

Covert URs: evidence from nasalization in Western Panjabi

Jonathan Charles Paramore and Ryan Bennett

University of California, Santa Cruz

jcparamo@ucsc.edu, rbennett@ucsc.edu

Abstract

Western Panjabi is an Indo-Aryan language spoken primarily in the Punjab province of Pakistan. Long vowels contrast in nasality in the language, but this contrast is neutralized before nasal consonants: $[\tilde{V}\tilde{V}N]$, $*[VVN]$. This paper presents air pressure data showing that predictably nasal vowels before nasal consonants in Panjabi, $[\tilde{V}\tilde{V}N]$, are phonetically identical to contrastive nasal vowels, $/\tilde{V}\tilde{V}/$ (Exp. 1). However, predictable $[\tilde{V}\tilde{V}N]$ and contrastive $/\tilde{V}\tilde{V}/$ vowels behave differently in their interaction with a nasal harmony process, in that only contrastive $/\tilde{V}\tilde{V}/$ vowels trigger nasal harmony (Exp. 2). This difference in phonological behavior can be straightforwardly analyzed by assuming (i) that pre-nasal vowels are underlyingly oral, even in morphemes which never show alternations in nasality, and (ii) that only underlyingly nasal vowels trigger harmony. These findings support a view of phonology in which abstract underlying representations (URs) can be completely covert, never surfacing unchanged in any allomorph. Surface-oriented approaches to phonology which reject covert URs — including usage-based models relying exclusively on observed exemplars to form lexical representations — have difficulty accounting for these facts.

1 Introduction

Since Halle (1959) and Chomsky & Halle (1968), research in generative phonology has traditionally assumed at least two distinct levels of representation for words and other lexical items. *Underlying representations* (URs) are the abstract forms that underlie all potential surface realizations of a morpheme or word. URs contain, at a minimum, any unpredictable aspects of the pronunciation of each morpheme or word; in particular, the basic phonemes which compose each morpheme. This information must simply be memorized. Chomsky & Halle (1968: Ch. 8) propose that URs contain *only* this information: they are the stored form of a lexical item after all predictable phonological and phonetic detail is stripped away (cf. Prince & Smolensky 1993/2004: Ch. 9). The UR for an English word like ‘price’ $[p^hɹaɪs]$ would thus be $/pɹaɪs/$, a sparse, redundancy-free, minimal specification of only those unpredictable aspects of the phonology and phonetics of ‘price’ which must be memorized.

Surface representations (SRs), on the other hand, are the forms of morphemes and words that emerge after URs are filtered through a grammar that applies predictable, language-specific, and context-specific phonological processes. The grammar, which is composed of rules or other formal devices, is thus responsible for transforming minimally-specified lexical representations

like /praɪs/ ‘price’ to SRs like [p^hjaɪs], which are full of predictable, redundant, contextually-determined phonetic detail.

In contrast to the sparse, redundancy-free, and phonetically abstract underlying representations encountered in generative phonology, a broad body of experimental literature has argued that lexical representations do, in fact, include phonetically-detailed, predictable, and context-dependent information about the pronunciation of words and morphemes (e.g. Johnson 1997, Bybee 2001, Pierrehumbert 2001, Johnson 2006, Pierrehumbert 2016, Goldrick & Cole 2023, and references therein). For example, a classic study by Goldinger (1998) explored the roles of talker familiarity and word frequency in a shadowing task and found that participants were faster to repeat words out loud if they had already heard that word previously uttered by the same speaker. He took this result as evidence that words are stored in memory as sets of *exemplars*, corresponding to phonetically-detailed memories of past experiences hearing each word, including detail about speaker identity.

Goldinger also found that speakers’ own productions sometimes imitated fine-grained phonetic properties of the prompts they heard. Imitation effects were stronger with low-frequency words than high-frequency words. Goldinger proposed that this effect can be explained if speech production targets are determined by some form of averaging over the stored exemplars for each word. Such averaging predicts that imitation effects should be more evident for low-frequency words: since low-frequency words have fewer stored exemplars in memory, any single exemplar should have a relatively outsized influence on how that word is pronounced. High-frequency words have many stored exemplars and so should have relatively stable production targets and be less liable to influence from recently heard productions than low-frequency words.

Exemplar-based models of lexical representation have also been used to model patterns of incomplete neutralization (see e.g. Kaplan 2017, Morrison 2021, Du & Durvasula 2025 for discussion and references). As a concrete example, we discuss the well-known case of word-final devoicing in German. In German, as in English and a range of other languages, vowels tend to be longer before voiced obstruents (e.g. Braunschweiler 1997). Consequently, the first vowel in [ʁɛ:d-ɐ] ‘wheels’ is slightly longer than the first vowel in [ʁɛ:t-ɐ] ‘councils’.

Voiced obstruents are devoiced word-finally in German. This produces phonological neutralization: both /ʁa:d/ ‘wheel’ and /ʁa:t/ ‘council’ are pronounced [ʁa:t]. However, Port & O’Dell (1985) found that this neutralization is phonetically incomplete: there is a small, but reliable acoustic difference in the length of preceding vowels, depending on the underlying voicing of the obstruent. Specifically, the pre-stop vowel in [ʁa:t] ‘wheel’ (from /ʁa:d/) tends to be phonetically longer than in [ʁa:t] ‘council’ (from /ʁa:t/). Phonetically, [ʁa:t] ‘wheel’ and [ʁa:t] ‘council’ are not homophones.

As discussed above, generative theories of phonology have traditionally assumed abstract, symbolic representations which are essentially devoid of phonetic detail.¹ This makes it difficult for such theories to account for incomplete neutralization, at least without ancillary assumptions or mechanisms (see Du & Durvasula 2025 for extensive discussion).

In contrast, usage-based, exemplar models of lexical representation can generate incomplete neutralization by leveraging connections between stored exemplars of individual morphemes, pooled across different contexts. For example, stored exemplars of the morpheme meaning ‘coun-

¹There are of course exceptions: a variety of research in generative phonology has assumed representations which include fine details of speech acoustics (e.g. Flemming 2001), speech production (e.g. Gafos 2002), and so on.

cil’ in German should always end in [t], as in plural [ʁɛ:t-ɐ] ‘councils’ and singular [ʁa:t] ‘council’. On the other hand, stored exemplars of the morpheme meaning ‘wheel’ could include *both* final [d] (as in [ʁɛ:d-ɐ] ‘wheels’) and final [t] (as in [ʁa:t] ‘wheel’). Stored tokens with final [d] will, as expected, have slightly longer vowels. Recall that production targets are determined by averaging over exemplars. Vowel length in stored exemplars like [ʁɛ:d] (from plural [ʁɛ:d-ɐ]) can thus induce vowel lengthening in singular [ʁa:t], as an indirect consequence of this averaging process. Incomplete neutralization then emerges from the basic dynamics of exemplar storage and interaction. Importantly, only phonetically-detailed surface forms are needed to produce this result: abstract, underlying forms play no role in this model of incomplete neutralization.

The preceding discussion provides only a small sample of the evidence which has been taken to support the existence of phonetically-detailed exemplar representations for words and morphemes. For a fuller picture, see e.g. Pierrehumbert (2016), Goldrick & Cole (2023) and research cited there. Models of lexical representation which assume phonetically-detailed, stored exemplars are sometimes called *usage-based* models, because lexical representations correspond in a fairly direct way to actual, physical experiences with words and other linguistics structures.

Exemplar models are not incompatible with abstract representations (e.g. Goldinger 1998, Goldrick & Cole 2023). Indeed, it has been argued that abstract phonemes and abstract lexical representations are needed to account for patterns of sound change and phonological productivity (e.g. Pierrehumbert 2016), for modeling spoken word recognition and phonological priming effects (e.g. Lahiri & Marslen-Wilson 1991, Gaskell & Marslen-Wilson 1996, LoCasto & Connine 2002, Jongman 2004, Lahiri & Reetz 2010, Cornell et al. 2011, Lahiri 2012, Ernestus 2014, Farris-Trimble & Tessier 2019, White et al. 2024), and for understanding phonological aphasias (e.g. Goldrick & Rapp 2007, Buchwald & Miozzo 2011), among various other experimental results (see also Krämer 2012). And of course, abstract lexical representations are justified to the extent that they are needed to model phonological alternations, phonotactics, and patterns of phonological generalization within individual languages (e.g. Kenstowicz & Kisseberth 1979, McCarthy 2007, Hyman 2018). These findings are largely compatible with the storage of phonetically-detailed episodic exemplars, provided that some type of abstract lexical representation is assumed to co-exist alongside them.

We thus assume that both abstract and exemplar representations are needed for an adequate understanding of phonological patterns and phonological behavior. A number of questions then arise. How abstract may lexical representations be? How far may they diverge from actual, observed surface forms? And how do language learners generate abstract, stored representations on the basis of the surface forms that they actually hear and produce?

In traditional generative frameworks, it is perfectly acceptable — and in some sense the norm — for stored lexical forms (URs) to differ extensively from surface forms (SRs) (e.g. Kenstowicz & Kisseberth 1979). More specifically, classical generative phonology holds the rather strong position that an underlying representation can be *completely covert*, with properties that never appear in *any* surface form or allomorph of that UR. For example, O’Hara (2017) discusses a stem-final [i]~[ø] alternation in Klamath, illustrated by forms like [ʔe:ɰ-a] ‘is deep’ ~ [ʔe:ɰi-ʔkʰ] ‘deep’. O’Hara argues that these alternations cannot be analyzed as the deletion or insertion of [i]. Instead, he proposes that stems like [ʔe:ɰ(i)-] contain an underlying final /e/, e.g. /ʔe:ɰe/. Stem-final /e/ deletes outside of initial syllables, but instead raises to [i] when deletion would produce an illicit consonant cluster. As a consequence, the underlying forms of stems like /ʔe:ɰe/ are *covert*: they contain a segment, /e/, which never appears in any surface allomorph (echoing a

long line of analyses of ‘yer’ vowels in Slavic languages, e.g. Gouskova 2012).² Analyses making use of covert URs involve a high level of abstraction: speakers must somehow arrive at a stored lexical entry for a stem like /ʔe:we/ which is qualitatively distinct from any surface form of that morpheme they have ever actually heard.

As various authors have argued, some form of ‘abstraction’ can result from statistical averaging over observed surface exemplars, as outlined above (e.g. Goldinger 1998, Goldrick & Cole 2023). However, it seems unlikely that any kind of computation over observed surface forms is sufficient, alone, to produce covert lexical representations like Klamath /ʔe:we/ → [ʔe:ɰ(i)-]. To the extent that covert underlying representations like /ʔe:we/ can be justified, they motivate richer, more articulated mechanisms for arriving at URs than can be produced by pooling over surface forms.

This paper provides phonetic and phonological evidence from Pakistani Panjabi which highlights the benefits of abstraction in lexical representation. We measure oro-nasal air pressure across three different types of vowels in Panjabi: oral vowels /VV/, contrastive nasal vowels /ṼṼ/, and vowels occurring before a tautomorphemic nasal consonant /VVN/. The results indicate that Panjabi speakers consistently produce /VVN/ vowels with nasalization identical to /ṼṼ/ vowels. However, nasal vowels in /VVN/ interact opaquely with a pattern of nasal harmony triggered by contrastive /ṼṼ/. For this reason, we argue that /VVN/ vowels should be analyzed as underlyingly oral, even though the vowels in question are non-alternating, and thus *always* phonetically nasal on the surface. Phonological opacity is crucial to our argumentation here: as already observed by Ettlinger (2008) and others, productive, general patterns of opacity are difficult to model without the assumption of URs that are categorically distinct from any single SR, or averaged set of SRs (see McCarthy 2007 for related discussion).

Criticisms of abstract underlying representations have also come from within the tradition of generative phonology. It has been proposed that the core functions of underlying representations – in particular, their role in determining unpredictable patterns of alternation across paradigms (Kenstowicz & Kisseberth 1979, McCarthy 2007) – can be subsumed by surface-oriented generalizations, including surface-to-surface relationships between allomorphs of the same morpheme (e.g. Burzio 1996, Scobbie et al. 1996, Hayes 1999, Albright 2002, Cole & Hualde 2011, Allen & Becker 2015, and various others). The pattern of opacity we discuss here cannot be straightforwardly modeled through the use of surface correspondence, because the key forms (derived /VVN/ → [ṼṼN]) are *morphologically simple*, and are also *non-alternating*: they have only a single surface form, and thus no possibility for producing opacity through relations between distinct surface forms (e.g. Benua 2000, McCarthy 2005, Mascaró 2016). To the extent that surface-oriented approaches *can* account for this data, they do so by implementing a formal distinction between predictable vs. unpredictable phonetic detail, thus recapitulating a core function of abstract URs (e.g. Cole & Kisseberth 1994, 1995). More to the point, we also argue that such approaches ultimately require the same covert underlying representations, /VVN/ → [ṼṼN], that we assume here.

²The URs we describe as ‘covert’ have typically been called ‘abstract’ (e.g. Kenstowicz & Kisseberth 1979, Albright 2002, Krämer 2012, O’Hara 2017, Hyman 2018). We opt for the term ‘covert’ here to specify more precisely exactly what type of abstract UR we are concerned with: namely, URs which contain segments that never surface faithfully, in any context.

2 Panjabi

Panjabi (alternatively spelled Punjabi) is an Indo-Aryan language spoken by more than 100 million people worldwide (Hussain et al. 2019). Most Panjabi speakers live in Pakistan (78 million), where Panjabi is by far the most widely spoken first language. Additionally, the Indian states of Punjab, Haryana, and Himachal Pradesh are home to about a quarter of native Panjabi speakers (33 million) (Bashir & Connors 2019). Panjabi is considered to possess a high degree of dialectal variety, with the varieties spoken in Pakistan often grouped together and termed Western Panjabi, and the varieties spoken in India termed Eastern Panjabi. Linguistic studies exploring Panjabi have largely focused on Eastern dialects, but a few descriptions of the Pakistani varieties have been published in recent years. The data used in this study focuses on Panjabi as spoken in Pakistan. As of the summer of 2023, all 36 speakers who participated in the two experiments described here live in one of the twin cities of Pakistan, Rawalpindi and Islamabad. A significant number (56%) were born within 25 miles of the twin cities, with the rest scattered throughout the Pakistani Punjab province, in and around the areas of Gujrat and Sialkot (19%), Faisalabad (14%), Lahore (6%), and Multan (5%). Though Western Panjabi dialects vary widely from town to town, apart from the two speakers from Lahore, the nasalization results from the present study displayed a remarkable uniformity across speakers.

The Panjabi vowel inventory is provided in Figure 1. As is evident, all long vowels contrast in nasality, but Bashir & Connors (2019: p.45) note that this contrast is confined to the morpheme-final syllable, as shown in (1).

(1) Contrastive vowel nasality in Panjabi

- a. [taa] ‘warmth, fever’ vs. [tãã] ‘that; so that’
b. [paavee] ‘cot leg’ vs. [tʃaavēē] ‘pumicestone’

Crucially, a coda nasal consonant neutralizes the nasality contrast such that no minimal pairs of the shape /CVVN/-/CVVN/ exist. This neutralization is likely due to the fact that all /V(V)N/ vowels are said to be characterized by a high degree of nasality (hence [tãã] ‘melody’ is a word, but not hypothetical *[taan]). Indeed, Zahid & Hussain (2012) conducted an acoustic study of vowel nasality in Panjabi. While their study consisted of only three speakers, each uttering 12 repetitions of various short and long vowels from both nasal vowel types, their results did not find a difference in nasality between contrastive /ṼṼ/ vowels and /VVN/ vowels in terms of their A1-P0 and A1-P1 values, measurements reported to correlate highly with nasality (Chen 1997, Styler 2017).

		Front			Central			Back		
		short	long	nasal	short	long	nasal	short	long	nasal
High	tense		ii	ĩĩ					uu	ũũ
	lax	ɪ						ʊ		
Mid	tense		ee	ẽẽ					oo	õõ
	lax		ɛɛ	ẽẽ	ə				ɔɔ	õõ
Low	tense									
	lax								aa	ãã

Figure 1: The Panjabi vowel inventory (Shackle 2003)

In addition to the nasality contrast on long vowels and the neutralization of this contrast before nasal consonants, Panjabi also employs a process of nasal harmony, in which nasality spreads leftward from a nasal vowel through all glides (/j/ and /v/) and vowels until a non-glide consonant blocks the spread. For example, words with the shape found in (2a-b), in which the word-final vowel is underlyingly / $\tilde{V}\tilde{V}$ /, are realized with nasality across the [VVG $\tilde{V}\tilde{V}$] sequence (Bhatia 1993, Bashir & Connors 2019). However, these descriptions of the nasal harmony process do not clarify whether the /VVN/ vowel in (2c-d) also instigates the spread of nasality, which serves as the central question of the second experiment discussed in the present work (section 4).

(2) Nasal harmony in Panjabi

- | | | | | |
|----|--------------|---|----------------------|---------------|
| a. | /saa-vãã/ | → | [sããvãã] | ‘breath-PL’ |
| b. | /matfii-jãã/ | → | [matfĩĩjãã] | ‘fish-PL’ |
| c. | /aav-aam/ | → | ?[ããvããm]/[aavããm] | ‘PL-citizen’ |
| d. | /siijaan/ | → | ?[sĩĩjããn]/[siijããn] | ‘recognition’ |

There are two possibilities with regard to the participation of /VVN/ vowels in nasal harmony. First, and perhaps most straightforwardly, if /VVN/ vowels are indeed phonetically identical to / $\tilde{V}\tilde{V}$ / vowels in Panjabi, it would be logical to assume that both vowel types instigate the nasal harmony process in a similar fashion. After all, no nasality difference exists on the surface that distinguishes the two vowels.

On the other hand, it is possible that, despite being phonetically identical to contrastive / $\tilde{V}\tilde{V}$ / vowels, non-alternating /VVN/ vowels do not instigate nasal harmony in the same way because they are phonologically distinct: specifically, because they are underlyingly oral, with predictable nasality provided on the surface by the phonological grammar.

This conclusion – which we argue in favor of – would support the existence of *covert URs*. Despite speakers consistently producing non-alternating /VVN/ vowels as categorically nasal [V $\tilde{V}\tilde{N}$], these vowels can be analyzed as oral in their URs, given that nasality is predictable in this context. This distinction between underlyingly nasal / $\tilde{V}\tilde{V}$ / and underlyingly oral /VVN/ → [V $\tilde{V}\tilde{N}$] can be leveraged to explain the uneven application of harmony: only underlying, contrastive nasal / $\tilde{V}\tilde{V}$ / is a harmony trigger. Without positing a difference in the underlying representations of contrastive / $\tilde{V}\tilde{V}$ / vowels and non-alternating /VVN/ nasal vowels in (2), it is difficult to explain why the two vowel types behave differently in their interaction with leftward nasal spread.

Crucially, /VVN/ vowels in words like those in (2c-d) do not exhibit morphological alternations in which the vowel occurs without the following nasal consonant. This means there are no surface forms (apart from those involving harmony) which directly and transparently demonstrate that these /VVN/ vowels are underlyingly distinct from contrastive / $\tilde{V}\tilde{V}$ / vowels (e.g. */aavaa=m/ - */aavaa=C/). As a result, the only direct, observed information speakers have regarding the nasality of pre-nasal /VVN/ vowels in words like ‘public’ and ‘recognition’ (2) is that they are consistently realized as [+NAS]. Given that speakers do not have any direct surface evidence that /VVN/ vowels are underlyingly oral, surface-oriented models like exemplar theory lack the necessary mechanisms to distinguish predictable pre-nasal vowels in /VVN/ from underlying contrastive nasal vowels / $\tilde{V}\tilde{V}$ /.

3 Experiment 1: How are /VVN/ vowels realized?

Experiment 1 measured the phonetic realization of Panjabi non-alternating /VVN/ vowels and compared their realization to contrastive / $\tilde{V}\tilde{V}$ / and oral /VV/ vowels in the language. As mentioned above, Zahid & Hussain (2012) analyzed /VVN/ vowels for three Panjabi speakers and did not discover a difference between their realization and the realization of contrastive / $\tilde{V}\tilde{V}$ / vowels with respect to acoustic correlates of vowel nasality. Thus, the first experiment aimed to replicate their findings using more speakers, and oro-nasal air pressure measurements rather than acoustic measurements.

3.1 Participants

Twenty native Panjabi speakers participated in the study (eleven men and nine women) ranging in age from 22-79 ($\mu = 39.7$; median = 33.5). All participants reported using Panjabi frequently throughout their daily lives in home, social, and academic settings. Participants were fluent in reading Panjabi in Shahmukhi, a Persio-Arabic script used by most Western Panjabi speakers. Besides Panjabi, all participants also spoke Urdu, the national language of Pakistan, and many were fluent in English. While almost all participants currently reside in Rawalpindi or Islamabad as of this writing, the cities in which they grew up span the Punjab Province. Ten speakers grew up in and around Rawalpindi, four in Gujrat, two in Sialkot, one in Sargodha, two near Toba Tek Singh, and one in the far south of Punjab in Sadiqabad. Despite the expansive distance between the hometowns of the participants, the results for experiment 1 were strikingly consistent across speakers.

