Phonology **36** (2019) 37–60. © Cambridge University Press 2019
doi:10.1017/S0952675719000034

*Phonotactic knowledge and phonetically unnatural classes: the plain uvular in Cochabamba Quechua**

Gillian Gallagher
New York University

While many phonological patterns target classes of sounds that can be defined phonetically, a large number of patterns in descriptive grammars involve sounds that cannot be easily characterised in phonetic terms. This finding suggests that phonological patterns themselves must be taken into account when learning phonological representations, and that phonological classes may emerge in learning from both phonetic factors (bottom-up) and phonological patterns (top-down). The current work presents a case of a phonetically unnatural class in South Bolivian Quechua that is active in the phonology of the language, and provides experimental support that this class is referred to by speakers' grammars. While many cases of phonetically unnatural classes have been documented in descriptions of language patterns, in most cases there is little or no evidence that these patterns or classes are represented by speakers as they are described by linguists.

1 Introduction

A defining property of phonological systems is that generalisations are stated over groups of sounds. Consequently, a central issue in phonological theory has been the determination of possible sound classes, with evidence coming from cross-linguistic phonological patterns and phonetic analysis. Distinctive feature theories (Jakobson *et al*. 1952, Chomsky & Halle 1968, Halle & Stevens 1971, Avery & Idsardi 2001), as well as theories of

* E-mail: GGILLIAN@NYU.EDU.
I am grateful to Ryan Bennett and Colin Wilson for discussion of this work, as well as audiences at Johns Hopkins University, New York University and the 2018 *Réseau Français de Phonologie*. I am also indebted to three anonymous reviewers and two associate editors at *Phonology* for their comments. This work would not have been possible without the consistent assistance of Gladys Camacho Rios, who facilitated all aspects of data collection in Bolivia. The materials and results for both experiments are available as online supplementary materials at https://doi.org/10.1017/S0952675719000034.

https://doi.org/10.1017/S0952675719000034 Published online by Cambridge University Press

phonetically based phonology (Steriade 1999, Flemming 2002, 2004, 2005), formalise and codify the observation that phonologically relevant sound classes are frequently defined by a common phonetic property.

Mielke (2008), drawing on previous work (Bach & Harms 1972, Anderson 1981), shows that, while most phonological patterns do hold of phonetically definable classes, not all do, calling into question the centrality of phonetic similarity to phonological representations. Mielke proposes instead that featural classes are learned from the phonological patterning of segments, as opposed to emerging exclusively from bottom-up analysis of phonetic commonalities.

A frequent source of phonetically unnatural classes is sound change or phonetic variability that makes it challenging to identify a stable phonetic correlate of class membership. This paper presents such a case in Bolivian Quechua. The etymological plain uvular stop /q/ is generally produced as a continuant (which I represent with the IPA symbol [ʁ]) in this Quechua variety, though the phonological patterning of the category remains consistently stop-like. This scenario provides a good test case for exploring the role of phonetic substance and phonological patterning in speakers' representations. I approach this question by testing speakers' sensitivity to phonotactic restrictions on [ʁ] as compared to their sensitivity to the same restrictions on the phonetically unambiguous stop [k]. Previous experimental work (Gallagher 2013, 2014, 2015, 2016b) has shown that Quechua speakers are sensitive to the restrictions on stop combinations, but only stimuli with phonetic stops have been tested. If featural representations and natural classes are restricted by phonetic similarity, then the continuant realisation of the plain uvular may obscure the phonotactic patterning of this segment to learners. Alternatively, if representations are informed by both phonetic similarity and phonological patterning, then speakers should be as sensitive to restrictions which are violated by the uvular as to those which are violated by other stops. The research in this paper could be framed as a search for evidence for traditional underlying representations. Is there evidence that speakers have learned an underlying category /q/, specified as [−continuant, −sonorant], despite the surface realisation as [ʁ]? Or does the surface production of this sound as [ʁ] lead to a representation as /ʁ/, specified as [+continuant]?

I present the results of two experiments that show that speakers have learned the phonotactic patterning of [ʁ] and support the inclusion of this segment in the broader class of stops. The findings support a theory where phonological representations are informed by both phonetic similarity and phonological patterns.

2 A phonetics–phonology mismatch

In the Cochabamba variety of South Bolivian Quechua, the plain uvular stop phoneme /q/ is produced primarily as a continuant [ʁ]. Figure 1 shows a representative example, taken from the recordings reported on

https://doi.org/10.1017/S0952675719000034 Published online by Cambridge University Press

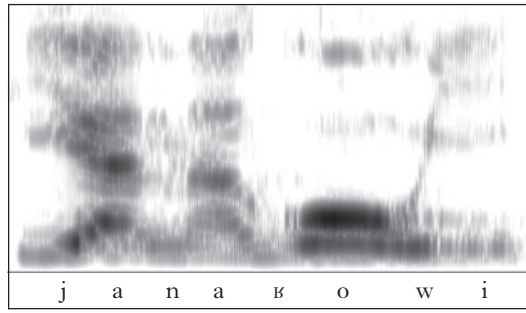


Figure 1
Production of /q/ as continuant [ʁ] in intervocalic position in [jana ʁowi] 'black guinea pig'.

in Gallagher & Whang (2014). While the continuant production is most common, productions as a voiceless stop [q] or, in post-nasal position, a voiced stop [ɢ], are also observed. Bills *et al*. (1969) also describe the continuant production as canonical, so this has been the pronunciation of this sound for at least the last fifty years.[1] While the plain stop is the focus of this paper, the other uvular stops, /qʰ/ and /q'/, are also undergoing changes in this variety of Quechua, though these are far from complete: /qʰ/ may be produced as [χ] and /q'/ as a voiceless implosive [ɢ̰]. This section begins by reviewing the phonological patterning of the plain uvular in §2.1, and continues with a brief acoustic analysis in §2.2.

2.1 The phonological patterning of /q/

The production of the etymological stop /q/ as a continuant [ʁ] creates a mismatch between phonetics and phonology in the synchronic system. There are at least three pieces of phonological evidence that the plain uvular belongs to the class of stops. First, there is the inventory structure of the language, which contrasts plain, aspirate and ejective stops at five places of articulation. If the plain uvular is considered a stop in terms of inventory structure, then the inventory is symmetrical and economical, in the sense of Clements (2003). If the plain uvular is considered a continuant (either a voiced fricative or a more heavily lenited approximant), the inventory would have a gap where the plain uvular stop should be. All three hypothetical positions for the plain uvular in the inventory are shown in (1), which shows the phonological position of the plain uvular as a stop (solid lines) and the alternative phonetic positions as either a fricative or an approximant (dashed lines).

[1] They transcribe the sound with the symbol [ɣ], though no sources to my knowledge explicitly suggest that the place of articulation of this category is in question. It behaves as a uvular in lowering and retracting surrounding vowels. The choice of [ɣ] likely reflects the influence of Spanish; speakers sometimes report this sound being similar to the Spanish 'g', which is lenited in many positions.

https://doi.org/10.1017/S0952675719000034 Published online by Cambridge University Press

(1)

	labial	dental	post-alveolar	velar	uvular	glottal
plain	p	t	ʧ	k	q	
aspirate	p^h	t^h	$ʧ^h$	k^h	q^h	
ejective	p’	t’	ʧ’	k’	q’	
fricative		s	ʃ	x	ʁ	h
nasal	m	n	ɲ			
approximant	w	l ɾ	ʎ j		ʁ	

Second, the plain uvular patterns as a stop in terms of syllable-structure constraints. Stops in southern Quechua varieties are only licit in prevocalic position, while fricatives and sonorants are also possible in preconsonantal positions (some, but not all, sonorants and fricatives are also permitted word-finally). This distribution is illustrated in the examples in (2).

(2) a. *Prevocalic position: stops, fricatives and sonorants*

p’aʁo	‘golden’	ʧiɾi	‘cold’	jut^hu	‘partridge’
takij	‘to sing’	ʁaʎu	‘tongue’	simi	‘mouth’

b. *Preconsonantal position: stops, fricatives and sonorants*

misk’i	‘delicious’	t’anta	‘bread’	waʎpa	‘chicken’
ʎaxta	‘town’	pampa	‘plain’	miɾk^ha	‘freckle’

c. *Preconsonantal position: no stops, including* [ʁ]

*mitk’i	*ʎakta	*t’aʁta

The third piece of phonological evidence for the stop status of the plain uvular is its participation in the laryngeal phonotactics of the language, shown in (3). Ejective and aspirate stops are permitted in either initial or medial position in roots, but they must be the leftmost stop. The plain uvular patterns with the other stops, in that it is never followed by an ejective or an aspirate.

