literature on Japanese linguistics (Kindaichi, 1942; Trubetskoy, 1969; Hibiya, 1995; Labrune, 2012; Ito and Mester, 1996, 2003), is exemplified in the complementary distribution of [g] and [ŋ] shown in the monomorphemic data in example (1) below, where the voiced oral velar stop is only permitted word-initially, and the velar nasal is only permitted word-medially.

(1) a. /kaŋami/ → [kaŋami]
"mirror"
b. /gimu/ → [gimu]
"obligation"

We assume throughout that non-alternating forms are stored surface-true as URs in the lexicon, in accordance with the phonological tradition of (Strong) Lexicon Optimization (Prince and Smolensky, 1993; Sanders, 2003). This stance is supported by psycholinguistic research on the contents of lexical representations, reviewed in section 4.1.

Japanese's extensive use of compounding in word-formation gives the opportunity for the phonotactic restriction to drive alternations, seen in examples (2)-(5) below. Here we see that when a /g/-initial morpheme is word-initial (either as a prosodically-free word, in examples (2)–(4), or as the first member (N1) of a compound, in example (5)[1]), it is realized with an initial [g], while when it occurs as the second member of a compound (N2) it is realized with initial [ŋ]. Critically for the current study, Ito and Mester (1996) observed that although in all cases the /g/-initial N2 *may* be realized word-medially with initial [ŋ], nasalization is optional when the N2 can stand on its own as a prosodically-free form (cf. the "b" series in (2)–(4) vs. (5c))—a case of optional paradigm uniformity.

(2) a. /hai + gan/ → [hai-ŋan] ~ [hai-gan]
lung cancer
"lung cancer"

b. /gan/→ [gan]
cancer
"cancer"

(3) a. /noo + geka/ →[noo-ŋeka] ~ [noo-geka]
brain surgery
"brain surgery"

b. /geka/→ [geka]
surgery
"surgery"

(4) a. /doku + ga/ → [doku-ŋa] ~ [doku-ga]
poison moth
"poison moth"

b. /ga/ → [ga], "moth" (a free-standing morpheme)

(5) a. /doku + ga/ → [doku-ŋa], *[doku-ga]
poison fang
"poison fang"

b. /ga + ʒoo/→ [ga-ʒoo]
fang castle
"main castle"

[1]We temporarily adopt here for the traditional assumption that the [g]-initial form of a free N2 is underlying, for expository ease and continuity with the previous literature. Our own proposal is laid out in section 4.

c. /ga/→ *[ga] (a bound morpheme)
fang
“fang”

Breiss et al. (2021b) examined this variation in a corpus derived from a pronouncing dictionary (NHK, 1993) and found that among compounds with free N2s, the two most prominent predictors of whether an item would be nasalized was the frequency of the N2’s free [g]-initial form, and the frequency of the whole compound. These effects ran in opposite directions: higher frequency compounds were more likely to be nasalized (the left facet of Figure 1); on the other hand, the more frequent the free N2, the less likely the nasalization was (the right facet of Figure 1).

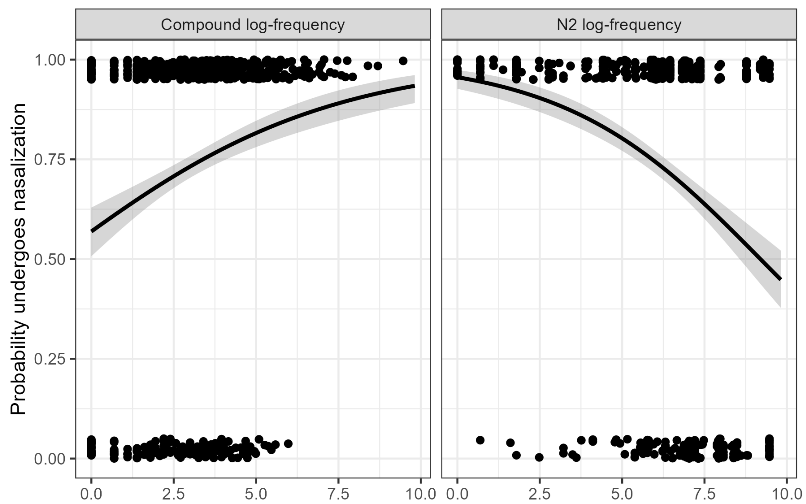

Figure 1: The effects of whole compound frequency (left facet) and N2 frequency (right facet) on the probability of nasalization (vertical axis), with binomial smooths in the corpus data. One dot represents one lexical item; vertical jitter has been added for readability. Figure and caption adapted with permission from Breiss et al. (*to appear*), data from Breiss et al. (2021b).

The corpus data was modeled as a case of probabilistic paradigm uniformity in Breiss et al. (2021a) using Output-Output Faithfulness constraints (Benua, 2000) indexed to items binned by the relative frequency of each compound and N2. The paper was limited, however, by the untested assumption of their model that the frequency-modulation of paradigm uniformity in their corpus data actually represents the synchronic knowledge of speakers. Additionally, their formal model was not explicitly informed by psycholinguistic considerations and thus its linking hypothesis between frequency (necessarily a lexical characteristic) and the phonological grammar was vulnerable to criticism on grounds of being stipulative—in other words, there was nothing in their model that prevented the opposite relation between frequency and paradigm uniformity from holding.

In this paper, we offer two improvements on the state of affairs in Breiss et al. (2021a). First, we model experimental data from Breiss et al. (*to appear*) (BKK) where the frequency-conditioning of the variable paradigm uniformity is reproduced in existing compounds and extended to novel ones. Second, we do this by extending the Voting Bases model of Breiss (2024) which is compatible with consensus understanding of the way lexical frequency is connected to the lexical representation and activation, and which offers an explicit linking hypothesis relating the real-time dynamics of the lexicon to the representation and computations of the phonological grammar.

3 BKK's Experiment 1

BKK carried out two experiments on Japanese nasalization. Their paper had the goal of seeing whether the corpus data was representative of speakers' generalizable knowledge, both in the aggregate and also at the level of the individual. They found that both individually and in aggregate, speakers' propensity to nasalize displayed sensitivity to the frequency of the free N2 and compound, in existing and novel compounds. In this paper, we focus our modeling efforts on the results of their Experiment 1, which we describe in some detail below.[2]

3.1 Stimulus selection

BKK chose stimuli that were roughly balanced between existing Japanese compounds of varying frequencies, and novel (that is, zero-frequency) semantically-compositional compounds. Both existing and novel stimuli had attested free N2s of a range of frequencies. Out of a desire to sample compounds with a wide range of frequencies that would likely be known to participants, existing compounds ranged from two to eight moras in length, while all novel compounds were four moras long.

3.2 Participants

BKK sought to recruit speakers of the phonologically-conservative Tōhoku dialect of Japanese, and used word-of-mouth and snowball sampling to find 20 speakers. In order to increase the precision of individual-level estimates of experimental manipulations, all but one speaker participated in the experiment twice, with the two sessions separated by a period ranging from a few weeks to a few months. Although the phonologically-conservative Yamanote dialect was

[2] They also sought to determine whether correlation between nasalization and the overall prosodic size of the compound, which is observed in the corpus (Breiss et al., 2021b) but is a typologically unusual pattern, was replicated in participants' online productions (Experiment 2). They actually found that there was no evidence of a direct relationship between nasalization and global prosodic length (cf. Jiang 2023). We therefore do not address this experimental data here, as our point is made in the simpler case of data from Experiment 1.

the subject of the corpus study of Breiss et al. (2021b) and of much of the previous linguistic analysis, this dialect is currently spoken only by the very elderly who might not be comfortable participating in a study online.