3.2 Stimuli and Design

Stimuli consisted of 63 monosyllabic tokens separated into three conditions: oral (CVV), contrastive nasal / $\tilde{V}\tilde{V}$ / (C $\tilde{V}\tilde{V}$), and pre-nasal /VVN/ (CVVN). The complete list of tokens is given in Figure A1 of the appendix. Tokens from the CVV condition were used to establish a baseline measurement of orality, so no nasal segments were included in these forms. Tokens from the C $\tilde{V}\tilde{V}$ condition had contrastive / $\tilde{V}\tilde{V}$ / vowels and served as a baseline for nasality; to avoid any coarticulatory effects, nasal consonants were excluded from all tokens in this condition. Finally, the CVVN condition consisted of monomorphemic tokens, all with a nasal coda. Nasal onsets were excluded from this condition to avoid the effects of carryover coarticulation. Whereas the oral CVV and / $\tilde{V}\tilde{V}$ / C $\tilde{V}\tilde{V}$ conditions provide clear baselines of orality and nasality, the CVVN condition allows for a comparison between the realization of /VVN/ vowels and the other two conditions. Whenever possible, near-minimal pairs or triplets across the three conditions were used. As is well known, differences in vowel height correlate with differences in nasality (Henderson 1984, Bell-Berti 1993, Delvaux et al. 2008). To control for this, a balanced range of vowel qualities was included in each condition as well.

3.3 Procedure

All participants were recorded individually in a sound-attenuated room at a university in Rawalpindi, Pakistan. A native Panjabi research consultant provided participants with a consent form in

Panjabi Shahmukhi script detailing what the experiment involved. They were told that the experiment measured air pressure variations from the nose and mouth to study the Panjabi sound system.

After signing the consent form and filling out a demographic questionnaire, the research consultant introduced participants to the equipment and experimental procedure. All instructions were given in Panjabi. A dual chamber oro-nasal air pressure mask connected to two separate transducers produced by Glottal Enterprises was used to measure air pressure from the nose and mouth separately at a sampling rate of 11,025 kHz. The mask was equipped with a handle which participants used to raise and lower the mask to and from their face during the recording. Before the recording began, participants completed a test trial, in which they uttered several words into the mask to ensure they both understood the procedure and held the mask firmly against their face during productions to prevent airflow leakage.

At this point, the recording started, and participants were presented with the 67 monosyllabic tokens in a randomized order using a Python Graphical User Interface. Participants read each item into the mask four times: twice slowly and twice using fast/rapid speech. The experiment included speech rate as a condition to test whether vowel duration modulates the degree of nasality on /VVN/ vowels (Solé 1992, 1995). Tokens were presented in the Shahmukhi script with English translations below and instructions repeated at the bottom of the screen. If a word was mispronounced, the research consultant instructed participants to repeat all four iterations of the token. Recording sessions lasted between 8 and 16 minutes.

3.4 Analysis

Recordings were hand-annotated with segment boundaries using Praat (Boersma & Weenink 2015). Details of the annotation procedure are outlined in the appendix.

After annotation, a Praat script iterated through the annotated files, converting the oral and nasal air pressure channels to intensity tracks. A series of R scripts extracted nasalance values ($A^n/A^n + A^o$) at 11 normalized time points for each segment from these intensity tracks. The values were then saved in a data frame for subsequent analysis.

3.5 Results

Data from nineteen of the twenty speakers that participated in the experiment is included in the results presented here, summing up to 4,761 tokens (63 tokens x 4 repetitions x 19 speakers - 27 low-quality tokens) in the analysis. One speaker, MA0, was removed from the analysis because they produced a large number of tokens in which contrastive / $\tilde{V}\tilde{V}$ / vowels and /VVN/ vowels were produced with abnormally low nasalance, essentially the same as in the oral condition.

Tokens were analyzed using R (RCoreTeam 2016). To visually compare the relative nasality between the three conditions, the vowel segment was extracted from each token, and nasalance – defined as the proportion of total air pressure present in the nasal cavity – was used to fit three generalized additive model (GAM) curves, one for each vowel type, to the data using ggplot2 (Wickham 2016). No difference was found between tokens uttered in slow speech and tokens uttered in fast speech, so all four repetitions of each token were lumped together in the model. This fact is evident from comparing the two oral-nasal waveforms in figure 2 for the word /k^háán/ (acute accent indicates high tone). The same speaker uttered both tokens in short succession, but

the words differed in speech rate; the vowel of the first repetition (#132) was 361 ms in duration, and the vowel of the second repetition (#133) was 216 ms in duration. Despite this significant durational difference, nasalization is categorically implemented across the entire span of both repetitions. This indicates that nasality on the /VVN/ vowel is intentionally implemented and controlled by the speaker rather than resulting from a coarticulatory effect due to the vowel's adjacency to the nasal consonant (Solé 1992, 1995).

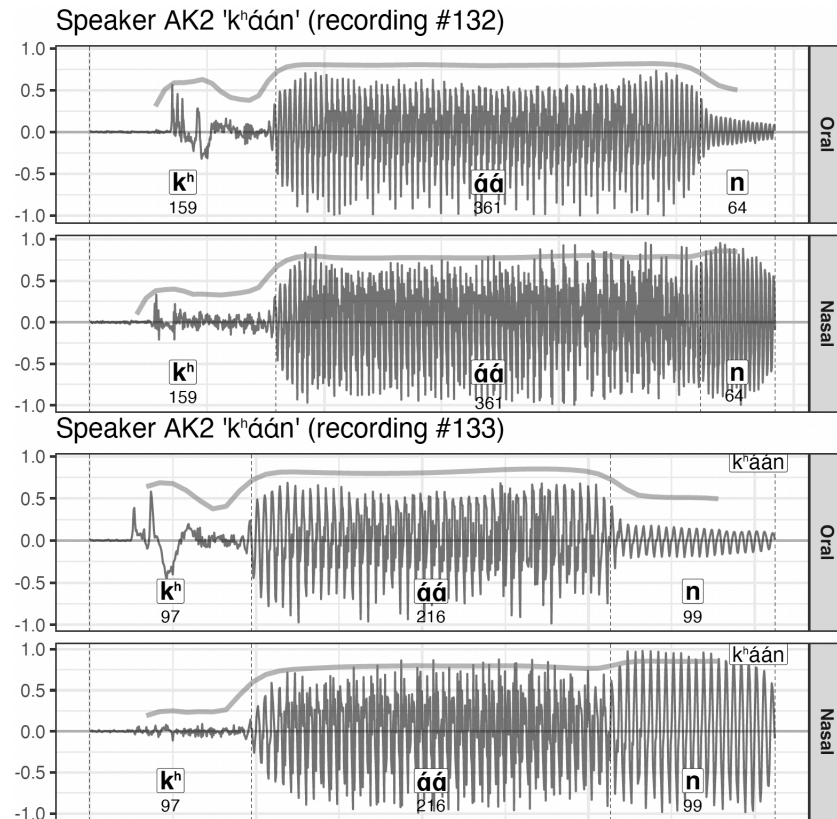


Figure 2: Oral and nasal waveforms stacked on each other for two different productions of /k'háán/ by the same speaker in quick succession that differ only in the speech rate. The waveforms show the normalized amplitude for each air pressure chamber, and the grey line indicates the normalized intensity contour.

To confirm that nasalance is not modulated by speech rate for /VVN/ vowels, we conducted a Pearson's correlation test to examine the relationship between the two variables. The results ($r(1287) = -.001$, $p = .97$), visualized by the scatterplot in figure 3, show no significant correlation between speech rate and mean nasalance for /VVN/ vowels.




Figure 3: Scatterplot showing the -.001 Pearson Correlation value between speech rate and mean nasalance for all CVVN vowels produced by nineteen speakers in experiment 1.

Figure 4 displays plots of the three GAM curves depicting mean nasalance across all tokens and speakers. As expected, nasalance for oral /VV/ vowels remains steady and noticeably lower than the other two vowel types at about 40% across the length of the oral vowel. For the contrastive / $\tilde{V}\tilde{V}$ / vowel condition (C $\tilde{V}\tilde{V}$), on the other hand, nasalance begins at about 45% at the onset of the vowel and rises steadily until reaching a plateau-like position at about 52% nasalance a third of the way through the vowel. This proportion of nasalance is maintained throughout the remainder of the vowel's duration. As for the nasalance of /VVN/ vowels in the CVVN condition, they follow a very similar pattern to contrastive / $\tilde{V}\tilde{V}$ / vowels, in which nasalance at the onset of the vowel sits at approximately 45%, followed by a steady rise in nasalance until a plateau is reached at about 52% nasalance which is maintained through the offset of the vowel.

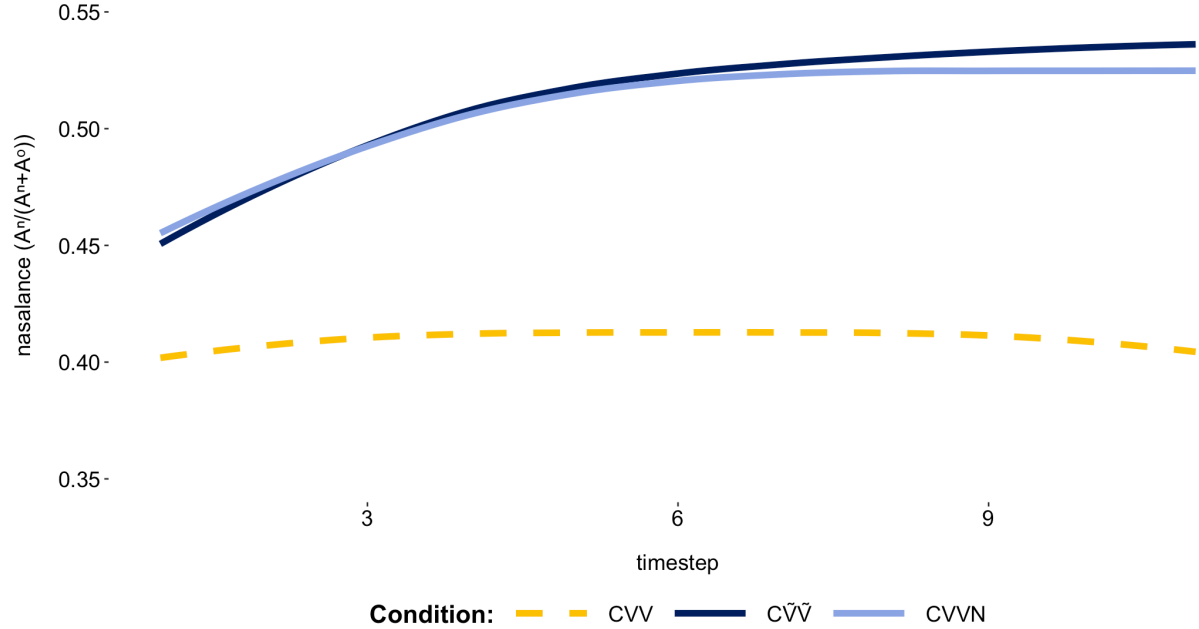


Figure 4: GAM curves separated by condition showing mean vowel nasalance across 11 normalized timesteps for all monosyllabic tokens from experiment 1 for nineteen speakers.

To quantitatively compare the nasalance between CŨŨ and CŨŨN vowels, a linear mixed-effects regression (LMER) model was fit to the data using the `lmer()` function from the `lme4` package (Bates et al. 2025) with mean vowel nasalance as the dependent variable and `CONDITION` (CŨŨ and CŨŨN), `TIMESTEP` (beginning, middle, final), the `CONDITION*TIMESTEP` interaction, and `VOWELQUALITY` included as the four fixed effects. Oral vowels were omitted as an additional level of `CONDITION` in the model to allow for a more straightforward comparison between CŨŨN and CŨŨ vowels. Moreover, rather than include all eleven steps at which nasalance was measured as separate levels of the `TIMESTEP` factor, the steps were binned into three different levels corresponding to the beginning third, middle third, and final third of the vowel. `VOWELQUALITY` was included as a predictor to account for the known differences in nasality across different vowel qualities mentioned above. A likelihood ratio test showed that the addition of `VOWELQUALITY` as a fixed effect significantly improved the fit ($\chi^2(5) = 36.42$, $p < .001$ ***), so it was retained in the model. A random intercept was included for word. By-speaker random slopes and intercepts were included for `CONDITION`.

The main effects of `CONDITION`, `TIMESTEP`, and `VOWELQUALITY` were effect coded, so the intercept (β_0) refers to the grand mean of nasalance averaged over all levels of the main effects, and the remaining coefficients — β_1 - β_{14} in table 1 — equal the difference between the grand mean and the listed factor level, controlling for the other predictors. The results of the LMER model are provided in table 1.

The `CONDITION*TIMESTEP` interaction was significant at the beginning and end of the vowel, which means that the degree of nasalance difference between CŨŨN and CŨŨ is not consistent across these timesteps. Nevertheless, the effect sizes of these interactions can be considered inconsequential since neither interaction deviates from the overall main effect by more than .002 (0.2%) in nasalance. Thus, it is safe to examine the main effect of `CONDITION` without worrying

about interaction effects clouding its interpretability.

Similarly, the main effect of VOWELQUALITY was significant for four of the six vowels included in the study, reaffirming that different vowels exhibit differing degrees of nasalization, controlling for other factors. However, because the effect sizes across vowel qualities were relatively minimal and vowel quality was balanced across the different nasality conditions, it is safe to set these effects aside as well.

Disregarding the small interaction and vowel quality effects, the primary takeaway from the model for our purposes is located in the first three rows. Specifically, neither level of the main effect of CONDITION (C $\tilde{V}\tilde{V}$ and CVVN) reach significance. That is, the null hypothesis that mean nasalance is identical for C $\tilde{V}\tilde{V}$ and CVVN vowels across their entire span cannot be rejected.

Predictor	β	95% CI	t	p
Intercept (grand μ)	.514	[.504, .525]	93.61	< .001 ***
C $\tilde{V}\tilde{V}$.002	[-.001, .005]	1.35	.18
CVVN	-.002	[-.005, .001]	-1.35	.18
Fin	.020	[.019, .020]	55.85	< .001 ***
Mid	.011	[.010, .011]	34.21	< .001 ***
Beg	-.030	[-.031, -.030]	-98.31	< .001 ***
uu	.010	[.005, .016]	3.54	.001 **
ii	.014	[.006, .023]	3.4	.002 **
oo	-.014	[-.020, -.007]	-3.98	< .001 ***
ɔɔ	-.002	[-.010, .006]	-.54	.59
ee	-.001	[-.014, .012]	-.17	.86
ɑɑ	-.008	[-.012, -.003]	-3.14	.004 **
CVVN:Fin	-.002	[-.003, -.002]	-6.9	< .001 ***
CVVN:Mid	.0004	[-.001, .001]	.03	.98
CVVN:Beg	.002	[.002, .003]	7.9	< .001 ***

Table 1: LMER model summarizing the effect of CONDITION (C $\tilde{V}\tilde{V}$, CVVN), TIMESTEP (beginning, middle, final), the CONDITION*TIMESTEP interaction, and VOWELQUALITY (ɑɑ, ɔɔ, oo, ee, uu, ii) on mean nasalance for nineteen speakers: $\text{lmer}(\text{nasalance} \sim \text{CONDITION} * \text{TIMESTEP} + \text{VOWELQUALITY} + (1 + \text{CONDITION} | \text{speaker}) + (1 | \text{word}))$

To directly compare whether or not C $\tilde{V}\tilde{V}$ and CVVN vowels are significantly different from each other in terms of nasalance, we ran an additional linear-mixed effects model which only included CONDITION (C $\tilde{V}\tilde{V}$, CVV, and CVVN) as a simple main effect. When exploring simple main effects, a single predictor variable (in this case, CONDITION) is included in the model, and one of the levels of that predictor serves as the intercept (in this case, C $\tilde{V}\tilde{V}$). The remaining coefficient values describe the amount that other levels of the predictor differ from the intercept. In essence, the model is dummy-coded.

As expected, predictions of the model in Table 2 show a significant difference between oral /VV/ and contrastive / $\tilde{V}\tilde{V}$ / vowels. β_0 (C $\tilde{V}\tilde{V}$) shows a mean nasalance of about .510, and β_1 (CVV) demonstrates a significant effect of CVV on nasalance of about -.093. These nasalance values are the established baselines for oral and nasal vowels in Panjabi.

Predictor	β	95% CI	t	p
Intercept (C $\tilde{V}\tilde{V}$)	.510	[.495, .526]	63.85	< .001 ***
CVV	-.093	[-.116, -.07]	-7.9	< .001 ***
CVVN	-.003	[-.022, .016]	-.29	.78

Table 2: Results of the LMER model predicting nasalance across the three vowel types: $\text{lmer}(\text{nasalance} \sim \text{CONDITION} + (1 + \text{CONDITION} | \text{speaker}) + (1 | \text{word}))$

Crucially, the effect of CVVN on nasalance did not reach significance, meaning that the null hypothesis that /VVN/ and / $\tilde{V}\tilde{V}$ / vowels are indistinguishable in terms of their nasalance values cannot be ruled out.

Nevertheless, while no effect of CONDITION was found to alter the mean nasalance of these two vowels, the result from the above LMER model does not allow us to conclude that the nasalance of / $\tilde{V}\tilde{V}$ / and /VVN/ are equivalent. Rather, an equivalence test must be conducted to determine whether we can accept the null hypothesis that the mean nasalance of the two vowel types is identical. Within the framework of classical statistics, a test of equivalence relies on determining a region of practical equivalence ($-h$ to h) for a specific research question and examining whether or not the confidence interval for the difference in means between two groups falls within that region (Wellek 2010, Lakens et al. 2018, Bonett 2021). If the confidence interval for the difference is wholly contained within that interval, the hypothesis that a difference exists between the two groups with respect to the dependent variable can be rejected. For this experiment, the dependent variable is mean nasalance, and the two groups of interest are the vowels in the C $\tilde{V}\tilde{V}$ and CVVN conditions.

To determine a region of practical equivalence (*RoPE*) of mean nasalance in Panjabi, we separated the speakers from experiment 1 into two groups, with group 1 comprised of the nine speakers in the study with the highest mean nasalance for / $\tilde{V}\tilde{V}$ / vowels and group 2 comprised of the ten speakers with the lowest mean nasalance for / $\tilde{V}\tilde{V}$ / vowels. We then ran a LMER model to compare the mean nasalance of / $\tilde{V}\tilde{V}$ / vowels between the two groups of speakers. Group membership was included as a simple main effect, and random intercepts for word were included as well. For our purposes, any differences in mean nasalance on / $\tilde{V}\tilde{V}$ / vowels between the two groups of speakers should be considered phonologically negligible, since all of these vowels are phonologically contrastive nasal / $\tilde{V}\tilde{V}$ /. In this way, the differences between the two groups can be used to establish a conservative value of h , and thus a region of practical equivalence within which a difference in nasalance can be disregarded as within-category phonetic variation, rather than a cross-category phonological difference.

Table 3 summarizes the outcome of the LMER model. Results show that the mean nasalance of Group 2 was significantly smaller than that of Group 1, with a 95% confidence interval of [-.033, -.030]. Nevertheless, despite a significant p-value between the two groups, the effect can be treated as meaningless because the vowels of both groups belong to the same / $\tilde{V}\tilde{V}$ / category in Panjabi. Taking the largest potential difference between these two groups as an effect size that can be ignored in terms of categorical nasality, a conservative window of h can be set at .033. Any difference within $\pm .033$ can be treated as insignificant because speakers abstract away from such small differences in nasalance when categorizing a vowel as [+NAS] or not.

Predictor	β	95% CI	t	p
Intercept (Group 1)	.527	[.507, .518]	179.45	< .001 ***
Group 2	-.031	[-.033, -.030]	-42.24	< .001 ***

Table 3: Results of the LMER model predicting C $\tilde{V}\tilde{V}$ nasalance across two groups of speakers. Group 1 is comprised of the nine speakers in the study with the highest mean nasalance on C $\tilde{V}\tilde{V}$ vowels, and Group 2 is comprised of the ten speakers in the study with the lowest mean nasalance on C $\tilde{V}\tilde{V}$ vowels.

Figure 5 provides additional evidence that a *RoPE* of $\pm.033$ for nasalance is a conservative value. The points on the plot represent the mean nasalance for / $\tilde{V}\tilde{V}$ / vowels for each speaker from experiment 1, with the tails indicating \pm two standard deviations (capturing 95% of productions) from that mean. When a listener normalizes expected nasalance levels for a given speaker, any nasalance value within two standard deviations of that speaker’s typical nasalance level for a / $\tilde{V}\tilde{V}$ / vowel can reliably be interpreted as nasal. Notably, the magnitude of the suggested *RoPE* is smaller than the range of nasalance values for all but two speakers from experiment 1. Thus, $\pm.033$ should be viewed as a relatively conservative threshold for nasality equivalence across vowel types.