(3) a. *Ejectives and aspirates in initial position*

t’anta	‘bread’	p^hawaj	‘to run’
k’aʧa	‘pretty’	q^haɾi	‘man’

b. *Ejectives and aspirates in medial position; no preceding stops*

ɾit’i	‘snow’	mosq^{h}oj	‘to dream’
saʧ’a	‘tree’	ʎimp^{h}i	‘colour’

c. *No ejectives or aspirates preceded by other stops, including* [ʁ]

*tant’a	*posq^{h}oj
*kaʧ’a	*ʧimp^{h}i
*ʁap’a	*ʁat^{h}a

Phonologically, then, while [ʁ] is phonetically continuant, it has the phonotactic distribution of a stop.

https://doi.org/10.1017/S0952675719000034 Published online by Cambridge University Press

2.2 The acoustics of /q/

While a full phonetic study of the plain uvular was not carried out for the current work, an available recording of a single speaker was analysed, with the aim of giving a slightly more detailed picture of the phonetic situation.

The speaker was a female, aged around 50, from Toro Toro, a rural Quechua-speaking community in the Cochabamba department of Bolivia. Recordings of both spontaneous and elicited speech from this speaker were generously shared with me by Gladys Camacho Rios. For the elicited speech, the speaker was presented with a Quechua verb in the infinitival form (marked by the suffix [-j]), and asked to produce an inflected form. The speaker produced three different inflected forms, adding the suffix /-ŋku/ (3PL), /-ɾqa/ (3SG PAST) or /-ni/ (1SG). For example, given the infinitive /saqij/ 'to leave', the speaker produced /saqiŋku/ 'they leave', /saqiɾqa/ 'he/she left' and /saqini/ 'I leave'. The verbs that are analysed in the current paper are given in their infinitival form in Table I, organised by segmental context. These words are a subset of stimuli originally designed for the examination of the effects of uvulars on surrounding vowels. Consequently, these forms do not represent the full range of contexts that a study designed to investigate uvular realisation itself might include. In particular, all of the plain uvulars are word-medial, preceding a stressed syllable; different prosodic positions may favour different surface realisations.

intervocalic	post-approximant	post-nasal	post-obstruent
t'uquj	ajqij	unquj	usqaj
saqij	q^{h}arquj	qanqij	musquj
uquj			p'usquj
puquj			
t'iqij			
t'uquj			

Table I

Target verbs in their infinitival form for the acoustic study, in phonemic transcription. Each verb was produced in three inflected forms.

The medial plain uvulars were inspected visually in Praat (Boersma & Weenink 2016), and were coded for their closure properties. The majority of tokens (34/36 = 94%) showed continuous voiced airflow and no burst. There were three post-nasal tokens that showed a voiced closure and a stop burst.

The spontaneous speech was in the form of a short interview (less than five minutes) about the process of making chicha, a beverage made from fermented corn. In spontaneous speech, the speaker produced seven instances of a plain uvular, all in intervocalic position. Of these seven productions, six were marked by continuous airflow, and one showed a short voiceless closure and release burst.

https://doi.org/10.1017/S0952675719000034 Published online by Cambridge University Press

While these data only represent analysis of a limited number of tokens from a single speaker, they show the dominance of the continuant realisation of the plain uvular. A full phonetic study of this sound would examine the rate of stop realisation across a larger range of speakers and phonological contexts, and would investigate the exact manner of articulation. The IPA does not distinguish between a voiced uvular fricative and a uvular approximant, as no language contrasts these manners of articulation at posterior places of articulation (Ladefoged & Johnson 2014: 188). I refer to this sound as a continuant throughout the remainder of the paper, remaining agnostic about whether it is a sonorant or a fricative, or variable between the two. Regardless of the precise phonetic realisation or the amount of variation, some degree of phonetics–phonology mismatch holds. A sound that patterns fully as a stop in the phonology of the language is not consistently realised as such in the phonetics. While the acoustic data here show that stop realisations of /q/ are possible, the stimuli in the experiments to follow crucially exposed participants only to a continuant realisation of this sound.

The following sections present two experiments that investigate speakers' awareness of the co-occurrence phonotactics in (3). If the phonetic identity of [ʁ] as a continuant is essential to its phonological representation (in featural terms, it is represented as [+continuant]), then speakers may treat [ʁ]–ejective or [ʁ]–aspirate combinations as grammatical, since they represent an attested phonetic combination. If the phonological patterning of [ʁ] is taken into account, and speakers have set up a more abstract phonological representation with a uvular stop /q/ (in featural terms, it is represented as [−continuant, −sonorant]), then [ʁ]–ejective and [ʁ]–aspirate should be treated like other stop–ejective and stop–aspirate combinations. Experiments 1 and 2 both focus on the stop–ejective and [ʁ]–ejective combinations; this choice was arbitrary, and the same results would be expected for stop–aspirate and [ʁ]–aspirate combinations.

3 Experiment 1: repetition

The first experimental paradigm was a repetition task, in which participants were asked to listen to and reproduce nonsense forms. To investigate the representation of [ʁ], responses to phonotactically illegal forms with [ʁ]–ejective and [k]–ejective pairs were compared to one another and to controls.

3.1 Participants

The participants were 22 Quechua–Spanish bilinguals, all of whom had grown up in Quechua-speaking communities in the Cochabamba department of Bolivia, and reported being equally comfortable in the two languages.

Data from three participants were removed from analysis, one because of a recording error, and two because they did not produce any ejectives. Of the remaining 19 participants, two were included with incomplete data;

https://doi.org/10.1017/S0952675719000034 Published online by Cambridge University Press

one participant skipped many items without repeating them (82% of the stimuli had responses) and for another participant there was a recording error (63% of the stimuli had responses).

The 19 participants were seven males and twelve females, aged 18–53 (one participant was 53, all others were between 18 and 23). All participants were affiliated with the applied linguistics programme at Universidad Mayor San Simón in Cochabamba, and were paid a small sum for participating. Participants were in their first semester of this programme, and had had limited exposure to linguistic analysis or linguistic descriptions of Quechua (moreover, descriptions of Quechua for pedagogical purposes do not mention the laryngeal restrictions, and even highly educated people are unaware of these restrictions).

3.2 Methods

3.2.1 *Stimuli*. The stimuli were 90 $(C_1)V(C)C_2V$ nonce words. There were three classes of target stimuli. Items in the CONTROL category ($n = 15$) had a fricative or sonorant in C_1 and an ejective in C_2; the ejective is phonotactically legal in these forms. Items in the VELAR–EJECTIVE category ($n = 15$) had [k] as C_1 and a medial ejective, and items in the UVULAR–EJECTIVE category ($n = 15$) had [ʁ] as C_1 and a medial ejective; items in both of these categories were phonotactically illegal. Filler items ($n = 45$) were phonotactically legal forms with varying segmental content. All target items and some sample filler items are shown in Table II. The uvular /q/ is represented as [ʁ], and the vowels /i u/ are transcribed as mid [e o] following a uvular, representing a regular process of vowel lowering (Bills *et al*. 1969, Gallagher 2016a, Holliday & Martin 2018).

control		velar–ejective		uvular–ejective		filler	
maʎʧ'i	ʎisp'u	*kaʎʧ'i	*kisp'u	*ʁaʎʧ'i	*ʁesp'u	aki	pisa
siɲʧ'u	musp'a	*kiɲʧ'u	*kusp'a	*ʁeɲʧ'u	*ʁosp'a	ampi	maʎtu
wajʧ'i	sunt'a	*kajʧ'i	*kunt'a	*ʁajʧ'i	*ʁont'a	ʧaɴʁa	piʎmi
miʎʧ'u	ʎast'i	*kiʎʧ'u	*kast'i	*ʁeʎʧ'u	*ʁast'i	ʧaxɾi	ɾaxma
ʎajʧ'u	wit'u	*kajʧ'u	*kit'u	*ʁajʧ'u	*ʁet'u	hasku	santu
wap'a	ʎat'a	*kap'a	*kat'a	*ʁap'a	*ʁat'a	k'apu	sapi
sump'i	maʎt'i	*kump'i	*kaʎt'i	*ʁomp'i	*ʁaʎt'i	kawki	juxpi
wap'i		*kap'i		*ʁap'i		ʎip^{h}u	

Table II

Target stimuli for the repetition task, and sample filler items.