Before participating in the experiment, each participant was screened with a short dialect questionnaire to ensure that their speech exhibited the allophonic distribution of word-initial [g] and word-medial [ŋ]. This was done because the literature addressing nasalization assumes that it is triggered in order to enforce compliance with this phonotactic restriction imposed on monomorphemic words; therefore it is important to base conclusions about the variability of nasalization on data from speakers who do exhibit the phonotactic in question.[3] For the purposes of the model which we develop, we will see that these monomorphemic words provide crucial evidence for the lower bound of the weight of the markedness constraint driving nasalization, since with data from compounds alone, it is not uniquely identified against the background of faithfulness constraints that the Voting Bases model uses (see section 5 for further details).

The dialect questionnaire consisted of a production task where speakers were asked to read aloud 10 monomorphemic words with word-initial [g] of varying frequencies, and 10 monomorphemic words with word-medial [ŋ]. The stimuli were written with *kanji* orthography, which does not distinguish between [g] and [ŋ]—this is also true of the main production experiment described below, so we follow BKK in assuming that the participants' production was not influenced by orthographic factors. The twenty words were shown to the participant in a random order, and their productions were recorded; only the eight participants who exhibited the target pattern of allophony in all monomorphemes were invited to participate in the main experiment.

3.3 Experiment structure

Each experimental session proceeded as follows. First, participants were asked to complete the dialect questionnaire; if they met the criterion discussed above, they were invited to complete the rest of the experiment. Before the production task, participants read out loud and indicated whether they knew half of the free N2s in the experiment—this was done in order to prime them, under the hypothesis that raising their resting frequency in the participant's lexicon would influence their phonological behavior (see further discussion in section section 4.4). After this knowledge check, participants saw each compound one at a time in a random order, and produced the form aloud while their speech was recorded. After the production task, participants produced and indicated knowledge of the other half of the free N2s in the experiment, as well as all of the compounds. The N2s that were primed were counter-balanced between participants across the two runs of the experiment.

[3]The question of how nasalization is produced or represented by speakers who entirely lack, or only variably enforce, the [g] ~[ŋ] allophony in monomorphemic words is not addressed by BKK, nor do we consider it here.

3.4 Results

BKK found that the participants reflected at an individual level the frequency-conditioned variability seen in the corpus study of Breiss et al. (2021b). In existing compounds (Figure 2), their productions were influenced by both the frequency of the compound (the left facet), for which higher values correlated with more nasalization, and by the frequency of the free N2 (the right facet), where higher values correlated with less nasalization.

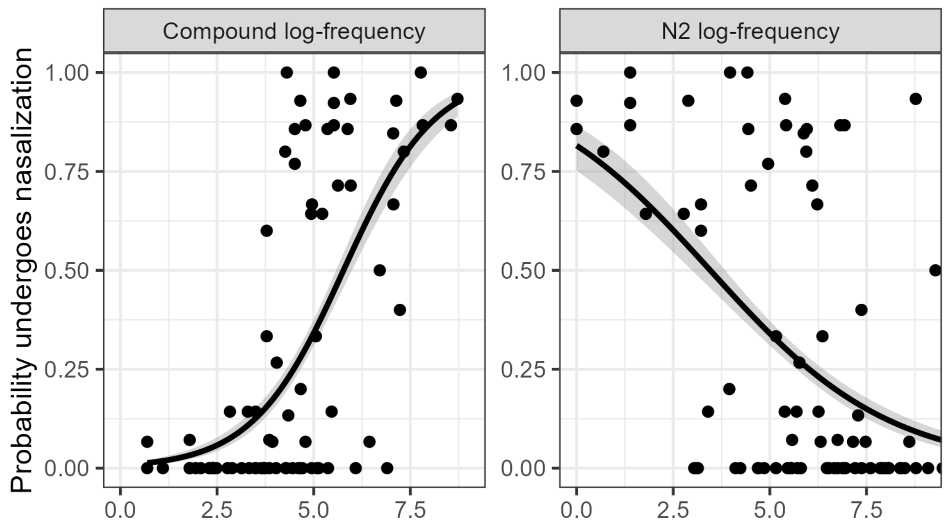

Figure 2: Probability of nasalization (the vertical axis) plotted against compound log-frequency (the left facet) and N2 log-frequency (the right facet), with binomial smooths for readability, in the experiment by BKK. Plot and caption reproduced with permission from Breiss et al. (*to appear*).

Figure 3 plots the same effect of N2 frequency in novel compounds: forms with higher-frequency N2s were less likely to undergo nasalization relative to those with lower-frequency N2s.

Finally, BKK found that the frequency effect was stable at the level of the individual, across existing and novel compounds, which is plotted in Figure 4. In this Figure, the horizontal axis plots the strength and direction of the effect of N2 log-frequency in novel compounds, and the vertical axis plots the strength and direction of the effect of N2 log-frequency in the existing compounds; see the caption of Figure 4 for further details. Although different participants were more or less sensitive to the frequency of a given N2, lying higher or lower on each axis, there was uniformity in this degree of sensitivity such that the two co-varied along a diagonal line through the center of the plot. BKK interpreted this correlation as evidence that morpheme usage frequency and phonological markedness have separable, distinct influences on speaker productions.

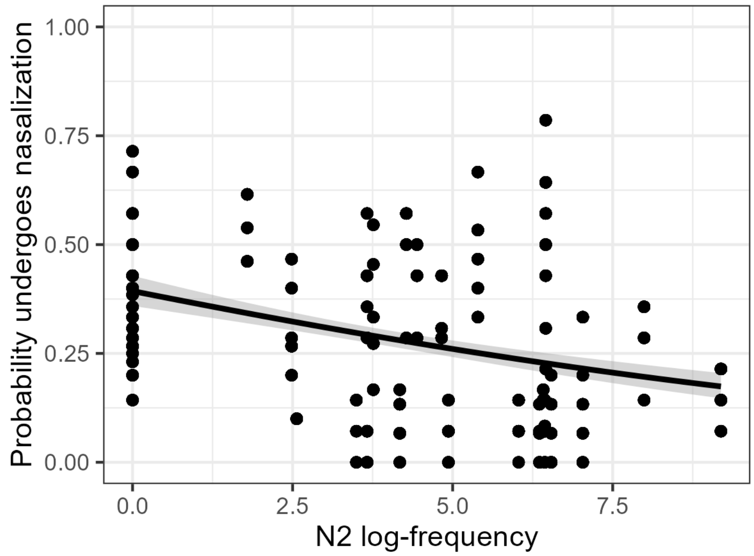

Figure 3: The probability of undergoing nasalization in novel compounds, plotted against N2 log-frequency, with a binomial smooth to aid readability. Plot and caption reproduced with permission from Breiss et al. (*to appear*).

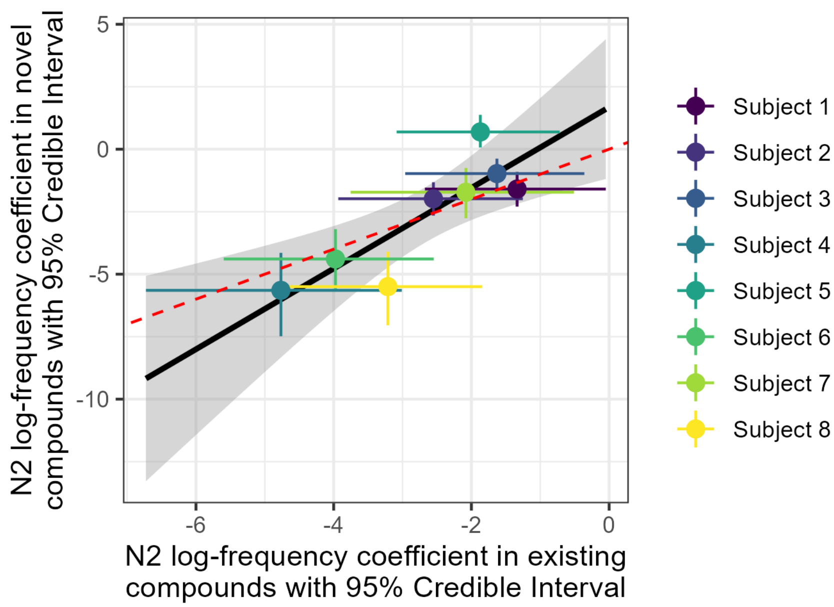

Figure 4: The coefficient of N2 log-frequency in novel compounds, derived from the model in Table 1 of Breiss et al. (*to appear*), is plotted on the horizontal axis, and the coefficient for N2 log-frequency in existing compounds, derived from the model summarised in Table 3 of Breiss et al. (*to appear*), is plotted on the vertical axis. Points represent median values of the posterior with ranges encompassing the 95% Bayesian Credible Intervals, colors represent speakers, and a linear smooth has been added for readability, with the line of slope 1 intersecting the origin in dotted red. Plot and caption adapted with permission from Breiss et al. (*to appear*).