Figure 5: Plot showing mean nasalance of / $\tilde{V}\tilde{V}$ / vowels for each speaker in experiment 1. Tails represent \pm two standard deviations around these mean nasalance values. Speakers are ordered from left to right by magnitude of standard deviation. The bolded point shows the mean nasalance across all speakers, and its tails represent the proposed *RoPE* of $\pm.033$. 17 of the 19 speakers exhibit a larger spread of / $\tilde{V}\tilde{V}$ / vowel nasalance than the *RoPE*.

Reexamining the 95% CI of the simple main effect of CVVN on nasalance in the LMER model in table 2, which compares nasalance in CVVN vowels to nasalance in C $\tilde{V}\tilde{V}$ vowels, we see that the entire confidence interval ([-.022, .016]) falls within the bounds of our established region of practical equivalence. Therefore, it seems appropriate to conclude that the mean nasalance of

the two vowel types is both statistically and practically equivalent. This allows us to assert that CVVN vowels surface as phonologically [+NAS] in the same way as C $\tilde{V}\tilde{V}$ vowels in Panjabi.

To sum up, the results of experiment 1 confirm and extend the findings from the acoustic analysis in Zahid & Hussain (2012): /VVN/ vowels are not meaningfully different from / $\tilde{V}\tilde{V}$ / vowels with respect to their phonetic nasality.³

4 Experiment 2: Do /VVN/ vowels trigger nasal harmony?

Experiment 1 found that contrastive / $\tilde{V}\tilde{V}$ / and non-alternating [$\tilde{V}\tilde{V}$ N] vowels are practically and statistically equivalent with respect to the acoustic realization of vowel nasality. Experiment 2 explores how these two vowel types interact with another phonological process in Panjabi: leftward nasal spread. Specifically, experiment 2 is designed to test if differences exist between contrastive / $\tilde{V}\tilde{V}$ / and non-alternating [$\tilde{V}\tilde{V}$ N] with respect to their ability to trigger nasal harmony in Panjabi.

The results of experiment 1 help frame two potential hypotheses regarding the interaction of surface [$\tilde{V}\tilde{V}$ N] vowels with nasal harmony. The first hypothesis asserts that, because contrastive / $\tilde{V}\tilde{V}$ / and non-alternating [$\tilde{V}\tilde{V}$ N] vowels appear to be identical in their phonetic nasality, they should *behave* identically in terms of triggering nasal harmony.

The abstract underlying representations provided by generative phonology also raise a second possibility. Vowel nasality in [$\tilde{V}\tilde{V}$ N] is predictable rather than contrastive in Panjabi. As such, surface [$\tilde{V}\tilde{V}$ N] could be analyzed as underlyingly /VVN/, with an oral vowel; predictable nasality would then be supplied by a rule or other formal grammatical device. This analysis produces a distinction between contrastively nasal vowels / $\tilde{V}\tilde{V}$ / – which would be underlyingly nasal – and underlyingly oral vowels in /VVN/.

Thus, a second hypothesis is that / $\tilde{V}\tilde{V}$ / and /VVN/ vowels might behave *differently* with respect to triggering harmony because they have a different underlying status. Specifically it could be that only underlying, contrastively nasal / $\tilde{V}\tilde{V}$ / triggers harmony, while underlying oral /VVN/ → [$\tilde{V}\tilde{V}$ N] does not.

Importantly, contextually nasal vowels /VVN/ → [$\tilde{V}\tilde{V}$ N] may be non-alternating. There is thus no surface evidence that these vowels are anything other than nasal. For this reason, hypothesis 2 is only stateable in frameworks which either (i) allow for covert URs which are qualitatively distinct from any observed surface form(s), or (ii) invoke other mechanisms for distinguishing between contrastive vs. predictable vowel nasality.

To the extent that surface-oriented frameworks like exemplar theory make predictions about the interaction of harmony with contextual vowel nasalization, we believe that they predict hypothesis 1: namely, that non-alternating [$\tilde{V}\tilde{V}$ N] should behave parallel to contrastive / $\tilde{V}\tilde{V}$ / in triggering harmony.

³The discussion here assumes that – because they are produced with identical nasality – / $\tilde{V}\tilde{V}$ / and /VVN/ are *perceptually* identical as well. Nevertheless, the production studies in this paper cannot answer questions about perception, so future research is needed to determine whether or not contrastive / $\tilde{V}\tilde{V}$ / vowels are perceptually distinct from /VVN/ vowels to listeners. See also Beddor & Krakow (1999).

4.1 Stimuli, design, and analysis

Sixteen additional native Panjabi speakers (eleven men and five women) who did not participate in experiment 1 took part in experiment 2. Ages ranged from 18 to 43 ($\mu = 28.2$; median = 27). Similar to the first experiment, stimuli were sourced from three main vowel types: oral /VV/, contrastive nasal / $\tilde{V}\tilde{V}$ /, and non-alternating pre-nasal /VVN/ vowels. Here, however, all words were either di- or trisyllabic, and the shape of the final two syllables was controlled for, such that every word contained a vowel–glide–vowel sequence (VVG \tilde{V} VV) across the last two syllables. This construction made it possible to examine whether the three vowel types differ in their interactions with nasal harmony.

Tokens from all three vowel conditions were produced with a CV postposition beginning with an oral stop (e.g. [aav \tilde{a} am=də] ‘public=GEN’). In addition to these three conditions, words ending in a contrastive / $\tilde{V}\tilde{V}$ / vowel followed by a nasal-initial NV postposition were included as an additional condition. In Panjabi, grammatical postpositions immediately follow the form they modify and act as enclitics, phonologically integrating with the preceding form and lacking independent stress or prosodic constituency (Bashir & Connors 2019: 244). By including a condition in which contrastive / $\tilde{V}\tilde{V}$ / vowels are immediately followed by a postposition beginning with a nasal consonant, the contextual environment of / $\tilde{V}\tilde{V}$ / more closely mimics that of /VVN/ vowels. This facilitates a more direct comparison of nasalance patterns between the two vowel types.

Altogether, four distinct conditions with 6-10 tokens each were included in the experiment, as shown in table 4. Twenty-seven filler items were also randomly dispersed throughout the token set. Complete glosses of all tokens included in the experiment are provided in Figure A2 in the Appendix.

VVG \tilde{V} V=CV		VVG \tilde{V} VN=CV	
/sət-aqjii də/	‘27’	/dʒəvaan də/	‘a youth’
/ɔɫii-jaə də/	‘pious men’	/ədʒəvɛɛn də/	‘seed’
/həvaə də/	‘air’	/siijaan də/	‘recognition’
/tʃoo-vii də/	‘24’	/teejaan də/	‘attention’
/paavee də/	‘cot leg’	/lɛɛ-jaan de/	‘bringing’
/beevaa də/	‘widow’	/aav-aam də/	‘citizens’
/dəvaə də/	‘medicine’		
/taavuu də/	‘uncle’		
/vək ^h aavee də/	‘boasting’		
/møk ^h iijaa də/	‘chief’		
VVG $\tilde{V}\tilde{V}$ =NV		VVG $\tilde{V}\tilde{V}$ =CV	
/tʃaavēē nɔ/	‘pumicestone’	/tʃaavēē də/	‘pumicestone’
/uu-vēē nɔ/	‘in that way’	/uu-vēē də/	‘in that way’
/doo-vāā nɔ/	‘both sides’	/doo-vāā paase/	‘both sides’
/kii-vēē nɔ/	‘somehow’	/kii-vēē də/	‘somehow’
/ɛɛ-vēē nɔ/	‘in vain’	/ɛɛ-vēē də/	‘in vain’
/matʃii-jāā nɪ/	‘fish-PL’	/matʃii-jāā də/	‘fish-PL’
/tiivīi nɪ/	‘woman’	/tiivīi də/	‘woman’
/saa-vāā nɪ/	‘breaths’	/saa-vāā də/	‘breaths’

Table 4: Experiment 2 tokens.

Both the contrastive oral /VV/ and contrastive nasal / $\tilde{V}\tilde{V}$ / conditions provide established baselines for an oral vowel with no nasality in the signal and a / $\tilde{V}\tilde{V}$ / vowel that triggers nasal harmony, respectively. The /VVN/ vowel condition, on the other hand, serves as the focus of the

experiment. Importantly, VVN sequences from the tokens in the /VVN/ condition for experiment 2 never straddle a morpheme boundary, so Panjabi speakers lack experience with alternations that reveal the status of these /VVN/ vowels when not followed by a nasal consonant. Thus, assuming similar results to experiment 1, in which /VVN/ vowels are phonetically identical to contrastive / $\tilde{V}\tilde{V}$ / vowels in terms of nasality, /VVN/ vowels in these tokens should always be realized as [+NAS]. Hypothesis 1 predicts that /VVN/ vowels should trigger nasal harmony at the same rate as contrastive / $\tilde{V}\tilde{V}$ / vowels because of their identical surface realizations. Hypothesis 2, in contrast, leaves open the possibility that the two vowels possess different underlying representations, which could result in the two vowels interacting with the nasal harmony process differently.

In the preceding section we suggested that surface-oriented frameworks like exemplar theory most naturally predict that allophonic vowel nasality in [$\tilde{V}\tilde{V}N$] should trigger harmony to the same extent as contrastive vowel nasality in / $\tilde{V}\tilde{V}$ /. In the event that [$\tilde{V}\tilde{V}N$] does *not* trigger harmony, there is at least one surface-oriented generalization which might explain that asymmetry. Assume that in both [$VV.G\tilde{V}\tilde{V}$] and [$VV.G\tilde{V}\tilde{V}N$] there is a *fixed period* of nasality anchored at the right edge of the word (see e.g. Beddor 2007). If so, we might expect nasality to begin later in [$\dots VV.G\tilde{V}\tilde{V}N$] than in [$\dots \tilde{V}\tilde{V}.G\tilde{V}\tilde{V}$], because the additional, word-final nasal segment in [$\dots VV.G\tilde{V}\tilde{V}N$] would cause the overall ‘window’ of nasality to shift rightward (see also Cole & Kisseberth 1994, 1995). This difference in the timing of phonetic nasality could give the impression of a difference in the application of harmony, without recourse to covert underlying forms.

One of our conditions – / $\tilde{V}\tilde{V}$ / vowels followed by an NV postposition, / $\tilde{V}\tilde{V}=NV$ / – may help assess this proposal. When a / $\tilde{V}\tilde{V}$ / vowel is followed by an NV postposition / $\tilde{V}\tilde{V}=NV$ /, two potential patterns of nasality arise. First, the / $\tilde{V}\tilde{V}=NV$ / condition could display a similar pattern to the / $\tilde{V}\tilde{V}=CV$ / condition. This would be the case if the underlying representation of contrastive / $\tilde{V}\tilde{V}$ / vowels drives nasal harmony. Alternatively, the contrastive / $\tilde{V}\tilde{V}$ / vowel in / $\tilde{V}\tilde{V}=NV$ / may behave similarly to a non-alternating /VVN/ vowel, such that the / $\tilde{V}\tilde{V}=NV$ / condition does *not* exhibit nasal harmony. The latter pattern may be expected if the nasal harmony process in the / $\tilde{V}\tilde{V}=CV$ / condition does not represent harmony at all, but instead a period of velar opening of fixed duration which shifts rightward in the presence of a following nasal consonant.

In short, if the presence of a nasal consonant shifts the nasalization window rightward, as outlined above, we would expect the nasalance traces in the / $\tilde{V}\tilde{V}=NV$ / condition to mimic that of the non-alternating /VVN=CV/ condition, rather than the contrastive / $\tilde{V}\tilde{V}=CV$ / condition.

4.2 Results

A total of 2,048 potential tokens (16 speakers x 4 repetitions x 32 words) resulted from the recordings. Unfortunately, the recording for one speaker (MI0) showed little to no differences between the contrastive oral /VV/ and nasal / $\tilde{V}\tilde{V}$ / conditions, so the data from this participant was removed. In addition, 50 more tokens across the remaining fifteen speakers were removed due to poor audio quality. After removing these words, a total of 1,870 tokens remained in the analysis.

All speakers in experiment 2 produced tokens from the oral /VV/ and the contrastive / $\tilde{V}\tilde{V}=CV$ / and / $\tilde{V}\tilde{V}=NV$ / conditions uniformly, but productions of the /VVN/ condition fell into two distinct patterns that require individual attention. All fifteen speakers lacked nasal harmony across the preceding VVG sequence in the /VVN/ condition (section 4.2.2.1). Thirteen speakers categorically nasalized the /VVN/ vowel immediately preceding the nasal consonant (section 4.2.2). However,

the remaining two speakers systematically lacked categorical nasality on all /VVN/ vowels, instead demonstrating a pattern of anticipatory coarticulation. We discuss these two speakers in section 4.2.2.3, where we speculate that their productions arise from dialectical differences.

Before examining the results in more detail, one other fact about nasal harmony in the / $\tilde{V}\tilde{V}$ =CV/ and / $\tilde{V}\tilde{V}$ =NV/ conditions must be considered. Across all speakers and words in the experiment, it was evident that nasal harmony on the penultimate vowel only sometimes applied; for 38% of contrastive / $\tilde{V}\tilde{V}$ / tokens (n=367), nasal harmony was not triggered by / $\tilde{V}\tilde{V}$ /. As no grammatical descriptions known to us report optionality or variability for nasal harmony in Panjabi, this result was unexpected.

Figure 6 illustrates with two productions of /uu-věẽ də/ by speaker AA1, uttered within one second of each other. One production of /dʒəvaan də/ is also included as a typical example of the /VVN/ condition for reference. Examining the nasal channels in the figure for both productions of /uu-věẽ də/, nasalance on the penultimate vowel exhibits categorical nasalization across the entire [VVGVV] sequence for recording #159 (top panel). For recording #160 (middle panel), on the other hand, nasalance remains at the level of an oral vowel for the entirety of the penultimate vowel's duration and then rises rapidly at the onset of the glide.

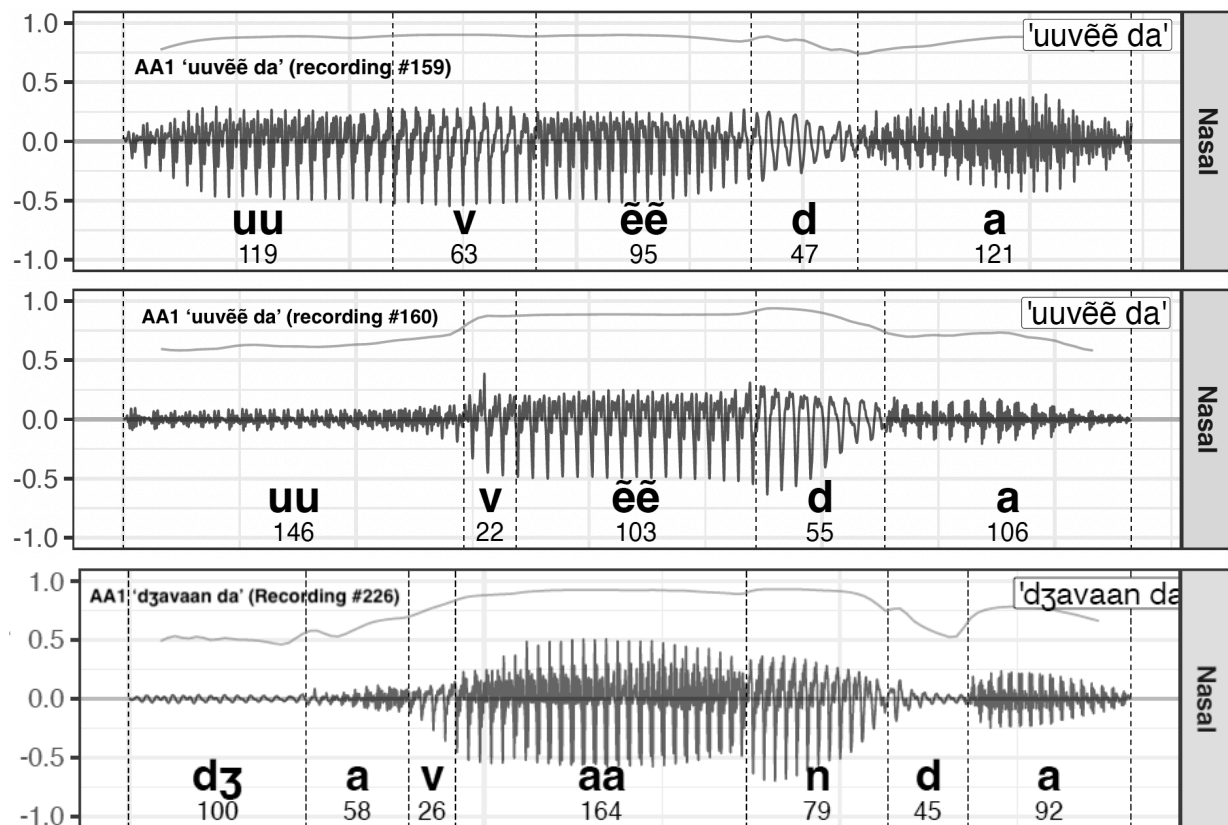


Figure 6: Nasal channel waveforms for two different productions of /uuvěẽ/ by the same speaker in quick succession that differ only in the application of the nasal harmony process. The final waveform shows a single production of /dʒəvaan/ by the same speaker. The y-axis depicts the normalized amplitude of the air pressure in the nasal cavity, and the light grey line shows the normalized intensity contour over time.

Importantly, this variable application of the nasal harmony process was present for all speak-

ers and words in both the contrastive $\tilde{V}\tilde{V}=CV/$ and $\tilde{V}\tilde{V}=NV/$ conditions. Moreover, the optional application of harmony was not present in the $/VVN/$ condition. These observations indicate that nasal harmony in Panjabi applies optionally when triggered by $\tilde{V}\tilde{V}/$.

Because experiment 2 aims to examine whether $/VVN/$ vowels differ from contrastive $\tilde{V}\tilde{V}/$ vowels in their ability to trigger nasal harmony in the language, all tokens that clearly lacked harmony in the $\tilde{V}\tilde{V}=CV/$ and $\tilde{V}\tilde{V}=NV/$ conditions were removed to prevent the obfuscation of the harmony pattern. This was accomplished for each speaker by removing tokens from the $\tilde{V}\tilde{V}=CV/$ and $\tilde{V}\tilde{V}=NV/$ conditions in which the penultimate vowel exhibited a mean nasalance clustering with the mean nasalance of penultimate vowels in the oral $/VV/$ condition for that speaker. For example, figure 7 shows a density plot depicting the concentration of mean nasalance for penultimate vowels in the $\tilde{V}\tilde{V}=CV/$ and $\tilde{V}\tilde{V}=NV/$ conditions for speaker AW0. As is evident, the distribution is bimodal, with the large majority of penultimate vowels exhibiting a mean nasalance of approximately .51-.53. Conversely, another smaller concentration of vowels cluster around a mean nasalance of roughly .42-.44. The vertical line shows the approximate cut in which tokens with mean nasalance to the left of the line were removed and tokens with mean nasalance to the right of the line retained. No tokens from the $/VVN/$ condition were removed in this process. Thus, the comparison between contrastive $\tilde{V}\tilde{V}/$ vowels and $/VVN/$ vowels in this study includes only those tokens in the $\tilde{V}\tilde{V}=CV/$ and $\tilde{V}\tilde{V}=NV/$ conditions in which the nasal harmony process *does* apply. That is, tokens like recording #160 (middle panel) in figure 6 above were removed.

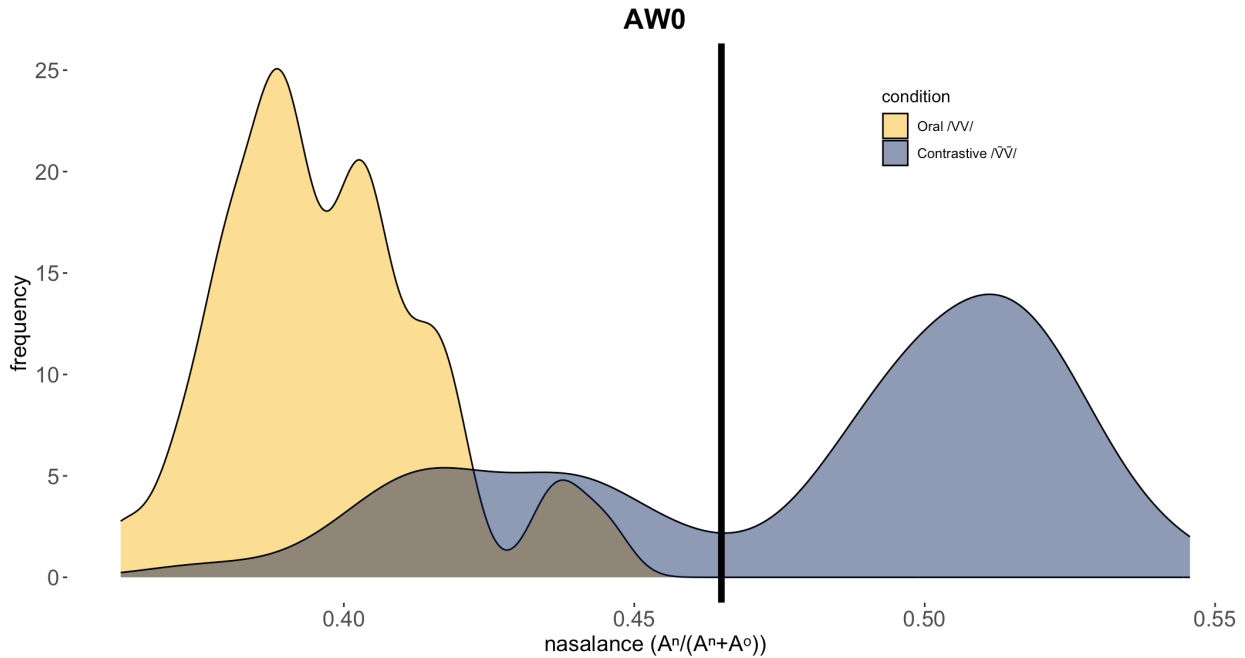


Figure 7: Density plot showing the distribution of mean nasalance for penultimate vowels for speaker AW0. Contrastive $\tilde{V}\tilde{V}/$ includes penultimate vowels from both the $\tilde{V}\tilde{V}=CV/$ and $\tilde{V}\tilde{V}=NV/$ conditions. The plot shows a bimodal distribution for contrastive $\tilde{V}\tilde{V}/$ vowels, with two distinct peaks around .43 and .52, respectively. The vertical line marks the cutoff point for speaker AW0, in which tokens with a mean nasalance to the left of the line were removed from the analysis. The multimodal distribution in the oral condition reflects inherent nasalance differences across vowel qualities (Bell-Berti 1993).