The stimuli were made from recordings of a native Quechua speaker from Cochabamba. The speaker was a 30-year-old female, equally proficient in both Quechua and Spanish. She was asked to read only phonotactically legal nonce forms, which were presented orthographically, in

https://doi.org/10.1017/S0952675719000034 Published online by Cambridge University Press

isolation. From these items, each stimulus item was made by cross-splicing together two different recordings during the closure portion of C_2. For example, a form like *[kap'i] was made by splicing [kapi] together with [wap'i], with the cut made during the labial stop closure. All items were normalised for intensity, but no other modifications were made.

3.2.2 *Procedure*. The stimuli were presented in random order, using PsyScope (http://psy.ck.sissa.it/). Participants were seated in front of a laptop computer, with a pair of Audio Technica noise-cancelling headphones. They were instructed to listen to the stimulus item, and then repeat it as precisely as possible. After repeating the item a single time, participants pressed the space bar to move on to the next trial. The experiment was self-paced, and there was no option to listen to a stimulus item more than once. Responses were recorded using a Marantz PMD660 recorder at 44k Hz, and an Audio Technica 831b lapel microphone. Instructions were given orally by the author, first in Quechua and then, if there was any confusion, in Spanish.

3.2.3 *Analysis*. Responses were transcribed phonetically by the author, relying on visual examination of waveforms and spectrograms where necessary, and coded for accuracy and type of error, if any. Accuracy and error types on the three categories of target items were then compared. For control items, errors were misproductions of a medial ejective as a plain stop, e.g. [musp'a] produced as [muspa]. For velar–ejective items, there were three types of errors: misproduction of a medial ejective as a plain stop, e.g. *[kat'i] produced as [kati] (coded as 'C_2'), misproduction of the form with an initial ejective instead of a medial ejective, e.g. *[kat'i] produced as [k'ati] (coded as 'reassociation'), or misproduction of both stops as ejective, e.g. *[kat'i] produced as *[k'at'i] (coded as 'double'; note that this last error does not repair the phonotactic violation). For uvular–ejective items, there were just two error types: misproduction of a medial ejective as a plain stop, e.g. *[ʁat'i] produced as [ʁati] (C_2), or misproduction of the form with an initial ejective instead of a medial ejective, e.g. *[ʁat'i] produced as [q'ati] (reassociation).

3.3 Results

Accuracy in repeating the three classes of target stimuli, uvular–ejective, velar–ejective and control, is shown in Fig. 2, along with the distribution of different error types (discussed below). On average, accuracy was high on controls (94%), and lower on the uvular–ejective and velar–ejective stimuli, showing a distinction between phonotactically legal and illegal structures. Accuracy for uvular–ejective forms (64%), where the initial segment is a phonetic continuant, was higher than accuracy for velar–ejective forms (46%), where the initial segment is a phonetic stop, showing that the phonetic distinction in manner of articulation between the uvular and velar also has an impact on accurate repetition.

https://doi.org/10.1017/S0952675719000034 Published online by Cambridge University Press

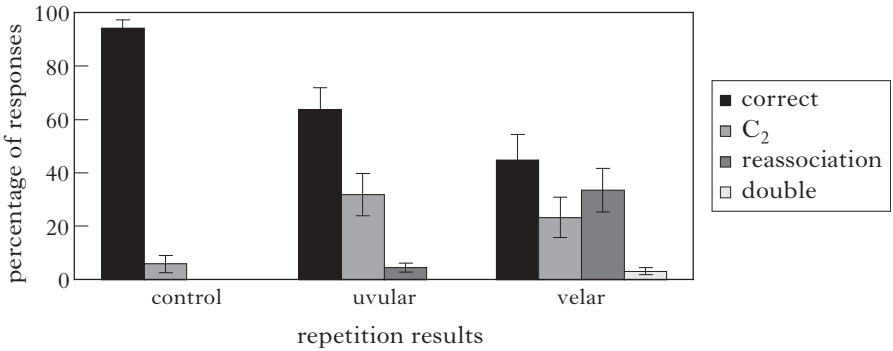


Figure 2

Responses on the repetition task, comparing correct responses as well as the three error types. Error bars represent standard error.

A binomial regression model with accuracy (correct or incorrect) as the dependent variable was fitted using the *lmer* function of the *lme4* package (Bates *et al*. 2015) in R (R Core Team 2014). The model had a single predictor of type, with uvular–ejective as the baseline, and random intercept and slopes by subject. The model finds that accuracy is significantly higher on control forms than on uvular–ejective forms ($\beta = 3.27$, SE = 1.00, $z = 3.27$, $p = 0.001$) and significantly lower on velar–ejective forms than on uvular–ejective forms ($\beta = -1.64$, SE = 0.64, $z = -2.55$, $p = 0.01$).

Among inaccurate responses, the frequency of error types differed between uvular–ejective and velar–ejective stimuli. For velar–ejective stimuli, the most common error (32% of all responses) was to produce the first stop as ejective rather than the second (*[kap'i] produced as [k'api]; reassociation in Fig. 2), though errors where there was no ejective at all (*[kap'i] produced as [kapi]; C_2) were also common (18%). Additionally, there were a small number of errors (4%) where both stops were produced as ejective (*[kap'i] produced as *[k'ap'i]; double); this change does not repair the phonotactic violation. For uvular–ejective stimuli, almost all errors (32% of responses) had no ejective (*[ʁap'i] produced as [ʁapi]). Errors where the initial stop rather than the second was ejective (*[ʁap'i] produced as [q'api]) were attested, but infrequent (4%), and there were no errors at all where both stops were produced as ejective.

3.4 Discussion

The higher error rate on both uvular–ejective and velar–ejective forms than on controls is consistent with these forms being phonotactically illegal, suggesting that Quechua speakers have learned restrictions on both [ʁ]–ejective and stop–ejective combinations. The repetition task

https://doi.org/10.1017/S0952675719000034 Published online by Cambridge University Press

also finds differences between uvular–ejective and velar–ejective forms, suggesting a role for the phonetic distinction in manner of articulation. There are at least two hypotheses about how phonetic and phonological knowledge interact in this task to produce the observed results.

On one hypothesis, both uvular–ejective and velar–ejective forms are penalised by a general phonotactic constraint on stop–ejective combinations, where the class of stops includes [ʁ]. Under this hypothesis, [ʁ] is specified phonologically as [−continuant, −sonorant], and the grammar contains a constraint on non-adjacent [−continuant, −sonorant]–[+constricted glottis] combinations. The difference between uvular–ejective and velar–ejective forms is not due to a difference in the phonotactic grammar, but rather is an effect that arises in either perception or production. The greater error rate on velar–ejective forms is a cumulative effect of both a phonotactic restriction and phonetic difficulty with perceiving and/or producing stop–ejective sequences (due to these sequences being unattested), while the error rate on uvular–ejective forms reflects only the phonotactic restriction. Given that stop–ejective sequences are unattested in the language, speakers may have trouble accurately identifying the ejective host when there are two possibilities and/or in coordinating a laryngeal gesture with the second of two stop closures. Since uvular–ejective forms are phonetically continuant–ejective, a phonetic combination speakers are familiar with, they do not pose the same perceptual or articulatory difficulties.

On a competing hypothesis, the difference between uvular–ejective and velar–ejective forms is due to the phonotactic grammar. The phonetic identity of [ʁ] as a continuant means that it cannot be grouped into the class of stops, and uvular–ejective forms are not penalised by a general constraint on stop–ejective combinations. Under this hypothesis, [ʁ] is specified phonologically as [+continuant], and is thus outside the scope of a phonotactic constraint on [−continuant, −sonorant]–[+constricted glottis] combinations. While velar–ejective forms are penalised by such a constraint, uvular–ejective forms must be penalised by a distinct, ad hoc, constraint on [ʁ]–ejective forms. This [ʁ]–ejective constraint is relatively weaker than the stop–ejective constraint, both because it is less general (has scope over fewer segmental combinations) and because it is phonetically and featurally arbitrary. While the repetition data are compatible with both hypotheses, the results of the forced-choice task presented in the next section support the first view.

The role of phonetic substance in the repetition task is further reflected in the distribution of error types in the repetition task. Reassociation of the ejective articulation to the initial stop is the preferred repair for velar–ejective forms (*[kap'i] is produced as [k'api] more often than as [kapi]), but complete removal is the preferred repair for uvular–ejective forms (*[ʁap'i] is produced as [ʁapi] more often than [q'api]). The perception study reported in Gallagher (2016b) found that reassociation of the ejective articulation is less perceptible to Quechua speakers than its complete removal from a form; for example, [kap'i]–[k'api] is a less distinct

https://doi.org/10.1017/S0952675719000034 Published online by Cambridge University Press

pair than [kap'i]–[kapi]. The greater frequency of reassociation errors than de-ejectivisation errors on velar–ejective forms correlates with this perceptual asymmetry: participants make the perceptually more minimal modification. In uvular–ejective forms, however, realising the initial consonant as ejective involves mapping a continuant [ʁ] to an ejective [q'], a perceptually larger change than mapping a stop [k] to an ejective [k'], which is only found in 4% of responses. These asymmetries in error type are consistent with the P-map hypothesis (Steriade 2009), which predicts that the likelihood of an unfaithful mapping will reflect the perceptual distance between faithful and unfaithful forms.

