

Speaker awareness of non-local ejective phonotactics in Cochabamba Quechua^{*#}

Gillian Gallagher, New York University

March, 2012

Abstract

This paper presents evidence that speakers of Cochabamba Quechua are aware of non-local restrictions on ejectives in their language. A repetition task was run to investigate the synchronic status of two restrictions in Quechua: the co-occurrence restriction on ejectives, which prohibits roots with two ejectives (e.g., *[k'ap'i]), and the ordering restriction on ejectives, which prohibits roots with an initial plain stop and a medial ejective (e.g., *[kap'i]). Medial ejectives are generally attested in the language, but only occur in roots with an initial fricative or sonorant (e.g., [mat'i] 'forehead'). Native Quechua speakers were asked to repeat a mixture of real and nonsense words with medial ejectives, where the nonsense words were either phonotactically legal but unattested roots or phonotactically illegal roots that violated either the co-occurrence restriction or the ordering restriction. Medial ejectives are accurately repeated significantly more often in nonce roots where the medial ejective is phonotactically legal than when it is illegal. There is variation among subjects as to whether accuracy differs greatly between the co-occurrence and ordering category targets. Additionally, there is variation in how roots that violate the ordering restriction are repaired: both de-ejectivization, e.g., target [kap'i] produced as [kapi], and movement of ejection, e.g., target [kap'i] produced as [k'api] are common. This variation in repair strategy has implications for the formal analysis of the restriction, which must predict all well-attested repairs.

0 Introduction

Non-local restrictions on ejectives and other laryngeally marked stops are found in a variety of genetically and geographically diverse languages (MacEachern 1999). It is common, for example, for a language with a contrast between ejectives and plain stops to systematically lack roots with multiple ejectives, e.g., roots like [k'ap'i]. This paper presents evidence from a repetition task that speakers of Cochabamba Quechua are aware of two non-local phonotactic restrictions on ejectives found in their language.

Like other varieties of Quechua spoken throughout Bolivia and Peru, Cochabamba Quechua (henceforth just “Quechua”) exhibits a range of long-distance restrictions on laryngeally marked segments, two of which are investigated here. Quechua prohibits roots with multiple ejectives, e.g., *[k'ap'i], *[p'int'u], and also disallows roots with an initial plain stop and a medial ejective, e.g., *[kap'i], *[pint'u], though medial ejectives are attested in roots with initial fricative or sonorant consonants, e.g., [suʎk'a] 'younger', [mat'i] 'forehead'.

* I am grateful to Lisa Davidson, Maria Gouskova and three anonymous reviewers for their comments on this work, as well as audiences at Stonybrook University, MIT, and *NELS 42* at the University of Toronto. Further thanks are due to Sean Martin and Neil Myler for help annotating the data and to Janine and Jaquin Hinojosa for helping to recruit participants in Cochabamba. A pilot study for this project was conducted with support from NSF Dissertation Improvement Grant #0950219. This work is IRB approved by NYU UCAIHS as HS #11-8357.

This is a pre-publication draft; the article is 'to appear' in *Natural Language and Linguistic Theory*.

The synchronic status of these two restrictions is investigated by asking Quechua speakers to repeat nonce words with medial ejectives. Medial ejectives are accurately repeated significantly more often in nonce roots where the medial ejective is phonotactically legal than when it is illegal, showing that Quechua speakers are aware of these restrictions on ejectives. In addition to this general result, there are two more points of interest. First, some subjects are equally accurate on words with a plain stop-ejective onset pair as on words with an ejective-ejective pair, while for other subjects accuracy is much higher for plain stop-ejective pairs. Second, there are two common repairs for targets with a plain stop-ejective sequence: de-ejectivization of the medial stop, e.g., target [kap'i] produced as [kapi], or movement of ejection, e.g., target [kap'i] produced as [k'api]. Both of these points of variation have implications for the formal analysis of the restrictions.

The rest of this paper is organized as follows. Section 1 provides background on the laryngeal restrictions of Quechua and Section 2 sets the theoretical background for the study. The similarities and differences between the two types of illegal sequences are laid out in §2.1, and the implications of the results for the formal analysis of the restriction against plain stop-ejective sequences are discussed in §2.2. Section 3 presents the methods and results of the experiment and Section 4 concludes.

1 Long-distance laryngeal restrictions in Quechua

The study in this paper investigates the synchronic knowledge of speakers of Cochabamba Quechua, a dialect of South Bolivian Quechua. The laryngeal phonotactics of this dialect are shared with other dialects of Bolivian and Peruvian Quechua, including Cuzco Quechua as described in Parker & Weber (1996), Parker (1997) and MacEachern (1999).

There are three contrastive series of stops and affricates: voiceless unaspirated or plain, aspirated and ejective. While MacEachern (1999) is the main source for the generalizations in this paper, the data have been augmented by the Ajacopa et al. (2007) dictionary of Bolivian Quechua as well as the author's own fieldwork in Cochabamba, Bolivia.² The consonantal inventory of the relevant dialects of Quechua is given in Table 1, from MacEachern (1999:29).

² Ajacopa et al. (2007) is a Bolivian Quechua-Spanish dictionary; all glosses have been translated into English by the author.

	labial	alveolar	postalveolar	palatal	velar	uvular	glottal
plain	p	t	tʃ		k	q	ʔ
aspirate	p ^h	t ^h	tʃ ^h		k ^h	q ^h	
ejective	p'	t'	tʃ'		k'	q'	
fricative		s	ʃ				h
nasal	m	n		ɲ			
liquid		l r		ʎ			
glide	w			j			

Table 1: Quechua consonant inventory.

Quechua has three phonemic vowels /i a u/ and two allophones [e o], which result from lowering of /i u/ adjacent to uvulars.³ Roots in Quechua are primarily CV(C)CV, where the optional coda is a continuant or nasal. [ʔ] is restricted to root initial position but is not contrastive as there are no vowel initial roots. As the language is strictly suffixing, root initial position is always also word initial position. Finally, stress is predictably penultimate; disyllabic roots thus have stress on the initial syllable in isolation, but stress may move onto the second syllable or off the root entirely when suffixes are added (e.g., [pitʃaj] ‘to sweep’, [pi'tʃanki] ‘you sweep’, [pitʃan'kitʃu] ‘are you sweeping?’).

Ejectives and aspirates occur only in the onset of a root syllable, and are subject to a range of long-distance restrictions, two of which are directly addressed in the present study. I will refer to two classes of restrictions, co-occurrence restrictions and ordering restrictions. Co-occurrence restrictions prohibit a root from containing two ejectives, two aspirates, or one ejective and one aspirate (1b). Ejectives and aspirates may co-occur with plain stops, and plain stops may occur in pairs (1a). The examples in this section are taken from Ajacopa et al. (2007).⁴

- (1) a. k'inti ‘a pair’
p'atʃa ‘clothes’
k^hastuj ‘to chew’
p^haskaj ‘to tie up’
katʃi ‘salt’
tapuj ‘to ask’

- b. *k'int'i
*k^hast^huj
*k'int^hi
*k^hint'i

³ Some researchers analyze all five vowels as phonemic in some varieties of Quechua, e.g., Cusihamán (1976) on Cuzco Quechua.

⁴ Final [j] is the infinitival suffix and is not part of the root.

Additional co-occurrence restrictions govern the combination of the glottal consonants [ʔ, h] with laryngeally marked stops. Glottal stop freely co-occurs with aspirates, and [h] freely co-occurs with ejectives (2a). Ejectives may not, however, occur in roots with initial glottal stop and aspirates may not co-occur with [h] (2b). This restriction against ejectives occurring in glottal stop initial roots is evidence that glottal stop is phonologically relevant in Quechua, despite being non-contrastive. If these roots are analyzed as phonologically vowel initial, a restriction against ejectives in vowel initial roots must be posited, as opposed to a restriction on the co-occurrence of the natural class of ejectives and glottal stop.

- (2) a. ʔantʰij ‘to moan’
 ʔaqʰa ‘fermented corn drink’
 harkʰaj ‘to protect’
 haytʰa ‘kick’
- b. *ʔantʰij
 *harkʰaj

The ordering restrictions in Quechua prohibit an ejective or aspirate from following a plain stop or affricate in a root. Ejectives and aspirates may occur in medial position in roots with initial fricatives or sonorants (3a), but never in a root with an initial stop or affricate (3b). Ejectives and aspirates may occur in initial position regardless of the medial consonant (3c).

- (3) a. rukʰiy ‘to pack tightly’
 sutʰi ‘clear, visible’
 rukʰu ‘decrepit’
 mapʰa ‘wax’
- b. *kapʰa
 *kapʰa
- c. kʰapa ‘cartilage’
 kʰapa ‘step’
 kʰiri ‘injury’
 qʰasa ‘ice’

The restrictions on ejectives and aspirates are summarized schematically in (4). The present study is concerned with two of these restrictions, boxed and bolded below: the co-occurrence restriction against pairs of ejectives and the ordering restriction on ejectives. In the 2300 roots in the Ajacopa et al. dictionary, there are no counterexamples to either of these restrictions.

(4)	<table style="border-collapse: collapse; margin: 0;"> <tr> <td style="padding: 2px 10px;">*K'-T'</td> </tr> <tr> <td style="padding: 2px 10px;">*K-T'</td> </tr> </table>	*K'-T'	*K-T'	*K ^h -T ^h	*K'-T ^h	*ʔ-T'
*K'-T'						
*K-T'						
		*K-T ^h	*K ^h -T'	*h-T ^h		

The current study is restricted to two of the eight unattested laryngeal patterns in order to make the project a manageable size. Investigation of the other patterns is left for future work. In particular, it should be noted that the co-occurrence and ordering restrictions are exactly parallel for the ejectives and aspirates. While the ejectives were chosen for the current study, this choice was arbitrary and the study could have just as easily been carried out with the aspirates.

2 Investigating phonotactic knowledge – long-distance restrictions in production

The present study asks Quechua speakers to produce unattested root structures, either roots with two ejectives (e.g., [k'ap'i]) or roots with a plain stop followed by an ejective (e.g., [kap'i]). Production of these two types of ungrammatical roots is compared to the production of roots with grammatical medial ejectives, either existing roots or unattested but phonotactically legal roots, in order to diagnose the synchronic status of these two restrictions in Quechua.

Several studies have shown that speakers of languages with static long-distance co-occurrence restrictions (i.e., generalizations that hold over the lexicon but are not supported by active, morphophonemic alternations), like the laryngeal restrictions in Quechua, are indeed aware of these restrictions. Berent and Shimron (1997) found that Hebrew speakers are sensitive to the location of identical consonants in a root via a rating task that compared words derived from nonce roots. In a triconsonantal Hebrew root, C₂ and C₃ may be identical, but roots with identical C₁ and C₂ are unattested. When asked to compare conjugated nonce roots, Hebrew speakers preferred roots with identical C₂ and C₃ over roots with identical C₁ and C₂.