3.5 Summary and goals for modeling

To summarize, the findings of BKK that are relevant for the modeling task of this paper are the following. Among those speakers for whom the phonotactic restriction enforcing [g]~[ŋ] allophony was exceptionless in monomorphemic words:

1. Phonotactically-driven nasalization is variable in compounds with prosodically-free N2s.
2. In these compounds, the probability of nasalization is increased by higher compound frequency, and decreased by higher N2 frequency.
3. The frequency effect is uniform within individuals across existing and novel compounds.

Below, we propose a formal model of these facts, using the Voting Bases model to relate a lexicon containing usage-frequency information to a phonological grammar couched in the Maximum Entropy (MaxEnt) framework.[4]

4 Modeling token frequency in the phonological grammar

Based on the facts laid out above, we seek a model of the phonological grammar that allows non-phonological properties of individual lexical items (here, frequency) to influence their participation in phonological processes (here, paradigm uniformity). Note that we specifically aim to model phonological and non-phonological influences on the outputs of the phonological grammar, rather than any possible morphological or paradigmatic effects on phonetic realization (see Purse et al. (2022) for a review), about which the Voting Bases model as laid out in Breiss (2024) makes no predictions.

4.1 The contents of a lexical entry

As prolegomena to the grammatical model, it will be important to establish some relevant context regarding the contents of the lexicon, because it is these representations that are at stake in discussions of token frequency. Psycholinguistic research has amassed a large body of evidence that the lexicon is richly structured, with numerous types of linked representations of various levels of detail grouped under the same lexical entry. We do not review this research in depth here, but simply highlight the findings relevant to developing the type of integrated phonological

[4]We do not attempt to model the frequency of the first compound member, N1, on the probability of nasalization in compounds, since this was not manipulated by BKK. Future work might profitably pursue this question experimentally and formally, since corpus data in Breiss et al. (2021b) suggests that higher N1 frequency may also independently lower the probability of nasalization; see Rebrus and Törkenczy (2017) for a similar finding of N1 frequency on compound coherence in Hungarian vowel harmony.

theory referenced above. For a thorough discussion and literature review on the (phonologically-relevant) contents of a lexical entry, see Pierrehumbert (2016); for more on how this information interacts with the Voting Bases theory in cases beyond those relevant for the nasalization, see Breiss (2021, 2024).

Since nasalization concerns paradigm uniformity, we assume the lexical entry for an existing word lists (among many other things) their allomorphs (cf. Strong Lexicon Optimization, Sanders 2006): for a non-alternating monomorpheme like [kaŋami] "mirror", this would be simply /kaŋami/; for a monomorpheme that can appear as an N2 and undergo nasalization, such as [ga]–[ŋa] "moth", the lexical entry would list both /ga/ and /ŋa/. Finally, we assume that existing compounds are stored whole, with nasalization applied so as to respect the phonotactic in the lexicon (Albright, 2008; Martin, 2007).

With regard to non-phonological characteristics of the lexicon, we follow a large body of evidence that lexical representations have differing degrees of salience or strength of encoding, which is often referred to as their *resting activation* (Morton, 1970). Following Breiss (2021, 2024), we take resting activation to correspond to the strength of a memory representation itself, not a number or rank stored in long-term memory as a characteristic of the lexical item. Thus, characteristics (long-term or dynamic) of lexical items like their frequency, and whether or not they were recently activated (for example, by priming), all contribute dynamically to an item's resting activation. Importantly, also following Breiss, we use the term "resting activation" as a stand-in for any scalar summary statistic that can be derived from an implemented model of lexical dynamics. We remain intentionally agnostic as to the specific model of these dynamics, simply stressing that so long as such a model can be used to drive a measure of relative salience influenced by the factors just mentioned, the Voting Bases model can make reference to it to scale faithfulness constraint violations. We discuss how resting activation is modeled as influencing the phonological grammar below in section 4.4.

4.2 The Voting Bases model

We now turn to a formal phonological model of the Japanese nasalization data. We use the Voting model of Base competition proposed in Breiss (2021, 2024). The Voting model has been used to model data in Lexical Conservatism in English and Spanish, and is broadly compatible with the view of the lexicon laid out above. Here, we extend the model to the probabilistic paradigm-uniformity found in Japanese nasalization.

The Voting Bases model has two parts: the first is that all listed stem allomorphs ("bases") in the lexicon exert an analogical pull on derivatives (operationalized using allomorph-specific faithfulness constraints), violations of which are scaled in proportion to the resting activation of the representation to which faithfulness is being assessed. The second part is that markedness constraints evaluate candidates in the standard way for constraint-based phonological models.

The Voting Bases model assumes a probabilistic, weighted-constraint phonological grammar; here, we use MaxEnt Harmonic Grammar (Smolensky, 1986; Goldwater and Johnson, 2003), but in principle we could also use another grammar formalism that has these characteristics, like Stochastic (or Noisy) Harmonic Grammar (Boersma and Pater, 2016). We use MaxEnt since it has various strengths; e.g. it directly relates Harmony to probability (Hayes, 2022), permits constraint cumulativity by default (Jäger and Rosenbach, 2006; Breiss, 2020), has a learning algorithm to set its weights, and is rooted in well-understood statistical techniques used widely outside linguistics (Jurafsky and Martin, 2009, ch. 5). We stress, however, that our analyses can be recast in terms of other stochastic constraint-based frameworks.

4.3 Constraints

In the analysis developed in this paper, we adopt the general approach of Ito and Mester (1996, 2003), following loosely Breiss et al. (2021a). We only use three constraints: a single markedness constraint to motivate nasalization (extending the spirit of the constraint *VgV from Ito and Mester (2003) to be compatible with nasal-final N1s, which pattern identically to vowel-final N1s), and a pair of faithfulness constraints which correspond to the second member's free form and to the analogical pull of the compound as a whole, if one exists. They are listed below.[5]

- ***Internal-[g]**: Assign one violation for each word-internal [g] in a candidate.
- **Id-[nasal]-N2**: Assign one violation for each segment in the listed allomorph for the free-standing N2 that does not match its corresponding segment in the feature [nasal].
- **Id-[nasal]-Compound**: Assign one violation for each segment in the listed allomorph for the full compound that does not match its corresponding segment in the candidate in the feature [nasal].

Note that the constraint definitions do not make reference to scaling or the contents of the lexicon; the proposal in the Voting Bases model is an architectural proposal about how psycholinguistic, "extra-grammatical" factors act within and beside the phonological grammar to influence certain variable phenomena.

[5]The first faithfulness constraint plays the same role as faithfulness to the Remote Base in an analysis of Lexical Conservatism. The second faithfulness constraint parallels faithfulness to the Local Base in a Lexical Conservatism analysis (Breiss, 2021, 2024). We use more transparent names here for the sake of clarity, since nothing in the Voting Bases model structurally prioritizes Remote Bases over Local ones.