4.2.1 Contrastive / $\tilde{V}\tilde{V}$ / and nasal harmony

Figure 8 displays GAM curves fit to nasalance at each of the 33 timesteps across the [VVG \tilde{V} VV] sequence for the oral /VV/ and contrastive / $\tilde{V}\tilde{V}$ / conditions, for all fifteen speakers.⁴ As expected, the oral /VV/ condition exhibits a relatively stable nasalance at about 41% across the entire [VVG \tilde{V} VV] sequence. Tokens from the contrastive / $\tilde{V}\tilde{V}$ / condition, on the other hand, show an increase in nasalance from the onset of the penultimate vowel until a plateau of nasalance is reached about a third of the way through the penultimate vowel and maintained through the end of the word. This pattern is qualitatively consistent with a process of nasal harmony triggered by the final / $\tilde{V}\tilde{V}$ / vowel and targeting both the glide and the penultimate vowel. The initial slope in nasalance across the first third of the penultimate vowel is due to the transition out of an oral onset consonant. In sum, the GAM curves highlight a clear and relatively stable distinction in nasalance between the oral /VV/ and contrastive / $\tilde{V}\tilde{V}$ / conditions across the length of the [VVG \tilde{V} VV] sequence.

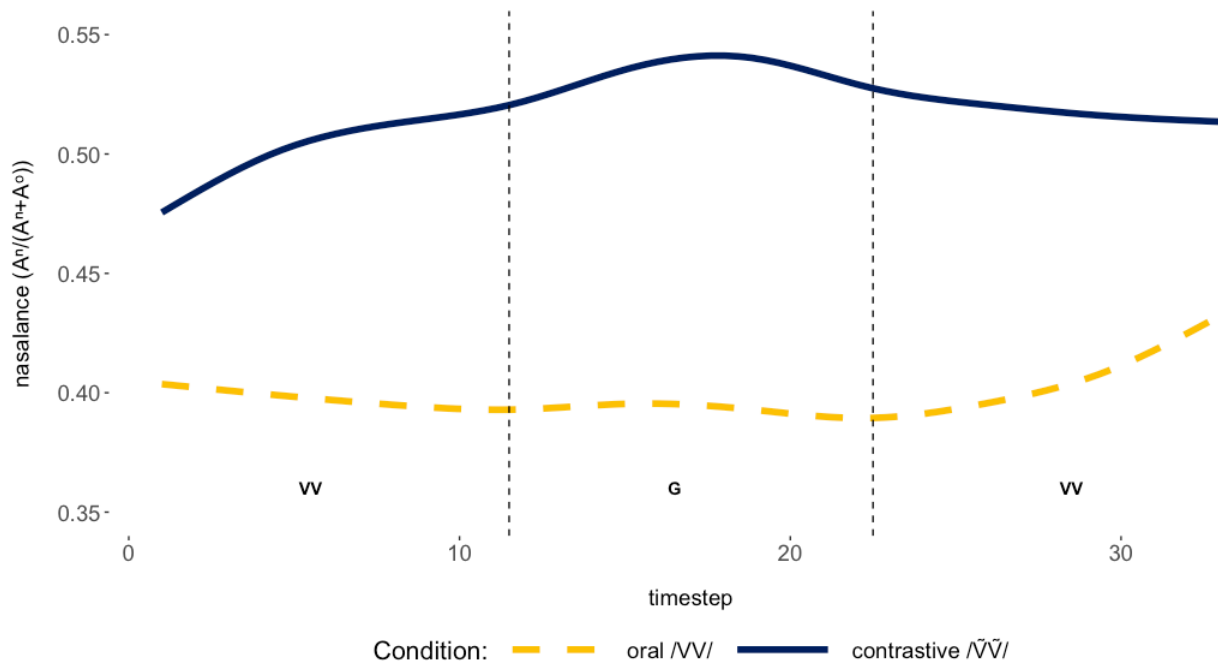


Figure 8: GAM curves for the oral /VV/ and contrastive / $\tilde{V}\tilde{V}$ / conditions in experiment 2. Curves visualize the mean vowel nasalance at 33 normalized timesteps across the [VVG \tilde{V} VV] sequence for all 15 speakers in the experiment. Vertical dashed lines represent segment boundaries.

4.2.1.1 Realization of contrastive / $\tilde{V}\tilde{V}$ / vowels

To quantitatively assess the difference in nasalance between the final oral /VV/ and contrastive / $\tilde{V}\tilde{V}$ / vowels, we ran an initial linear mixed effects model, with CONDITION (2 levels), TIMESTEP (3

⁴Because nasalance measurements were taken at 11 equidistant timesteps for each segment, figure 8 and other similar figures reporting nasality in [VVG \tilde{V} VV] involve some distortion in time. For instance, glides were typically much shorter than either of the vowels in any given token, and so the time steps during [G] in figure 8 are not strictly comparable to the time steps for the adjacent vowels [VV].

levels), the CONDITION*TIMESTEP interaction, and VOWELQUALITY (4 levels) included as fixed effects. The three levels of the TIMESTEP variable roughly correspond to the first third (Beg), middle third (Mid), and final third (Fin) of the final vowel in the [VVG VV] sequence. By-speaker random slopes and intercepts were included for condition and random intercepts for word. A likelihood ratio test comparing the full model described above to a model without VOWELQUALITY as a fixed effect showed no significant difference in fit ($\chi^2(3) = 3.25$, $p = .35$). As such, VOWELQUALITY was removed from this model to reduce model complexity.

Results of the LMER are provided in table 5. The main effect for all levels of CONDITION and TIMESTEP are included in the model output for clarity. The interaction between CONDITION and TIMESTEP is significant at the beginning third and final third of the vowel, meaning that the effect of CONDITION on nasalance is not stable across the vowel’s timespan. Nevertheless, the effect size (β) of these interactions is marginal at less than 1% difference in nasalance, so they can be safely set aside to examine the main effects.

The interesting portion of the model concerns the main effect of CONDITION on nasalance. As is evident from the first three rows, the mean nasalance for word-final vowels in the oral /VV/ and contrastive / $\tilde{V}\tilde{V}$ / conditions diverged in opposite directions by 8% nasalance from the grand mean. This result is unsurprising, as the two vowel types exhibit a phonological contrast in nasality. However, the outcome does provide an approximate magnitude of the nasalance difference we should expect between vowels with categorically distinct nasality values in Panjabi.

Predictor	β	95% CI	t	p
Intercept (grand μ)	.461	[.441, .480]	45.66	< .001 ***
Contrast. / $\tilde{V}\tilde{V}$ /	.079	[.064, .093]	10.64	< .001 ***
Oral /VV/	-.079	[-.093, -.064]	-10.64	< .001 ***
Beg	-.003	[-.004, -.002]	-5.79	< .001 ***
Mid	-.0002	[-.001, .001]	.44	.66
Fin	.003	[.002, .004]	4.69	< .001 ***
Contrast. / $\tilde{V}\tilde{V}$ /:Beg	.006	[.005, .008]	11.54	< .001 ***
Contrast. / $\tilde{V}\tilde{V}$ /:Mid	-.0002	[-.001, .001]	-.42	.67
Contrast. / $\tilde{V}\tilde{V}$ /:Fin	-.006	[-.007, -.005]	-9.76	< .001 ***

Table 5: Results of the LMER model – $\text{lmer}(\text{nasalance} \mid \text{CONDITION} * \text{TIMESTEP} + (1 + \text{CONDITION} \mid \text{speaker}) + (1 \mid \text{word}))$ – predicting nasalance across the final vowel for all 15 speakers in experiment 2. Effect-coded fixed effects include CONDITION (oral /VV/ and contrastive / $\tilde{V}\tilde{V}$ /), TIMESTEP (beginning, middle, and final), and their interaction.

4.2.1.2 Contrastive / $\tilde{V}\tilde{V}$ / vowels trigger nasal harmony

Having determined the approximate nasalance values for contrastive / $\tilde{V}\tilde{V}$ / vowels, we can now assess whether the segments in the preceding VVG sequence exhibit similar nasalance levels for this condition. If so, this provides quantitative evidence attesting to a process of regressive nasal harmony triggered by contrastive / $\tilde{V}\tilde{V}$ / vowels.

To accomplish this, we fit an additional LMER model, comparing the degree of nasalance of the preceding glide and penultimate vowel to that of the final vowel in the contrastive / $\tilde{V}\tilde{V}$ / condition. As before, by-speaker random slopes and intercepts for segment (penult, glide, final) and random intercepts for word were included in the model. Here, however, the three segments

in the [VVG \tilde{V} V] sequence were dummy coded to evaluate the simple main effect of SEGMENT on nasalance.

The output of the model is given in table 6. The intercept corresponds to the mean nasalance of the contrastive / $\tilde{V}\tilde{V}$ / vowel in the final syllable of words like /saavũã/ ‘breaths’ across all tokens and speakers. Results indicate that the mean nasalance of the final vowel is approximately 52%. Both the preceding glide and the penultimate vowel were found to be statistically different in nasalance from the final vowel. However, we calculated the nasalance *RoPE* for the speakers in experiment 2 using the same method as in experiment 1 and found the region of practical equivalence to be $\pm .039$. As is evident, the 95% CIs for both the glide and the penultimate vowel fall entirely within this equivalence region. The upshot of these results is that the degree of nasalance across the [VVG \tilde{V} V] sequence in the contrastive / $\tilde{V}\tilde{V}$ / condition confirms a process of nasal harmony triggered by the word-final / $\tilde{V}\tilde{V}$ / vowel.

Predictor	β	95% CI	t	p
Intercept (final / $\tilde{V}\tilde{V}$ /)	.515	[.498, .532]	60.43	< .001 ***
Glide	.015	[.008, .022]	4.02	= .001 **
Penult /VV/	-.014	[-.024, -.003]	-2.59	= .02 *

Table 6: Results of the LMER model – `lmer(nasalance SEGMENT + (1 + SEGMENT|speaker) + (1|word)` – predicting nasalance across the [VVG \tilde{V} V] sequence of the contrastive / $\tilde{V}\tilde{V}$ / condition for all 15 speakers in experiment 2. The main effect of SEGMENT was dummy-coded, with the final / $\tilde{V}\tilde{V}$ / as the intercept and the glide and penultimate vowels as two additional factor levels.

4.2.2 Nasality of /VVN/ vowels

The crucial condition in this experiment is /VVN/. The GAM curves in Figure 9 show that nasalance on the penultimate vowel does not diverge from the oral /VV/ condition with respect to nasalance until the latter portion of the penultimate vowel, at which point a cline-like increase in nasalance begins and continues through the onset of the /VVN/ vowel. Strikingly, nasalance on the final vowel itself in /VVN/ resembles the level of nasalance found on the contrastive / $\tilde{V}\tilde{V}$ / vowel through most of its duration.

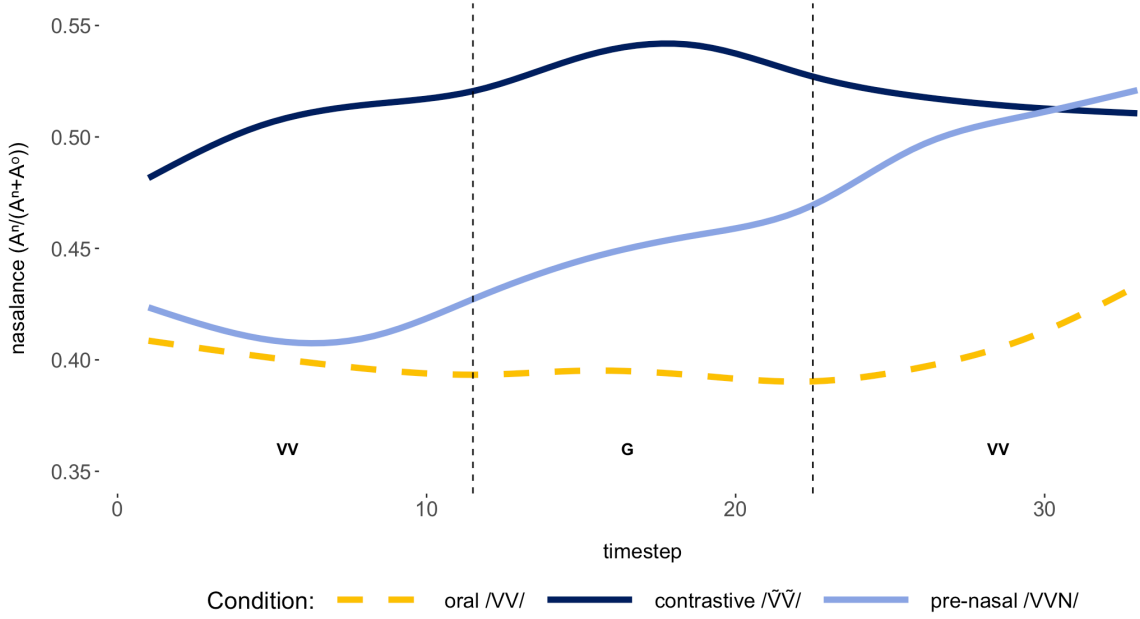


Figure 9: GAM curves for the oral /VV/, contrastive / $\tilde{V}\tilde{V}$ =CV/, and pre-nasal /VVN/ conditions in experiment 2. Curves visualize the mean vowel nasalance at 33 normalized timesteps across the [VVG \tilde{V} VV] sequence for the 13 speakers in the experiment who categorically nasalized the word-final /VVN/ vowel. Vertical dashed lines represent segment boundaries.

To confirm the results from experiment 1 that these /VVN/ vowels are identical in nasalance to contrastive nasal vowels, we first ran a linear mixed effects model (table 7) comparing nasalance on the final vowel between the / $\tilde{V}\tilde{V}$ =CV/ and /VVN/ conditions. CONDITION (2 levels) was included as a main effect, and – due to the visual change in nasalance over time for both target vowels in figure 9 – TIMESTEP (3 levels) and its interaction with CONDITION were also included as fixed effects. As before, the three levels of the TIMESTEP effect roughly correspond to the first third, middle third, and final third of the final vowel in the [VVG \tilde{V} VV] sequence. In addition, including a main effect of VOWELQUALITY (4 levels) did not result in a significantly better fit to the data ($\chi^2(3) = 7.76$, $p > .05$), so VOWELQUALITY was removed from the model. By-speaker random slopes and intercepts were included for condition and random intercepts for word. All levels of the main effects – CONDITION, TIMESTEP – are included in the model output displayed in table 7 for clarity.

As is evident from the final three rows in table 7, the interaction between CONDITION and TIMESTEP is significant, with the effect of CONDITION at each level of TIMESTEP diverging significantly in opposite directions. Still, the effect size at each timestep remains relatively small, so the absence of a significant difference between the nasalance of /VVN/ and contrastive / $\tilde{V}\tilde{V}$ / vowels observed in the first three rows is likely reliable.

Predictor	β	95% CI	t	p
Intercept (grand μ)	.506	[.492, .519]	74.85	< .001 ***
Contrast. / $\tilde{V}\tilde{V}$ /	.009	[-.002, .019]	1.65	.11
/VVN/	-.009	[-.019, .002]	-1.65	.11
Beg	-.008	[-.009, -.007]	-18.17	< .001 ***
Mid	.002	[.001, .003]	4.79	< .001 ***
Fin	.006	[.005, .007]	11.82	< .001 ***
Contrast. / $\tilde{V}\tilde{V}$ /:Beg	.011	[.010, .012]	25.46	< .001 ***
Contrast. / $\tilde{V}\tilde{V}$ /:Mid	-.002	[-.003, -.001]	-4.78	< .001 ***
Contrast. / $\tilde{V}\tilde{V}$ /:Fin	-.009	[-.010, -.008]	-18.27	< .001 ***

Table 7: Results of the LMER model – $\text{lmer}(\text{nasalance } \text{CONDITION} * \text{TIMESTEP} + (1 + \text{CONDITION} | \text{speaker}) + (1 | \text{word})$ – predicting nasalance across the final vowel for 13 speakers who categorically nasalize /VVN/ vowels. Effect-coded fixed effects include CONDITION, TIMESTEP, and CONDITION*TIMESTEP.

Nevertheless, to compare nasalance values between the two conditions more closely, three separate models for each third of the vowel were run with the CONDITION variable dummy-coded to show the simple main effect of CONDITION at each of the three timepoints. The results of the models are summarized in table 8.⁵ The first two pairwise comparisons show that no significant difference in nasalance was found on the final and middle portions of the final vowel between the / $\tilde{V}\tilde{V}$ =CV/ and /VVN/ conditions. Furthermore, the 95% CI for the difference in means for both the middle and final thirds of the vowel falls entirely inside the bounds of the interval of practical equivalence established for experiment 2 ($\pm .039$). This confirms, therefore, that the two conditions are equivalent in the final two-thirds of the target vowel.

Predictor	β	95% CI	t	p
Intercept (/ $\tilde{V}\tilde{V}$ / Fin 1/3)	.511	[.493, .528]	58.65	< .001 ***
/VVN/ Fin 1/3	.003	[-.015, .022]	.37	.72
Intercept (/ $\tilde{V}\tilde{V}$ / Mid 1/3)	.508	[.494, .521]	73.95	< .001 ***
/VVN/ Mid 1/3	-.007	[-.017, .004]	-1.22	.23
Intercept (/ $\tilde{V}\tilde{V}$ / Beg 1/3)	.518	[.502, .535]	61.79	< .001 ***
/VVN/ Beg 1/3	-.041	[-.065, -.016]	-3.19	< .005 **

Table 8: Results of three LMER models predicting nasalance across the final, middle, and beginning positions of the final vowel for 13 speakers who categorically nasalize /VVN/ vowels. CONDITION (/ $\tilde{V}\tilde{V}$ / and /VVN/) was included as a dummy-coded fixed effect.

The results for the first third of the vowel show a significant difference in nasalance between the two conditions. As we argue in the next section, this is likely a consequence of the fact that

⁵Because three models are being run simultaneously, the reader may want to interpret results with a Bonferroni-corrected 95% alpha value of .017. Nevertheless, because the research question at hand is attempting to demonstrate that there is no significant difference between the two vowel types, interpreting the results without a Bonferroni correction is more conservative. That is, a higher alpha value of .05 makes the task of proving equivalence more stringent, biasing the results against our hypothesis. We leave it to the reader to determine which avenue of interpretation to pursue.

nasal harmony applies in contrastive /VVG $\tilde{V}\tilde{V}$ / \rightarrow [$\tilde{V}\tilde{V}.\tilde{G}\tilde{V}\tilde{V}$], but not in pre-nasal /VVG $\tilde{V}\tilde{V}$ N/ \rightarrow [VV.G $\tilde{V}\tilde{V}$ N]. If the glide is fully nasalized in [$\tilde{V}\tilde{V}.\tilde{G}\tilde{V}\tilde{V}$] but not in [VV.G $\tilde{V}\tilde{V}$ N], then we expect to find evidence of an oral-to-nasal transition at the beginning of the final vowel in [VV.G $\tilde{V}\tilde{V}$ N]. Taking that effect into account, it seems reasonable to conclude that contrastive / $\tilde{V}\tilde{V}$ / vowels and pre-nasal /VVN/ vowels are equivalent in terms of their surface nasalance, not only in monosyllabic tokens (experiments 1) but in longer words as well. Apparent differences between them arise from different patterns of local, anticipatory nasal coarticulation with the preceding glide.