4 Experiment 2: forced choice

The second experimental paradigm was a forced-choice task, where participants were asked to listen to pairs of nonsense words and choose which one sounded more natural as a hypothetical word of Quechua. This task involves perception, but not production, and also asks for a more abstract response from participants than the repetition task. As in the repetition task, treatment of [ʁ]–ejective and [k]–ejective pairs were compared both to one another and to controls.

4.1 Participants

The participants were 26 Quechua–Spanish bilinguals, 19 of whom had also completed the repetition task. For participants who completed both tasks, the order in which they did the two tasks was balanced. There were 16 males and ten females, aged 18–53 (one participant was 53, all others were between 18 and 23). They were paid a small sum for participating.

4.2 Methods

4.2.1 *Stimuli.* The stimuli were 60 $C_xV(C)C_yV$ nonce word pairs, where the two items in the pair differed in the position of the two onset consonants (e.g. [sunt'a]–[t'unsa]). There were three classes of target pairs. Pairs in the CONTROL category ($n = 12$) contrasted the position of an ejective and a fricative/sonorant; both members of the pair were phonotactically legal. Pairs in the VELAR–EJECTIVE ($n = 12$) category contrasted the position of an ejective and [k], and pairs in the UVULAR–EJECTIVE ($n = 12$) category contrasted the position of an ejective and [ʁ]; in these pairs, one member was phonotactically legal and one was phonotactically illegal. There were also 24 filler pairs that did not include ejectives. Sample items from each category are shown in Table III. Nasal codas reflect place assimilation, and vowel height reflects the position of the uvular, if present.

As for the repetition task, the stimuli were made from recordings of a native Quechua speaker from Cochabamba producing phonotactically legal items. From these recordings, each stimulus item was made by

https://doi.org/10.1017/S0952675719000034 Published online by Cambridge University Press

control	velar–ejective	uvular–ejective	filler
maʎʧ'i–ʧ'aʎmi	*kaʎʧ'i–ʧ'aʎki	*ʁaʎʧ'i–ʧ'aʎʁe	maʎtu–taʎmu
siɲʧ'u–ʧ'insu	*kiɲʧ'u–ʧ'iŋku	*ʁeɲʧ'a–ʧ'eɴʁa	ʧaxɾi–ɾaxʧi
sump'i–p'unsi	*kump'i–p'uŋki	*ʁomp'i–p'uɴʁe	ʎatʃu–ʧaʎu
wap'i–p'awi	*kap'i–p'aki	*ʁap'i–p'aʁe	sapi–pisa
ʎast'i–t'asʎi	*kast'i–t'aski	*ʁast'i–t'asʁe	ʧinsu–siɲʧu
sunt'a–t'unsa	*kunt'a–t'uŋka	*ʁont'a–t'oɴʁa	sumpi–punsi

Table III
Sample stimuli for the forced-choice task.

cross-splicing together two different recordings during the closure portion of C_2. All items were normalised for intensity, and no other modifications were made.

4.2.2 *Procedure.* Stimulus pairs were presented in random order, using PsyScope. The order of stimuli within a pair was also randomised. Participants were seated in front of a laptop computer with headphones. They listened to the two stimulus items in a pair, and were then asked to choose which one of the two sounded more natural as a hypothetical word of Quechua. Stimulus items were not presented orthographically. A choice for the first word was indicated by pressing the '1' key on the computer, and a choice for the second word by pressing the '0' key. The experiment was self-paced, and there was no option to listen to the stimuli more than once. Instructions were given orally by the author, in both Quechua and Spanish.

4.3 Results

Responses to filler items were not analysed. Responses to target items were coded for whether participants chose the item with the ejective in initial position or not. Figure 3 shows that in control pairs (e.g. [sunt'a] *vs.* [t'unsa]), where both the initial and medial ejective are phonotactically legal, participants preferred the medial ejective overall. This preference was reversed for uvular–ejective (e.g. *[ʁast'i] *vs.* [t'asʁe]) and velar–ejective pairs (e.g. *[kast'i] *vs.* [t'aski]), where only the initial ejective was phonotactically legal.

A binomial regression model with initial ejective preference (chose initial ejective or not) as dependent variable was fitted, again using the *lmer* function in the *lme4* package. The model had a single predictor of type, with uvular–ejective as the baseline, and random intercept and slopes by subject. The model finds that the preference for the initial ejective is lower in control pairs than in uvular–ejective pairs ($\beta = -0.59$, SE = 0.22, $z = -2.68$, $p = 0.007$), but that the preference for the initial ejective

https://doi.org/10.1017/S0952675719000034 Published online by Cambridge University Press

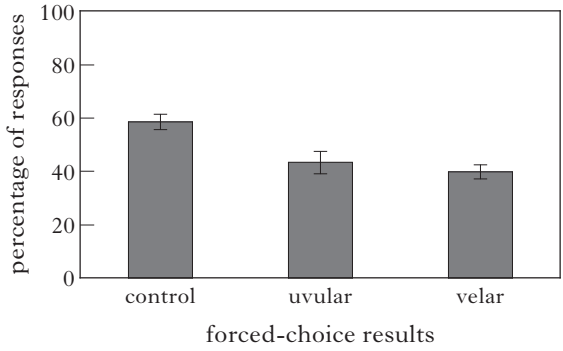


Figure 3
Percentage of responses that preferred the form with the medial ejective in the forced-choice task. Error bars indicate standard error.

does not differ between uvular–ejective and velar–ejective pairs ($\beta = 0.10$, SE = 0.19, $z = 0.53$, $p = 0.60$).

4.4 Discussion

The results of the forced-choice task echo those of the repetition task: both provide support for the presence of phonotactic restrictions on both stops and [ʁ] in speakers' synchronic grammars. Unlike in the repetition task, the forced-choice task did not find any difference between uvular–ejective and velar–ejective forms.

This difference between the two tasks is consistent with a theory in which [ʁ]–ejective and [k]–ejective forms are both penalised by a general phonotactic constraint on stop–ejective combinations, and [ʁ] is specified phonologically as [−continuant, −sonorant], like the other stops. These two structures are not distinguished by the phonotactic grammar, but they do pose different challenges in speech production that show up in the repetition task. When participants are asked to perceive and judge nonce forms, as in the forced-choice task, their responses reflect the phonotactic grammar, in which both ungrammatical structures are penalised equally. When participants are asked to perceive and reproduce nonce forms, as in the repetition task, both phonotactic well-formedness and phonetic details are reflected in error rates. Velar–ejective forms, which are phonetically stop–ejective, likely pose gestural coordination challenges in production that uvular–ejective forms, which are phonetically sonorant–ejective, do not. From an articulatory perspective, Quechua speakers have learned to time glottal gestures with the initial oral stop closure in a word, a pressure that is present for velar–ejective forms, but not for uvular–ejective forms. Furthermore, as discussed in §3.4, the repetition task involves mapping from one form to another, while the forced-choice task does

https://doi.org/10.1017/S0952675719000034 Published online by Cambridge University Press

not. The phonetic differences between uvular–ejective and velar–ejective forms may also influence mappings in the repetition task, due to considerations of perceptual similarity.

The difference between the two tasks is harder to explain under a view where the grammar contains a general constraint on stop–ejective combinations and a separate, ad hoc constraint on [ʁ]–ejective combinations. Under this hypothesis, the repetition task would need to be considered a more sensitive task, reflecting smaller grammatical distinctions among ungrammatical forms, while the forced-choice task reflects only the larger differences between grammatical and ungrammatical forms. While this is possible, the existing literature on metalinguistic tasks suggests that these tasks are quite sensitive (Daland *et al*. 2011, Kager & Pater 2012), as they often reveal subtle differences between ungrammatical structures.