4.4 Modeling resting activation

The discussion in 4.1 above left open how a specific numerical value for resting activation might be calculated on the basis of the psycholinguistic characteristics of item's lexical entry. Here, we take the approach of modeling it as function of log-frequency of the frequency of the allomorph, which is passed through the sigmoid function $\frac{1}{1+e^{-logfreq}}$ that translates the linear predictor (i.e. $-logfreq$) into the bounded interval of {0,1}, which will be the scaling factor applied to faithfulness violations. This is illustrated in Figure 5. The effect of this non-linear transformation will be to preserve the idea that it is less penalized to be unfaithful to low-frequency lexical items compared to higher-frequency ones, while damping down the difference between extreme values of the scale and rendering it bounded. The final move we make here is rather than using *raw* log-frequencies, we use *scaled and centered* log-frequencies, following the statistical analysis in BKK. This corresponds to the notion that it is not so much the *absolute* frequency of each item that is important, but how frequent it is relative to the other competitor items in the lexicon (here approximated by the population of items in the experiment), which is in line with previous work on morphological decomposition in stored forms (Hay, 2001). Finally, in the analysis that we develop below, we do not model the priming of N2, since BKK did not find substantial evidence that it affected their experimental data.[6]

4.5 Schematic illustrations

Before modeling the experimental data itself, it will be useful to work with some toy data to get a feel for how resting-activation-scaled faithfulness violations interact with the dynamics of a MaxEnt grammar. First, let us consider the case of novel compounds, since they are the simplest case to lay out the workings of the analysis. Recall the empirical pattern: here, although the frequency of the compound is zero, we nevertheless find that nasalization is modulated by the frequency of N2. Now, consider the case of two hypothetical novel compounds, one with a higher-frequency N2, and one with a lower-frequency N2, such that when the sigmoid transformation is applied to their frequencies the higher-frequency form scales its violations of Id-[nasal]-N2 by 0.7, and the lower scales its own violations of the same constraint by 0.3 (these specific numbers are chosen purely for the sake of illustration). Using the constraints defined in section 4.3 above, we can define the tableaux below in Figure 6.

[6]The Voting Bases framework is easily extensible to multiple predictors of resting activation: to incorporate priming, one could simply treat the term passed into the sigmoid as itself a log-linear model, adding a coefficient (weight) for the effect of priming, in addition to a coefficient for the effect of lexical frequency. This is beyond the current scope of this paper, however, and so we simply assume a fixed coefficient for lexical frequency, since there being only one predictor in the log-linear model for resting activation would make the coefficient of frequency redundant with the weight of the faithfulness constraint being scaled. Similarly extensions of the Voting Bases model could also model by-participant variability in the priming effect using a hierarchical model structure.

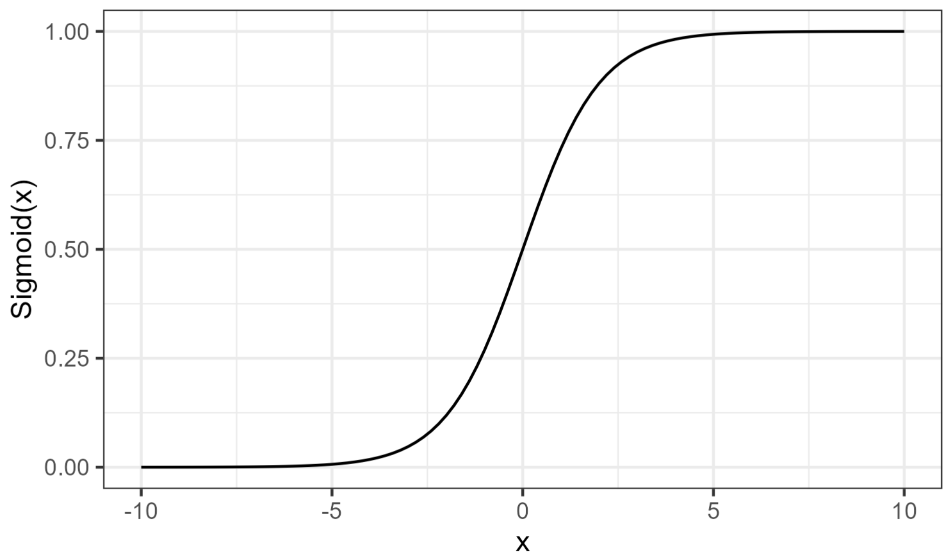

Figure 5: Sigmoid function that translates the (centered) frequencies into the scaling factors. See text for details.

$/.../_{N1}$, $/g.../_{High-freq.N2}$	*INTERNAL-[g]	ID-[nas]$_{N2}$		
Weight:	2	1	H	*p*
a. [...g ...]	1		2	.21
b. [...ŋ ...]		.7	.7	.79
$/.../_{N1}$, $/g.../_{Low-freq.N2}$	*INTERNAL-[g]	ID-[nas]$_{N2}$		
Weight:	2	1	H	*p*
c. [...g ...]	1		2	.15
d. [...ŋ ...]		.3	.3	.85

Figure 6: Schematic application of the Voting model of Base Competition to the formation of a novel compound in the *wug*-test.

We can see that the pull of faithfulness to the N2 with higher frequency is stronger than the one with lower frequency, though both are relatively marginal outcomes since the weight of *INTERNAL-[g] dominates the distribution of probabilities in this scenario.

Moving on to existing compounds, we now must add another item to the lexical entry we are considering in our left-hand input cell to our tableaux, shown in Figure 7. For the sake of minimal contrasts, we assume that the frequency of both N2s are equal and medial relative to the examples in Figure 6 above, allowing us to examine the effect of compound frequency holding N2 frequency constant. However, in our analysis of the actual data, both scaling factors are independently set on a per-item basis.

Here we see that the scaling of the compound again depends on frequency, but because of the

/…/$_{N1}$, /g…/$_{N2}$, /…ŋ…/$_{High-freq.compound}$ Weight:	*INTERNAL-[g] 2	ID-[nas]$_{N2}$ 1	ID-[nas]$_{Compound}$ 1	H	*p*
e. […g …]	1		.7	2.7	.09
f. […ŋ …]		.5		.5	.91
/…/$_{N1}$, /g…/$_{N2}$, /…ŋ…/$_{Low-freq.compound}$ Weight:	*INTERNAL-[g] 2	ID-[nas]$_{N2}$ 1	ID-[nas]$_{Compound}$ 1	H	*p*
g. […g …]	1		.3	2.3	.14
h. […ŋ …]		.5		.5	.86

Figure 7: Schematic application of the Voting model of Base Competition to the formation of an existing compound in the *wug*-test.

assumption we made about the listed form of the compound—specifically, that phonologically well-formed words are preferentially the target of lexicalization (Albright, 2008; Martin, 2007)—we find that the faithfulness to the compound's UR penalizes the candidate that does not exhibit nasalization and violates markedness.

Finally, we lay out the case where the competition between candidates is driven primarily by faithfulness. Above, where markedness had a high weight, the candidate that satisfied markedness had a higher probability than the one which violated it, and the effects of the faithfulnesss constraints were on the probability of the minority candidate. In the scenario where markedness is low and the weights of the faithfulness constraints are dominant, the majority candidate is the one that satisfies faithfulness to the whole compound, and the presence of the N2 is the main reason that the unfaithful (but markedness-satisfying) candidate gets appreciable probability; this is a type of "analogical" effect where markedness has little role, as in Figure 8, in which the markedness constraint is assigned a very low weight (here, arbitrarily set as 0.1).

5 The model in action

Moving on to the analysis itself, we first fit models separately to the existing and novel compound data, to demonstrate the suitability of the Voting model in each context, and also to allow better comparison to the statistical models fit to the experimental data above. We then consider how nasalization might be modeled more comprehensively by incorporating information from non-alternating monomorphemes where the complementary distribution between [g] and [ŋ] is enforced, and also by incorporating our knowledge that the free form of N2s surface non-alternatingly with initial [g], despite the presence of an [ŋ]-initial stem allomorph.