To sum up, pre-nasal /VVN/ and contrastive / $\tilde{V}\tilde{V}$ / vowels do not show a statistically significant difference in nasality for the middle and final thirds of the final vowel. Furthermore, while the beginning third of / $\tilde{V}\tilde{V}$ / vowels was significantly more nasalized than the same portion of /VVN/ vowels, the nasality vs. orality of the preceding glide in [$\tilde{V}\tilde{V}.\tilde{G}\tilde{V}\tilde{V}$] vs. [VV.G $\tilde{V}\tilde{V}$ N] explains this effect.

4.2.2.1 Nasal Harmony

Having confirmed that non-alternating /VVN/ vowels and contrastive / $\tilde{V}\tilde{V}$ / vowels exhibit similar nasality in the tokens used in experiment 2, we can now quantitatively examine whether the two vowel types differ in how they trigger nasal harmony.

From a cursory examination of the GAM curves in figure 9 above, it seems clear that the two vowel types – non-alternating, pre-nasal /VVN/ and underlying contrastive / $\tilde{V}\tilde{V}$ / – interact in a qualitatively distinct manner with the nasal harmony process. To corroborate these facts quantitatively, we ran two separate models to examine nasalance on the glides and the penultimate vowels in the [VVG $\tilde{V}\tilde{V}$] sequence, respectively.

The first model shows a pairwise comparison between the glides in the /VVG $\tilde{V}\tilde{V}$ N=CV/ and contrastive /VVG $\tilde{V}\tilde{V}$ =CV/ conditions. As will be demonstrated, glides in the /VVG $\tilde{V}\tilde{V}$ N=CV/ condition exhibit a significantly smaller degree of nasalance than glides in the contrastive /VVG $\tilde{V}\tilde{V}$ =CV/ condition, consistent with our claims in the preceding section.

The second model provides a pairwise comparison between the penultimate vowels in the oral /VVG $\tilde{V}\tilde{V}$ =CV/ and pre-nasal /VVG $\tilde{V}\tilde{V}$ N=CV/ conditions. While penultimate vowels in the /VVG $\tilde{V}\tilde{V}$ N=CV/ condition statistically differ from penultimate vowels in the oral /VV/ condition, the magnitude of the effect is much smaller than the difference between penultimate vowels in the oral /VV/ and contrastive / $\tilde{V}\tilde{V}$ / conditions. Taken together, the results of these two models suggest that pre-nasal /VVN/ vowels do not trigger nasal harmony on either the penultimate vowel or the preceding glide.

The results of the model comparing glides in the contrastive /VVG $\tilde{V}\tilde{V}$ =CV/ and pre-nasal /VVG $\tilde{V}\tilde{V}$ N=CV/ conditions are given in table 9. The data in this model is limited to the thirteen speakers who categorically nasalize pre-nasal /VVN/ vowels in experiment 2; see section 4.2.2.3 for discussion of the remaining two speakers.

The model's intercept is the final / $\tilde{V}\tilde{V}$ / vowel in the / $\tilde{V}\tilde{V}$ =CV/ condition. Contrastive / $\tilde{V}\tilde{V}$ / vowels were included in the model in order to compare the glides in the two conditions to a segment that is clearly categorically [+NAS]. Doing so provides a clearer picture as to whether /VVN/ vowels are triggering nasal harmony.

The glide in the contrastive /VVG $\tilde{V}\tilde{V}$ =CV/ condition is significantly more nasalized than the contrastive / $\tilde{V}\tilde{V}$ / vowel that follows it (as is evident in figure 11). However, the 95% confidence

interval of the effect size falls within the region of practical equivalence of $\pm 3.9\%$.⁶ This suggests a practical equivalence in the degree of nasalance on the glide and following / $\tilde{V}\tilde{V}$ / in the /VVG $\tilde{V}\tilde{V}$ =CV/ condition.

The slightly increased nasalance on the glide in contrastive [$\tilde{V}\tilde{V}$.G $\tilde{V}\tilde{V}$] relative to nasal vowels is likely due to the fact that glides have narrower constrictions in the oral cavity than vowels. All else being equal, narrowed oral constrictions will force a greater proportion of the overall airflow through the nasal cavity (e.g. Bell-Berti 1993, Delvaux et al. 2008: 595). In any case, these results provide clear support for the claim that nasal harmony in /VVG $\tilde{V}\tilde{V}$ =CV/ \rightarrow [$\tilde{V}\tilde{V}$.G $\tilde{V}\tilde{V}$ =CV] strings nasalizes the preceding glide.

For glides in the pre-nasal /VVG $\tilde{V}\tilde{V}$ N=CV/ \rightarrow [VV.G $\tilde{V}\tilde{V}$ N=CV] condition, on the other hand, mean nasalance falls significantly below the mean nasalance of the contrastive / $\tilde{V}\tilde{V}$ / vowels in the /VVG $\tilde{V}\tilde{V}$ N=CV/ condition (by 7.0%). The 95% confidence interval for the difference lands entirely outside the *RoPE* established for experiment 2. We conclude that, unlike glides preceding contrastive / $\tilde{V}\tilde{V}$ / vowels, glides preceding a non-alternating /VVN/ vowel are not targets of nasal harmony.

Predictor	β	95% CI	t	p
Intercept (Final / $\tilde{V}\tilde{V}$ / in / $\tilde{V}\tilde{V}$ =CV/)	.516	[.498, .533]	59.06	< .001 ***
glide (/ $\tilde{V}\tilde{V}$ =CV/ condition)	.015	[.007, .023]	3.88	.002 **
glide (/VVN=CV/ condition)	-.070	[-.100, -.041]	-4.63	< .001 ***

Table 9: Results of the LMER model predicting nasalance across the glide in the [VVG $\tilde{V}\tilde{V}$] sequence for 13 speakers. CONDITION was included as a dummy-coded fixed effect. The final / $\tilde{V}\tilde{V}$ / in the / $\tilde{V}\tilde{V}$ =CV/ condition serves as the intercept in the model.

We now turn to a model comparing nasalance of the penultimate vowels in the pre-nasal /VVG $\tilde{V}\tilde{V}$ N=CV/, oral /VVG $\tilde{V}\tilde{V}$ =CV/, and contrastive /VVG $\tilde{V}\tilde{V}$ =CV/ conditions. The results of this model are summarized in table 10. Here, penultimate oral vowels /VV/ serve as the intercept: the goal is to determine whether penultimate vowels in the pre-nasal /VVG $\tilde{V}\tilde{V}$ N=CV/ condition significantly differ from those in the oral /VVG $\tilde{V}\tilde{V}$ =CV/ condition.

The model results indicates that the penultimate vowel in the contrastive /VVG $\tilde{V}\tilde{V}$ =CV/ condition is about 15.2% more nasalized than those in the oral /VVG $\tilde{V}\tilde{V}$ =CV/ condition. The penultimate vowel in the pre-nasal /VVG $\tilde{V}\tilde{V}$ N=CV/ condition is also significantly different from the oral /VVG $\tilde{V}\tilde{V}$ =CV/ condition, but only by approximately 3.3%. So while there is some evidence of nasalization on penultimate vowels in pre-nasal /VVG $\tilde{V}\tilde{V}$ N=CV/, the amount of nasalization relative to oral /VVG $\tilde{V}\tilde{V}$ =CV/ is much smaller than the amount of nasalization found for contrastive /VVG $\tilde{V}\tilde{V}$ =CV/ \rightarrow [$\tilde{V}\tilde{V}$.G $\tilde{V}\tilde{V}$ =CV].

Predictor	β	95% CI	t	p
Intercept (penult oral VV)	.373	[.344, .402]	25.31	< .001 ***
penult VV (/VVN=CV/)	.033	[.006, .06]	2.42	.02 *
penult VV (/ $\tilde{V}\tilde{V}$ =CV/)	.152	[.124, .179]	10.69	< .001 ***

⁶We assume here that the nasalance *RoPE* for nasalized glides is equivalent to that of nasal vowels, but more research is needed to confirm this assumption.

Table 10: Results of the LMER model predicting nasalance across the penultimate vowel in the [VVG \tilde{V} V] sequence for 13 speakers who categorically nasalize pre-nasal /VVN/ vowels in the [VVG \tilde{V} VN=CV] condition. CONDITION (penultimate vowel (oral), penultimate vowel (/VVN/), and penultimate vowel [+nas]) was included as a dummy-coded fixed effect. The intercept is the penultimate vowel in the oral CV condition.

To summarize, non-alternating, pre-nasal /VVN/ vowels and contrastive / $\tilde{V}\tilde{V}$ / vowels were not found to be significantly different in terms of nasality when uttered in the final position of a [VVG \tilde{V} V] sequence. Despite the similar surface nasality on the stem-final vowels in /VVN/ and contrastive / $\tilde{V}\tilde{V}$ / conditions, pre-nasal /VVN/ and contrastive / $\tilde{V}\tilde{V}$ / triggered distinct patterns of nasalization across the preceding [VVG] sequence. The segments preceding the contrastive / $\tilde{V}\tilde{V}$ / vowel were significantly nasalized across the entire span, to a similar degree as the / $\tilde{V}\tilde{V}$ / vowel itself. This confirms previous impressionistic descriptions that contrastive / $\tilde{V}\tilde{V}$ / vowels in Panjabi trigger leftward nasal harmony across [VVG] sequences, affecting both preceding vowels and glides (albeit optionally).

Conversely, [VVG] sequences preceding /VVN/ exhibited a different pattern of nasality. The immediately preceding glide in [VVG \tilde{V} VN] was significantly less nasalized than contrastive / $\tilde{V}\tilde{V}$ /; in contrast, glides preceding the final contrastive / $\tilde{V}\tilde{V}$ / in [$\tilde{V}\tilde{V}$ G $\tilde{V}\tilde{V}$] exhibited similar nasalance to contrastive / $\tilde{V}\tilde{V}$ /. This is consistent with our claim that glides are nasalized by harmony when they precede contrastive / $\tilde{V}\tilde{V}$ /, but not when they precede pre-nasal /VVN/ \rightarrow [$\tilde{V}\tilde{V}$ N].

For the pre-nasal [VVG \tilde{V} VN] condition, nasalance is low on the glide [G], and apparently even lower on the penultimate vowel (figure 9). In other words, nasality steadily declines in strength for segments which are further away from the [$\tilde{V}\tilde{V}$ N] sequence which we analyze as the only phonologically nasalized segments in the [VVG \tilde{V} VN] string.

Nasalance on the penultimate vowel in [VVG \tilde{V} VN] was found to be significantly different from nasalance on the penultimate, pre-glide vowel in the oral condition /VVG \tilde{V} V=CV/. However, this difference is small in magnitude, only approximately 3.3%. These results suggest that the pre-nasal vowel in /VVN/ \rightarrow [$\tilde{V}\tilde{V}$ N] is itself phonologically [+NAS], but does not trigger nasal harmony across the preceding glide and penultimate vowel in [VVG \tilde{V} VN]. Any vestigial nasality on the glide and penultimate vowel appears to reflect lower-level, local coarticulation for nasality rather than nasal harmony as such.

4.2.2.2 Harmony is not an artifact of stress

A potential confound in our interpretation of nasalance values in contrastive [$\tilde{V}\tilde{V}$ G $\tilde{V}\tilde{V}$] vs. pre-nasal [VVG \tilde{V} VN] vowels comes from stress placement. Stress has been described as weight sensitive in Panjabi, though in somewhat inconsistent terms, as we discuss below. Both syllables in the contrastive [$\tilde{V}\tilde{V}$ G $\tilde{V}\tilde{V}$] condition are bimoraic, whereas, for [VVG \tilde{V} VN], the penultimate syllable is bimoraic, and the final syllable at least plausibly trimoraic. These differences in syllable weight could affect stress placement across conditions. If so, based on some prior descriptions of weight-sensitive stress in Panjabi (e.g. Bashir & Connors 2019), we might expect penultimate stress on [$\tilde{V}\tilde{V}$.G $\tilde{V}\tilde{V}$], and final stress on [VV.G $\tilde{V}\tilde{V}$ N].

Stress has also been found to correlate with increased nasality in some contexts and languages (e.g. Krakow 1989, Avelino et al. 2020, Zellou 2022 and references there). This at least raises the possibility that the increased nasalance on the penultimate vowel and glide in the contrastive [$\tilde{V}\tilde{V}$ G $\tilde{V}\tilde{V}$] condition may not represent nasal harmony, but could instead result from long-distance

coarticulation for nasality between the stressed penultimate syllable and the final contrastive nasal vowel.⁷

We can try to assess this claim by comparing the mean duration of the penultimate and final vowels across the two conditions of most interest, contrastive [ṼṼG̃ṼṼ] vs. pre-nasal [VVG̃ṼVN]. The penultimate and final vowels in the [ṼṼG̃ṼṼ] condition were 103 ms and 100 ms in mean duration, respectively, compared to a mean duration of 60 ms for the penultimate vowel and 141 ms for the final vowel in the pre-nasal [VVG̃ṼVN] condition. These measurements are at least consistent with the claim that contrastive [ṼṼG̃ṼṼ] and pre-nasal [VVG̃ṼVN] differ with respect to stress placement, stress being penultimate in the former case, and final in the latter.

Nevertheless, there are several reasons why the potential confound of stress is unlikely to be the cause of the nasality differences between contrastive [ṼṼG̃ṼṼ] and pre-nasal [VVG̃ṼVN]. First, no phonetic analyses of the acoustic correlates of stress exist for Panjabi, so it is unclear whether duration is a good indicator of stress location. In fact, Bhatia (1993) states that “vowel quality and length remain unaffected by any difference in stress placement”, which indicates, at least impressionistically, that duration is *not* a phonetic correlate to stress in Panjabi. Furthermore, of five different impressionistic descriptions of Panjabi stress (Bhatia 1993, Shackle 2003, Dhillon 2007, Bashir & Connors 2019, Hussain et al. 2019), no two descriptions provide the same generalizations. To cite one relevant example, while Bashir & Connors (2019: 50), Hussain et al. (2019: 293-294), and Dhillon (2007: 84-86) place stress on the final syllable in [VV.VVC] words, Bhatia (1993: 343) and Shackle (2003: 592) describe stress as falling on the penultimate syllable. Suffice it to say that if stress is a confounding factor in the present study, its perceptual salience seems to be rather weak, which also hints that its correlates may be as well.

In at least some languages where stress conditions greater nasal coarticulation, it has been reported that there is a positive correlation between the duration of the stressed vowel and the degree of nasalance on that vowel (e.g. discussion in Zellou 2022: 21, 35). Alternatively, if the nasal harmony pattern in the contrastive [ṼṼG̃ṼṼ] condition reflects a nasal coarticulation window of approximately fixed duration (sections 4.1, 4.2.3), we would expect to find a negative correlation between the duration of the penultimate vowel and the amount of nasalization on that vowel in [ṼṼG̃ṼṼ] (Solé 1992, 1995): as the penultimate vowel gets longer, it should be less affected by nasal coarticulation with the final vowel.

However, no such correlation, positive or negative, is present in our data. Considering all forms in the contrastive [ṼṼG̃ṼṼ] conditions (including those that do not exhibit categorical nasal harmony), a Pearson’s correlation analysis ($r(871) = .061$, $p = .07$) between penultimate vowel duration and mean nasalance on the penultimate vowel did not detect a significant effect (figure 10).

⁷To our knowledge, reported effects of stress on nasal coarticulation involve local interactions between [NV] and [VN] sequences, rather than the kind of longer-distance effects we are entertaining here. We nonetheless consider this scenario in the interest of ruling it out as a possible counter-explanation, particularly since nasal coarticulation is known to occur over a wide timespan in some languages (e.g. Pouplier et al. 2024).




Figure 10: Scatterplot showing the estimated .061 Pearson Correlation value between penultimate vowel duration and mean nasalance in the contrastive $[\tilde{V}\tilde{V}]$ conditions.

If the nasal harmony pattern observed in the $[\tilde{V}\tilde{V}\tilde{G}\tilde{V}\tilde{V}]$ conditions reflected long-distance coarticulation, we would expect at least *some* association between nasalance and vowel duration, be it positive or negative. The lack of such a correlation argues against treating nasal harmony as a phonetic process of nasal coarticulation, rather than a phonological process of nasal spreading. (We return to this point in section 4.2.3.)

In sum, the incongruent descriptions of Panjabi stress in prior literature, the uncertainty about the phonetic correlates of stress in this language, and the finding that nasalance on the penultimate vowel in $[\tilde{V}\tilde{V}\tilde{G}\tilde{V}\tilde{V}]$ is insensitive to duration all argue against re-analyzing nasal harmony as a stress-sensitive process of long-distance nasal coarticulation. While further work on the phonetics of stress in Panjabi would of course be welcome, at present we see no reason to treat nasality on the penults of contrastive $[\tilde{V}\tilde{V}\tilde{G}\tilde{V}\tilde{V}]$ as anything other than the result of a categorical process of phonological nasal harmony (see also section 4.2.3).

4.2.2.3 Nasalization in Lahori Panjabi

In contrast to the above results, GAM curves across the $[VVG\tilde{V}\tilde{V}]$ sequence for the two speakers that produced $/VVN/$ vowels with anticipatory coarticulation rather than categorical nasality are displayed in figure 11. As can be seen by the curves, nasalance across the penultimate vowel and glide for the $/VVN/$ condition is proportionally similar to the oral $/VV/$ condition, but a cline-like increase in nasalance begins near the onset of the $/VVN/$ vowel, only nearing the nasalance level of the $[\tilde{V}\tilde{V}=CV]$ condition at the offset of the $[VVG\tilde{V}\tilde{V}]$ sequence. This pattern suggests that the $/VVN/$ vowel is not intentionally targeted as $[+NAS]$ by these speakers, which is quite distinct from all other speakers in either experiment 1 or experiment 2. Crucially, however, the two speakers that exhibited this pattern grew up in Lahore, a large city of about 11 million people

located on the eastern border of Pakistan, and no other speakers in either experiment grew up in this region. While further research is needed, we speculate that the dialect of Panjabi in Lahore displays a distinct pattern of nasality on /VVN/ vowels, such that they do not surface as [+NAS] vowels but instead are coarticulatorily nasalized as a result of the following nasal consonant. We leave it to future research to explore the veracity of the assumption that Lahori Panjabi does indeed represent a distinct dialect in terms of its nasality processes.




Figure 11: GAM curves for the oral /VV/, contrastive / $\tilde{V}\tilde{V}$ =CV/, and /VVN/ conditions. Curves visualize the mean vowel nasalance at 33 normalized timesteps across the [VVG $\tilde{V}\tilde{V}$] sequence for the two speakers that exhibit coarticulatory nasality on /VVN/ vowels.

4.2.3 Against a surface-oriented alternative analysis: fixed nasalization windows

We have argued that nasality spreads leftward to the penult in contrastive /VVG $\tilde{V}\tilde{V}$ / \rightarrow [ṼṼG̃ṼṼ] in Panjabi, but not in pre-nasal /VVG $\tilde{V}\tilde{V}\tilde{N}$ / \rightarrow [VVG̃ṼṼN]. We have further argued that spreading is best construed as a phonological harmony process which propagates the abstract feature [+NAS] leftward from contrastively nasal vowels in / $\tilde{V}\tilde{V}$ /, but not from predictable, contextually nasal vowels in /VVN/ \rightarrow [ṼṼN].

As described in sections 4.1 and 4.2.3, an alternative, surface-oriented analysis of this data might be possible in which nasalization in contrastive /VVG $\tilde{V}\tilde{V}$ / \rightarrow [ṼṼG̃ṼṼ] owes to a coarticulatory window in which the velum is lowered for a more or less fixed period of time. In pre-nasal /VVG $\tilde{V}\tilde{V}\tilde{N}$ / \rightarrow [VVG̃ṼṼN], that window would shift rightward to align with the right edge of the nasal consonant; as a consequence, nasalization could also shift off of the preceding glide and penultimate vowel.

There are at least three arguments against this analysis. The first comes from the optionality of nasal harmony in Panjabi (section 4.2). Leftward nasal spreading in contrastive /VVG $\tilde{V}\tilde{V}$ / often yields nasalized [ṼṼG̃ṼṼ], but outcomes in which harmony fails to apply [VVG̃ṼṼ] are also

frequent across all items and speakers. If there is a fixed period of velar lowering in all words that end in a nasal segment, then the non-application of nasal harmony in [VVG \tilde{V} \tilde{V}] is quite puzzling. On the other hand, if nasal harmony is an abstract phonological process, then categorical variation between nasalized [V \tilde{V} G \tilde{V} \tilde{V}] and non-nasalized [VVG \tilde{V} \tilde{V}] can be straightforwardly analyzed as optional rule application (see also Plug et al. 2019).

Second, in section 4.2.2.2 we observed that there is no correlation between the duration of the penultimate vowel in nasalized [V \tilde{V} G \tilde{V} \tilde{V}] and the degree of nasalance on that vowel. This result is again surprising if harmony owes to a coarticulatory window of approximately fixed duration: in that scenario, we should see a decrease in nasalance as the penultimate vowel grows longer, and more distant from the phonetic source of nasality at the right edge of the word in [V \tilde{V} G \tilde{V} \tilde{V}].