Two other comments, which do not bear on the main point of the paper, can be made about the results of the forced-choice task. First, the preference in control pairs is for the form with the medial ejective, not the initial ejective. This is surprising, given that initial ejectives are almost twice as frequent as medial ejectives in the language (in a list of 1083 roots in Laime Ajacopa 2007, there are 212 initial ejectives and 130 medial ejectives). One explanation for the preference for medial ejectives could be that speakers prefer a cluster with a stop in C_2 (cf. syllable-contact effects; Gouskova 2004) – clusters with a stop in C_2 make up 75% of the 467 clusters in the roots list. A second potentially surprising result is that the preference for the grammatical form in the velar–ejective and uvular–ejective categories is somewhat small. Here, both forms have a cluster with a stop in C_2, and participants still only choose the form with an attested onset combination and an initial ejective 57–60% of the time. It is not clear why the preference is not stronger.

5 Interim summary and comparison with previous work

Taken together, the results of the repetition and forced choice tasks favour an interpretation where [ʁ] is grouped with the stops and [ʁ]–ejective combinations are grammatically restricted, as part of a single phonotactic constraint on stop–ejective combinations. In terms of distinctive features, there is a level of representation at which [ʁ] is specified as [−continuant, −sonorant], and thus is subject to a phonotactic constraint on [−continuant, −sonorant]–[+constricted glottis] combinations. This interpretation is further supported by a comparison of the current results with the results of similar prior studies. Previous experimental work with Quechua speakers has found that while unattested stop combinations consistently correlate with repetition and perception errors, not all unattested structures do.

In addition to the restrictions on stop–ejective and stop–aspirate combinations, [h]–aspirate combinations are also categorically unattested in

https://doi.org/10.1017/S0952675719000034 Published online by Cambridge University Press

Quechua.[2] Gallagher (2015) compares speaker behaviour on nonce words with aspirate–aspirate and [h]–aspirate combinations for evidence of what kinds of featural classes the laryngeal restrictions in Quechua refer to. Both [h] and aspirates are [+spread glottis] segments ([h] is a fricative and thus [+continuant], while the aspirated stops are [−continuant]), but speakers treat aspirate–aspirate and [h]–aspirate combinations differently. On a repetition task, nonce words with pairs of aspirates (e.g. *[k^{h}ap^{h}i]) are repeated with errors that repair the phonotactic violation a significant portion of the time, while nonce words with [h]–aspirate combinations (e.g. *[hap^{h}i]) are repeated accurately, and do not differ from controls ([map^{h}i]). Repetition errors on all unattested combinations of stops are also found in Gallagher (2013, 2014, 2016b). This pattern of errors leads Gallagher (2015) to propose that speakers' grammars contain a markedness constraint that refers to the feature [+long VOT], which picks out aspirates (and ejectives), to the exclusion of [h]. Under this view, the restrictions on stops could be enforced by a markedness constraint on non-adjacent [−continuant, −sonorant]–[+long VOT] combinations.

The distinction between aspirate–aspirate forms on the one hand and [h]–aspirate forms on the other is relevant to the interpretation of the current results in two ways. First, it shows that errors on a repetition task do not track whether the target structure is attested or unattested, but rather can differentiate between unattested structures. Errors on a repetition task, like those found for both uvular–ejective and velar–ejective forms in the current study, can thus be considered strong evidence for the phonotactic illegality of a structure. From another angle, the current results suggest that the repetition task also does not simply track putative perceptual or articulatory challenges incurred by forms with multiple stops. Given that [ʁ]–ejective combinations correlate with significant repetition errors, the claim in previous work that repetition errors on unattested stop combinations have a phonological source is supported.

Second, the null finding for unattested [h]–aspirate combinations contrasts with the repetition errors found for uvular–ejective forms. The substantial error rate on uvular–ejective forms in the current study is unlikely to have arisen simply because uvular–ejective combinations are unattested or because of a segment-specific restriction on [ʁ]–ejective pairs. If [ʁ] is a member of the class of stops and [ʁ]–ejective combinations are penalised by a general stop–ejective constraint, then the different effects for [ʁ]–ejective compared to [h]–aspirate can be explained: the more general stop–ejective restriction is stronger than the more specific [h]–aspirate restriction (e.g. they have a higher weight in a weighted constraint grammar). If [ʁ]–ejective combinations are penalised by a distinct, ad hoc, constraint, however, this difference is harder to explain. Putative restrictions on [h]–aspirate and [ʁ]–ejective combinations have scope over the same number of

[2] [ʔ]–ejective combinations are absent as well. However, [ʔ] is non-phonemic in Quechua, though it is impressionistically reported to be present before vowel-initial words. See Gallagher (2015) for supporting acoustic evidence for a glottal stop in certain phrasal positions.

https://doi.org/10.1017/S0952675719000034 Published online by Cambridge University Press

segments, and have similar formal structure from the perspective of a statistical learner. From a phonetic standpoint, [ʁ] is unrelated to ejectives, while [h] shares a laryngeal constriction with aspirates. The [h]–aspirate restriction is thus featurally natural, so it may be expected to be stronger than a featurally unnatural [ʁ]–ejective constraint. The finding that [ʁ]–ejective combinations trigger repetition errors while [h]–aspirate combinations do not runs counter to this prediction, and further supports a representation where [ʁ] belongs to the class of stops.

To summarise, the experimental results support a synchronic restriction on [ʁ]–ejective combinations. I have further argued that the results support a general restriction on stop–ejective pairs, with the class of stops including [ʁ]. That is, the experiments are interpreted as supporting a constraint *[−continuant, −sonorant]–[+long VOT], where [−continuant, −sonorant] includes [ʁ].[3] The remainder of this paper looks at how Quechua speakers could learn that the surface continuant [ʁ] is grouped with the stops.

6 The role of phonological and phonetic evidence for sound classes

Unnatural classes like the class of 'stops' in Quechua motivate the proposal that phonetic and phonological information interact in constructing phonological representations. While some sound classes may be learned from shared phonetic properties, the phonological patterning of segmental categories must also be evaluated and considered in constructing phonological representations (as proposed in Mielke 2008). This section sketches how phonetic and phonological information could inform the learning of the phonetically unnatural class of stops in Quechua, based on what is known about early phonological learning.

The ubiquity of accidental gaps in natural languages means that phonological patterns do not emerge clearly from looking at the distribution of individual segments. Distinguishing between accidental and systematic gaps has long been recognised as a major problem in learning a phonotactic grammar (Fischer-Jørgensen 1952, Chomsky & Halle 1968, Iverson & Salmons 2005, Hayes & White 2013), and the classic method for determining whether a gap is systematic or accidental is to reference natural classes. When an absent structure can be captured as part of a broader restriction on phonetically definable classes of segments, it is considered a systematic gap, and if the absent structure must be stipulated at the segmental level, it is typically considered an accidental gap. Using natural classes to guide phonotactic learning is a critical component of the success of inductive

[3] The grammar could equally well be said to contain two constraints, *[−continuant, −sonorant]–[+constricted glottis] (against stop–ejective combinations) and *[−continuant, −sonorant]–[+spread glottis, −continuant] (against stop–aspirate combinations); using [+long VOT] allows for a featurally simple way of distinguishing the laryngeally marked stops from [h].

https://doi.org/10.1017/S0952675719000034 Published online by Cambridge University Press

C_1	attested C_2	unattested C_2
p	p t ʧ k ʁ s m n ɲ l ʎ ɾ w	p^h p' t^h t' $ʧ^h$ ʧ' k^h k' q^h q' **j**
p^h	p t ʧ k s m ɲ ʎ ɾ w j	p^h p' t^h t' $ʧ^h$ ʧ' k^h k' **ʁ** q^h q' **n l**
p'	p t ʧ k ʁ ɲ l ɾ w	p^h p' t^h t' $ʧ^h$ ʧ' k^h k' q^h q' **s m n ʎ j**
t	p t k ʁ s m n ʎ ɾ w j	p^h p' t^h t' **ʧ** $ʧ^h$ ʧ' k^h k' q^h q' **ɲ l**
t^h	t k ʁ s m n ɲ l ʎ ɾ w j	**p** p^h p' t^h t' **ʧ** $ʧ^h$ ʧ' k^h k' q^h q' **ɾ j**
t'	p t k ʁ s n l ʎ ɾ w j	p^h p' t^h t' **ʧ** $ʧ^h$ ʧ' k^h k' q^h q' **m ɲ**
ʧ	p t ʧ k ʁ m n l ʎ ɾ w j	p^h p' t^h t' $ʧ^h$ ʧ' k^h k' q^h q' **s ɲ**
$ʧ^h$	p t ʧ k ʁ s m l ʎ ɾ	p^h p' t^h t' $ʧ^h$ ʧ' k^h k' q^h q' **n ɲ w j**
ʧ'	p ʧ k ʁ s ɲ m l ʎ ɾ w j	p^h p' **t** t^h t' $ʧ^h$ ʧ' k^h k' q^h q' **n**
k	p t ʧ k s m n l ʎ ɾ w j	p^h p' t^h t' $ʧ^h$ ʧ' k^h k' **ʁ** q^h q' **ɲ**
k^h	p t ʧ k m n ɲ l ʎ ɾ w j	p^h p' t^h t' $ʧ^h$ ʧ' k^h k' **ʁ** q^h q' **s l ʎ**
k'	p t ʧ k s m n ɲ ʎ ɾ w j	p^h p' t^h t' $ʧ^h$ ʧ' k^h k' **ʁ** q^h q' **l**
ʁ	p t ʧ ʁ s m n ɲ ʎ ɾ w	p^h p' t^h t' $ʧ^h$ ʧ' **k** k^h k' q^h q' **l j**
q^h	p t ʧ ʁ s ɲ ʎ ɾ w j	p^h p' t^h t' $ʧ^h$ ʧ' **k** k^h k' q^h q' **m n l**
q'	p t ʧ s m ɲ l ʎ ɾ w j	p^h p' t^h t' $ʧ^h$ ʧ' **k** k^h k' **ʁ** q^h q' **n**