In all cases, we fit the MaxEnt models using the *Solver()* function in Microsoft Excel (Fylstra

/…/$_{N1}$, /g…/$_{N2}$, /…ŋ…/$_{High-freq.compound}$ Weight:	*Internal-[g] 0.1	Id-[nas]$_{N2}$ 1	Id-[nas]$_{Compound}$ 2	H	*p*
i. […g …]	1		.7	1.5	.27
j. […ŋ …]		.5		0.5	.73
/…/$_{N1}$, /g…/$_{N2}$, /…ŋ…/$_{Low-freq.compound}$ Weight:	*Internal-[g] 0.1	Id-[nas]$_{N2}$ 1	Id-[nas]$_{Compound}$ 2	H	*p*
k. […g …]	1		.3	0.7	.45
l. […ŋ …]		.5		0.5	.55

Figure 8: Schematic application of the Voting model of Base Competition to the formation of an existing compound in the *wug*-test, in a regime where faithfulness is strong and markedness weak.

et al., 1998), and used a relatively weak Gaussian prior of Normal(0,10) on constraint weights, which has the effect of allowing weights to vary in response to values that best fit the data, while making extreme values (here, above twenty or so) less appealing. For more on priors on weights in MaxEnt phonological models, see Wilson (2006) and White (2017). All models fit in this paper are provided in the supplementary materials.

5.1 Existing compounds

We first applied the analysis sketched in section 4.5 to data from existing compounds. Recall that in these forms, compounds with higher-frequency N2s are more likely to resist nasalization than those with lower-frequency N2s, but that compound frequency itself also influences nasalization, with higher-frequency compounds favoring the surface-realization of their underlying [ŋ]. We take as our data to model the counts of compounds produced having undergone nasalization or not, in cases where speakers know both the compound and the N2 in question.

The best-fitting constraint weights for existing compounds are 6.93 for Id-[nasal]-Compound, and 7.22 for Id-[nasal]-N2, with *Internal-[g] receiving a weight of zero. The weights of the two faithfulness constraints were not significantly different from one another, as assessed via a likelihood ratio test: Δlog-likelihood $= 1.3$, $p = 0.11$; a similar conclusion was suggested by the near-zero difference in the sample-size corrected AIC of the two models: ΔAICc $= 1.8$. AICc differences greater than 10 are typically taken to indicate strong support for the model with the lower AICc value; for more on model-comparison in statistical models and phonological grammars, see Shih (2017) and Wilson and Obdeyn (2009). This result suggests that the attractive influence of both Bases is critical in driving the alternation in attested forms; the zero weight of the markedness constraint *Internal-[g] indicates that in existing compounds, analogical faithfulness is doing all the work here, despite the assumption in the literature that the alternation is markedness-driven. We will revisit the role of markedness when evaluating the joint model of novel and existing compounds below in section 5.3.

We also compared the full model to one where the two faithfulness constraints were allowed to take on different values but were not scaled by frequency. As one might expect, since low- and high-frequency forms have the same violation profiles in the phonological grammar, a grammar without access to frequency information can only predict one rate of nasalization across all forms; this model fits the data dramatically less well (Δlog-likelihood = 264.87, $p < .001$ with one degree of freedom, ΔAICc = 527.15).

Finally, we evaluate the absolute performance of the model by examining how well it fits the data: although the two models have different internal structures, we can ask whether the theoretically-informed MaxEnt model here does as good a job in explaining the data patterns as the theory-neutral mixed effects logistic regression model reported by BKK when assessing the statistical robustness of the experimental results.[7] We do this using the measure of R^2, which ranges from zero to one, and can be thought of as the proportion of the variation in the dependent variable (here, whether nasalization applies or not) explained by the collection of independent variables (the phonological and lexical characteristics of interest).

We used the *r2_bayes()* function from the *performance* package (Lüdecke et al., 2021) to obtain the marginal R^2 of the statistical model—that is, the amount of variance in the data explained by the fixed effects—and compared it to the R^2 for the MaxEnt model.[8] Since the statistical model is Bayesian, we obtain a median and 95% Credible Interval for our R^2: 0.48 and [0.31, 0.56], respectively. This is lower, though still relatively comparable, to the MaxEnt models R^2 of 0.63, for which we have only a point estimate. Although the two are relatively close, the point value for the marginal R^2 of the MaxEnt model is outside the 95% Credible Interval of the statistical model; this comparison suggests that the theoretically-structured model out-performs the theory-blind statistical one. While we find this result to be encouraging, this conclusion is tentative, however—since the MaxEnt model does not capture variation at the level of the speaker, it may be overfitting the data somewhat, attributing to the population grammar variance that should more conservatively be attributed to speaker-level idiosyncrasies.

[7]The model specification in BKK was as follows: `Nasalization ~ 1 + LogN2Freq*N2Primed + LogCompoundFreq + NasalFinalN1 + (1 + LogN2Freq*N2Primed + LogCompoundFreq + NasalFinalN1 | Speaker) + (1 + N2Primed | Compound)`; see BKK section 3.3 and 3.4.1 for details.

[8]We used marginal R^2, which makes reference to fixed effects only, since the conditional R^2 that takes into account the variance explained by both fixed and random effects has no direct comparison in the MaxEnt model we fit. For more on the relationship between mixed effects models and hierarchical structures in linguistic data, see Zymet (2019).

5.2 Novel compounds

Turning to entirely novel compounds, we fit the analysis sketched in section 4.5 to the data obtained by BKK. Here, the best fit weight for the constraints gives only Id-[nasal]-N2 a nonzero weight, at 2.18; faithfulness to existing compounds receives no weight (as expected, since novel compounds by definition do not have an existing compound to be faithful to), and also the markedness constraint is weighted zero. This is also to be expected, since the one degree of freedom in the model is actually the harmony *difference* between faithful and repairing compounds, and so since faithfulness receives a non-zero positive weight, the Gaussian prior prefers the other constraint to remain at zero.

As above, we compared the fit of the MaxEnt model of novel compounds to one that did not allow access to frequency information, which fit the data much less well relative to a model that does allow the lexicon to scale violations of faithfulness constraints (Δlog-likelihood $= 11.12$, $p < .001$ with one degree of freedom, ΔAICc $= 22.19$). We also compared the absolute fit of our theoretically-motivated MaxEnt model to the purely statistical model fit by BKK[9], and find that the R^2 of our model, 0.11, falls within the 95% Credible Interval of the median of that of the statistical model, 0.06 [0.00, 0.18]. While low in absolute terms, it is reassuring that it is in line with the statistical model, suggesting that our grammatical model is not overfit to the data. We suspect that the cause of the poor model fit may be that there is greater between-individual variation in novel compounds than in existing ones, while the model we fit to the novel data has fewer parameters, and thus is less expressive, than that fit to existing data.

5.3 Fitting a joint model

To get a more holistic picture of nasalization, we fit a joint model to both existing and novel compounds, and integrated the fact that the participants were included in the experiment on the basis of exhibiting complementary distribution of [g] and [ŋ] in monomorphemes. Therefore, in addition to both sets of compound data, the model included the monomorphemes used in the dialect questionnaire to screen participants for inclusion in the experiment, including frequency-based scaling of their faithfulness violations. Since we assume lexicon optimization, we cannot accurately assess the weight of the markedness constraint *Internal-[g] that drives the language-wide phonotactic because the number of monomorphemes that we surveyed was relatively small in comparison to all the data in the Japanese lexicon that exhibits this complementary distribution (Ito and Mester, 2003), and thus contributes to the actual weight of *Internal-[g] in speakers' grammars. However, we can find a lower bound on its weight by constraining the sets of weights

[9]Model specification: `Nasalization ∼ 1 + LogN2Freq*N2Primed + (1 + LogN2Freq*N2Primed | Speaker) + (1 + N2Primed | Compound)`; see BKK section 3.3 and 3.4.2 for details.

we consider to those that maximize the likelihood of the compound data, while simultaneously preserving allophony in monomorphemes (operationalized as having 95% or greater probability of faithful realization). The final model yielded weights listed in Table 9, and predictions plotted in Figure 10.