A third argument against this approach comes from the analysis of a condition in our data which we have not yet discussed: contrastive nasal vowels followed by a nasal-initial postposition, /VGG \tilde{V} \tilde{V} =NV/.

If contrastive nasal vowels / \tilde{V} \tilde{V} / are responsible for triggering nasal harmony — understood as a phonological process — the presence of a following /=NV/ postposition should not affect the application of harmony in any way in /VVG \tilde{V} \tilde{V} /. Conversely, if the consistent nasalance in the /VVG \tilde{V} \tilde{V} =CV/ condition is due to the anchoring of a fixed window of nasalization to the right edge of the rightmost nasal segment, we would expect this window to shift to the right to align with the initial nasal consonant of the postposition. In this scenario, nasalization in /VGG \tilde{V} \tilde{V} =NV/ should mirror the nasalance pattern of the /VVG \tilde{V} \tilde{V} N/ condition: there should be no evidence of harmony.

Each item in the /VGG \tilde{V} \tilde{V} =CV/ condition was also produced in a corresponding /VGG \tilde{V} \tilde{V} =NV/ context in our study. This allows for a direct comparison of nasality across the two conditions. The GAM curve for all fifteen speakers for the /VVG \tilde{V} \tilde{V} =NV/ condition is provided in figure 12, alongside the baseline oral /VVG \tilde{V} \tilde{V} / and /VVG \tilde{V} \tilde{V} =CV/ conditions for comparison.

As is evident from figure 12, the /VVG \tilde{V} \tilde{V} =NV/ condition shows a high degree of nasalance across the entire [VVG \tilde{V} \tilde{V}] sequence, closely resembling the nasalance contour for the /VVG \tilde{V} \tilde{V} =CV/ condition (oral postposition). (The slight difference in nasality at the very end of /VVG \tilde{V} \tilde{V} =CV/ vs. /VVG \tilde{V} \tilde{V} =NV/ undoubtedly reflects local coarticulation with the following /C/ in /VVG \tilde{V} \tilde{V} =CV/ vs. /N/ in /VVG \tilde{V} \tilde{V} =NV/.)

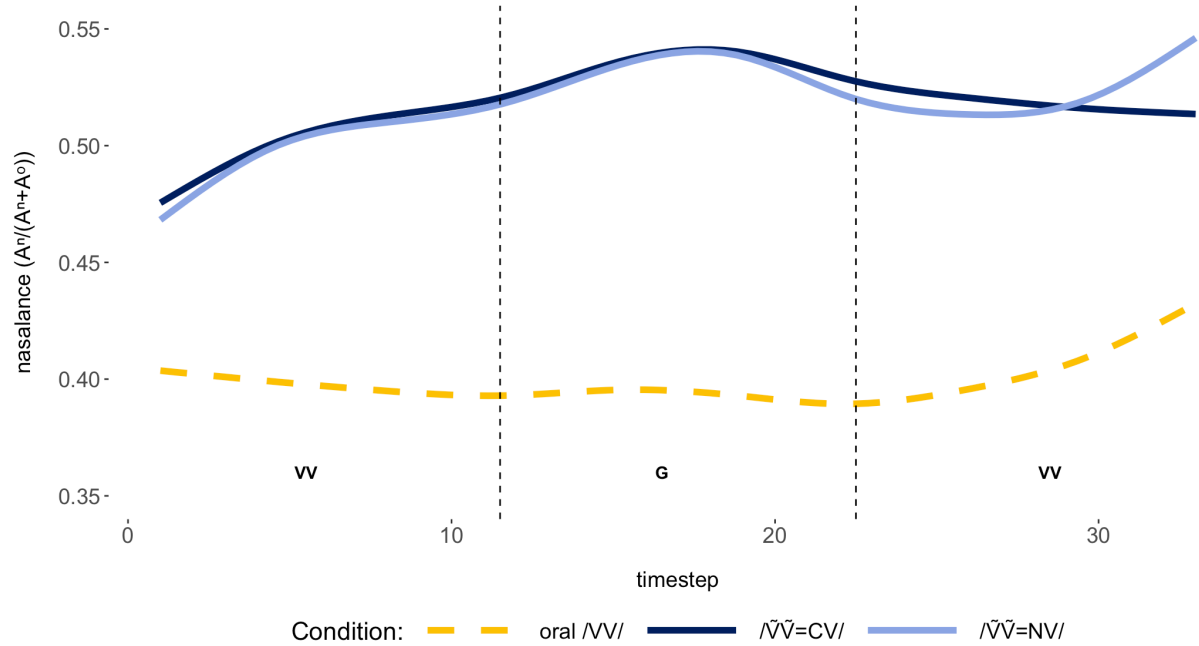


Figure 12: GAM curves comparing the $\tilde{V}\tilde{V}=N\tilde{V}$ condition to the oral $/VV/$ and $\tilde{V}\tilde{V}=CV$ conditions. Curves visualize the mean vowel nasalance at 33 normalized timesteps across the $[VVG\tilde{V}\tilde{V}]$ sequence for each condition for all fifteen speakers.

A pairwise comparison between the penultimate vowels in the $\tilde{V}\tilde{V}=CV$ and $\tilde{V}\tilde{V}=N\tilde{V}$ conditions in table 11 shows that the mean nasalance of the penultimate vowel in the $\tilde{V}\tilde{V}=N\tilde{V}$ condition is not significantly different from the mean nasalance of the penultimate vowel in $\tilde{V}\tilde{V}=CV$. Further, the estimated (but non-significant) difference in nasalance between these conditions is practically insignificant because the 95% CI around the estimate falls within the region of practical equivalence for nasality ($\pm .039$) established for experiment 2.

Predictor	β	95% CI	t	p
Intercept (Penult VV in $\tilde{V}\tilde{V}=CV$)	.499	[.480, .518]	51.05	< .001 ***
Penult VV in $\tilde{V}\tilde{V}=N\tilde{V}$	-.004	[-.022, .014]	-.42	.68

Table 11: Results of the LMER model predicting nasalance across the penultimate vowel in the $[VVG\tilde{V}\tilde{V}]$ sequence for 15 speakers. CONDITION (penultimate vowel ($\tilde{V}\tilde{V}=CV$) and penultimate vowel ($\tilde{V}\tilde{V}=N\tilde{V}$)) was included as a dummy-coded fixed effect. The intercept is the penultimate vowel in the $\tilde{V}\tilde{V}=CV$ condition.

We conclude that the results of experiment 2 cannot be analyzed by assuming that a nasalization window of approximately fixed length is anchored at the right edge of the rightmost nasal segment. The results are most straightforwardly consistent with a phonologically-controlled, optional harmony process which spreads nasality leftward from contrastive nasal vowels $\tilde{V}\tilde{V}$, but not from derived, predictable pre-nasal vowels in $/VVN/ \rightarrow [\tilde{V}\tilde{V}N]$.

4.3 Summary of experiment 2

This experiment sought to uncover how /VVN/ vowels interact with nasal harmony in Panjabi. While the results in this experiment and experiment 1 confirm that /VVN/ vowels show identical nasality to underlying / $\tilde{V}\tilde{V}$ / for most speakers, section 4.2.1 showed that – in support of Hypothesis 2 mentioned at the beginning of this section – they do not trigger nasal harmony like contrastive / $\tilde{V}\tilde{V}$ / vowels. Section 4.2.3 introduced nasal-initial postpositions after contrastive / $\tilde{V}\tilde{V}$ / vowels to examine whether the presence of a following nasal consonant would alter the purported nasal harmony pattern. Findings demonstrated that contrastive / $\tilde{V}\tilde{V}$ / vowels consistently triggered nasal harmony regardless of whether the following postposition was oral- or nasal-initial, suggesting that an underlying / $\tilde{V}\tilde{V}$ / vowel’s ability to induce nasal harmony is not affected by the presence of a following nasal consonant.

To sum up the results in terms of the two proposed hypotheses at the beginning of this section, despite the same surface realization of non-alternating /VVN/ and contrastive / $\tilde{V}\tilde{V}$ / vowels, the two vowel types act distinctly in terms of their phonology. Hypothesis 1, which posits that contrastive / $\tilde{V}\tilde{V}$ / and /VVN/ vowels should be phonologically identical because of their surface realizations, does not align with these results. Conversely, Hypothesis 2, which permits a divergence in the phonological behavior of contrastive / $\tilde{V}\tilde{V}$ / and /VVN/ vowels via distinct URs, provides a sensible explanation of the findings.

5 Discussion

5.1 Analyzing /VVN/ vowels in Panjabi with abstract URs

The above experiments established several facts about non-alternating /VVN/ vowels in Panjabi. First, these vowels are phonologically nasal on the surface. Their realization mirrors the categorical nasality of contrastive / $\tilde{V}\tilde{V}$ / vowels in both experiments, and changes in speech rate do not impact the degree of nasalization on /VVN/ \rightarrow [$\tilde{V}\tilde{V}N$]. Nonetheless, derived nasal vowels in /VVN/ \rightarrow [$\tilde{V}\tilde{V}N$] differ from contrastive nasal vowels in / $\tilde{V}\tilde{V}$ / in failing to trigger leftward nasal harmony: /VVG $\tilde{V}\tilde{V}$ / \rightarrow [$\tilde{V}\tilde{V}.\tilde{G}\tilde{V}\tilde{V}$], but /VVG $\tilde{V}\tilde{V}N$ / \rightarrow [$\tilde{V}\tilde{V}.\tilde{G}\tilde{V}\tilde{V}N$], *[$\tilde{V}\tilde{V}.\tilde{G}\tilde{V}\tilde{V}N$] (experiment 2).

The interaction between nasal harmony and /VVN/ vowels can be handled straightforwardly within any framework that (i) permits covert URs, and (ii) has mechanisms for generating counterfeeding opacity. Here, we provide an analysis using ordered rewrite rules (e.g. Chomsky & Halle 1968), but various other approaches are also viable. We gloss over the optionality of nasal harmony, but this simplification has no consequences for our overall analysis.

The derivation in (3) produces an opaque interaction between nasal harmony and the contextual nasalization of pre-nasal vowels in /VVN/ \rightarrow [$\tilde{V}\tilde{V}N$] (which we dub ‘assimilation’ in (3)). In (3a), the word-final nasal vowel / $\tilde{a}\tilde{a}$ / is present underlyingly. Harmony is the first rule to apply in the derivation, so it can only be sensitive to underlying structure. As such, harmony is triggered by the final / $\tilde{a}\tilde{a}$ / in /saa $\tilde{u}\tilde{a}\tilde{a}$ / \rightarrow [saa $\tilde{u}\tilde{a}\tilde{a}$], as desired.

- (3) Counterfeeding between nasal harmony and nasal assimilation in Panjabi

		a. /saavãã/	b. /aavaam/
Harmony:	$[-CONS] \rightarrow [+NAS]/_X_{0[-CONS]}V_{[+NAS]}$	sããvãã	-
Assimilation:	$V \rightarrow [+NAS]/_C_{[+NAS]}$	-	aavããm
		[sããvãã]	[aavããm]

In (3b), the /VVN/ vowel is underlyingly oral, /aavaam/. Since harmony is the first rule to apply, it cannot ‘know’ that the final vowel of /aavaam/ will eventually be nasalized by the assimilation rule in (3) – harmony is only sensitive to underlying nasality, not derived, predictable nasality. As such, harmony fails to apply to /aavaam/. The underlying form /aavaam/ subsequently undergoes assimilation, producing the surface form [aavããm]. This is a textbook case of counterfeeding opacity: assimilation applies ‘too late’ in the derivation to feed nasal harmony (e.g. Kiparsky 1976, Baković 2011).

To reiterate a crucial aspect of our analysis, underlying representations like /aavaam/ are *covert* in Panjabi, because pre-nasal vowels in non-alternating morphemes never actually surface as oral. The use of covert URs is a centerpiece of the analysis in (3): it is the key to distinguishing between derived (predictable) and contrastive (unpredictable) nasality for the purposes of determining when nasal harmony applies.

5.2 Panjabi nasal opacity: a challenge for surface-oriented frameworks

5.2.1 Opacity in exemplar models

A primary concern of generative phonology has been the modeling of productive, general alternations in phonological systems. Consider the famous ‘Bach test’ of Halle et al. (1978): asked to pluralize the proper name *Bach* [bax], English speakers will uniformly provide [bax-s], as opposed to any of the other plural forms conceivably suggested by the vocabulary of English (*[bax-z], *[bax-ɪz], *[baɣ-z], *[bax-en], etc.). This is a PRODUCTIVE pattern of generalization inasmuch as (i) speakers may have never heard, or themselves produced, the plural of *Bach* before, and (ii) the velar fricative [x] does not normally occur in English, so speakers have little or no direct evidence about how to pluralize forms ending in [x]. Traditionally, productive phonological generalization of this type has been modeled with abstract, feature based rules or other formal mechanisms (e.g. the voicing assimilation rule $[-SON] \rightarrow [\alpha VOI]/[-SON, \alpha VOI]_\#$; Halle et al. 1978).

Usage-based models like exemplar theory have not been equally concerned with modeling productive phonological generalizations and alternations of this type, though work in this vein does exist (see e.g. Pierrehumbert 2006, Kirchner et al. 2010, Kirchner 2011, Kaplan 2015, Pierrehumbert 2016, Goldrick & Cole 2023).⁸ This is not to suggest that exemplar models lack mechanisms for productive generalization (as Pierrehumbert 2016 emphasizes). We are simply noting that the issue has received less attention in the exemplar theory literature than in generative phonology. Relatedly, the precise mechanisms that determine productive generalization in exemplar theoretic models remain a somewhat open question (e.g. how, exactly, ‘similarity’ is calculated between novel and stored forms for the purposes of pattern extension; how statistical

⁸An infamous exception is the ‘English past-tense debate’; see e.g. Pinker (2002), Seidenberg & Plaut (2014). We note that this is a debate about morphological patterning (with phonological conditioning, Albright & Hayes 2003), where here we are concerned with purely phonological patterns.

regularities are actually computed for the purposes of learning phonotactics, etc. (e.g. Albright & Hayes 2003, Becker et al. 2011, Daland et al. 2011, Becker et al. 2012, Gorman 2013, Gouskova & Becker 2013, Becker et al. 2017).

Perhaps more to the point, we are unaware of any attempt within the exemplar theory literature to model regular, opaque interactions like the counterfeeding interaction between nasal assimilation and nasal harmony in Panjabi discussed in (3) (see also Kawahara 2001, Ettlinger 2008, 2009). For that reason, our assessment of exemplar theory and related models in this section is based on our own understanding of the framework, rather than any specific, previous proposals. Our focus is on ‘eliminativist’ exemplar theories – that is, theories which dispense with discrete, symbolic abstraction of the speech stream (Goldinger 1998, Goldrick & Cole 2023).

Several mechanisms have been proposed for producing some form of abstraction in exemplar theory. The first mechanism is at the very core of the theory: the storage of many, phonetically-detailed memory traces for individual lexical items (section 1). If production and perception are mediated by some form of averaging over these traces, the result is abstraction: any single production target, or perceptual classification, will reflect the aggregate structure of the exemplar cloud, rather than any single, specific, stored concrete exemplar (e.g. Goldinger 1998, Pierrehumbert 2001, 2002, 2003, Wedel 2004, 2006). This is a species of abstraction.

As we have argued above, surface, phonetic nasality is functionally equivalent in contrastive / $\tilde{V}\tilde{V}$ / and pre-nasal /VVN/ \rightarrow [$\tilde{V}\tilde{V}N$] in Panjabi. Consequently, no form of averaging over stored exemplars will produce an abstract difference between contrastive / $\tilde{V}\tilde{V}$ / and pre-nasal /VVN/ \rightarrow [$\tilde{V}\tilde{V}N$] as such. We conclude that this first mechanism for abstraction in exemplar theory – averaging over stored exemplar clouds – doesn’t provide any insight into why contrastive / $\tilde{V}\tilde{V}$ / triggers harmony, but pre-nasal /VVN/ \rightarrow [$\tilde{V}\tilde{V}N$] does not.

A second mechanism for producing abstraction in exemplar theory comes from phonetic influence between morphologically and/or semantically related forms. For example, in section 1 we discussed incomplete neutralization in German: the vowel in / ʁa:d / \rightarrow [ʁa:t] ‘wheel’ is slightly longer, phonetically, than the vowel in / ʁa:t / \rightarrow [ʁa:t] ‘council’. In exemplar models, the slight lengthening of the vowel in [ʁa:t] ‘wheel’ can be attributed to influence from the plural form [ʁɛ:d-ɐ] ‘wheels’, where the vowel is transparently long as a result of a general process of vowel lengthening before voiced obstruents.

Incomplete neutralization is a kind of phonetic opacity: singular [ʁa:t] ‘wheel’ opaquely inherits the vowel length of plural [ʁɛ:d-ɐ] ‘wheels’, where lengthening is phonetically transparent. (This bears a close resemblance to the generative mechanism of OUTPUT-OUTPUT FAITHFULNESS, discussed in section 5.2.2.) The question naturally arises whether connections between morphologically related forms can also be leveraged to model the opaque non-application of harmony in forms ending in /VVN/ \rightarrow [$\tilde{V}\tilde{V}N$].

We believe the answer is no: the requisite morphological relations are simply not present for the /VVN/ \rightarrow [$\tilde{V}\tilde{V}N$] items in our study. Words like [ədʒəʋāāŋ] ‘omum seed’ are monomorphemic roots. To force the opaque non-application of harmony in [ədʒəʋāāŋ] through phonological or phonetic connections with morphologically related forms, there must be an item like [ədʒəʋaa], in which harmony *transparently* fails to apply, because there is no nasal vowel to trigger it. However, no such form exists, because [ədʒəʋāāŋ] is a non-decomposable root: it always ends in [$\tilde{V}\tilde{V}N$] on the surface. (We elaborate on this point in more detail in section 5.2.2.)

For suffixed forms like [ləɛ-jāāŋ] ‘bring-NOML’, one could appeal to influence from a non-nasalized, related form of the stem vowel ([ləɛ-jaa] ‘bring-IMP’) to explain why harmony fails to

apply in [lɛɛ-jǎǎn]. However, this explanation also falls short, because it equally — and again, incorrectly — predicts that nasal harmony should be inhibited in suffixed forms containing a *contrastive* nasalized vowel in the suffix, such as [mætʃii-vaq] ‘fisherman’ vs. [mætʃii-jǎǎ], *[mætʃii-jǎǎ] ‘fish-PLURAL’.

A third potential mechanism for producing opacity in /VVGVVN/ → [VV.GǎǎVN], *[ǎǎ.GǎǎVN] is to invoke influence between *phonologically* related forms. Research in exemplar theory has sometimes assumed that phonologically similar forms can influence each other, through their gestalt similarity, regardless of their morphological (non-)relatedness (e.g., Bybee 2001: p.22). Concretely, if pre-nasal /VVGVVN/ → [VV.GǎǎVN] is influenced by *oral* [VV.GVV], that influence could inhibit harmony in the penult of [VV.GǎǎVN], in order to maintain consistency in the nasality of the penultimate vowel.

The obvious question here is why pre-nasal /VVGVVN/ → [VV.GǎǎVN] should be influenced by oral [VV.GVV], specifically. Indeed, forms with contrastive nasal vowels, [ǎǎ.GǎǎV], are at least equally plausible candidates for influencing pre-nasal [VV.GǎǎVN] through some notion of overall phonological similarity. Privileging oral [VV.GVV] for this role is not well-motivated. Furthermore, it simply sneaks covert underlying representations in through the back door, by implicitly admitting that pre-nasal /VVGVVN/ → [VV.GǎǎVN] is in some sense more closely related to oral [VV.GVV] than contrastively nasal [ǎǎ.GǎǎV] (see also (McCarthy 1999) on sympathy theory).

Fourth, it is possible that the distinction between pre-nasal /VVGVVN/ → [VV.GǎǎVN] and contrastive [ǎǎ.GǎǎV] is simply an artifact of historical change. Presumably, pre-nasal assimilation developed *after* the emergence of nasal harmony, thus yielding an opaque synchronic interaction between the two processes. Speakers might simply memorize forms like [mætʃii-jǎǎ] ‘fish-PLURAL’ vs. [aav-ǎǎm] ‘PLURAL-citizen’. This would absolve the theory of any need to model the interaction between nasal harmony and local nasal assimilation.

The problem with this analysis is a familiar one: it fails to account for the productivity and generality of both nasal harmony and local nasal assimilation. As far as we are aware, these are completely regular processes in the varieties of Panjabi analyzed here. Further, they are both triggered by productive affixes, such as plural [-jǎǎ] or the nominalizer [-jǎǎn] (e.g., [lɛɛ-jǎǎn] ‘bring-NOML’). While we have not carried out wug-testing, we feel confident in assuming that these suffixes would evince the same phonological behavior even when attached to completely novel stems. Something must ensure this; simple rote memorization will not do the job.