In Semitic triconsonantal roots, homorganic consonants are generally underattested or unattested. Frisch and Zawaydeh (2001) asked Arabic speakers to judge conjugated nonce verbal roots on a well-formedness scale from 1-7, and found that speakers gave lower ratings to nonce roots with a homorganic pair of consonants than to those without.

Rose & King (2007) also investigated co-occurrence restrictions in Semitic triconsonantal roots in two Ethiopian languages, Amharic and Chaha. As in other Semitic languages, homorganic consonants are underattested in roots in Amharic and Chaha. Additionally, these languages exhibit laryngeal harmony, whereby stops that differ in laryngeal features are underattested (note that this is the opposite pattern from what is seen in Quechua and many other languages, where stops are prohibited from having the same laryngeal features). When asked to read nonsense quadrisyllabic tongue twisters, subjects made more errors when the stimuli contained sequences of syllables that violated either phonotactic restriction than when there was no violation. The highest error rate was found in stimuli that violated both restrictions. The work of Rose & King employs a speech error elicitation paradigm; the current study follows this work as well as that of Goldrick & Larson (2008) in interpreting elicited production errors as a window into the phonological grammar.

Coetzee (2005) uses a phoneme identification task to show that English speakers are aware of the restriction on words of the shape sC_iVC_i , like *[spap] and *[skak] (Fudge 1969; Davis 1991). Given a $sCVC$ word where the final C is perceptually ambiguous, English speakers are more likely to identify the ambiguous C as belonging to the category that results in a grammatical structure, e.g., a stimulus with VC transitions ambiguous between [k] and [p] in words like [skak] and [skap] is more likely to be interpreted as [p], yielding the grammatical [skap] as opposed to the ungrammatical *[skak].

One aim of the current study is to contribute further evidence about speakers' synchronic phonological knowledge, particularly of static, long-distance phonotactics. The study has two further goals, discussed in more detail in §2.1 and §2.2 below: to investigate any potential differences between the two types of phonotactically illegal sequences, and to apply the types of errors made in the production of plain stop-ejective sequences to the analysis of ordering restrictions.

2.1 Differences between co-occurrence and ordering restrictions

This section begins by looking at the Quechua lexicon, showing that there is ample evidence for the systematic absence of ejective-ejective as well as plain stop-ejective onset pairs, suggesting that both restrictions should be learned as part of the synchronic grammar, possibly even as a single restriction. Cross-linguistically, however, co-occurrence restrictions are far more common than ordering restrictions, as discussed in §2.1.2. The phonetic grounding of the two restrictions is summarized in §2.1.3 and how these observations relate to previous research and motivate hypotheses about the performance of Quechua speakers is discussed in §2.1.4.

2.1.1 Generalizations over the lexicon

Looking at generalizations over Quechua roots alone, there is no reason to suspect a difference in speakers' accuracy in repeating ejective-ejective sequences and plain stop-ejective sequences. Both the co-occurrence restriction and the ordering restriction are exceptionless in the language, and, moreover, there is comparable evidence that both ejective-ejective and plain stop-ejective sequences are systematically absent, as opposed to accidentally absent due to the relative rarity of a given consonant type in a given position. Ejectives and plain stops are both common in initial and medial positions in Quechua. In a word list of 2300 $CV(C)CV$ roots compiled from the Ajacopa et al. (2007) dictionary, there are 457 roots with initial ejectives, 533 roots with initial plain stops and 241 roots with medial ejectives.⁵ If consonants combined at chance, a list of 2300 roots should have 48 ejective-ejective pairs ($(457*241)/2300$) and 56 plain stop-ejective pairs ($(533*241)/2300$). It is thus unlikely that the absence of either of these root types would be attributed to a sparsity of data.

⁵ Here and throughout the paper “medial” position is used to refer to onset of the second syllable in a disyllabic root.

Another potential source of a difference between the two restrictions could be a formal or analytic bias in learning (Moreton 2002, 2008; Saffran & Thiessen 2003; Wilson 2003). One shape such a bias could take is a preference for hypothesizing and learning formal generalizations that have low complexity and a wide scope over the data. Neither of these metrics distinguishes co-occurrence and ordering restrictions in Quechua, however. Both restrictions require similarly complex grammatical statements, with complexity defined both as the number of features that must be referred to, as well as any abstract structure that must be assumed (e.g., autosegmental structure or syllabic positions), and both restrictions rule out the same number of onset combinations.

A single statement can capture both restrictions: *[-continuant, -sonorant]...[+constricted glottis] (where “...” indicates intervening material between co-occurring onsets), a constraint that disallows a sequence of any stop (specified by [-continuant, -sonorant]) followed by an ejective (specified by [+constricted glottis]). Alternatively, two statements could be postulated, *[+cg]...[+cg] for the co-occurrence restriction and *[-cont, -son, -cg, -sg]...[+cg] for the ordering restriction.⁶ While more features are required to pick out the class of plain stops than ejectives, [-cont, -son, -cg, -sg] vs. [+cg], suggesting that ordering restrictions requires a more complex generalization than co-occurrence restrictions, the degree of abstraction required for the two generalizations is the same. Both restrictions are long-distance, and thus both *[+cg]...[+cg] and *[-cont, -son, -cg, -sg]...[+cg] require either some autosegmental structure (cf. Hayes & Wilson 2008), or a mechanism for ruling out non-adjacent sequences.

Thus, given the Quechua lexicon and the structure of the two restrictions alone, Quechua speakers should treat the two types of illegal structures similarly.

2.1.2 Typology

Co-occurrence restrictions on multiple instances of the same feature are relatively common cross-linguistically, both for laryngeal features and other types of features.⁷ This type of restriction is commonly referred to as the Obligatory Contour Principle (OCP) (Leben 1973; Goldsmith 1976; McCarthy 1986, 1988). In particular, languages with ejectives often permit only one ejective per root. Such a restriction is attested in a genetically and geographically diverse range of languages, including dialects of Quechua (Quechuan) and Aymara (Aymaran), Shuswap (Salishan), Tz’utujil (Mayan), Old Georgian (Kartvelian) and Hausa (Afro-Asiatic), as described in detail in MacEachern (1999).

In addition to ejectives, co-occurrence restrictions are found for aspirates in Sanskrit (Indo-Aryan) and Ancient Greek (Indo-European) (‘Grassman’s Law’), Gojri (Indo-Aryan), dialects of

⁶ The statement *[-cont, -son]...[+cg] captures both the co-occurrence and ordering restrictions on ejectives; further constraints are needed to capture the remaining laryngeal restrictions in the language. The ordering and co-occurrence restrictions on aspirates requires *[-cont, -son]...[+sg] as well, and the restrictions involving glottal segments require further constraints. The restriction against ejective-aspirate pairs could be captured by the combined work of *[-cont, -son]...[+cg] and *[-cont, -son]...[+sg].

⁷ Only restrictions on consonants are considered here, though vowels also commonly interact non-locally in vowel harmony systems.

Aymara and Quechua, Ofo (Siouan), Souletin Basque (isolate), and some North East Bantu languages (Ohala 1981) ('Dahl's Law'). Pairs of implosives are absent in Hausa and Muna (Austronesian) (van den Berg & Sidu 1996), as are pairs of voiced stops in Japanese (Ito & Mester 2003 and references there; Kawahara 2008). Pairs of pre-nasals are restricted in Muna and some North East Bantu languages (Meinhof 1932; Herbert 1977) ('Meinhof's Law' or the 'Ganda Rule').⁸ Restrictions on homorganic segments are well established in Semitic, Muna and many other languages (Greenberg 1960; McCarthy 1986; Frisch, Broe & Pierrehumbert 2004; Coetzee & Pater 2008).

In contrast, an ordering restriction on ejectives is only found in the Andean region of South America, in dialects of Quechua and Aymara, according to MacEachern's (1999) survey. Examination by the current author of the grammars and dictionaries of a further 29 languages with an ejection contrast in multiple root positions reveals no more instances of a restriction of this type.

In addition to the restrictions on ejectives in Quechua and Aymara, ordering restrictions are attested for aspirates in dialects of Quechua and Aymara as well as Souletin Basque. Outside of the domain of laryngeal features, ordering restrictions are found in a limited number of cases of restrictions on coronal consonants. In Aymara, Kera and Pengo, Hansson (2001) reports that an alveolar and postalveolar may only co-occur in one order: roots with postalveolar-alveolar sequences are attested (tʃ...t), but alveolar-postalveolar sequences are absent (*t...tʃ).

2.1.3 *Phonetic grounding*

A perceptual basis for co-occurrence restrictions on ejectives (and aspirates) is identified and argued for in Gallagher (2010a,b). English speakers participated in an AX discrimination task where they were asked to distinguish between stimuli with different numbers of ejectives and plain stops. The stimuli were made from recordings of Bolivian Quechua; further details can be found in Gallagher (2010a,b). The experiment compared performance on the two categories of "different" pairs schematized in (5). Stimulus pairs in the 1 vs. 0 category in (5a) contrast a form with a single ejective with a form with two plain stops. Stimulus pairs in the 2 vs. 1 category in (5b) also differ in whether a single stop is ejective or plain, but in this category the contrast is between a form with two ejectives and a form with one ejective and one plain stop.

- (5) a. 1 vs. 0 *k'api-kapi* or *kap'i-kapi*
 b. 2 vs. 1 *k'ap'i-kap'i* or *k'ap'i-k'api*

Performance on the 2 vs. 1 category is significantly worse than performance on the 1 vs. 0 category (64% correct vs. 75% correct). Subjects incorrectly respond "same" to pairs contrasting 2 vs. 1 ejectives more often than to pairs contrasting 1 vs. 0 ejectives.

⁸ Many thanks to Andries Coetzee and Joe Pater for sharing their root list and O/E calculations for Muna from the van den Berg & Sidu dictionary. Additional thanks to an anonymous reviewer for pointing out some of these cases of dissimilation.

These results show that a contrast in ejection is less perceptible in the context of another ejective than in the context of another plain stop. The relevance of this result to co-occurrence restrictions on ejectives is that the perceptually difficult contrast between 2 vs. 1 ejectives is exactly the contrast that is absent in a language with a co-occurrence restriction. A co-occurrence restriction on ejectives is thus perceptually grounded because it eliminates a perceptually difficult contrast.

Ohala (1981, 1993) also proposes that dissimilatory co-occurrence restrictions, as seen for ejectives in Quechua, are perceptually natural. Ohala's account of long-distance laryngeal restrictions is based on misperception caused by the affect of laryngeal features like ejection on adjacent segments. Ohala hypothesizes that glottalization on a vowel between two ejectives may be attributed by the listener to just one of the two ejectives, and that systematic misperception of this kind results in dissimilatory patterns.