Constraint	Weight
*Internal-[g]	0.16
Id-[nasal]-Compound	5.12
Id-[nasal]-N2	4.25

Figure 9: Best-fitting weights for the experimental data, existing and novel compounds combined, that preserves the allophony in monomorphemes.

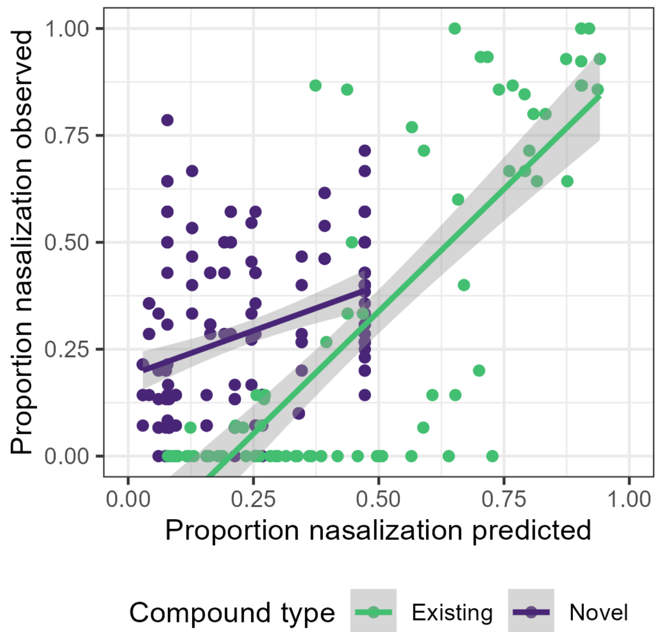

Figure 10: Predicted (vertical axis) vs. observed (horizontal axis) rates of nasalization for categories existing (green) and novel (purple) compounds under the combined model (weights in Table 9).

Although the integration of non-alternating monomorphemic words was intended to give a better picture of the weight of the markedness constraint, their inclusion ended up leaving the weights of the Ident-[nasal] constraints relatively unchanged, with the markedness constraint only getting a modest weight. We take this to indicate that much of the variability seen in the experimental data is actually driven by a strong analogical effect of faithfulness, rather than paradigm-uniformity being parasitic on markedness. This is consistent with the weights obtained when fitting the individual datasets above; when forced to account for the non-alternation of monomorphemes, a modest weight for *Internal-[g] suffices, relative to the stronger weights of faithfulness required to drive the paradigm-uniformity effect.

Finally, we compare the joint model to one where we force either the weights of both faithfulness constraints to be the same (Δlog-likelihood = 28.56, $p < .001$ with one degree of freedom, ΔAICc = 56.54, favoring the model with independently-weighted faithfulness constraints), or one where faithfulness is not scaled by frequency (Δlog-likelihood = 761.88, $p < .001$ with one degree of freedom, ΔAICc = 1,466.66, favoring the model where faithfulness is scaled by frequency). In both cases, the more complex model presented here fits the data significantly better than the alternatives, supporting two basic tenets of the Voting Bases model: multiple faithfulness constraints active in the grammar, each scaled by the resting activation of the lexical items they refer to (with frequency being the current proxy for resting activation). The R^2 values for the combined model components are within 0.01 of their values for the separately-fit models, demonstrating no compromise in absolute model fit.

6 Discussion

This paper has proposed a model of variable voiced velar nasalization in Japanese, drawing on experimental data published in Breiss et al. (*to appear*). The model integrates grammatical and functional determinants of variation, drawing on the Voting Bases framework of lexicon-grammar interaction, which was originally developed to model an entirely separate phonological phenomenon, Lexical Conservatism in English and Spanish (Breiss, 2024). Here, we address several major issues that the model raises, notably about whether the proposed system can be learned from the actual Japanese lexicon (section 6.1), about the competence-performance distinction (section 6.2), and about how the Voting Bases model's mechanism for integrating usage frequency and formal grammar compares to other propositions in the literature (section 6.3). Finally, we close the paper with a more general discussion about how we might understand the broader empirical landscape of frequency effects in phonological patterning in light of the proposal in this paper.

6.1 Whence the weights? Evidence in the lexicon

Having observed that there is robust frequency-conditioning of nasalization in both existing and novel compounds, we can ask what the source of this frequency-conditioning might be. By hypothesis, the relationship between frequency and resting activation is one that is automatic and not overtly learned. However, we find that the model performs significantly better when allowed to set the weights of faithfulness constraints referencing different allomorphs to different weights. This result suggests that, setting aside the relationship between frequency and activation, the speakers must be able to attribute different amounts of influence to different faithfulness constraint violations depending on which base the violation is assessed against. Put another way, the learner needs to be able to figure out how analogically-driven her lexicon is. Here, we present a preliminary investigation of what kind of evidence might exist in the Japanese lexicon that could allow speakers to assign different weights to Id-[nasal]-Compound and Id-[nasal]-N2.

We fit a grammar with the constraints in section 4.3 and frequency-driven scaling of faithfulness violations to the set of compounds in the corpus analyzed by Breiss et al. (2021b) that had a free N2. We found that the optimal weights of the grammar were zero for both *Internal-[g] and Id-[nasal]-Compound, and 1.08 for Id-[nasal]-N2. We had anticipated there being little to no weight assigned to the markedness constraint in this dataset for the same reasons discussed above in section 5.3, but we also found that instead of a tension between faithfulness to the compound itself and faithfulness to the N2, the grammar instead left it to the paradigm uniformity effect alone to perturb the otherwise at-chance distribution of variation (at chance because the weight of Id-[nasal]-Compound was at zero, indicating, all else equal, that the alternating and non-alternating candidates were equiprobable). This is qualitatively the same finding as for the novel compounds.

We compared the model fit to the corpus data to one where the grammar was forced to assign the same weight to Id-[nasal]-Compound and Id-[nasal]-N2, and found that it was significantly out-performed by the model that allowed the grammar to allot differing weights to different faithfulness constraints to different bases (Δlog-likelihood $= 45.3$, $p < .001$ with one degree of freedom). We take this as tentative evidence that there is an empirical basis in the lexicon for assigning different degrees of faithfulness to different bases.

6.2 Competence, performance, and formal modeling

This paper has proposed a model of Japanese nasalization that integrates token frequency into the workings of the phonological grammar. Since the prospect of integrating a putatively performance-related factor like token frequency into a formal phonological model is not an uncontroversial one, below we directly address some possible criticisms of this approach. We certainly do not think that these are the last words on the topic, but we do feel that by explicitly discussing what we are doing and our motivations for doing it, we take a first step towards a clearer understanding of the stakes and consequences of the choices made in modeling information about usage jointly with the phonological grammar.

One initial objection to formally modeling the frequency-conditioned variation in nasalization might be that there is nothing competence-related to model here at all—the variation is solely driven by "performance" factors (Chomsky, 1965). We respond that this cannot be true of Japanese nasalization: the fact that only compounds whose N2 is morphologically free exhibit frequency-sensitive variation, despite the existence of bound morphemes with [g]- and [ŋ]-initial forms like [ga]/[ŋa] "fang", as shown by the examples in (5), requires an explanation that makes reference to grammatical structures. Further afield, cases like Lexical Conservatism much more strongly blur the line between the contents of the lexicon and the phonological grammar and are well-modeled by a framework like Voting Bases. The fact that this paper demonstrates both paradigm uniformity and Lexical Conservatism emerge as special cases of the same theory speaks to the theoretical insight that can be gained by jointly modeling "performance-related" and "competence-related" influences on the phonological grammar.