Table IV

Attested and unattested onset–onset pairs in Quechua roots, with a stop as C_1. Bolding indicates unattested structures that do not violate the laryngeal co-occurrence constraints.

phonotactic models (e.g. Hayes & Wilson 2008, Futrell *et al*. 2017), and Wilson & Gallagher (2018) show that models without natural classes fail to distinguish systematic and accidental gaps.

The trouble with establishing the class of stops on the basis of the distribution of individual segments in Quechua is illustrated below. Table IV gives the attested and unattested onset–onset combinations in a list of 1083 roots from Laime Ajacopa (2007), in the subset of cases where C_1 is a stop (this discussion leaves aside the question of how a learner discovers that non-adjacent onset consonants interact). Unattested combinations that do not fall under the purview of the laryngeal restrictions are given in bold. Table IV shows that the class of stops does not emerge straightforwardly from looking at the distribution of the individual segments that comprise this class: not all segments can be followed by the same set of C_2's. One reason for distributional differences is the co-occurrence restriction on velars and uvulars; while most stops can be followed by all other plain stops, [k k^h k'] cannot be followed by [ʁ], and [ʁ q^h q'] cannot be followed by [k]. A second reason that the distributions are not exactly the same is because there are accidental gaps, e.g. [p'] is not followed by [m] or [ʎ]. The simple exercise of listing co-occurrence patterns based on segments reveals the fundamental problem with deriving classes from segmental distributions: given the prevalence of accidental

https://doi.org/10.1017/S0952675719000034 Published online by Cambridge University Press

gaps in natural language, not all members of a class will have the same distribution.

The expected co-occurrences for a subset of the unattested onset combinations are presented in Table V. The expected values take into account the frequency of individual segments, and are an additional metric for determining whether an absent structure is systematically or accidentally unattested. Wilson & Gallagher (2018) demonstrate that a statistical calculation and class representations are *both* necessary to distinguish systematic and accidental gaps, and expected values are used as a reference point in several inductive learning models (Albright & Hayes 2003, Hayes & Wilson 2008). Table Va shows the unattested plain–ejective and plain–aspirate combinations, while Table Vb shows the expected values for a subset of the accidental gaps.

The key point in comparing Tables Va and b is that the expected values in Table Va, which correspond to systematic gaps, are not uniformly higher than the expected values in Table Vb, which correspond to accidental gaps. Unattested stop–ejective or stop–aspirate combinations have expected values ranging from 0.12 for [t^h–t^h] to 2.68 for [ʧ'–k']. No accidental gaps have values so low, the lowest being 0.39 for [k^h–l], and some have values that are quite high, e.g. 3.43 for [ʧ–s]. Even under a more statistically sensitive view of the distribution of individual segments, the systematically absent combinations do not emerge as a distinct group, and thus cannot form the basis for identifying the class of stops as a set of sounds with the same distribution.

From a learning perspective, then, it is not clear how looking at the distribution of segments would lead to the observation that all stops have the same distribution, and thus form a class, nor is it clear how a learner would distinguish systematic restrictions on laryngeal features from accidental gaps. Distinguishing the accidental and systematic gaps in Quechua, as in other languages, requires reference to classes. Indeed, as stated above, the existence of a distinction between systematic and accidental gaps at all is largely contingent upon a class-based evaluation. The co-occurrence restrictions on stops followed by ejectives and aspirates, or on uvular–velar combinations, are considered systematic gaps, and thus part of the phonotactic grammar, because they can be stated over classes of segments and hold of all members of those classes. In contrast, an unattested combination like [p'–m] is considered an accidental gap, because it is just this individual segment pair that is absent: the other labial stops, [p p^h], can combine with [m]. Behavioural and neurological studies have shown that speakers distinguish between systematic and accidental gaps defined in this way (Frisch & Zawaydeh 2001, Rose & King 2007, Gwilliams & Marantz 2015).

To acquire phonological classes, I suggest that learners must have access both to individual segmental distributions and to the distribution of phonetically defined sound classes. In Quechua, an evaluation of the distribution of phonetically definable laryngeal and manner classes will show that stop–ejective and stop–aspirate combinations are unattested, while

https://doi.org/10.1017/S0952675719000034 Published online by Cambridge University Press

(a)

	pʼ	tʼ	ʧʼ	kʼ	qʼ	pʰ	tʰ	ʧʰ	kʰ	qʰ
p	0.72	1.27	0.89	1.87	0.76	0.59	0.38	0.25	0.72	0.55
t	0.63	1.11	0.78	1.63	0.66	0.52	0.33	0.22	0.63	0.48
ʧ	0.83	1.47	1.03	2.15	0.88	0.69	0.44	0.29	0.83	0.64
k	0.67	1.19	0.83	1.75	0.71	0.56	0.36	0.24	0.67	0.52
ʁ	0.52	0.91	0.64	1.34	0.55	0.43	0.27	0.18	0.52	0.40
pʼ	0.42	0.75	0.52	1.1	0.45	0.35	0.22	0.15	0.42	0.32
tʼ	0.61	1.08	0.76	1.58	0.65	0.50	0.32	0.22	0.61	0.47
ʧʼ	1.04	1.83	1.28	2.68	1.10	0.85	0.55	0.37	1.04	0.79
kʼ	0.77	1.36	0.95	1.99	0.81	0.63	0.41	0.27	0.77	0.59
qʼ	0.49	0.86	0.6	1.26	0.52	0.40	0.26	0.17	0.49	0.37
pʰ	0.47	0.83	0.58	1.22	0.5	0.39	0.25	0.17	0.47	0.36
tʰ	0.22	0.39	0.27	0.57	0.23	0.18	0.12	0.08	0.22	0.17
ʧʰ	0.35	0.61	0.43	0.89	0.37	0.28	0.18	0.12	0.35	0.26
kʰ	0.39	0.69	0.48	1.02	0.42	0.32	0.21	0.14	0.39	0.30
qʰ	0.41	0.72	0.50	1.06	0.43	0.34	0.22	0.14	0.41	0.31

(b)

p–j	1.36	t–ʧ	3.03	tʼ–ɲ	0.90	k–ʁ	2.94	ʁ–l	0.52
pʰ–ʁ	2.05	t–ɲ	0.92	ʧ–s	3.43	k–ɲ	0.99	ʁ–j	0.98
pʰ–n	1.02	t–l	0.63	ʧ–ɲ	1.22	kʰ–ʁ	1.71	qʰ–k	2.74
pʰ–l	0.47	tʰ–p	1.28	ʧʰ–n	1.81	kʰ–s	1.62	qʰ–m	1.20
pʼ–s	1.75	tʰ–ʧ	1.06	ʧʰ–ɲ	1.22	kʰ–l	0.39	qʰ–n	0.89
pʼ–m	1.25	tʰ–ɾ	1.16	ʧʰ–w	2.64	kʰ–ʎ	1.48	qʰ–l	0.41
pʼ–n	0.92	tʰ–j	0.41	ʧʰ–j	1.57	kʼ–ʁ	3.35	qʼ–k	3.26
pʼ–ʎ	1.60	tʼ–ʧ	2.95	ʧʼ–t	4.94	kʼ–l	0.77	qʼ–ʁ	2.12
pʼ–j	0.80	tʼ–m	1.80	ʧʼ–n	2.25	ʁ–k	3.47	qʼ–n	1.06

Table V

Expected values for unattested onset combinations with a stop in C_1 and (a) an ejective or aspirate in C_2 (systematic gaps); (b) a plain stop, fricative or sonorant in C_2 (accidental gaps).

fricative–ejective, sonorant–ejective, fricative–aspirate and sonorant–aspirate combinations are attested. The distribution of [ʁ] could then be noticed as anomalous, given the distribution of the other sonorants and fricatives.