Co-occurrence restrictions and ordering restrictions also differ in their articulatory grounding. Ejectives are articulatory complex, consisting of both an oral constriction and raising and constriction of the glottis that must be precisely timed with the closure and release of the oral constriction (Lindau 1984; Kingston 1985). Languages may avoid roots with multiple ejectives because such roots would contain two articulatorily difficult sounds in close proximity to one another. This is the line of analysis taken by Walter (2007) with respect to place co-occurrence restrictions. Walter argues that avoiding repetition of relatively taxing articulatory gestures is the basis for dissimilatory phonological restrictions.

Ordering restrictions are also phonetically grounded in that they favor ejection in initial position, a position that is cross-linguistically linked to cue strengthening for laryngeal contrasts (Pierrehumbert & Talkin 1991; Keating et al. 2003) and is generally a favored prosodic position for contrasts (Beckman 1998; Zoll 1998). While non-initial ejectives are generally attested in Quechua, a phonetically natural preference for initial ejectives can be seen just in roots with two stop onsets. I don't know of a phonetic explanation for why this preference is observable only in roots with two stops; that is, while a general preference for initial ejectives is phonetically natural, the prohibition on medial ejectives only in roots with a preceding plain stop is not clearly phonetically based.⁹

2.1.4 Predictions and discussion

The preceding observations about co-occurrence and ordering restrictions, as well as formal analyses in the literature, support two competing hypotheses about Quechua speakers' accuracy in repeating these two types of illegal roots: that accuracy will be equivalent on both types of roots, or that accuracy will differ. The results of the experiment, discussed below in §3, show variation among subjects; for some subjects, there is a stark difference in accuracy for co-occurrence and ordering restriction violating roots, while for other subjects there is no difference.

⁹ An anonymous reviewer points out that evidence for the relative markedness of initial vs. medial ejectives could come from language acquisition. If Quechua children acquire initial ejectives before medial ejectives, this would support the idea that ejectives are more natural in initial position.

One hypothesis is that, since speakers have no experience with either ejective-ejective or plain stop-ejective sequences, and have comparable evidence that both sequences are systematically absent, there should be no difference in accuracy between the two categories. This is also the prediction made if co-occurrence and ordering restrictions are analyzed formally as a single restriction, as proposed by Mackenzie (2009). Mackenzie analyzes similar facts in Aymara as the effect of a single markedness constraint ruling out ejectives that are preceded by a stop (either an ejective as in the co-occurrence restriction, or a plain stop as in the ordering restriction) in the same root.

The analysis of Mackenzie contrasts with the analyses of MacEachern (1999) and Gallagher (2010a, 2011), which treat co-occurrence and ordering restrictions as independent restrictions and thus predict that accuracy on the two types of roots could differ. The available evidence from the cross-linguistic typology suggests that ejective-ejective pairs are more dispreferred than plain stop-ejective sequences. Several studies have shown that cross-linguistic asymmetries of this kind may influence behavior on experimental tasks with non-native structures.

While speakers generally have trouble with non-native phonemes, sequences of phonemes, or position of phonemes, several studies have found that subjects are not equally inaccurate on all non-native structures. Broselow & Finer (1991) found that Japanese and Korean speakers are more accurate when producing the onset cluster [pr] than when producing [br], and are in turn even less accurate on [fr], even though all three onset clusters are unattested in both languages. When producing English voiced and voiceless coda stops, Broselow, Chen & Wang (1998) find that native speakers of Mandarin, a language with no stop codas, make more errors on voiced stop codas than on voiceless stop codas, and sometimes devoice voiced targets. Hansen (2004) examined errors made by Vietnamese speakers when producing words with English codas and found that speakers were most accurate with /s/ and /f/, followed by /v/ and /l/, and had the most difficulty with /ʃ/. While all of these target codas are phonemes of Vietnamese, none are allowed in coda position in that language. In a repetition task, Davidson (2006) finds that English speakers are more accurate on non-native word-initial clusters when C_1 is /f/ than when it is /z/ and are least accurate when C_1 is /v/. Moreton (2002) further finds that English listeners have a stronger perceptual bias against onset [dl] than onset [bw], despite the fact that both are absent in English. These results lend some plausibility to the idea that Quechua speakers may not have equal amounts of trouble with ejective-ejective sequences and plain stop-ejective sequences, even though both are unattested in the language.

One explanation for differences among unattested structures is the relative markedness of these structures, as assessed by their cross-linguistic frequency. Broselow & Finer (1991) argue that better performance on [pr] than on [br] is due to the greater sonority difference between C_1 and C_2 in [pr] than in [br], and thus the greater markedness of [br]. Broselow, Chen & Wang (1998) similarly claim that the higher error rate on voiced codas than voiceless codas, and the tendency to produce voiced codas as voiceless, is due to the relative markedness of voiced codas.

This particular hypothesis is pursued explicitly by the investigations of sonority sequencing in clusters in Berent et al. (2007) and Berent & Lennertz (2010), which find that speakers

distinguish between onset sequences based on sonority even among onset sequences that are all unattested in their native language. Across languages, onset clusters with a larger increase in sonority between C_1 and C_2 (e.g., [bl]) are more common than those with a smaller increase (e.g., [bn]) which are in turn more common than onsets with a sonority plateau (e.g., [bd]), and onsets that decrease in sonority are least common of all (e.g., [lb]). English speakers are more likely to report that a nonce word with a falling sonority onset is disyllabic, e.g., [lbif] is perceived or classified as [ləbif], than that a nonce word with a sonority plateau is disyllabic, e.g., [bdif] is perceived or classified as [bədif], even though both types of onsets are unattested in English.

2.2 Error type and the analysis of the ordering restriction

The ordering restriction in Quechua disallows a sequence of a plain stop followed by an ejective, e.g., *[pat'i]. Such an ungrammatical sequence could be repaired in one of (at least) two ways, either by de-ejectivizing the medial stop, producing target [pat'i] as [pati], or by moving ejection to grammatical initial position, producing [pat'i] as [p'ati]. These two different mappings are predicted by different formal analyses, and thus the actual errors made when producing plain stop-ejective sequences bear on the formal analysis of the restriction.

The main challenge for an analysis of the ordering restriction is to distinguish medial ejectives in roots with an initial stop, e.g., *[pat'i], from medial ejectives in roots with an initial continuant or sonorant, e.g., [mat'i] 'forehead'. An ordering restriction is not a straightforward case of positional licensing or faithfulness (Beckman 1998; Zoll 1998); medial ejectives exist in Quechua, but whether they are grammatical or not is dependent on the identity of the root-initial consonant. In autosegmental terms, this pattern can be analyzed as predictable association between a [+cg] autosegment and a segmental position; if a [+cg] autosegment is present, it associates to the leftmost stop. In Optimality Theory, the principle of Richness of the Base requires that an analysis must map any input to some grammatical output. Thus, the grammar must map an ungrammatical input like /pat'i/ to a grammatical alternative, either [p'ati] or [pati], while mapping an input like /mat'i/ faithfully to [mat'i]. There are at least two approaches to distinguishing grammatical and ungrammatical medial ejectives. In one approach, a markedness constraint that penalizes medial ejectives interacts with multiple faithfulness constraints such that the markedness constraint may be satisfied by moving ejection, making the mapping /pat'i/ → [p'ati] optimal, but not by deleting it, making /mat'i/ → [mat'i] optimal. Another approach is that the operative markedness constraint specifically penalizes medial ejectives that are preceded by plain stops; /mat'i/ surfaces faithfully because it doesn't incur a violation of markedness while /pat'i/ is modified because it does.

The first approach assumes a general markedness constraint, like a categorical constraint from the ALIGN family (McCarthy & Prince 1993) or a constraint against all non-initial ejectives. This constraint is phonetically grounded, in that it favors ejection where it is perceptually

strongest. The analysis sketched here is roughly that of McCarthy (2003) and Gallagher (2011) and assumes the categorical ALIGN constraint in (6).¹⁰

- (6) ALIGN([+cg]; RT, L) Every instance of [+constricted glottis] is aligned with the left edge of a root.

The constraint in (6) is violated by all non-root-initial ejectives; [pat'i] and [mat'i] both incur a violation mark. The different behavior of the two types of roots can be accounted for with the relative ranking of ALIGN with the two faithfulness constraints in (7).

- (7) MAX([+cg]) Every instance of [+cg] in the input is present in the output.

- NOFLOP([+cg]) Given a [+cg] autosegment in the input, if that [+cg] autosegment has an output correspondent, then it is associated to the same segment in the input and the output.

The constraint MAX([+cg]) prohibits deletion of a [+constricted glottis] autosegment from input to output, but is not violated by changes of segmental association, while NOFLOP([+cg]) prohibits changes in segmental association but not deletion. If MAX([+cg]) outranks ALIGN, but ALIGN outranks NOFLOP([+cg]), then ejection will reassociate in order to satisfy ALIGN but will not delete. In a form with two stops like /pat'i/, reassociation is possible and results in the output [p'ati] being optimal (8a). In a form with only one stop like /mat'i/, however, reassociation is not possible and the ranking of MAX([+cg]) over ALIGN means that a surface medial ejective is optimal over deletion of ejection (8b).¹¹

(8) a.

/pat'i/	MAX([+cg])	ALIGN([+cg]; RT, L)	NOFLOP([+cg])
i. pat'i		* !	
→ ii. p'ati			*
iii. pati	* !		

b.

/mat'i/	MAX([+cg])	ALIGN([+cg]; RT, L)	NOFLOP([+cg])
→ i. mat'i		*	
ii. mati	* !		

¹⁰ Categorical alignment constraints are also known as COINCIDE constraints (McCarthy 2003, Zoll 1998).

¹¹ The full analysis of inputs like [mat'i] requires additional constraints to be high-ranked. Non-stops must both be prohibited from hosting a [+cg] feature (*[m'ati]) and blocked from turning into stops in order to host [+cg] (*[mat'i/ → [p'ati]).

The above analysis relies on the relative ranking of faithfulness constraints to differentiate between medial ejectives in roots with initial stops and those in roots with initial non-stops. Another line of analysis, pursued in different forms by Parker (1997), MacEachern (1999) and Mackenzie (2009), is to propose a markedness constraint that specifically penalizes ejectives that are preceded by other stops. The constraint in (9) is an example of such a constraint; it is violated by a form like [pat'i] but not by a form like [mat'i].

- (9) *[-cont, -son]...[+cg] A [-continuant, -sonorant] segment is not followed by a [+constricted glottis] segment in the root.

The constraint in (9) penalizes a root-internal sequence of non-adjacent segments with certain feature values and is the type of constraint proposed by Mackenzie (2009). A gradient alignment constraint, which requires ejection to occur on the leftmost stop in the root, is proposed by Parker (1997) and MacEachern (1999) and makes the same predictions as the constraint in (9). With a markedness constraint like that in (9), mapping /pat'i/ to [pati] is optimal if *[-cont, -son]...[+cg] and NOFLOP([+cg]) outrank MAX([+cg]), as in (10a).