Another objection to the grammatical model that may be put forward in the current paper is that by incorporating both resting activation (a psycholinguistic quantity) and phonological markedness (a grammatical quantity) into the same formal model, we compromise the distinction between competence and performance to the extent that it is not clear what we can actually say the model is a model *of*. If true, this would indeed be a flaw of the approach; however, a virtue of the Voting Bases model is that lexical influence on the grammar is clearly delimited: the model only allows the lexicon to scale the weights of faithfulness constraints to corresponding lexical representations. Manipulating the resting activation of a given UR has identifiable, localized influences on the computations of the phonological grammar, and instantiates a linking hypothesis consistent with a consensus view of the basic structure of the lexicon. This mechanism can be seen as one way of implementing the idea of "grammar dominance" put forth, for example, by Coetzee (2016) and Coetzee and Kawahara (2013). The "core" phonological grammar—weighted constraints which can assess violations of novel candidates—can be recovered by simply ignoring the influence of the lexicon on constraint violations, and can be studied in novel contexts like *wug*-tests, where there is no relevant lexical representation to bear on the grammar.

A final objection that we consider is that the very act of jointly modeling usage frequency and the phonological grammar risks leading the analyst to think of fundamentally performance-related factors as in fact competence-related, thus undercutting the goal of researchers whose focus is only understanding linguistic competence. We contend that this is simply false, and in fact, the reverse is true: for a researcher who *only* cares about linguistic competence, modeling usage factors jointly with theories of competence is vital. When confronting data derived from language use (that is, modeling corpus data as in Breiss et al. (2021a), or experimental data where stimuli are existing morphemes of the language as in Breiss et al. (*to appear*), a joint model will better expose the true influence of competence-related factors on the data under study, with the performance-related parts of the model accounting for the otherwise-distorting influence of these factors. Simply ignoring performance-related factors in a formal model makes the strong claim that they have no effect, an assumption which is untenable in the cases examined here, and, we suggest, is also false in many (if not all) types of linguistic data that speakers might have prior usage-based experience with (Arnon and Snider, 2010; Smith and Moore-Cantwell, 2017; Zymet, 2018; Morgan and Levy, 2016, 2023). Rather, an integrated approach that jointly models grammar and usage is essential to disentangle competence from performance factors, if this is the goal of the analysis.

6.3 Comparison with other models

The Voting Bases model is one of several approaches in the literature that propose to model the interaction of usage frequency and phonological grammar. In particular, it is similar to the methods proposed in Coetzee and Pater (2008) and Coetzee and Kawahara (2013) which directly scale the weight of faithfulness constraints by the frequency of the form they make reference to, and that of Baird (2021) where a simulated perception-production loop comes to the same result via online learning. This family of approaches involves lowering the weight of faithfulness constraints to high-frequency forms relative to lower-frequency forms which enables them to model data like coronal stop deletion in English (Coetzee and Kawahara, 2013), where higher-frequency monomorphemes (like *just*) tend to get produced more often with a deleted coronal stop than phonologically-similar words (like *jest*). Common to these models is that they assume that the underlying form is /t/-ful, and thus the task of their model must relate higher frequency to therefore have lower constraint weights for it.

A weakness of these models is that, with the possible exception of Baird (2021), the directionality between frequency and constraint weight is arbitrary—the primary goal set in these studies was to fit the data, which is better than the alternative which does not model the effects of lexical frequencies at all, but they suffered somewhat for the lack of clear functional grounding the relation.

By contrast, the frequency-faithfulness relation that Voting Bases model adopts runs in the opposite direction—more frequent forms exact a greater penalty for unfaithful realizations relative to less frequent forms; constraint violations are less severe for low-frequency vs. high-frequency forms. This allows the model to fit a similar range of data, but with a linking hypothesis that is explicitly rooted in resting activation, a construct that is externally justified by a large body of work in psycholinguistics, as reviewed in Breiss (2021, 2024). Lexical items with higher resting activation are more insistent on faithfulness to themselves, corresponding to their increased salience in the language processing system. The main contribution of the Voting Bases model in modeling this phenomenon is that the influence of the lexicon on the grammar should be, in principle, derivable without reference to any facts about the experiment in question; given some independently-established computationally-implemented model of lexical dynamics that represents a scalar quantity of resting activation (or similar construct), the strong prediction of the Voting Bases model is that that quantity should be able to be a fully adequate scaling factor for faithfulness constraint violations. The specific mechanism that is used in this paper—scaling the weights by the sigmoidal transformation of the resting activation—is used since it represents, to us, a reasonable first stab, but the linking function may need to be revised in light of future findings.

In summary, we suggest that the Voting Bases model, because of its functional grounding of frequency effects in externally-motivated psycholinguistic phenomena, is on firmer footing than theories that have alternative linking functions between frequency and grammar, which are arguably arbitrary.

6.4 Towards a unified picture of token frequency in phonology

In this section, we broaden our view of token frequency effects in phonology, and discuss how considering the varying functional roles of frequency can reconcile some seemingly-contradictory bodies of evidence.

First, there is evidence that higher token-frequency leads to more markedness-reducing alternations. Coetzee and Kawahara (2013) found that higher-frequency lexical items were more likely to undergo phonological processes of simplification and (markedness-)reduction: high-frequency English words like *jus(t)* underwent an optional process of coronal stop deletion at a higher rate than low-frequency words like *jes(t)*, and high-frequency Japanese words like [baggu] "bag" underwent geminate devoicing more often than low-frequency words like [budda] "Buddha" (Kawahara and Sano, 2013). Zuraw (2007) examines frequency-conditioned application of markedness-reducing phonological processes in a corpus of written Tagalog, and likewise finds higher rates of repair within higher-frequency units (words, clitic groups, etc), subject to the markedness principles of the language.

On the other hand, there is also evidence to show that higher-frequency forms are more likely to be exceptional, and thus marked with regard to the overall properties of the grammar. Smith and Moore-Cantwell (2017) found that higher-frequency comparative constructions are more likely to flout grammar-wide trends driven by markedness. In a similar vein, Anttila (2006) and Mayer (2021) found that higher-frequency morphologically-complex forms were more likely to behave opaquely with respect to grammar-wide phonological processes.

We can compare these effects to the ones observed in Breiss et al. (2021b) and Breiss et al. (*to appear*): higher-frequency N2s act as stronger attractors, yielding *more* faithfulness to their preserved surface [g] resulting in lower rates of nasalization, whereas higher compound frequency as a whole yielded higher rates of nasalization. Thus it seems that for compounds, higher frequency is correlated with more phonological-process application and markedness-reduction; this is broadly in line with the findings of Coetzee and Kawahara (2013) where higher-frequency words undergo more phonological alternations. However, we found that at the same time, in compounds with free N2s, higher free N2 frequency is related to *less* process application, with higher-frequency supporting the retention of a marked structure (word-medial [g]).

We suggest we can resolve this tension by distinguishing between the processes that token frequency can impact: one is whether to set up an independent lexical representation for a surface allomorph, and the other is influencing the strength of that representation in the lexicon of the speaker.

If a form is exceptional and high-frequency, it may be more economical for a speaker to pay a one-time "cost" of encoding the exception as a listed form that is not derived by the grammar, thus relieving the phonology of the difficulty of having to generate the exceptional or idiosyncratic form on each of the many frequent occasions of use (cf. Adaptor Grammars (Johnson et al., 2007, *et seq.*) or Fragment Grammars (O'Donnell, 2015) which offer computationally-explicit implementations of this general idea). For lower-frequency exceptional forms, the likelihood of listing is less since the price trades off less favorably with the amount of times it is used; thus lower-frequency forms are more susceptible to change and regularization to the dominant grammatical trends over time compared to higher-frequency forms. Another aspect of this trade-off is the emergence of Lexicon Optimization (Prince and Smolensky, 1993; Sanders, 2003, 2006); even if a form is not particularly exceptionful, if a UR almost always surfaces with a phonological process applied to it, with sufficient frequency it becomes less costly to just store the form with phonological process applied—that is, create a separate allomorph that is specific to the environment that would trigger the phonological rule. This, similarly, relieves the grammar of the job of having to repair the form every time. Thus, we find Lexicon Optimization targeting forms like *jus(t)* over forms like *jest*, making these forms restructured to automatically have the phonological alternation applied, thus giving the appearance of having undergone a markedness-improving repair in the grammar, but actually the frequency of the form has resulted in restructuring to the lexicon (see Breiss and Wilson (2020) for an initial attempt at a computational model of the phonological grammar and lexicon that exhibits this property).