Lastly, the patterns of nasal harmony described in this paper could be accounted for with a more complex co-occurrence restriction on nasality, *[ǎǎVN], which bans a nasalized glide before a nasalized (long) vowel, followed by a nasal consonant. The trigram constraint *[ǎǎVN] straightforwardly excludes unattested *[ǎǎǎVN] sequences while allowing [ǎǎǎVN]. Furthermore, it does so with reference to surface forms alone: this renders the lack of *[ǎǎǎVN] sequences transparent, rather than opaque, and thus vitiates the need for abstract underlying representations. Within an exemplar-theoretic framework, a constraint like *[ǎǎVN] could plausibly emerge from distributional learning over the Panjabi lexicon, based on surface forms alone.

However, two central issues have been raised in the literature concerning the viability of such complex co-occurrence constraints. The first is computational: the cost of learning constraints over segmental sequences quickly becomes intractable as the length of those sequences increases (Hayes & Wilson 2008: 391-392). Constraints are stated over windows of length *n*, where *n* is some number of segments or other symbols (e.g. *[g#], *[VtV], etc.). As *n* increases, the hypothesis

space of possible constraints expands exponentially (the exact rate of growth depends on how many natural classes constraints may refer to). This makes it increasingly difficult for a learner to efficiently navigate the constraint space as constraints expand their segmental scope. For example, Hayes & Wilson (2008) argue that constraints on segment combinations must generally be limited to $n = 2$; they relax this to $n = 3$ for constraints which refer only to the major class features [syllabic], [consonantal], and [sonorant]. The constraint *[$\check{G}\check{V}\check{V}N$] of course goes beyond these limitations.

The second, and perhaps more pressing, concern is empirical. Allowing the learner to posit highly specific co-occurrence restrictions based solely on attested patterns in the lexicon makes it difficult—if not impossible—to distinguish between principled phonotactic gaps and accidental ones (Chomsky & Halle 1965, Hayes & White 2013, Wilson & Gallagher 2018). Principled gaps are those which native speakers judge to be ill-formed, whereas accidental gaps are unattested but fully acceptable to speakers. A statistical learner, lacking a mechanism for making this distinction, would be just as likely to form constraints that exclude acceptable but unattested forms as it would be to discover true phonological generalizations.

To test the extent to which such overgeneration might occur in Panjabi, we examined phonotactic patterns in the Pakistani Panjabi corpus, *pnbTenTen* (Jakubíček et al. 2013), a 2.8 million-word corpus compiled from 5,351 webpages in 2017 written in Shahmukhi script, the standard orthography for Pakistani Panjabi. The corpus contains approximately 100,000 unique words of varying frequencies. Using a custom Python script, we extracted the set of 3-gram sequences absent from the corpus in the following manner. First, we extracted the complete set of unique Shahmukhi characters from the corpus and transliterated those characters into their corresponding natural class category: nasal consonant (N), liquid (L), obstruent (O), long oral vowel (VV), or long nasal vowel ($\check{V}\check{V}$). Several complications involved in transliterating Shahmukhi to broad natural classes led us to exclude glides and short vowels from the analysis. First, Shahmukhi is a partial abjad, meaning that short vowels are written using optional diacritics, which are generally omitted from everyday writing and inferred from context. As a result, consistent transliteration of short vowels is not possible. Furthermore, glides in Shahmukhi are not represented by a unique orthographic symbol but are instead written with the same symbol as their corresponding vowel. For example, ج may represent /j/ or /i:/ and و may represent /v/, /u:/, /o:/, or /ɔ:/. The exact context for each symbol-to-sound mapping is complex and not well documented for Shahmukhi. Consequently, we chose to transcribe these symbols consistently as vowels and omit all 3-gram sequences containing glides or short vowels in our analysis of potential phonotactic gaps in the corpus.

After transliterating the corpus into broad natural classes, we generated all possible 3-gram sequences for the five target natural classes mentioned above and identified which sequences are unattested. As expected, many missing sequences reflect principled phonotactic restrictions – e.g. the absence of [#NN] sequences, which reflects an illicit initial consonant cluster of two nasals. However, other gaps are less easily explained. One particularly striking case is the complete absence of [$L\check{V}\check{V}L$] sequences (liquid–nasal vowel–liquid). Panjabi readily allows [LVVL] sequences with oral vowels – for example, [laɖl] ‘red, ruby’ appears 634 times in the corpus – and native speaker consultants uniformly judged [lããl] to be a fully acceptable, though unattested, word. Additionally, smaller bigram $L\check{V}\check{V}$ and $\check{V}\check{V}L$ sequences are attested in the corpus, so less complex constraints based on simpler bigram co-occurrence restrictions cannot account for the gap.

While future work on an annotated corpus of Pakistani Punjabi is needed to confirm these results, the analysis here shows that a statistical learner limited to making generalizations over observed forms might posit a constraint like $*L\tilde{V}\tilde{V}L$ simply because no such sequence occurs in the lexicon. This is clearly an undesirable outcome – an example of a complex co-occurrence constraint banning a form that is well-formed for native speakers, despite being absent from the lexicon. While it is theoretically possible to model Panjabi’s nasal harmony patterns using surface-based generalizations alone, doing so comes at a significant cost. The computational demands of learning high-order co-occurrence restrictions are substantial, and, more importantly, this approach runs the risk of overfitting the lexicon, thereby failing to distinguish between accidental and principled gaps and resulting in the mischaracterization of speakers’ phonological competence.

In sum, in order to model the Panjabi data presented here, exemplar theory must be augmented with some mechanism(s) that allow speakers to arrive at covert underlying representations. We believe this is a matter of principle, rather than a consequence of the particular assumptions discussed above: it follows, more or less logically, from the surface-oriented character of all contemporary versions of exemplar theory.

5.2.2 Opacity in surface-oriented generative models

In the generative tradition, it has sometimes been suggested that underlying representations can be dispensed with entirely, their functions being subsumed by surface-to-surface relationships between morphologically related surface allomorphs, or other purely surface-oriented generalizations (e.g. Burzio 1996, Scobbie et al. 1996, Hayes 1999, Albright 2002, Cole & Hualde 2011, Allen & Becker 2015, and various others; cf. Krämer 2012, Bermúdez-Otero 2018, Hyman 2018 and references there). We illustrate with a case of counterfeeding opacity in Berbice Dutch Creole (Dow 2014).

In Berbice Dutch Creole, nasal consonants generally assimilate in place of articulation to following oral stops. Additionally, stem-final vowels may delete when followed by a suffix. These two processes interact opaquely, in a counterfeeding relationship, as can be seen from derivations like $/nimi-t\epsilon/ \rightarrow [nim-t\epsilon]$ ‘know-ANTERIOR’ (4).

(4) Counterfeeding opacity in Berbice Dutch Creole

		$/nimi-t\epsilon/$
Place assimilation:	$C[+NASAL] \rightarrow [\alpha PL]/_ [\alpha PL, -CONT, -SON]$	—
Vowel deletion:	$V \rightarrow \emptyset / _]_{STEM} -_{SUFFIX} [$	$[nim-t\epsilon]$
		$[nim-t\epsilon]$

As in the rule-based analysis of Panjabi nasal harmony (3), the rule-based analysis of Berbice Dutch Creole hinges on a distinction between underlying $/N/+stop$ clusters (which undergo place assimilation), and derived $/N/+stop$ clusters (which do not assimilate).⁹

⁹The Berbice Dutch Creole facts are more nuanced than the text here implies. For example, nasal place assimilation is in fact optional in forms like $[nim-t\epsilon] \sim [nin-t\epsilon]$ ‘know-ANTERIOR’, and stem-final vowel deletion shows evidence of morphological idiosyncrasy. We present a simplified version of the Berbice Dutch Creole data for expository purposes; see Dow (2014) for more details.

An alternative, surface-oriented analysis of these facts is also possible, which does not make crucial reference to underlying representations. Following Benua (2000), we may assume that OUTPUT-OUTPUT FAITHFULNESS constraints enforce similarity between morphological bases and complex forms built on those bases. OO-FAITH constraints evaluate similarity along specific phonological dimensions: the constraint OO-IDENT-PLACE (5b) specifically bans changes in place of articulation between a morphological base like [nimi] and a corresponding inflected form like [nim-tɛ] (5). Importantly, OO-FAITH constraints refer only to the *surface* forms of bases and their morphological derivatives.

The simplified tableau in (5c) shows how output-oriented faithfulness relations can produce counterfeeding opacity without drawing a distinction between underlying and derived forms. In short, the regular process of nasal place assimilation is inhibited in [nim-tɛ] because it would produce a phonological discrepancy between the base [nimi] and the derived form *[nin-tɛ]. The input to phonological evaluation is a concatenation of independently observed surface allomorphs, [nimi] and [-tɛ], rather than an abstract UR.

(5) Counterfeeding opacity in Berbice Dutch Creole as output-oriented faithfulness

- a. NT PLACE ASSIMILATION: assign one violation for every nasal consonant followed by an oral stop, such that the nasal and stop differ in their place of articulation.
- b. OO-IDENT-PLACE: assign one violation for every morphologically complex form $\mathbb{B}\text{-}\mathbb{A}$, where \mathbb{B} is a morphological base and \mathbb{A} an affix, in which the \mathbb{B} substring of $\mathbb{B}\text{-}\mathbb{A}$ contains a segment which differs in place of articulation from the corresponding segment in the base form \mathbb{B} .

[nimi-tɛ] / BASE: [nimi]		OO-IDENT-PLACE	NT PLACE ASSIMILATION
c.	a. nin-tɛ	*! W	L
	b. nim-tɛ		*

However, it is well-known that output-oriented faithfulness does not provide a general analysis of phonological opacity (e.g. Wolf 2008, Bermúdez-Otero 2011, Trommer 2013, Mascaró 2016). OO-FAITH constraints can only produce opacity in morphologically complex forms, by demanding that phonological similarity to a morphological base take precedence over general phonotactic constraints (5).

In morphologically simple forms — that is, roots — there is no base to be faithful to, and thus no mechanism for generating opacity. This is precisely the situation presented by the opaque non-application of nasal harmony in forms like [dʒəvãɳ], *[dʒãvãɳ] ‘public’ in Panjabi. The root [dʒəvãɳ] ‘youth, soldier’ has no morphological base to be faithful to; hence, output-oriented mechanisms for producing opacity fall short of explaining why [dʒəvãɳ] fails to undergo harmony. This ‘missing base’ problem is not specific to OO-FAITH: it generalizes to all surface-oriented phonological theories known to us which attempt to generate opaque phonological interactions by means of surface relationships between morphologically related forms.

A conceptually distinct, surface-oriented approach to opacity in nasal spreading is provided by OPTIMAL DOMAINS THEORY (ODT) (Cole & Kisseberth 1994, 1995). Briefly summarizing, in ODT, harmony domains are specified in surface, output forms. Surface-oriented constraints require all segments in a harmony domain to share the same value for a particular feature. The distributions of harmony domains themselves are also specified by surface-oriented constraints.

Cole & Kisseberth (1995) provide an analysis of nasal harmony in Máíhiki (formerly, Orejón) which is strikingly similar to the pattern of nasal harmony we describe for Panjabi (though see Sylak-Glassman et al. 2014 for a different description of the core facts). Nasality spreads rightward from nasal vowels in the initial syllable, the only position in which vowel nasality is contrastive. Harmony may spread through the glide [j], which we assume is nasalized, following Cole & Kisseberth (1995). Harmony is illustrated in (6a), in which we indicate nasal harmony domains by means of parentheses.

(6) Nasal harmony in Máíhiki

- a. Long-distance rightward nasal spreading from contrastive / \tilde{V} /:
/sěje/ → (sějě) ‘species of bird’
- b. Only local rightward nasal spreading from nasal consonants:
/naji/ → (nā)ji ‘grandson’

However, words beginning in a nasal stop show only local nasalization of the following vowel, without longer-distance harmony (6b).¹⁰ Cole & Kisseberth (1995) analyze this contrast by assuming that contrastively nasal vowels initiate harmony domains that must spread as far to the right as possible (6a), while nasal consonants initiate harmony domains that only need spread as far as the rightmost syllable boundary (6b).

Crucially, Cole & Kisseberth (1995) must *also* invoke covert underlying representations for this analysis to function as intended. Roots like [nāji] ‘grandson’ are non-alternating: they are always realized with an initial nasal vowel. Lexical representations like /nāji/, which are faithful to the invariant surface realization [nāji], incorrectly predict nasal harmony in [nāji], *[nājĩ] on analogy with [sějě], in which nasal harmony is initiated by an underlying nasal vowel (6a). The only tractable lexical representation for ‘grandson’ must therefore be the abstract, covert UR /naji/ – despite the fact that this morpheme never surfaces as such (i.e. the initial vowel is always nasal on the surface). We conclude that optimal domains theory also requires covert underlying representations.

5.2.3 The learnability of abstract URs

A major objection to permitting abstract URs in linguistic analysis is the challenge they pose for acquisition. Specifically, if abstract URs are to be considered, the space of potential URs a learner must evaluate to determine the UR that maximizes the likelihood of observing the data grows substantially. Additionally, lexicon optimization – where UR → SR disparities are introduced only to the extent necessary – becomes more difficult to enforce when both a minimally abstract and an overly abstract UR account equally well for the data.

Instead of rejecting abstract URs due to these concerns, a substantial body of research has proposed strategies to mitigate them. For instance, Tesar (2014, 2016) shows that the space of potential URs need not be searched exhaustively if it is efficiently organized in such a way that allows broad swaths of implausible URs to be ruled out in one fell swoop. In addition, several approaches (e.g., Pater et al. 2012, O’Hara 2017, Rasin & Katzir 2018, Paramore 2025, Wang & Hayes 2025) offer distinct mechanisms for minimizing abstraction in UR learning, thereby aiding

¹⁰Sylak-Glassman et al. (2014) argue that post-nasal vowels are oral rather than nasal in Máíhiki; this does not affect our discussion.

lexicon optimization. We refer interested readers to these works and the references therein for further discussion of the learnability of abstract URs.

6 Conclusion

We have presented original phonetic data from Panjabi to establish the novel descriptive generalization that nasal harmony interacts opaquely with local nasal assimilation in /VVN/ sequences. We have further argued that modeling this opaque interaction requires reference to abstract underlying forms that never phonetically surface — forms we dub *covert URs*.

Covert URs allow for a straightforward distinction between underlying nasal vowels (which trigger harmony) and phonologically derived nasal vowels (which do not). Drawing this distinction *without* covert URs proves difficult. Many pre-nasal vowels in /VVN/ → [ṼṼN] sequences are in non-alternating roots. This means that the vowels in question are *always* nasal [ṼṼN] on the surface. The lack of alternations in nasality presents a challenge for distinguishing non-alternating, contrastive nasal vowels /ṼṼ/ (which are unpredictably nasal) from non-alternating, derived nasal vowels [ṼṼN] (which are predictably nasal) by reference to surface forms alone.

We are broadly sympathetic to the goals of exemplar theory, and the empirical findings that motivate exemplar theory and related frameworks. The question is not whether surface-oriented models of phonological computation are justified; the question is whether they are *sufficient*. We contend that an adequate theory of phonology will require mechanisms for learners to arrive at abstract lexical representations that depart from observed surface forms in filtering out at least some surface-predictable information. This conclusion, while traditional, strikes us as surprising in a modern context in which grammatical abstraction has been downplayed in favor of sophisticated, surface-oriented models of phonological learning and generalization.

References

- Albright, Adam & Bruce Hayes. 2003. Rules vs. analogy in English past tenses: a computational/experimental study. *Cognition* 90(2). 119–161.
- Albright, Adam C. 2002. *The identification of bases in morphological paradigms*. UCLA dissertation.
- Allen, Blake & Michael Becker. 2015. *Learning alternations from surface forms with sublexical phonology*. Ms., University of British Columbia and Stony Brook University. <https://ling.auf.net/lingbuzz/002503>.
- Avelino, Heriberto, Roberto Zariquiey & Jorge Iván Pérez-Silva. 2020. Nasal coarticulation and prosody in Kakataibo. *Phonetica* 77. 29–54.
- Baković, Eric. 2011. Opacity and Ordering. In *The Handbook of Phonological Theory*. John A. Goldsmith, Jason Riggle & Alan C. Yu (eds.). Blackwell Publishing Ltd.
- Bashir, Elena & Thomas J. Connors. 2019. *A descriptive grammar of Hindko, Panjabi, and Saraiki* (Mouton-CASL Grammar Series). De Gruyter Mouton.
- Bates, Douglas, Martin Maechler, Ben Bolker & Steven Walker. 2025. *Linear mixed-effects models using ‘Eigen’ and S4*. [Computer software program] Available from <https://cran.r-project.org/package=lme4>.

- Becker, Michael, Lauren Eby Clemens & Andrew Nevins. 2017. Generalization of French and Portuguese plural alternations and initial syllable protection. *Natural Language & Linguistic Theory* 35(2). 299–345.
- Becker, Michael, Nihan Ketrez & Andrew Nevins. 2011. The surfeit of the stimulus: analytic biases filter lexical statistics in Turkish laryngeal alternations. *Language* 87(1). 84–125.
- Becker, Michael, Andrew Nevins & Jonathan Levine. 2012. Asymmetries in generalizing alternations to and from initial syllables. *Language* 88(2). 231–268.
- Beddor, Patrice Speeter. 2007. Nasals and nasalization: the relation between segmental and coarticulatory timing. In *Proceedings of the XVIth International Congress of Phonetic Sciences*, 249–254. Available online at <http://www.icphs2007.de/conference/Papers/1728/>. Saarbrücken, Germany: Universität des Saarlandes.
- Beddor, Patrice Speeter & Rena Arens Krakow. 1999. Perception of coarticulatory nasalization by speakers of English and Thai: evidence for partial compensation. *Journal of the Acoustical Society of America* 106(5). 2868–2887.
- Bell-Berti, Fredericka. 1993. Understanding velic motor control: studies of segmental context. In *Nasals, Nasalization, and the Velum*. Marie K. Huffman & Rena A. Krakow (eds.). Vol. 5. (Phonetics and Phonology). Academic Press.
- Benua, Laura. 2000. *Phonological relations between words*. New York: Garland.
- Bermúdez-Otero, Ricardo. 2011. Cyclicity. In Marc van Oostendorp, Colin Ewen, Elizabeth Hume & Keren Rice (eds.), *The Blackwell companion to phonology*, 2019–2048. Malden, MA: Wiley-Blackwell. <https://doi.org/10.1002/9781444335262.wbctp0085>.
- Bermúdez-Otero, Ricardo. 2018. In defence of underlying representations: Latin rhotacism, French liaison, Romanian palatalization. *Probus* 30(2). 171–214.
- Bhatia, Tej K. 1993. *Punjabi* (Descriptive Grammars). Routledge.
- Boersma, Paul & David Weenink. 2015. *Praat: Doing phonetics by computer (version 5.4.21)*. <http://www.fon.hum.uva.nl/praat/>.
- Bonett, Douglas. 2021. Statistical methods for psychologists: introduction to statistical inference and experimental design.
- Braunschweiler, Norbert. 1997. Integrated cues of voicing and vowel length in German: a production study. *Language and Speech* 40(4). 353–376.
- Buchwald, Adam & Michele Miozzo. 2011. Finding levels of abstraction in speech production: evidence from sound-production impairment. *Psychological Science* 22(9). 1113–1119.
- Burzio, Luigi. 1996. Surface constraints versus underlying representations. In *Current trends in phonology: Models and methods*. Jacques Durand & Bernard Laks (eds.). Salford: European Studies Research Institute, University of Salford. 118–136.
- Bybee, Joan. 2001. *Phonology and language use*. Cambridge: Cambridge University Press.
- Chen, Marilyn Y. 1997. Acoustic correlates of English and French nasalized vowels. *Journal of the Acoustical Society of America* 102(4). 2360–2370.
- Chomsky, Noam & Morris Halle. 1965. Some controversial questions in phonological theory. *Journal of Linguistics* 1. 97–138.
- Chomsky, Noam & Morris Halle. 1968. *The sound pattern of English*. New York: Harper & Row.
- Cole, Jennifer & José Hualde. 2011. Underlying representations. In Marc van Oostendorp, Colin Ewen, Elizabeth Hume & Keren Rice (eds.), *The Blackwell companion to phonology*, 1–26. Malden, MA: Wiley-Blackwell.