(10) a.

/pat'i/	NOFLOP([+cg])	*[-cont, -son]...[+cg]	MAX([+cg])
i. pat'i		* !	
ii. p'ati	* !		
→ iii. pati			*

b.

/mat'i/	NOFLOP([+cg])	*[-cont, -son]...[+cg]	MAX([+cg])
→ i. mat'i			
ii. mati			* !

It is also possible to derive the mapping /pat'i/→[p'ati] if *[-cont, -son]...[+cg] and MAX([+cg]) outrank NOFLOP([+cg]), as in the analysis in (8) above with categorical alignment.¹²

¹² In addition to the difference in repairs predicted by the two types of markedness constraints, the constraints also make different predictions about allowable sequences in longer roots. Given a constraint ranking where a categorical alignment constraint can only be satisfied by moving ejection, a hypothetical root like [matik'u] should be grammatical. While [matik'u] has an ejective following a plain stop, moving ejection leftwards ([mat'iku]) does not alleviate the violation of categorical markedness. Given a gradient alignment constraint, or a constraint on a plain stop-ejective sequence, a root like [matik'u] should be impossible. The limited number of trisyllabic roots in Quechua makes it difficult to test these two hypotheses, though there are roots of the shape [mat'iku] and none like [matik'u]. Arguing for categorical alignment, McCarthy (2003) attributes the absence of [matik'u] like roots to an accidental gap, while Parker (1997) argues that these roots are systematically absent due to a gradient alignment constraint. The experiment presented in this paper does not contribute to this debate, as only disyllabic roots are investigated.

(11) a.

/pat'i/	MAX([+cg])	*[-cont, -son]...[+cg]	NOFLOP([+cg])
i. pat'i		* !	
→ ii. p'ati			*
iii. pati	* !		

b.

/mat'i/	MAX([+cg])	*[-cont, -son]...[+cg]	NOFLOP([+cg])
→ i. mat'i			
ii. mati	* !		

Since the ordering restriction in Quechua does not trigger any alternations, the various analyses presented above are all equally consistent with the patterns in the lexicon. The repetition task explicitly asks Quechua speakers to produce forms like [pat'i]. If these forms are modified when repeated, the way in which they are modified could help adjudicate between the competing analyses. Particularly, if targets like [pat'i] are produced as [pati], with de-ejectivization, this is inconsistent with an analysis based on a categorical alignment constraint. We will see that both de-ejectivization and movement of ejection are attested repairs for targets that violate the ordering restriction, supporting an analysis with *[-cont, -son]...[+cg] that more easily accounts for both types of repairs.

3 The experiment – repetition task

A repetition task compares the accuracy of producing medial ejectives when phonotactically legal and illegal. Targets with phonotactically legal medial ejectives are either existing roots or nonce roots where a medial ejective occurs with an initial fricative or sonorant (e.g., [jak'i]). These two phonotactically legal categories of targets are referred to as “real” and “gap”, respectively. Targets with phonotactically illegal medial ejectives have either an initial ejective (e.g., [k'at'u]), violating the co-occurrence restriction on ejectives, or an initial plain stop (e.g., [kat'u]), violating the ordering restriction on ejectives. These two categories of phonotactically illegal targets are referred to as “co-occurrence” and “ordering”.

The null hypothesis is that speakers will do equally well on all four categories of words: real, gap, ordering and co-occurrence. If there is an affect of attestation (familiarity), speakers may be more accurate in repeating the real words as opposed to the nonsense words in the gap categories. Real words serve as a baseline for the general error rate in the task.

The main purpose of the experiment is to test whether phonotactic restrictions play a role in accuracy on the task, and thus are part of speakers' grammar. Again, the null hypothesis is that medial ejectives will be repeated equally accurately among all categories. If phonotactic restrictions play a roll in accuracy, however, the phonotactically illegal medial ejectives in the co-occurrence and ordering categories should be repeated less accurately than the legal medial ejectives in the real and gap categories.

The current experiment also tests for a difference in accuracy between the co-occurrence and ordering categories. If speakers have synchronic knowledge of universal markedness asymmetries, like the relative typological rarity of restrictions on plain stop-ejective pairs as compared to restrictions on ejective-ejective pairs, then more errors are expected on items in the co-occurrence category than in the ordering category.

3.1 Stimuli

3.1.1 Stimuli creation

The stimuli for the repetition task consist of real and nonsense words with medial ejectives. The words fall into one of four categories: real words (real), accidental gaps (gap), co-occurrence restriction violators (co-occurrence) and ordering restriction violators (ordering). The medial ejectives in the real and gap categories are phonotactically legal while the medial ejectives in the co-occurrence and ordering categories are phonotactically illegal. There are 15 items in each category, for a total of 60 items, shown in IPA in Table 2.

real		gap	
haʎp'a 'earth'	mast'a 'rug'	hutʃa	rik'u
hutʃuj 'small'	mat'i 'forehead'	hup'i	sanq'e
ʎank'aj 'to work'	misk'i 'sweet'	limp'u	saq'o
ʎant'a 'wood'	ruk'ana 'finger'	lip'u	jak'i
ʎimp'ij 'clean'	satʃa 'tree'	luk'a	jank'u
ʎutʃu 'naked'	suʎk'a 'younger'	mant'u	jap'i
matʃi 'nice'	wirp'a 'lip'	nump'ij	jump'a
map'a 'dirtiness'		rank'i	
co-occurrence		ordering	
tʃanq'o	q'amp'a	tʃink'aj	pitʃu
tʃup'a	q'antʃa	tʃuk'i	pint'a
k'ant'i	q'ap'u	kap'ij	pit'u
k'ap'u	t'ap'u	kat'u	puk'aj
k'at'u	t'ik'aj	kitʃi	tik'aj
p'ank'u	t'ink'aj	kuntʃa	tip'a
p'intʃa	t'uk'i	kunt'a	tump'ij
p'int'i		pank'u	

Table 2: Stimulus items by category.

The real words were designed to serve as a baseline for error rate on the task, to which errors on the gap, co-occurrence and ordering categories would be compared. No other fillers were included. This stimulus set is appropriate for testing the relative error rate on medial ejectives, as

affected by attestation (real words vs. gap, co-occurrence and ordering words) and phonotactic legality (real and gap words vs. co-occurrence and ordering words). The stimulus set does not allow for an assessment of error rates for ejectives in initial position, or for any other error type.

Fairly common words that the author had heard used in speech were selected for the real category to increase the likelihood that they would be recognizable to participants. Words in the gap category are nonsense words that do not violate any known phonotactic restriction, either long-distance or local (e.g., the sequence *ji is unattested in Quechua), and whose absence from the lexicon is likely accidental. Nonce roots with onset combinations that were similar were avoided, as these could be due to a non-categorical restriction on similar consonants that speakers may be aware of. For example, n...t' is an unattested onset sequence, but as both consonants are coronal this consonant pair was not used to construct a word in the gap category. There are 4 unattested onset pairs with an ejective second member that satisfy the criteria for a “gap” (l...p', n...p', j...k', j...p' each occur 0 times), and these pairs were used to construct 7 of the 15 stimuli in the gap category. Three more stimuli contain onset pairs that occur about ½ as often as predicted by chance (l...k' and s...q'). The remaining 5 gap stimuli contain well-attested (occurring about as often as predicted by chance) onset combinations (h...tʃ', h...p', m...t', r...k').¹³

Within each category, an attempt was made to balance between CVCV and CVCCV words. In the co-occurrence category, there are 9 CVCV words and 6 CVCCV words; this imbalance resulted from the need to also balance for frequency measures across categories (discussed in §3.2.2 below). In the nonce words, the coda consonant was always a nasal, though it is variable in the real words.

All stimuli were spliced together from recordings of a male, middle-aged, native speaker of Cochabamba Quechua, taken on-site in Cochabamba using a Marantz PMD-660 solid-state portable recorder and an Audio Technica AT831b lavalier microphone. The speaker was asked to read a list of both real and phonotactically legal nonsense words from a computer screen, repeating each word twice. The stimuli were made by splicing together initial and medial syllables from these recordings. Whenever possible, the first repetition from the recordings was used for splicing stimuli; the second repetition was only used if the first repetition was mispronounced or disfluent. Stimuli were spliced together during the closure of the second stop, keeping VC transitions or CC transitions intact. For example, the stimulus *pank'u* was made by splicing together the initial syllable of *panku* and the final syllable of *yank'u*, making the cut during the closure of [k]/[k']. Initial and final syllables in the stimuli were always spliced from original initial and final syllables in the recorded speech.

3.1.2 Frequency measures

The stimuli in the gap, co-occurrence and ordering categories are balanced as much as possible

¹³ The degree of attestation of an onset pair C₁-C₂ is determined by the Observed/Expected (O/E) ratio (Pierrehumbert 1993; Frisch et al. 2004), calculated by multiplying the probability of C₁ as the word-initial consonant by the probability of C₂ as the medial syllable onset by the total number of roots.

for two frequency measures: expected co-occurrence of the onset consonants and number of neighbors. Expected co-occurrence of a C_1 - C_2 pair is computed as the number of occurrences of C_1 in initial position multiplied by the number of occurrences of C_2 in medial position, divided by the total number of roots. For example, the expected co-occurrence of r-k' is the number of times [r] occurs as C_1 in a root (89) multiplied by the number of times [k'] occurs as C_2 in a root (71), divided by the total number of roots (2300) = 2.75. Much research has shown that frequency can affect performance on a range of tasks involving nonce words. In production studies, subjects respond more quickly when asked to repeat nonce words containing high probability sequences than those with lower probability (Vitevitch & Luce 1999, 2005; Vitevitch et al. 1997). Similarly, in a word-recognition task subjects respond more quickly to higher probability nonce words (Frisch, Large & Pisoni 2000) and in word-likeness rating studies higher probability nonce words are rated higher than lower probability words (Bailey & Hahn 2001; Coleman & Pierrehumbert 1997; Frisch, Large & Pisoni 2000; Munson 2001; Vitevitch et al. 1997). Stimuli in each of the three nonce categories are balanced with respect to probability to ensure that differences in accuracy between the categories result from the phonological properties of the stimuli as opposed to being an artifact of a preference for higher probability items.