As reviewed above, lexical frequency also influences the resting activation of a lexical item once it is listed in the lexicon. In the Voting Bases model, higher resting activation leads to the listed form exerting a stronger pull on the surface realization of a related form; where this pressure goes against the broader principle of markedness in the grammar, as in cases of paradigm uniformity, we find that marked structures with high-frequency output-bases are preserved; in cases where the listed form coincides with the output of the markedness-reducing process, as in many cases of Lexical Conservatism (Steriade, 1997; Steriade and Stanton, 2020; Breiss, 2021), then the higher-frequency form promotes an unmarked surface form.

Recent work by Jarosz et al. (2024) has laid out a class of models which exhibit characteristics that align favorably with the dynamics of frequency laid out here, suggesting that an integrated, implemented model that can jointly account for the variety of frequency effects reviewed in this section is perhaps quite close at hand. Future work may profitably explore how well these models can provide converging evidence from computational learning simulations to support the psycholinguistic, experimental, and diachronic evidence for the contents of the lexicon that the Voting Bases theory relies on. In sum, the broader landscape of token frequency in phonology is compatible with the functional grounding given to frequency under the Voting Base model, though much empirical and formal work remains to be done to further support the predictions of the framework more broadly as a candidate for a general theory of the influence of the dynamic lexicon on the probabilistic grammar.

References

Albright, A. (2008). Cumulative violations and complexity thresholds. Unpublished ms, MIT.

Anttila, A. (2006). Variation and opacity. *Natural Language & Linguistic Theory*, 24(4):893–944.

Arnon, I. and Snider, N. (2010). More than words: Frequency effects for multi-word phrases. *Journal of memory and language*, 62(1):67–82.

Baird, M. (2021). Deriving frequency effects from biases in learning. *Proceedings of the Linguistic Society of America*, 6(1):514–525.

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

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

Breiss, C. (2020). Constraint cumulativity in phonotactics: evidence from artificial grammar learning studies. *Phonology*, 37(4):551–576.

Breiss, C. (2024). When bases compete: A voting model of lexical conservatism. *Language*, 100(2):308–358.

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

Breiss, C., Katsuda, H., and Kawahara, S. (2021b). A quantitative study of voiced velar nasalization in Japanese. In *University of Pennsylvania Working Papers in Linguistics*, volume 27.

Breiss, C., Katsuda, H., and Kawahara, S. (accepted). Token frequency modulates optional paradigm uniformity in japanese voiced velar nasalization. *Phonology*.

Breiss, C. and Wilson, C. (2020). Extending adaptor grammars to learn phonological alternations. *Proceedings of the Society for Computation in Linguistics*, 3(1):480–483.

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

Chomsky, N. (1965). *Aspects of the Theory of Syntax.* The MIT Press.

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.

Coetzee, A. W. and Pater, J. (2008). Weighted constraints and gradient restrictions on place co-occurrence in muna and arabic. *Natural Language & Linguistic Theory*, 26(2):289–337.

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.

Hay, J. (2001). Lexical frequency in morphology: Is everything relative? *Linguistics*, 39(6):1041–1070.

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

Hibiya, J. (1995). The velar nasal in Tokyo Japanese: A case of diffusion from above. *Language Variation and Change*, 7:139–152.

Ito, J. and Mester, A. (1996). *Correspondence and compositionality: The ga-gyō variation in Japanese phonology*. Clarendon Press Oxford.

Ito, J. and Mester, A. (2003). *Japanese morphophonemics: Markedness and word structure*, volume 41. MIT Press.

Jäger, G. and Rosenbach, A. (2006). The winner takes it all—almost: Cumulativity in grammatical variation. *Linguistics*, 44(5):937–971.

Jarosz, G., Hughes, C., Lamont, A., Prickett, B., Baird, M., Kim, S., and Nelson, M. (2024). Type and token frequency jointly drive learning of morphology. *ms*.

Jiang, H. (2023). Learnability of length-referencing alternations. MA Thesis, University College London.

Johnson, M., Griffiths, T. L., Goldwater, S., et al. (2007). Adaptor grammars: A framework for specifying compositional nonparametric bayesian models. *Advances in neural information processing systems*, 19:641.

Jurafsky, D. and Martin, J. H. (2009). *Speech and language processing: An introduction to natural language processing, computational linguistics, and speech recognition, 2nd Edition*. Prentice Hall.

Kawahara, S. and Sano, S.-i. (2013). A corpus-based study of geminate devoicing in japanese: Linguistic factors. *Language Sciences*, 40:300–307.

Kindaichi, H. (1942). Ga-gyo bionron. *Onin no Kenkyû (Reprinted in Kindaichi 1967)*.

Labrune, L. (2012). *The phonology of Japanese*. Oxford University Press.

Lüdecke, D., Ben-Shachar, M. S., Patil, I., Waggoner, P., and Makowski, D. (2021). performance: An R package for assessment, comparison and testing of statistical models. *Journal of Open Source Software*, 6(60):3139.

Martin, A. T. (2007). *The evolving lexicon*. PhD thesis, UCLA.

Mayer, C. (2021). *Issues in Uyghur backness harmony: Corpus, experimental, and computational studies*. University of California, Los Angeles.

Morgan, E. and Levy, R. (2016). Abstract knowledge versus direct experience in processing of binomial expressions. *Cognition*, 157:384–402.

Morgan, E. and Levy, R. P. (2023). Productive knowledge and item-specific knowledge trade off as a function of frequency in multiword expression processing.

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

NHK (1993). Nhk broadcasting culture research institute.

O'Donnell, T. J. (2015). *Productivity and reuse in language: A theory of linguistic computation and storage*. MIT Press.

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

Prince, A. and Smolensky, P. (1993). Optimality Theory: Constraint interaction in generative grammar. *Optimality Theory in phonology*.

Purse, R., Fruehwald, J., and Tamminga, M. (2022). Frequency and morphological complexity in variation. *Glossa: a journal of general linguistics*, 7(1).

Rebrus, P. and Törkenczy, M. (2017). Gradient harmonicity in compounds. In *Proceedings of the*

Annual Meetings on Phonology, volume 4.
Sanders, N. (2006). Strong lexicon optimization. *University of Massachusetts Amherst Phonology Group. http://sanders. phonologist. org/Papers/sanders-umass. pdf.*
Sanders, R. N. (2003). *Opacity and sound change in the Polish lexicon.* University of California, Santa Cruz.
Shih, S. S. (2017). Constraint conjunction in weighted probabilistic grammar. *Phonology*, 34(2):243–268.
Smith, B. W. and Moore-Cantwell, C. (2017). Emergent idiosyncrasy in English comparatives. In *NELS 47: Proceedings of the 47th meeting of the North East Linguistic Society. Amherst: Graduate Linguistic Student Association.*
Smolensky, P. (1986). Information processing in dynamical systems: Foundations of harmony theory. Technical report, Colorado Univ at Boulder Dept of Computer Science.
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. and Stanton, J. (2020). Productive pseudo-cyclicity and its significance. Talk at LabPhon 17.
Trubetskoy, N. (1969). Principles of phonology [translated from the Russian by C. Baltaxe]. *Berkeley: U. of California Press.*
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.
Wilson, C. and Obdeyn, M. (2009). Simplifying subsidiary theory: statistical evidence from Arabic, Muna, Shona, and Wargamay. Unpublished manuscript, Johns Hopkins University.
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.
Zymet, J. (2018). *Lexical propensities in phonology: corpus and experimental evidence, grammar, and learning*. PhD thesis, UCLA.
Zymet, J. (2019). Learning a frequency-matching grammar together with lexical idiosyncrasy: Maxent versus hierarchical regression. In *Proceedings of the Annual Meetings on Phonology*, volume 7.