Both segmental categories and phonetically based sound classes are plausibly learned from statistical analysis of the acoustic properties of sounds (Maye *et al*. 2002, Vallabha *et al*. 2007, Lin & Mielke 2008, McMurray *et al*. 2009), and experimental work has shown that infants can learn artificial language patterns based both on segments and classes (Saffran *et al*. 1996, Saffran & Thiessen 2003, Chambers *et al*. 2003, Cristià *et al*. 2011). Cristià & Peperkamp (2012) further show that

https://doi.org/10.1017/S0952675719000034 Published online by Cambridge University Press

infants may learn patterns over phonetic classes without encoding the distribution of individual segments, and additional studies with both infants and adults find that patterns are easier to learn when stated over phonetically based classes (Saffran & Thiessen 2003, Pycha *et al.* 2003, Wilson 2003), underscoring the statistical argument above that phonetic classes form the basis of phonotactic learning.

7 Discussion and conclusion

The experimental evidence suggests that Quechua speakers have learned a phonotactic generalisation over a phonetically unnatural class, grouping [ʁ] with the other phonetic stops and affricates, and thereby providing support for the proposal that phonological representations are themselves influenced by phonological patterns. A theory of features emerging from phonological patterns is typically supported on the basis of descriptions of phonological patterns that reference unnatural classes, as in Quechua, but there is little evidence that such classes and the patterns that reference them are learned by speakers and are part of their synchronic grammars. The current results thus provide important evidence that, at least in this one case, speakers do actually learn a phonetically unnatural class.

I have not attempted to outline or implement a formal model of how sound classes are learned, but have instead shown that there is support for both phonetic and phonological information being included in a model of representational learning. To make the contributions of the various types of phonetic and phonological evidence to learned representations precise, future work needs to investigate speakers' knowledge of phonetically unnatural classes in a diverse sample of phonological systems. The Quechua case shows several properties that may (or may not) be crucial to speakers treating this sound as a stop in behavioural tasks. First, /q/ is produced as a continuant the majority of the time, but it is variably produced as a stop. Second, there are two categorical phonological patterns – coda restrictions and laryngeal restrictions – that provide top-down evidence that [ʁ] belongs in the class of stops. Third, [ʁ] fills the slot of the plain uvular in the inventory, which supports its representation as a stop under traditional phonemic analysis (Sapir 1925) and under the theory of the contrastive hierarchy (Dresher 2009, Mackenzie 2013), where features are used minimally to define the contrasts in the inventory. Comparison with other languages is needed to show which of these properties is crucial for learning a phonetically unnatural class.

The Quechua case is just one example of what is a pervasive aspect of sound systems. While phonetic descriptions of languages are characterised by variability, phonological patterns are frequently stated over a phonemic or discrete representation that abstracts over the surface variability in sound categories. Bennett *et al.* (2016) discuss another case of a laryngeal co-occurrence restriction that targets a class that is challenging to define phonetically. In Kaqchikel, pairs of ejectives are restricted from co-

https://doi.org/10.1017/S0952675719000034 Published online by Cambridge University Press

occurring. This generalisation holds of the consistently ejective [t' ʦ' ʧ' k'] as well as /q'/, which is variably produced as ejective [q'] or implosive [ɢ̩]. The glottalised labial is also variably produced as ejective [p'] or implosive [ɓ], but is not subject to the co-occurrence restriction, and may combine with other ejectives. Bennett *et al*. examine multiple acoustic correlates of these sounds, including VOT, burst amplitude and voice quality on the surrounding vowel, and find that none of them reliably distinguishes the class /t' ʦ' ʧ' k' q'/, which are targeted by the co-occurrence restriction, from /ɓ/, which is not. Bennett *et al*. propose that the articulatory alignment of the glottal gesture with the closure or release phase of the oral stop gesture is a potential correlate of this distinction, which learners may use to set up phonological representations. An additional factor in learning the relevant class of segments is likely their phonological patterning. The phonological literature is rife with other similar examples. In French, the rhotic phoneme patterns as a liquid, but is produced in some contexts as a voiceless fricative [χ] (Fougeron & Smith 1993). The lateral sonorant [l] must be included in the class of voiced stops in Setswana, where it is a contextual variant of /d/ (Zsiga & Boyer 2017, Zsiga 2018), and in Taiwanese, where etymological /d/ has undergone a full change in all contexts, resulting in the class [b l g] (Pan 2004, Wang 2018); see Mielke (2005) for more on the ambiguous phonological status of laterals.

Inductive learning models of sound categories (Vallabha *et al*. 2007, McMurray *et al*. 2009, Dillon *et al*. 2013, Feldman *et al*. 2013, Martin *et al*. 2013) and phonological patterns (Albright & Hayes 2003, Hayes & Wilson 2008, Peperkamp *et al*. 2006, Calamaro & Jarosz 2015) have shown that the input data to learners contains a wealth of information about the formal structure of the language, but the link between these two learning procedures has yet to be worked out in full. A pressing task for future work is to develop an explicit, implemented model of how phonological sound classes emerge from the phonetic signal to form the basis of, and to be informed by, phonological patterns.

REFERENCES

Albright, Adam & Bruce Hayes (2003). Rules *vs*. analogy in English past tenses: a computational/experimental study. *Cognition* **90**. 119–161.

Anderson, Stephen R. (1981). Why phonology isn't 'natural'. *LI* **12**. 493–539.

Avery, Peter & William J. Idsardi (2001). Laryngeal dimensions, completion and enhancement. In T. Alan Hall (ed.) *Distinctive feature theory*. Berlin & New York: Mouton de Gruyter. 41–70.

Bach, Emmon & Robert T. Harms (1972). How do languages get crazy rules? In Robert P. Stockwell & Ronald K. S. Macaulay (eds.) *Linguistic change and generative theory*. Bloomington: Indiana University Press. 1–21.

Bates, Douglas, Martin Mächler, Benjamin M. Bolker & Steven C. Walker (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software* **67**. 1–48.

https://doi.org/10.1017/S0952675719000034 Published online by Cambridge University Press

Bennett, Ryan, Kevin Tang & Juan Ajsivinac Sian (2016). Against phonetic realism as the source of root co-occurrence restrictions. Paper presented at the 2016 Annual Meeting on Phonology, University of California, Los Angeles. Slides available (November 2018) at https://www.kevintang.org/projects/kaqchikel_phonotactics/Files/Slides/Bennett_Tang_AMP2016.pdf.
Bills, Garland D., Bernardo Vallejo C. & Rudolph C. Troike (1969). *An introduction to spoken Bolivian Quechua*. Austin: University of Texas Press.
Boersma, Paul & David Weenink (2016). *Praat: doing phonetics by computer* (version 6.0.21). http://www.praat.org.
Calamaro, Shira & Gaja Jarosz (2015). Learning general phonological rules from distributional information: a computational model. *Cognitive Science* **39**. 647–666.
Chambers, Kyle E., Kristine H. Onishi & Cynthia Fisher (2003). Infants learn phonotactic regularities from brief auditory experience. *Cognition* **87**. B69–B77.
Chomsky, Noam & Morris Halle (1968). *The sound pattern of English*. New York: Harper & Row.
Clements, G. N. (2003). Feature economy in sound systems. *Phonology* **20**. 287–333.
Cristià, Alejandrina, Grant L. McGuire, Amanda Seidl & Alexander L. Francis (2011). Effects of the distribution of acoustic cues on infants' perception of sibilants. *JPh* **39**. 388–402.
Cristià, Alejandrina & Sharon Peperkamp (2012). Generalizing without encoding specifics: infants infer phonotactic patterns on sound classes. In Alia K. Biller, Esther Y. Chung & Amelia E. Kimball (eds.) *Proceedings of the 36th Annual Boston University Conference on Language Development*. Somerville, Mass.: Cascadilla. 126–138.
Daland, Robert, Bruce Hayes, James White, Marc Garellek, Andrea Davis & Ingrid Norrmann (2011). Explaining sonority projection effects. *Phonology* **28**. 197–234.
Dillon, Brian, Ewan Dunbar & William Idsardi (2013). A single-stage approach to learning phonological categories: insights from Inuktitut. *Cognitive Science* **37**. 344–377.
Dresher, B. Elan (2009). *The contrastive hierarchy in phonology*. Cambridge: Cambridge University Press.
Feldman, Naomi H., Emily B. Myers, Katherine S. White, Thomas L. Griffiths & James L. Morgan (2013). Word-level information influences phonetic learning in adults and infants. *Cognition* **127**. 427–438.
Fischer-Jørgensen, Eli (1952). On the definition of phoneme categories on a distributional basis. *Acta Linguistica Hafniensia* **7**. 8–39.
Flemming, Edward (2002). *Auditory representations in phonology*. London & New York: Routledge.
Flemming, Edward (2004). Contrast and perceptual distinctiveness. In Bruce Hayes, Robert Kirchner & Donca Steriade (eds.) *Phonetically based phonology*. Cambridge: Cambridge University Press. 232–276.
Flemming, Edward (2005). Deriving natural classes in phonology. *Lingua* **115**. 287–309.
Fougeron, Cécile & Caroline L. Smith (1993). Illustrations of the IPA: French. *Journal of the International Phonetic Association* **23**. 73–76.
Frisch, Stefan A. & Bushra Zawaydeh (2001). The psychological reality of OCP-place in Arabic. *Lg* **77**. 91–106.
Futrell, Richard, Adam Albright, Peter Graff & Timothy J. O'Donnell (2017). A generative model of phonotactics. *Transactions of the Association for Computational Linguistics* **5**. 73–86.
Gallagher, Gillian (2013). Speaker awareness of non-local laryngeal phonotactics in Cochabamba Quechua. *NLLT* **31**. 1067–1099.
Gallagher, Gillian (2014). An identity bias in phonotactics: evidence from Cochabamba Quechua. *Laboratory Phonology* **5**. 337–378.