Neighbors were counted as the number of attested roots that were one change away from the stimulus, where a single change is a change in either onset consonant, a change in either vowel, a change of coda, the insertion/deletion of a coda or the deletion of the second onset. For example, the nonsense word *jump'a* has three neighbors: *lump'a* 'purity', *hump'a* 'to fill the mouth with water', and *juma* 'semen'. Vowel change and coda neighbors only existed for those words in the gap category that contained an attested onset consonant pair. The expected co-occurrence values and number of neighbors for stimuli in the gap, co-occurrence and ordering categories are given in Table 3. Both frequency measures were included in the analysis of the results of the experiment, to test whether any differences between the three categories on these measures had a significant effect on subjects' repetitions.

gap	exp.	neigh.	coocc.	exp.	neigh.	ord.	exp.	neigh.
nump'ij	.30	2	q'ap'u	1.37	4	kap'ij	1.55	6
jump'a	.73	3	q'amp'a	1.37	4	tump'ij	1.88	2
jap'i	.73	4	p'intʃ'a	1.51	2	tip'a	1.88	5
limp'u	1.11	0	t'ap'u	1.58	2	kuntʃ'a	1.97	1
lip'u	1.11	1	p'int'a	1.63	0	kitʃ'i	1.97	6
jank'u	1.30	2	k'ap'u	1.63	4	kunt'a	2.13	2
jak'i	1.30	3	q'antʃ'a	2.08	1	kat'u	2.13	6
luk'a	1.98	5	p'ank'u	2.10	4	pitʃ'u	2.57	9
hup'i	2.42	9	tʃup'a	2.16	10	pint'a	2.77	3
rik'u	2.75	4	k'ant'i	2.25	1	pit'u	2.77	8
rank'i	2.75	6	k'at'u	2.25	5	tik'aj	3.33	4
hutʃ'a	3.08	8	t'ink'aj	2.81	2	pank'u	3.58	4
sanq'e	3.63	4	t'ik'aj	2.81	6	puk'ay	3.58	10
saq'o	3.63	9	t'uk'i	2.81	6	tʃink'aj	3.67	3
mant'u	3.80	5	tʃanq'o	3.83	2	tʃuk'i	3.67	11
average	2.01	4.33		2.15	3.53		2.63	5.33

Table 3: Expected co-occurrence of onset combination and neighbors for nonce word stimuli.

A caveat with regards to frequency measures is that the measures here are based on a word list of 2300 roots taken from the Ajacopa et al. (2007) dictionary of Bolivian Quechua. There are likely many roots that are not included in the dictionary, as well as roots that are not known to all speakers or are pronounced slightly differently. While this is surely true of any language corpus, the small size of the word list here increases the importance of each individual form.

3.1.3 Acoustic properties of the stimuli

The acoustic properties of ejectives have been shown to vary widely, both between and within languages (Lindau 1984; Kingston 1985; Wright et al. 2002). The acoustic study of ejectives in Bolivian Quechua in Gallagher (2010a) (see also Parker's (2007) acoustic analysis of Cuzco Quechua) finds that the main correlates to ejection are a long VOT and a large burst amplitude; follow-up measurements in Gallagher (2011) finds that pitch in the first 30 milliseconds of the following vowel is depressed in ejectives relative to other stops.

In order to compare the results of the repetition task for phonotactically legal and illegal medial ejectives, it is important to verify that the cues to ejection in the stimuli are comparable across categories. To this end, VOT in all non-affricate stimuli was normalized to 130 ms, a relatively long VOT value within the normal range, to ensure that this cue to ejection was both clear and constant across stimuli. In stop ejectives, the period of VOT is silent, making it easy to delineate and manipulate VOT length by cutting or pasting silence in Praat. In affricate ejectives, however, the friction period is followed by a rapid decrease in intensity of friction, and often a

period of silence, but it is not straightforward to delineate the VOT period from the frication period. Consequently, VOT in affricate ejectives was left unmodified, though the entire period from burst to vowel was always at least 130 ms.

For all nonsense stimuli, burst amplitude, intensity difference between burst and VOT period, and pitch/glottalization in the following vowel were measured and compared across stimulus categories. Burst amplitude is measured in arbitrary units directly off the waveform in Praat by subtracting the minimum amplitude from the maximum. Additionally, the difference between maximum and minimum intensity in dB is taken for the period from the beginning of the burst until the onset of the following vowel, measured as the beginning of the first glottal pulse. This measure doubles as a second indicator of burst amplitude, and is particularly useful for affricates, where an isolated burst is not present. One of two measures of voice quality in the following vowel was taken. When available, pitch in the first 30 ms of the vowel was taken. In cases of extreme glottalization, pitch is often not resolved in Praat; in these tokens, the duration of glottalization was measured. Figure 1 gives an example of a stimulus item in the co-occurrence category, with the relevant periods for acoustic measures marked.




Figure 1: Spectrogram and waveform for stimulus [p'int'i], with relevant landmarks for acoustic measures of burst, VOT and vowel quality marked.

For each measure, a one-way ANOVA was run with category (gap, co-occurrence or ordering) as a factor. There was no effect of category for any measure, suggesting that differing error rates on the three categories are not due to acoustic differences in the stimuli. The acoustic measurements for each stimulus are reported in Appendix A.

3.2 Procedure

Eight native speakers of Quechua participated in the experiment. Seven of the subjects (3 male and 4 female, age 30-65) were from the Cochabamba area of Bolivia and participated in the experiment in and around the city of Cochabamba. The eighth subject, recruited in New York City, was a male from Cuzco, Peru in his mid-forties. While there are dialectal differences between Cochabamba and Cuzco Quechua, the phonotactic restrictions on ejectives (and other laryngeal features) are the same in the two dialects. All the subjects were bilingual in Quechua and Spanish; Spanish was the metalanguage used to conduct the experiment.

Subjects listened to two repetitions of each target word, 400 ms apart, through Audio Technica headphones played on an Ipod, and were instructed to repeat what they heard as precisely as possible. Instructions were given just once, at the beginning of the experiment. Subjects had as long as they needed to respond after hearing the second stimulus; once they had responded, the experimenter played the next stimulus. Responses were recorded using a Marantz PMD-660 solid-state portable recorder and an Audio Technica AT831b lavalier microphone. The experiment consisted of a total of 60 stimulus items, presented in random order for each subject, and took about 5 minutes.

3.3 Analysis

Each response was transcribed, relying on examination of spectrograms and waveforms to determine the ejective or non-ejective status of a given target, and coded for correctness and type of modification (if any). Two transcribers were used, both native English speakers with some knowledge of Quechua and experience listening to Quechua ejectives. Where the two transcribers disagreed, a third transcriber, also a native English speaker, was consulted. Additionally, each response was measured acoustically.

Responses were coded as correct if the medial segment was produced as an ejective, and no other changes were made that changed the category of the target item. Responses were excluded from the correct/incorrect analysis if the response was not recognizable as the target item or the three transcribers all differed in their transcription (7/480 = 1.5% of tokens), if there was excessive background noise making the response inaudible (13/480 = 2.7% of tokens), or if the target was produced in a way that changed the phonotactic status of the medial ejective (e.g., [pitʃu] produced as [mitʃu]) (9/480 = 1.9% of tokens). A total of 29 responses were removed, 6% of total responses. Errors that did not change the phonotactic status of the word were ignored for the purposes of the correct/incorrect analysis. These were predominately place of articulation errors (e.g., [kat'u] produced as [kap'u]), insertion of a coda, either medially or finally (e.g., [kat'u] produced as [kant'u] or [kat'un]), or a change in vowel. The number of responses in each category not included in the correct/incorrect analysis are given by subject in Table 4.

subject	gap	real	ordering	co-occurrence	total
1	0	1	1	0	2
2	0	0	2	2	4
3	0	1	1	2	4
4	2	0	5	3	10
5	1	0	2	3	6
6	0	0	0	0	0
7	1	0	0	1	2
8	0	1	0	0	1
total	4	3	11	11	29

Table 4: Tokens excluded from the correct/incorrect analysis, by subject and category.

The remaining 451 responses were coded as “correct” (ejective) or “incorrect” (non-ejective) based on inspection of burst amplitude, intensity difference and length of VOT.

For responses to targets in the co-occurrence and ordering categories, each response was coded for the type of error, if any. Common errors included de-ejectivizing a stop (e.g., [kant’u] repeated as [kantu]), moving ejection (e.g., [kant’u] repeated as [k’antu]), changing place of articulation, inserting/deleting a coda, and changing the manner of C₁ (e.g., [kant’u] repeated as [hant’u]).

In addition to error data, acoustic measurements of the medial stop were taken for each response. VOT, burst amplitude, intensity difference, and pitch and glottalization in the following vowel were measured. Measurements were taken following the same procedure as described for measuring the stimuli.

3.4 Results

3.4.1 Accuracy

As predicted by the main hypothesis, subjects perform nearly at ceiling in accurately repeating the phonotactically legal medial ejectives in the real and gap categories, and are overall around chance at accurately repeating the phonotactically illegal medial ejectives in the co-occurrence and ordering categories. Percent correct by category is shown in a barplot in Figure 2, averaged across subjects; errors bars indicate 95% confidence interval between subjects. The raw numbers of incorrect and correct tokens are given in Table 6 below.

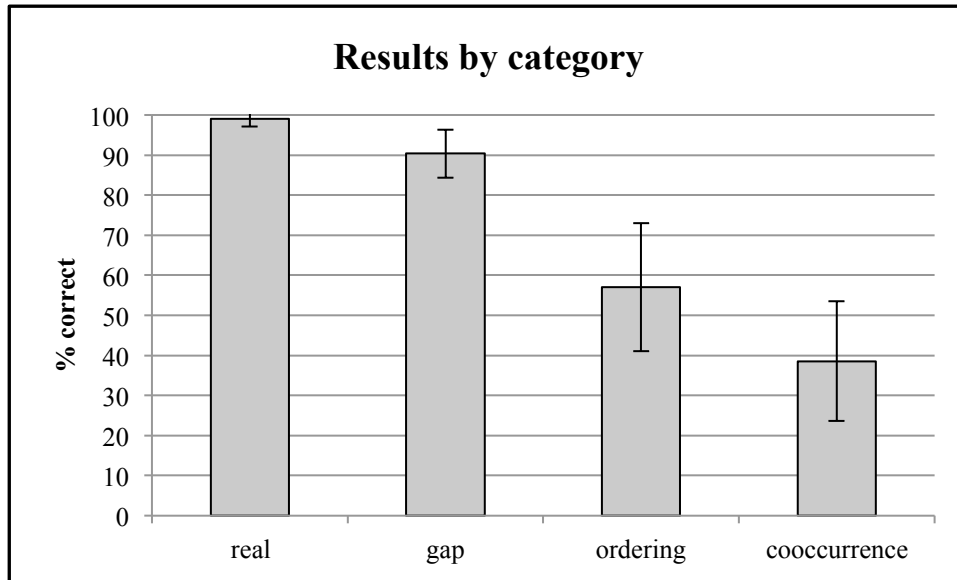


Figure 2: Percent correct by category, averaged across subjects. Error bars indicate 95% confidence interval between subjects.