- Cole, Jennifer & Charles Kisseberth. 1994. An optimal domains theory of harmony. *Studies in the Linguistic Sciences* 24(2).
- Cole, Jennifer & Charles Kisseberth. 1995. Nasal harmony in optimal domains theory. *WECOL* 7. 44–58.
- Cornell, Sonia A., Aditi Lahiri & Carsten Eulitz. 2011. “What you encode is not necessarily what you store”: evidence for sparse feature representations from mismatch negativity. *Brain research* 1394. 79–89.
- Daland, Robert, Bruce Hayes, James White, Marc Garellek, Andrea Davis & Ingrid Norrmann. 2011. Explaining sonority projection effects. *Phonology* 28(2). 197–234.
- Delvaux, Véronique, Didier Demolin, Bernard Harmegnies & Alain Soquet. 2008. The aerodynamics of nasalization in French. *Journal of Phonetics* 36. 578–606.
- Desmeules-Trudel, Félix & Marc Brunelle. 2018. Phonotactic restrictions condition the realization of vowel nasality and nasal coarticulation: duration and airflow measurements in Québécois French and Brazilian Portuguese. *Journal of Phonetics* 69. 43–61. <https://doi.org/10.1016/j.wocn.2018.05.001>.
- Dhillon, Rajdip. 2007. Stress in punjabi. In *Berkeley linguistics society*, vol. 33, 84–95.
- Dow, Michael. 2014. Hybrid opacity in Berbice Dutch Creole. In John Kingston, Claire Moore-Cantwell, Joe Pater & Robert Staubs (eds.), *Proceedings of the Annual Meetings on Phonology*, vol. 2. Washington, DC: Linguistic Society of America. <https://doi.org/10.3765/amp.v1i1.28>.
- Du, Naiyan & Karthik Durvasula. 2025. Psycholinguistics and phonology: the forgotten foundations of generative phonology. *Cambridge Elements in Phonology*. <https://doi.org/10.1017/9781009347631>.
- Ernestus, Mirjam. 2014. Acoustic reduction and the roles of abstractions and exemplars in speech processing. *Lingua* 142. 27–41. <https://doi.org/10.1016/j.lingua.2012.12.006>.
- Ettlinger, Marc. 2008. *Input-driven opacity*. University of California, Berkeley dissertation.
- Ettlinger, Marc. 2009. Variation as a window into opacity. In Malcolm Elliott, James Kirby, Osamu Sawada, Eleni Staraki & Suwon Yoon (eds.), *Chicago Linguistics Society (CLS)* 43. Chicago: University of Chicago, Chicago Linguistic Society.
- Farris-Trimble, Ashley & Anne-Michelle Tessier. 2019. The effect of allophonic processes on word recognition: eye-tracking evidence from Canadian raising. *Language* 95(1). e136–e160. <https://doi.org/10.1353/lan.2019.0023>.
- Flemming, Edward. 2001. Scalar and categorical phenomena in a unified model of phonetics and phonology. *Phonology* 18(1). 7–44.
- Gafos, Adamantios. 2002. A grammar of gestural coordination. *Natural Language & Linguistic Theory* 20(2). 269–337.
- Gaskell, M. Gareth & William Marslen-Wilson. 1996. Phonological variation and inference in lexical access. *Journal of Experimental Psychology: Human perception and performance* 22(1). 144–158.
- Goldinger, Stephen D. 1998. Echoes of echoes? an episodic theory of lexical access. *Psychological Review* 105(2). 251–279.
- Goldrick, Matthew & Jennifer Cole. 2023. Advancement of phonetics in the 21st century: exemplar models of speech production. *Journal of Phonetics* 99.
- Goldrick, Matthew & Brenda Rapp. 2007. Lexical and post-lexical phonological representations in spoken production. *Cognition* 102(2). 219–260.

- Gorman, Kyle. 2013. *Generative phonotactics*. Philadelphia, PA: University of Pennsylvania dissertation.
- Gouskova, Maria. 2012. Unexceptional segments. *Natural Language & Linguistic Theory* 30(1). 79–133.
- Gouskova, Maria & Michael Becker. 2013. Nonce words show that Russian yer alternations are governed by the grammar. *Natural Language & Linguistic Theory* 31(3). 735–765.
- Halle, Morris. 1959. *The sound pattern of Russian*. The Hague: Mouton.
- Knowledge unlearned and untaught: What speakers know about the sounds of their language. 1978. In Morris Halle, Joan Bresnan & George A. Miller (eds.), *Linguistic theory and psychological reality*, 294–304. Cambridge, MA: MIT Press.
- Hayes, Bruce. 1999. Phonological restructuring in Yidin^y and its theoretical consequences. In Ben Hermans & Marc van Oostendorp (eds.), *The derivational residue in phonological Optimality Theory*, 175–205. Amsterdam: John Benjamins Publishing.
- Hayes, Bruce & James White. 2013. Phonological naturalness and phonotactic learning. *Linguistic Inquiry* 44. 45–75.
- Hayes, Bruce & Colin Wilson. 2008. A maximum entropy model of phonotactics and phonotactic learning. *Linguistic Inquiry* 39(3). 379–440.
- Henderson, Janette Barclay. 1984. *Velopharyngeal function in oral and nasal vowels: a cross-language study*. University of Connecticut dissertation.
- Hussain, Qandeel, Michael Proctor, Mark Harvey & Katherine Demuth. 2019. Punjabi (Lyallpuri variety). *Journal of the International Phonetic Association* 50(2). 282–297.
- Hyman, Larry. 2018. Why underlying representations? *Journal of Linguistics* 54. 591–610.
- Jakubiček, Miloš, Adam Kilgarriř, Vojtěch Kovář, Pavel Rychlý & Vít Suchomel. 2013. The TenTen Corpus family. In *7th international corpus linguistics conference cl*, 125–127.
- Johnson, Keith. 1997. Speech Perception without Speaker Normalization. In *Talker Variability in Speech Processing*. Keith Johnson & John W. Mullennix (eds.). San Diego: Academic Press. Chap. 8, 145–165.
- Johnson, Keith. 2006. Resonance in exemplar-based lexicon: the emergence of social identity and phonology. *Journal of Phonetics* 34. 485–499.
- Jongman, Allard. 2004. Phonological and phonetic representations: the case of neutralization. In Augustine Agwuele, Willis Warren & Sang-Hoon Park (eds.), *Proceedings of the 2003 Texas Linguistics Society Conference*, 9–16. Somerville, MA: Cascadilla Press.
- Kaplan, Abby. 2015. Positional neutralization in an exemplar model: the role of unique inflectional bases. In Adam Albright & Michelle A. Fullwood (eds.), *Proceedings of the Annual Meetings on Phonology*, vol. 3. Washington, DC: Linguistic Society of America. <https://doi.org/10.3765/amp.v2i0.3745>.
- Kaplan, Abby. 2017. Incomplete neutralization and the (a)symmetry of paradigm uniformity. In Aaron Kaplan, Abby Kaplan, Miranda K. McCarvel & Edward J. Rubin (eds.), *West Coast Conference on Formal Linguistics (WCCFL) 34*, 319–328. Available online at <http://www.lingref.com/document/#3326>. Somerville, MA: Cascadilla Proceedings Project.
- Kawahara, Shigeto. 2001. *Similarity among Variants: Output-Variant Correspondence*. Senior Thesis, International Christian University. Available online as ROA-497, Rutgers Optimality Archive, <http://roa.rutgers.edu/>.
- Kenstowicz, Michael & Charles Kisseberth. 1979. *Generative phonology: description and theory*. New York: Academic Press.

- Kiparsky, Paul. 1976. Abstractness, opacity, and global rules. In *The Application and Ordering of Grammatical Rules*. Andreas Koutsoudas (ed.). The Hague: Mouton.
- Kirchner, Robert. 2011. Preliminary thoughts on “phonologisation” within an exemplar-based speech processing system. *UCLA Working Papers in Linguistics* 6.
- Kirchner, Robert, Roger K. Moore & Tsung-Ying Chen. 2010. Computing phonological generalization over real speech exemplars. *Journal of Phonetics* 38(4). 540–547.
- Krakow, Rena A. 1989. *The articulatory organization of syllables: a kinematic analysis of labial and velar gestures*. Yale University dissertation.
- Krämer, Martin. 2012. *Underlying representations*. Cambridge: Cambridge University Press.
- Lahiri, Aditi. 2012. Asymmetric phonological representation of words in the mental lexicon. In Abigail Cohn, Cécile Fougeron & Marie Huffman (eds.), *The Oxford Handbook of Laboratory Phonology*, 146–162. Oxford, UK: Oxford University Press.
- Lahiri, Aditi & William Marslen-Wilson. 1991. The mental representation of lexical form: a phonological approach to the recognition lexicon. *Cognition* 38. 245–294.
- Lahiri, Aditi & Henning Reetz. 2010. Distinctive features: phonological underspecification in representation and processing. *Journal of Phonetics* 38(1). 44–59.
- Lakens, Daniel, Anne M. Scheel & Peder M. Isager. 2018. Equivalence testing for psychological research: a tutorial. *Advances in Methods and Practices in Psychological Science* 1(2). 259–269.
- Lapierre, Myriam. 2023. Two types of [NT]s in Panāra: evidence for temporally ordered subsegmental units. *Glossa* 8(1).
- LoCasto, Paul & Cynthia Connine. 2002. Rule-governed missing information in spoken word recognition: schwa vowel deletion. *Perception & Psychophysics* 64(2). 208–219.
- Mascaró, Joan. 2016. Morphological exceptions to vowel reduction in Central Catalan and the problem of the missing base. *Catalan Journal of Linguistics* 15. 27–51. <https://doi.org/10.5565/rev/catj1.184>.
- McCarthy, John J. 1999. Sympathy and phonological opacity. *Phonology* 16. 331–399.
- McCarthy, John J. 2005. Optimal Paradigms. In Laura Downing, Tracy Alan Hall & Renate Rafelsiefen (eds.), *Paradigms in phonological theory*, 170–210. Oxford, UK: Oxford University Press.
- McCarthy, John J. 2007. Derivations and levels of representation. In Paul de Lacy (ed.), *The Cambridge handbook of phonology*, 99–117. Cambridge, UK: Cambridge University Press.
- Morrison, Donald Alasdair. 2021. Vowel nasalisation in Scottish Gaelic: no evidence for incomplete neutralisation in initial mutation. *Morphology* 31(2). 121–146.
- O’Hara, Charlie. 2017. How abstract is more abstract? learning abstract underlying representations. *Phonology* 34. 325–345.
- Paramore, Jonathan Charles. 2025. Learning covert URs via disparity minimization. In *Society for Computation in Linguistics*, vol. 8, 29.
- Pater, Joe, Robert Staubs, Karen Jesney & Brian Smith. 2012. Learning probabilities over underlying representations. In *Proceedings of the twelfth meeting of the special interest group on computational morphology and phonology*, 62–71.
- Pierrehumbert, Janet. 2001. Exemplar dynamics: Word frequency, lenition and contrast. In *Frequency and the emergence of linguistic structure*. Joan Bybee & Paul Hopper (eds.). Amsterdam: John Benjamins. 137–157.

- Pierrehumbert, Janet. 2002. Word-specific phonetics. In Carlos Gussenhoven & Natasha Warner (eds.), *Papers in laboratory phonology VII*, 101–140. Berlin: Mouton de Gruyter. <https://doi.org/10.1515/9783110197105.101>.
- Pierrehumbert, Janet. 2003. Phonetic diversity, statistical learning, and acquisition of phonology. *Language and speech* 46(2-3). 115–154.
- Pierrehumbert, Janet. 2006. The statistical basis of an unnatural alternation. In Louis Goldstein, D.H. Whalen & Catherine Best (eds.), *Papers in laboratory phonology VIII*, 81–107. Berlin: Mouton de Gruyter.
- Pierrehumbert, Janet. 2016. Phonological representation: beyond abstract versus episodic. *Annual Review of Linguistics* 2. 33–52.
- Pinker, Steven. 2002. Four decades of rules and associations, or whatever happened to the past tense debate? In *Language, brain, and cognitive development: essays in honor of Jacques Mehler*. Cambridge, MA: MIT Press.
- Plug, Leendert, Abdurraouf Shitaw & Barry Heselwood. 2019. Inter-consonantal intervals in Tripolitanian Libyan Arabic: accounting for variable epenthesis. *Laboratory Phonology* 10(1). 5. <https://doi.org/10.5334/labphon.122>.
- Port, Robert F. & Michael L. O'Dell. 1985. Neutralization of syllable-final voicing in German. *Journal of Phonetics* 13. 455–471.
- Poupplier, Marianne, Francesco Rodriguez, Justin J.H. Lo, Roy Alderton, Bronwen G. Evans, Eva Reinisch & Christopher Carignan. 2024. Language-specific and individual variation in anticipatory nasal coarticulation: a comparative study of American English, French, and German. *Journal of Phonetics* 107. 101365. <https://doi.org/10.1016/j.wocn.2024.101365>.
- Prince, Alan & Paul Smolensky. 1993/2004. *Optimality Theory: constraint interaction in Generative grammar*. Malden, Mass: Blackwell Publishers.
- Rasin, Ezer & Roni Katzir. 2018. Learning abstract underlying representations from distributional evidence. In Sherry Hucklebridge & Max Nelson (eds.), *North east linguistic society*, vol. 48. MA: GLSA.
- RCoreTeam. 2016. *R: A language and environment for statistical computing*. Vienna: R Foundation for Statistical Computing. <https://www.R-project.org/>.
- Scobbie, James M., John S. Coleman & Steven Bird. 1996. Key aspects of declarative phonology. In Jacques Durand & Bernard Laks (eds.), *Current trends in phonology: models and methods*, vol. 2, 685–709. Manchester, UK: University of Salford, European Studies Research Institute.
- Seidenberg, Mark S. & David C. Plaut. 2014. Quasiregularity and its discontents: the legacy of the past tense debate. *Cognitive science* 38(6). 1190–1228.
- Shackle, Christopher. 2003. Panjabi. In *The Indo-Aryan Languages*. George Cardona & Dhanesh Jain (eds.). (Routledge Language Family Series). Routledge. Chap. 16, 581–621.
- Solé, Maria-Josep. 1992. Phonetic and phonological processes: the case of nasalization. *Language and Speech* 35(1). 29–43.
- Solé, Maria-Josep. 1995. Spatio-temporal patterns of velopharyngeal action in phonetic and phonological nasalization. *Language and Speech* 38(1). 1–23.
- Styler, Will. 2017. On the acoustical features of vowel nasality in English and French. *Journal of the Acoustical Society of America* 142(4). 2469–2482.
- Sylak-Glassman, John, Stephanie Farmer & Lev Michael. 2014. *An Agreement by Correspondence analysis of Máihiki nasalization harmony*. Presentation at ABC↔ conference. <http://www.>

- linguistics.berkeley.edu/phonlab/documents/2014/ABCC/Sylak-Glassman_etal.pdf.
- Tesar, Bruce. 2014. *Output-Driven Phonology*. Cambridge: Cambridge University Press.
- Tesar, Bruce. 2016. Phonological learning with output-driven maps. *Language Acquisition* 24(2). 148–167.
- Trommer, Jochen. 2013. Stress uniformity in Albanian: morphological arguments for cyclicity. *Linguistic Inquiry* 44(1). 109–143.
- Wang, Yang & Bruce Hayes. 2025. Learning phonological underlying representations: the role of abstractness. *Linguistic Inquiry*.
- Wedel, Andrew. 2004. *Self-organization and categorical behavior in phonology*. Santa Cruz, CA: University of California, Santa Cruz dissertation.
- Wedel, Andrew. 2006. Exemplar models, evolution and language change. *Linguistic Review* 23. 247–274.
- Wellek, Stefan. 2010. *Testing statistical hypotheses of equivalence and noninferiority*. 2nd edn. Taylor & Francis Group.
- White, Yosiane, David Embick & Meredith Tamminga. 2024. Affix priming with variable ING in English: implications for unique vs. dual representation. *Journal of Memory and Language* 138. 104535.
- Wickham, Hadley. 2016. *ggplot2: Elegant graphics for data analysis*. Springer-Verlag New York. <https://ggplot2.tidyverse.org>.
- Wilson, Colin & Gillian Gallagher. 2018. Accidental gaps and surface-based phonotactic learning: a case study of South Bolivian Quechua. *Linguistic Inquiry* 49(3). 610–623.
- Wolf, Matthew. 2008. *Optimal interleaving: serial phonology-morphology interaction in a constraint-based model*. Amherst, MA: University of Massachusetts, Amherst dissertation.
- Zahid, Saira & Sarmad Hussain. 2012. An acoustic study of vowel nasalization in Punjabi. *Language & Technology* 61.
- Zellou, Georgia. 2022. Coarticulation in phonology. *Cambridge Elements in Phonology*. <https://doi.org/10.1017/9781009082488>.

7 Appendix

Condition 1 - CVV		Condition 2 - CŨŨ		Condition 3 - CVVN	
baat	‘story’	bãã	‘arm’	paan	‘Betel nut chewing’
bii	‘seed’	bããk	‘early teen’	biin	‘a musical instrument’
boo	‘smell’				
piid	‘pain’	pĩĩg	‘swing’		
poo	‘12th Panjabi month’	põõtʃ	‘arrival’	koon	‘who’
daa	‘strategy (wrestling)’			daam	‘price, money’
daak	‘post’	dããg	‘big stick’		
deed	‘one and a half’				
doo	‘two’			doon	‘ropes on a knit bed’
dzii	‘family member’				
		dzõõ	‘barley’		
dzoog	‘coming together’	dzõõk	‘leach’		
		dzũũ	‘louse’	dzuun	‘June’
taa	‘warmth, fever’	tãã	‘that; so that’	taan	‘melody’
t ^h aa	‘was’	t ^h ãã	‘place, room’	t ^h aan	‘a bolt of cloth’
tɔɔk	‘yoke, iron collar’			tɔɔŋ	‘neck’
tʃaa	‘desire, wish’	tʃ ^h ãã	‘shade’	tʃ ^h aan	‘non-refined flour’
ruup	‘beauty’	rũũ	‘cotton’	luun	‘salt’
saa	‘breath’				
seek	‘warmth’	sẽẽk	‘termite’		
sii	‘sew.IMP’			siin	‘ploughing’
sɔɔ	‘hundred’	sõõ	‘sleep.IMP’	sɔɔŋ	‘rainy season’
suud	‘profit, interest’	tũũ	‘you’		
		vããg	‘like, similar’	vaan	‘strings of a knit bed’
gaal	‘abuse’	gãã	‘cow’		
kaadz	‘button-hole’	kãã	‘crow’	k ^h áán	‘a proverb’
kuuk	‘melodious sound’	kũũdz	‘black sparrow’		
k ^h uu	‘a well’	k ^h ũũdz	‘a corner, angle’	xuun	‘blood’
xaab	‘dream’			xaantʃaa	‘corner, trick’

Figure A1: Experiment 1 tokens, consisting of 63 monosyllabic words separated into three groups: oral vowels (CVV), nasal vowels (CŨŨ), and /VVN/ vowels (CVVN).

7.1 Annotation Procedure & Excrescent nasals

The first author and two undergraduate research assistants completed the annotations. The first author verified the quality of all annotations. Boundaries between oral consonants and vowels were marked at the onset/offset of high amplitude in the oral channel, along with the onset of increased formant activity in the spectrogram. The boundary between vowels and the onset of nasal consonants was marked at the attenuation of higher formant activity coupled with a rapid

Oral Condition - CVVGVV			
sət-aqjii	seven-twenty	ɔɭii-jaa	pious man-PL
həvaɑ	air	tʃoo-vii	four-twenty
paavee	cot leg	beevaa	a widow
dəvaɑ	medicine, remedy	taavuu	father's oldest brother
mək ^h iijaa	chief	vək ^h aavee	boasting

Nasal Condition - CVVGṼṼ			
tʃãũṽẽẽ	pumicestone	ũũ-ṽẽẽ	that-way
dõõ-ṽãã	two-way	kĩĩ-ṽẽẽ	what-way
ẽẽ-ṽẽẽ	this-way	mətʃĩĩ-jãã	fish-PL
sãã-ṽãã	breath-PL	tĩĩṽĩĩ	a woman, wife

/VVN/ Condition - CVVGVVN			
dʒəṽããŋ	youth, soldier	ədʒəṽããŋ	omum seed
siijããŋ	recognition	teejããŋ	attention
lɛɛ-jããŋ	to bring-NOML	aaṽ-ããm	PL-citizen

Figure A2: Experiment 2 tokens, consisting of 25 di- and trisyllabic words separated into three groups: oral /VV/, contrastive nasal /ṼṼ/, and pre-nasal vowels /VVN/.

decline in waveform amplitude in the oral channel. A sample annotation is provided in figure A3.




Figure A3: Spectrogram and waveform showing a sample annotation for /paan/ for speaker AR0. The top waveform shows amplitude for the oral channel, and the bottom shows amplitude for the nasal channel.

For almost all of the CṼṼ tokens with oral stops in the coda position, oral occlusion for the coda consonant was achieved before the offset of articulation of nasality for the /ṼṼ/ vowel. This resulted in the physical presence of an intrusive nasal stop for a short period between the gestures of the vowel and the coda (see e.g. Desmeules-Trudel & Brunelle 2018, Lapierre 2023). Due to the minimal amount of oral air expelled during the production of nasal consonants, nasality increased rapidly at the offset of the contrastive /ṼṼ/ vowel when these intrusive nasal conso-

nants were present. To avoid corrupting the reported nasality of the / $\tilde{V}\tilde{V}$ / vowel, intrusive nasals were annotated as part of the consonant segment whenever present. Importantly, it is clear from other phonotactic restrictions in Panjabi that this nasal stop is indeed intrusive and not present in the phonology. Specifically, Bhatia (1993) reports that NC coda clusters are illicit in the language when following a long vowel. A representative example of these intrusive nasal consonants is evidenced by the spectrogram and waveform in figure A4.




Figure A4: Spectrogram and waveform showing the intrusive [ŋ] in /bāāk/ for speaker AK2.