https://doi.org/10.1017/S0952675719000034 Published online by Cambridge University Press

Gallagher, Gillian (2015). Natural classes in cooccurrence constraints. *Lingua* **166**. 80–98.
Gallagher, Gillian (2016a). Asymmetries in the representation of categorical phonotactics. *Lg* **92**. 557–590.
Gallagher, Gillian (2016b). Vowel height allophony and dorsal place contrasts in Cochabamba Quechua. *Phonetica* **73**. 101–119.
Gallagher, Gillian & James Whang (2014). An acoustic study of trans-vocalic ejective pairs in Cochabamba Quechua. *Journal of the International Phonetic Association* **44**. 133–154.
Gouskova, Maria (2004). Relational hierarchies in Optimality Theory: the case of syllable contact. *Phonology* **21**. 201–250.
Gwilliams, L. & A. Marantz (2015). Non-linear processing of a linear speech stream: the influence of morphological structure on the recognition of spoken Arabic words. *Brain and Language* **147**. 1–13.
Halle, Morris & Kenneth N. Stevens (1971). A note on laryngeal features. *MIT Quarterly Progress Report* **101**. 198–212.
Hayes, Bruce & James White (2013). Phonological naturalness and phonotactic learning. *LI* **44**. 45–75.
Hayes, Bruce & Colin Wilson (2008). A maximum entropy model of phonotactics and phonotactic learning. *LI* **39**. 379–440.
Holliday, Nicole & Sean Martin (2018). Vowel categories and allophonic lowering among Bolivian Quechua–Spanish bilinguals. *Journal of the International Phonetic Association* **48**. 199–222.
Iverson, Gregory K. & Joseph C. Salmons (2005). Filling the gap. *Journal of English Linguistics* **33**. 207–221.
Jakobson, Roman, C. Gunnar, M. Fant & Morris Halle (1952). *Preliminaries to speech analysis: the distinctive features and their correlates*. Cambridge, Mass.: MIT Press.
Kager, René & Joe Pater (2012). Phonotactics as phonology: knowledge of a complex restriction in Dutch. *Phonology* **29**. 81–111.
Ladefoged, Pater & Keith Johnson (2014). *A course in phonetics*. 7th edn. Stamford, Conn.: Cengage Learning.
Laime Ajacopa, Teofila, Efraín Cazazola & Félix Layme Pairumani (2007). *Diccionario bilingüe: iskay simipi yuyayk'ancha* (*Quechua–Castellano Castellano–Quechua*). 2nd edn. La Paz.
Lin, Ying & Jeff Mielke (2008). Discovering place and manner features: what can be learned from acoustic and articulatory data? *Penn Working Papers in Linguistics* **14**. 241–254. Available (November 2018) at http://repository.upenn.edu/pwpl/vol14/iss1/19/.
Mackenzie, Sara (2013). Laryngeal co-occurrence restrictions in Aymara: contrastive representations and constraint interaction. *Phonology* **30**. 297–345.
McMurray, Bob, Richard N. Aslin & Joseph C. Toscano (2009). Statistical learning of phonetic categories: insights from a computational approach. *Developmental Science* **12**. 369–378.
Martin, Andrew, Sharon Peperkamp & Emmanuel Dupoux (2013). Learning phonemes with a proto-lexicon. *Cognitive Science* **37**. 103–124.
Maye, Jessica, Janet F. Werker & LouAnn Gerken (2002). Infant sensitivity to distributional information can affect phonetic discrimination. *Cognition* **82**. B101–B111.
Mielke, Jeff (2005). Ambivalence and ambiguity in laterals and nasals. *Phonology* **22**. 169–203.
Mielke, Jeff (2008). *The emergence of distinctive features*. Oxford: Oxford University Press.
Pan, Ho-hsien (2004). Nasality in Taiwanese. *Language and Speech* **47**. 267–296.
Peperkamp, Sharon, Rozenn Le Calvez, Jean-Pierre Nadal & Emmanuel Dupoux (2006). The acquisition of allophonic rules: statistical learning with linguistic constraints. *Cognition* **101**. B31–B41.

https://doi.org/10.1017/S0952675719000034 Published online by Cambridge University Press

Pycha, Anne, Pawel Nowak, Eurie Shin & Ryan Shosted (2003). Phonological rule-learning and its implications for a theory of vowel harmony. *WCCFL* **22**. 423–435.
R Core Team (2014). *R: a language and environment for statistical computing*. Vienna: R Foundation for Statistical Computing. http://www.r-project.org.
Rose, Sharon & Lisa King (2007). Speech error elicitation and co-occurrence restrictions in two Ethiopian Semitic languages. *Language and Speech* **50**. 451–504.
Saffran, Jenny R., Richard N. Aslin & Elissa L. Newport (1996). Statistical learning by 8-month-old infants. *Science* **274**. 1926–1928.
Saffran, Jenny R. & Erik D. Thiessen (2003). Pattern induction by infant language learners. *Developmental Psychology* **39**. 484–494.
Sapir, Edward (1925). Sound patterns in language. *Lg* **1**. 37–51.
Steriade, Donca (1999). Phonetics in phonology: the case of laryngeal neutralization. *UCLA Working Papers in Linguistics: Papers in Phonology* **3**. 25–145.
Steriade, Donca (2009). The phonology of perceptibility effects: the P-map and its consequences for constraint organization. In Kristin Hanson & Sharon Inkelas (eds.) *The nature of the word: studies in honor of Paul Kiparsky*. Cambridge, Mass.: MIT Press. 151–179.
Vallabha, Gautam K., James L. McClelland, Ferran Pons, Janet F. Werker & Shigeaki Amano (2007). Unsupervised learning of vowel categories from infant-directed speech. *Proceedings of the National Academy of Sciences* **104**. 13273–13278.
Wang, Sheng-Fu (2018). Phonotactics and phonetic context in the perception of onset nasality in Taiwanese. In Gillian Gallagher, Maria Gouskova & Sora Heng Yin (eds.) *Supplemental proceedings of the 2017 Annual Meeting on Phonology*. http://dx.doi.org/10.3765/amp.v5i0.4231.
Wilson, Colin (2003). Experimental investigation of phonological naturalness. *WCCFL* **22**. 533–546.
Wilson, Colin & Gillian Gallagher (2018). Accidental gaps and surface-based phonotactic learning: a case study of South Bolivian Quechua. *LI* **49**. 610–623.
Zsiga, Elizabeth (2018). In favor of [fortis]: evidence from Setswana and Sebirwa. In Gillian Gallagher, Maria Gouskova & Sora Heng Yin (eds.) *Proceedings of the 2017 Annual Meeting on Phonology*. http://dx.doi.org/10.3765/amp.v5i0.4252.
Zsiga, Elizabeth C. & One Tlale Boyer (2017). Sebirwa in contact with Setswana: a natural experiment in learning an unnatural alternation. In Jason Kandybowicz & Harold Torrence (eds.) *Africa's endangered languages: documentary and theoretical approaches*. Oxford: Oxford University Press. 343–366.

https://doi.org/10.1017/S0952675719000034 Published online by Cambridge University Press