To test for the factors that influence subjects' accuracy, a Mixed Logit Model with response (correct or incorrect) as the dependent variable was run using *lmer()* in the *lme4* package (Bates & Maechler 2010) in the R software. Three user-coder predictors were included in the model, testing for (i) a difference between legal (gap and real) and illegal (co-occurrence and ordering categories), (ii) a difference between co-occurrence and ordering categories and (iii) a difference between gap and real categories. In addition, the continuous factors of expected frequency and number of neighbors were also included as predictors in the model. Random intercepts and slopes were fit for both subject and stimulus item. The significance of each factor's contribution to the model was determined through model comparison.¹⁴ Two of the binary predictors based on category are significant: the difference between legal and illegal categories, and the difference between co-occurrence and ordering categories. There is no significant difference between real and gap categories, nor any effect of expected frequency or neighborhood density. Table 5 gives the final model.¹⁵

¹⁴ A full model with all predictors is fit to the data, and then predictors are removed one-by-one and the resulting simpler model is compared to the full model using the *anova()* function. If a model without a given predictor differs significantly (at a level of .05) from a model with that predictor, then the predictor is significant. This process of removing predictors one-by-one is continued until the model contains only predictors whose removal is significant.

¹⁵ The Wald's z value is the estimate divided by the standard error; it expresses the effect of a factor relative to the error for that factor.

	Estimate	Standard Error	Wald's z	p
intercept	2.68	0.45	5.94	< .0001
<i>legal vs. illegal</i>	4.52	0.56	8.08	< .0001
<i>co-occurrence vs. ordering</i>	-1.37	0.45	-3.02	< .03

Table 5: Results of final Mixed Logit Model with “correct” as the dependent variable.

Examination of individual subjects shows that all subjects make more errors on phonotactically illegal ejectives than legal ones. Within the illegal ejectives, some subjects make more errors on targets in the co-occurrence category, while others make more errors on the ordering category. Table 6 shows percent correct by category for each subject. Since the number of analyzable tokens differs for each subject, the number of correct items and total items in each category is also given to the left of the percentage.

subject	co-occurrence		ordering		gap		real	
1	12/15	80%	13/14	93%	15/15	100%	13/14	93%
2	4/13	31%	6/13	46%	12/15	80%	15/15	100%
3	1/12	8%	4/15	27%	10/14	71%	15/15	100%
4	5/12	42%	4/10	40%	12/13	92%	15/15	100%
5	7/12	58%	10/13	77%	14/14	100%	15/15	100%
6	2/15	13%	8/15	53%	14/15	93%	15/15	100%
7	6/14	43%	5/15	33%	12/14	86%	15/15	100%
8	5/15	33%	13/15	87%	15/15	100%	14/14	100%
total	42/108	39%	63/110	57%	104/115	90%	117/118	99%

Table 6: Percent correct by category and subject.

The barplot in Figure 3 shows accuracy on the co-occurrence and ordering categories for each subject. As can be seen, overall accuracy on illegal categories varies greatly between subjects, as does relative accuracy on the co-occurrence and ordering categories. Some subjects perform better on the co-occurrence category than on the ordering category, while others show the opposite pattern. Additionally, subjects differ in the size of the difference in accuracy between the two categories.

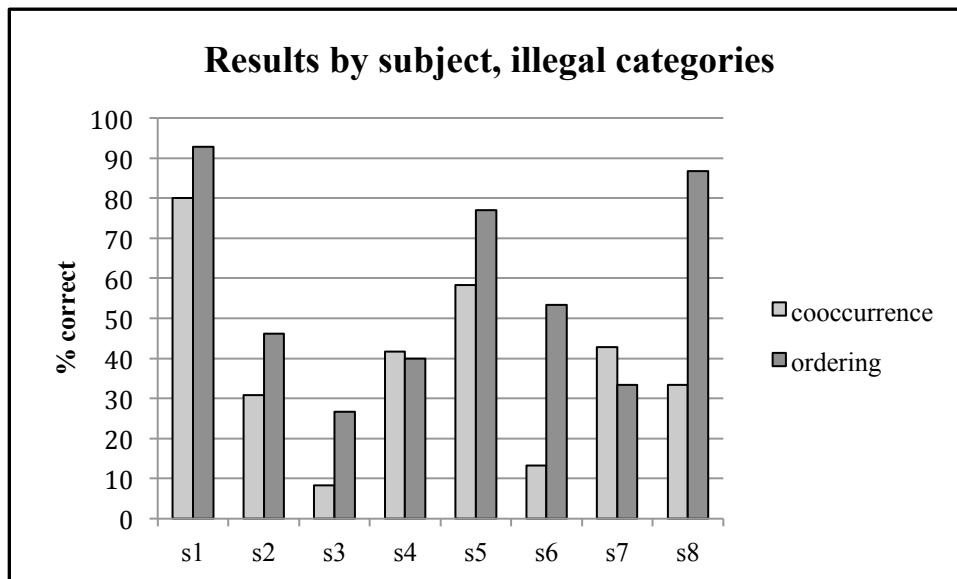


Figure 3: Percent correct by subject, ordering and co-occurrence categories.

3.4.2 Acoustic measures of responses

Acoustic measurements were taken for all medial target ejectives. For correct responses with medial ejectives, the acoustic properties of the ejective can be compared across all four categories, allowing for an analysis of the effect of attestation and phonotactic legality on the realization of an ejective. For each acoustic measure, a Helmert coded, linear mixed model was run with a random intercept and slope for subject and item. The Helmert coded factors made three comparisons: (i) co-occurrence vs. ordering categories (ii) gap vs. co-occurrence and ordering categories (iii) real vs. gap, co-occurrence and ordering categories. For each acoustic measure, the same method of model comparison as described for the accuracy results above was used to determine the significance of each predictor.

For medial ejectives, VOT in the real category differs significantly from the other three categories (Est = 15.06, SE = 6, $t = 2.51$, $p < .02$); no other differences are significant. Neither burst amplitude nor intensity difference differs between categories. Pitch in the following vowel differs between the co-occurrence and ordering categories (Est. = 27.88, SE = 11.36, $t = 2.45$, $p < .05$) and glottalization differs marginally between the real category and the others (Est. = 13.93, SE = 6.98, $t = 1.1$, $p = .057$). Figure 4 shows the average values by category for the significant measures. Average values for the other measures are given in Appendix B, along with the average values for each measure for target ejectives produced as plain stops.

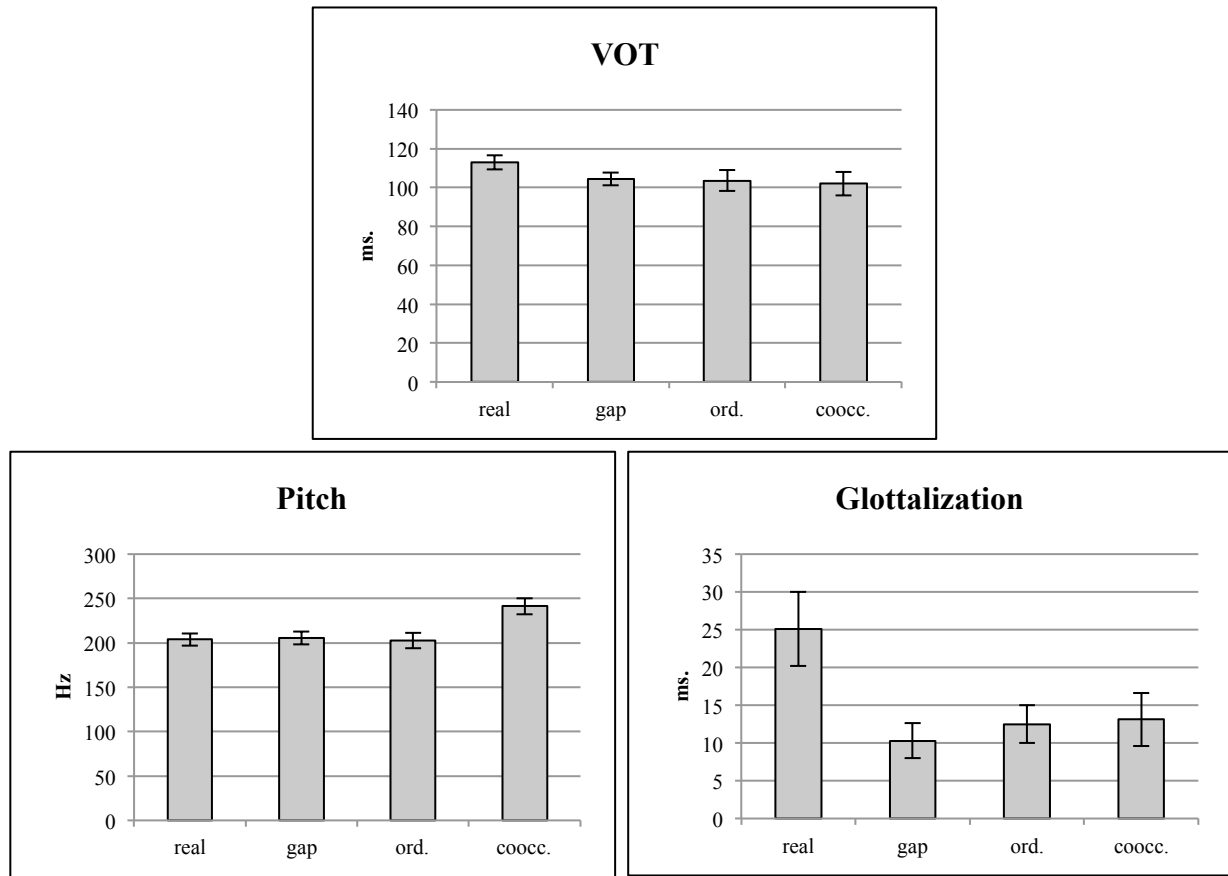


Figure 4: Average values by category for VOT, pitch and glottalization; error bars indicate Standard Error.

These measures show that phonotactic legality does not effect the production of an ejective, with the exception of pitch for ejectives in cooccurrence category stimuli. Ejectives in attested words show a longer VOT and more glottalization, an articulatorily more extreme or prototypical production, than the legal and illegal nonce words. Within the nonce words, however, the articulation of ejection does not differ between the phonotactically legal words in the gap category and illegal words in the ordering and cooccurrence categories. The only exception is the pitch of ejectives in the cooccurrence category, which is higher than for the other categories.

3.4.3 Error types

For 58% of the responses, no error of any kind was recorded. For the remaining 42% of responses that did contain some error, subjects' repetitions differed from the target in a variety of ways. In 11% of responses, there was some mismatch in place of articulation (e.g., target [pit'u] produced as [pik'u]), insertion/deletion of a coda (e.g., target [ʎuntʃ'u] produced as [ʎutʃ'u]), or a change in a vowel (e.g., target [mant'u] produced as [mant'i]). These errors do not change the

status of a word as either grammatical or ungrammatical and will not be discussed further, as they do not bear on the main question of how ungrammatical targets are repaired.

A handful of errors (2% of total responses) changed an ungrammatical target into a different ungrammatical target by doubling ejection on ordering targets (target [kant'u] produced as [k'ant'u]). Doubling of ejection was particularly common for subject 5, but was also attested for subjects 1 and 7. Another .5% of errors changed a grammatical target into an ungrammatical structure by changing the manner of C₁ to a stop. There are two errors of this type, both for the same word: target [map'a] is produced as [pap'a] for two subjects.

The majority of errors were repairs (28.5%), changing an ungrammatical target to a grammatical item in a variety of ways. By far the most common repair was to de-ejectivize a medial ejective (e.g., target [t'ap'u] produced as [t'apu]), occurring in 22% of responses. The next most common repair was to move ejection from medial to initial position in ordering category targets (e.g., target [tap'u] produced as [t'apu]), occurring in 5% of responses. The remaining 1.5% changed the manner of C₁, thus rendering the medial ejective grammatically licit (e.g., target [kant'u] produced as [hant'u]).

For the co-occurrence category, de-ejectivization of the medial ejective is the only commonly attested repair strategy. For the ordering category, however, there is some variation between de-ejectivization of the medial ejective and movement of ejection from the medial to the initial stop. Figure 5 shows the percentage of responses in the ordering category that are either de-ejectivization or movement of ejection. Table 7 gives the number of errors of each type for each subject.




Figure 5: Percent of responses on the ordering category that were de-ejectivizing C₂ or moving ejection from C₂ to C₁, by subject.

subject	deletion	movement
1	1/14 = 7%	0/14 = 0%
2	6/13 = 46%	1/13 = 7.5%
3	5/15 = 33%	7/15 = 47%
4	4/10 = 40%	2/10 = 20%
5	1/13 = 8%	3/13 = 23%
6	4/15 = 27%	3/15 = 20%
7	4/15 = 27%	6/15 = 40%
8	2/15 = 13%	0/15 = 0%
total	15/37 = 41%	15/37 = 41%

Table 7: Number of errors on the ordering category that were either de-ejectivizing C_2 or moving ejection from C_2 to C_1 , by subject.

As can be seen from Figure 4 and Table 7, there is substantial variation between subjects as to whether de-ejectivization, movement, or both were employed as repairs.

3.5 Discussion

3.5.1 Accuracy and acoustics

The accuracy results provide experimental support for the synchronic reality of both co-occurrence and ordering restrictions on ejectives, the latter of which is cross-linguistically quite rare. Subjects make more errors on the ungrammatical medial ejectives in the co-occurrence and ordering categories than on the grammatical ejectives in either the gap or real categories. This result shows that the static phonotactic restrictions on ejective-ejective pairs and plain stop-ejective pairs are synchronically present in the grammars of Quechua speakers.

The relative accuracy on the co-occurrence and ordering categories differs substantially across subjects. Two subjects are substantially more accurate on the ordering category than on the co-occurrence category (Subjects 6 and 8), though this asymmetry between the two categories is not echoed in all the other subjects. This result does not clearly favor one of the two analyses of the restrictions over the other. Subjects with a small distinction in accuracy between co-occurrence and ordering categories are compatible with an analysis like Mackenzie's (2009), where co-occurrence and ordering restrictions are treated as a single restriction, while the two subjects that show larger distinctions between co-occurrence and ordering categories suggest that the two restrictions must have distinct formal representations. What is clear from the current data is that all subjects have learned both the co-occurrence and ordering restrictions on ejectives; data from more subjects is needed to reliably quantify the range of variation between subjects and to make conclusions about the formal representation of the restrictions.

Errors on the repetition task could have multiple sources, resulting from misperception, misproduction, or more abstract grammatical change. Each of these possibilities is discussed in

more detail below. While it is not possible to decisively locate the source of errors on the repetition task given the available data, the general result that Quechua speakers distinguish between legal and illegal structures holds regardless of the source of these errors. It is well known that perception of phonological structures is strongly influenced by the phonological grammar; speakers are attuned to contrasts that are attested in their language and are biased towards perceiving phonotactically legal structures (Werker & Tees 1984; Hallé et al. 1998; Pitt 1998; Dupoux et al. 1999; Moreton 2002; Hallé & Best 2007). Similarly, articulatory coordination is learned and language specific (Bradley 2002; Browman & Goldstein 1990; Byrd 1995, 1996; Gafos 2002; Hall 2003; Kochetov 2002; Zsiga 2003). Thus, the distinction between phonotactically legal and illegal ejectives in the repetition task is a grammatical distinction, regardless of whether the source of the errors is in the perception grammar, the production grammar, or a more abstract, phonological grammar.

It could be that Quechua speakers misperceive illegal ejectives as plain stops, precisely because they are illegal. If Quechua speakers are biased towards perceiving phonotactically legal structures, they may be less accurate in perceiving the difference between an ejective and a plain stop in environments where this contrast does not occur in the native lexicon. Since roots like [k'ap'i] and [kap'i] are unattested in the language, but roots like [k'api] and [kapi] are attested, speakers may be biased to perceiving a medial plain stop as opposed to an ejective. In this scenario, Quechua speakers are misperceiving the target item, but are accurately repeating their intended target. A follow up perception study testing Quechua speakers' accuracy on discriminating between the illegal targets and the common, legal repairs would test whether errors on the repetition task could primarily result from systematic misperception.

An articulatory explanation for the repetition errors would hypothesize that speakers accurately hear the illegal targets, and intend to reproduce them, but fail. This explanation is particularly unlikely for the targets in the ordering category, since it is not obvious why it would be more difficult to produce a medial ejective when the initial segment of the word is a stop than when it is a sonorant or fricative. It is possible, however, that targets in the co-occurrence category pose some articulatory difficulty, since they require producing two articulatorily complex sounds relatively close together. If errors on the co-occurrence category are a result of articulatory difficulty, then accurate repetition of multiple ejectives should increase with the intervening distance. This could be tested, for example, by seeing if Quechua speakers are more accurate on hypothetical trisyllabic roots like [k'amit'u] than on [k'at'u], when a syllable intervenes between the two ejectives.

While the current data do not include a direct test of the articulatory hypothesis, there is reason to be doubtful that articulatory difficulty plays a role. First, a comparison can be made between CVCV co-occurrence targets and CVCCV co-occurrence targets. If articulatory difficulty is at play, it might be expected that subjects are more accurate on CVCCV co-occurrence targets than on CVCV co-occurrence targets, since a coda lengthens the interval between the two target ejectives. This does not turn out to be the case, however. Measuring from the onset of V₁ until the closure of medial stop, duration does differ significantly between CVCV

(129 ms.) and CVCCV (276 ms.) words (Generalized Linear Model: Est. = 147.2, SE = 8, $z = 18.5$, $p < .0001$), but there is no difference in accuracy between the two types of stimuli (Generalized Linear Model: Est. = 0.3, SE = 0.42, $z = 0.73$, $p = 0.47$).

Second, as shown in §3.4.2, the acoustic measures of medial ejectives do not show any consistent differences between legal and illegal ejectives. There is thus also no evidence from the acoustic data that phonotactically illegal ejectives are articulatorily more difficult than legal ejectives for Quechua speakers.

Finally, it should be noted that sequences of ejectives in adjacent syllables do occur across morpheme boundaries in Quechua. For example, in an adjective-noun phrase, an adjective with a medial ejective and a noun with an initial ejective would create a sequence of ejectives across just a single vowel (e.g., [misk'i t'anta] 'sweet bread'). Examples of this type are given in the Bills et al.'s (1969) description of Cochabamba Quechua, and are not reported to be modified in any way, though the actual production of phrases of this type has not been directly investigated.

3.5.2 Error types and ordering restrictions

Different formal analyses consistent with the ordering restriction data predict different input-output mappings. Under some analyses, the position of ejection is switched and an input with a plain stop-ejective sequence maps to an ejective-plain stop sequence, e.g., /pat'i/ → [p'ati]; under other analyses, the medial stop is de-ejectivized from an input with a plain stop-ejective sequence, e.g., /pat'i/ → [pati]. Instead of choosing between these two analyses, the repetition data support both mappings. When repeating targets in the ordering category, both movement of ejection and de-ejectivization are common repairs, though de-ejectivization seems more common.

One explanation for variation in error type is that different types of errors result from different components of the grammar. For example, one error type could be a result of misperception, while another is a result of a phonological change or misproduction. Alternatively, both error types could be generated in a single grammatical component and the choice of repair could be truly variable. As discussed above, the current data is not sufficient to localize the source of the errors in the repetition task, and thus firm conclusions about the implications of the variable repairs for the formal analysis of ordering restrictions cannot be made.

If, however, both types of errors occur in a single grammatical component (i.e., both are due to misperception, misproduction, or phonological change), then the variability in repairs favors an analysis of the ordering restriction as the effect of a markedness constraint explicitly disallowing a plain stop-ejective sequence, as opposed to a categorical ALIGN constraint. Recall the three analyses sketched in §2.2 above, summarized in (12).

- (12) a. MAX[*cg*] >> ALIGN[*cg*] >> NOFLOP[*cg*] /pat'i/ → [p'ati]
 b. *[-cont, -son]...[+cg], MAX[*cg*] >> NOFLOP[*cg*] /pat'i/ → [p'ati]
 c. *[-cont, -son]...[+cg], NOFLOP[*cg*] >> MAX[*cg*] /pat'i/ → [pati]

If the operative markedness constraint in the ordering restriction is ALIGN, then the only possible repair for a plain stop-ejective sequence is movement of ejection. The ALIGN based analysis crucially relies on the ranking of MAX[*cg*] over ALIGN in order to account for the grammaticality of medial ejectives that are not preceded by plain stops; if ALIGN outranks MAX[*cg*], then deletion of ejection is predicted both for inputs with ungrammatical medial ejectives like /pat'i/ and for inputs with grammatical medial ejectives like /mat'i/. Given the markedness constraint *[-cont, -son]...[+cg], however, only medial ejectives preceded by plain stops are penalized and faithfulness is not needed to protect medial ejectives preceded by non-stops, as in /mat'i/. As shown in (12b,c), the relative ranking of MAX[*cg*] and NOFLOP[*cg*] determines whether deletion or movement of ejection is the predicted repair for a violation of *[-cont, -son]...[+cg]. While both markedness constraints can account for the de-ejectivization repair, the variation between de-ejectivization and movement is easily captured with *[-cont, -son]...[+cg] by assuming that MAX[*cg*] and NOFLOP[*cg*] are variably or stochastically ranked with respect to one another.

In a stochastic grammar (Boersma 1997, 1998; Boersma & Hayes 2001; Coetzee & Pater 2009), constraints are weighted probabilistically, as opposed to standing in a fixed ranking as in classic Optimality Theory (Prince & Smolensky 1993/2004). Each constraint occupies a probability distribution of weights, and the distributions of constraints can be overlapping. If two constraints A and B overlap, then for any given input-output mapping there is a chance that A will be weighted higher than B and that B will be weighted higher than A. The probability of each relative weighting scenario is determined by the amount of overlap between the two constraints. For variable repairs to the ordering restriction to be analyzed in this way, the probabilistic weighting distribution of *[-cont, -son]...[+cg] would be higher and non-overlapping with the weighting distributions of MAX and NOFLOP, while the distributions of these two constraints would overlap to some degree.

Finally, in addition to the errors that repair phonotactic violations, 2% of targets were actually changed from one ungrammatical form into a different ungrammatical form by doubling ejection in ordering targets (e.g., producing [kap'i] as [k'ap'i]). While I have no conclusive analysis of these errors, the overall number of errors of this type is much smaller than the number of errors that repair a phonotactic violation. Moreover, the majority of changes from one ungrammatical form to another are produced by a single subject (subject 5), while repairs to phonotactically illegal targets are attested for all subjects.

Section 4 Conclusion

Quechua speakers have more difficulty producing medial ejectives when those medial ejectives violate a phonotactic restriction in their language than when they do not, supporting the synchronic status of these restrictions. Roots with two ejectives, e.g., [k'ap'i], and roots with a plain stop followed by an ejective, e.g., [kap'i], are both unattested in Quechua, and when asked to repeat roots with these structures Quechua speakers are likely to modify them to conform to existing structures in the language.

The relative error rate for targets that violate the co-occurrence and ordering restrictions varies between subjects. Two subjects show a large difference in accuracy between co-occurrence and ordering categories, while the remaining subjects show small or no differences in accuracy. The type of error that is made on targets in the ordering category is also highly variable within and among subjects; de-ejectivization of the medial stop (e.g., [kap'i] produced as [kapi]) and movement of ejection to initial position are both attested (e.g., [kap'i] produced as [k'api]). More data from more subjects is needed to reliably quantify and draw conclusions from the amount of variation, both with respect to accuracy on the co-occurrence and ordering categories and with respect to repairs.

A further topic for future research is localizing the source of errors on the repetition task. It is a priori possible that errors result from misperception, misproduction, or from abstract modification. Further experiments must investigate Quechua speakers' perception and articulation of legal and illegal structures more directly.

References

- Ajacopa, Teofila Laime, et al. 2007. *Diccionario Bilingüe, Iskay Simipi Yuyayk'anacha: Quechua – Castellano Castellano – Quechua*. La Paz, Bolivia.
- Bailey, Todd and Ulrike Hahn. 2001. Determinants of word likeness: phonotactics or lexical neighborhoods? *Journal of Memory and Language* 44:568-591.
- Bates, Doug and Martin Maechler. 2010. lme4: Linear mixed effects models using S4 classes. <<http://CRAN.Rproject.org/package=lme4>>
- Beckman, Jill. 1998. Positional faithfulness. Doctoral Dissertation, University of Massachusetts, Amherst.
- Berent, Iris and Joseph Shimron. 1997. The representation of Hebrew words: evidence from the obligatory contour principle. *Cognition* 64:39-72.
- Berent, Iris and Tracy Lennertz. 2010. Universal constraints on the sound structure of language: Phonological or acoustic? *Journal of Experimental Psychology: Human Perception and Performance* 36(1):212-223.
- Berent, Iris, Donca Steriade, Tracy Lennertz and Vered Vaknin. 2007. What we know about what we have never heard: Evidence from perceptual illusions. *Cognition* 104:591-630.
- Bills, Garland and Rudolph C Troike; Bernardo Vallejo. 1969. *An Introduction to Spoken Bolivian Quechua*. Austin:University of Texas Press.
- Boersma, Paul and Bruce Hayes. 2001. Empirical tests of the Gradual Learning Algorithm. *Linguistic Inquiry* 32:45-86.
- Boersma, Paul. 1997. How we learn variation, optionality and probability. *Proceedings of the Institute of Phonetic Sciences of the University of Amsterdam* 21:43-58.
- Boersma, Paul. 1998. *Functional Phonology: Formalizing the interactions between articulatory and perceptual drives*. Doctoral dissertation, University of Amsterdam. The Hague: Holland Academic Graphics.
- Bradley, Travis. 2002. Gestural timing and derived environment effects in Norwegian clusters. In L. Mikkelsen & Chris Potts, eds., *Proceedings of WCCFL 21*, 43-56. Somerville, MA: Cascadilla Press.
- Broselow, E., & Finer, D. 1991. Parameter setting in second language phonology and syntax.

- Second Language Research, 7(1), 35–59.
- Broselow, E., Chen, S., & Wang, C. 1998. The emergence of the unmarked in second language phonology. *Studies in Second Language Acquisition*, 20, 261–280.
- Browman, C., & Goldstein, L. 1990. Tiers in articulatory phonology, with some implications for casual speech. In J. Kingston, & M. Beckman (Eds.), *Papers in laboratory phonology I: Between the grammar and physics of speech*. Cambridge: Cambridge University Press.
- Byrd, Dani. 1995. C-centers revisited. *Phonetica* 52:285-306.
- Byrd, Dani. 1996. Influences on articulatory timing in consonant sequences. *Journal of Phonetics* 24:209-244.
- Coetzee, Andries and Joe Pater. 2008. Weighted constraints and gradient restrictions on place co-occurrence in Muna and Arabic. *Natural Language and Linguistic Theory* 26:289-337.
- Coetzee, Andries and Joe Pater. 2009. The place of variation in phonological theory. In John Goldsmith, Jason Riggle and Alan Yu (eds.), *The Handbook of Phonological Theory*, (2nd ed). Cambridge: Blackwell.
- Coetzee, Andries. 2005. The Obligatory Contour Principle in the perception of English. In Sónia Frota, Marina Vigária and Maria João Freitas, eds., *Prosodies*, 223-245. Berlin: Mouton de Gruyter.
- Coleman, John and Janet Pierrehumbert. 1997. Stochastic phonological grammars and acceptability. *Proceedings of the third meeting of the Association for Computational Linguistics special interest group in computational phonology*. Somerset, NJ: Association for Computational Linguistics.
- Cusihuamán, Antonio. 1976. *Diccionario Quechua: Cuzco-Collao*. Lima: Ministerio de Educación and Instituto de Estudios Peruanos.
- Davidson, Lisa. 2006. Phonology, phonetics of frequency: Influences on the production of non-native sequences. *Journal of Phonetics* 34(1):104-137.
- Davis, Stuart. 1991. Coronals and the phonotactics of nonadjacent consonants in English. In Carole Paradis and Jean-Francois Prunet (eds.), *The special status of coronals: Internal and external evidence*, 49-60. San Diego: Academic Press.
- Dupoux, Emmanuel, Yuki Hirose, Kazuhiko Kakehi, Christophe Pallier and Jacque Mehler. 1999. Epenthetic vowels in Japanese: A perceptual illusion? *Journal of Experimental Psychology: Human Perception and Performance* 25:1568-1578.
- Frisch, Stefan and Bushra Zawaydeh. 2001. The psychological reality of OCP-place in Arabic. *Language* 77:91-106.
- Frisch, Stefan, Janet Pierrehumbert, and Michael Broe. 2004. Similarity avoidance and the OCP. *Natural Language and Linguistic Theory* 22: 179-228.
- Frisch, Stefan, Nathan Large and David Pisoni. 2001. Perception of wordlikeness: effects of segment probability and length on the processing of nonwords. *Journal of Memory and Language* 42:481-496.
- Fudge, Erik. 1969. Syllables. *Journal of Linguistics* 5:253-287.
- Gafos, Adamantios. 2002. A grammar of gestural coordination. *Natural Language and Linguistic Theory* 20:269-337.
- Gallagher, Gillian. 2010a. The perceptual basis of long-distance laryngeal restrictions. Doctoral Dissertation, Massachusetts Institute of Technology.
- Gallagher, Gillian. 2010b. Perceptual distinctness and long-distance laryngeal restrictions. *Phonology* 27:435-480.

- Gallagher, Gillian. 2011. Auditory features in Phonology – The case for [long VOT]. *The Linguistic Review* 28:281-313.
- Goldrick, Matthew and Larson, Meredith. (2008). Phonotactic probability influences speech production. *Cognition* 107: 1155-1164.
- Goldsmith, John. 1976. *Autosegmental Phonology*. Doctoral Dissertation, Massachusetts Institute of Technology.
- Greenberg, Joseph. 1960. The patterning of root morphemes in Semitic. *Word* 6:162-181.
- Hall, Nancy. 2003. *Gestures and segments: Vowel intrusion as overlap*. Doctoral dissertation, University of Massachusetts, Amherst.
- Hallé, Pierre and Catherine Best. 2007. Dental-to-velar perceptual assimilation: A cross-linguistic study of the perception of dental stop+*/l/* clusters. *Journal of the Acoustical Society of America* 121:2899-2914.
- Hallé, Pierre, Juan Segui, Uli Frauenfelder and Christine Meunier. 1998. Processing of illegal consonant clusters: A case of perceptual assimilation? *Journal of Experimental Psychology: Human Perception and Performance* 24(2):592-608.
- Hansen, J. 2004. Developmental sequences in the acquisition of English L2 syllable codas. *Studies in Second Language Acquisition*, 26, 85–124.
- Hansson, Gunnar Olafur. 2001. Theoretical and typological issues in consonant harmony. Doctoral Dissertation, University of California, Berkeley.
- Hayes, Bruce & Colin Wilson. 2008. A maximum entropy model of phonotactics and phonotactic learning. *Linguistic Inquiry* 39:379-440.
- Herbert, R.K. 1977. Phonetic analysis in phonological description: prenasalized consonants and Meinhof's Rule. *Lingua* 43:339-373.
- Kawahara, Shigeto. 2008. Phonetic naturalness and unnaturalness in Japanese loanword phonology. *Journal of East Asian Linguistics* 17:317-330.
- Keating, Patricia, Taehong Cho, Cécile Fougeron and Chai-Shune Hsu. 2003. Domain-initial articulatory strengthening in four languages. In John Local, Richard Ogden and Rosalind Temple, eds., *Phonetic interpretation: Papers in laboratory phonology VI*, 145-163. Cambridge: Cambridge University Press.
- Kingston, John. 1985. The phonetics and phonology of the timing of oral and glottal events. Doctoral Dissertation, University of California, Berkeley.
- Kochetov, Alexei. 2002. *Production, perception and emergent phonotactic patterns: A case of contrastive palatalization*. New York: Routledge.
- Leben, W. 1973. *Suprasegmental Phonology*. Doctoral Dissertation, Massachusetts Institute of Technology.
- Lindau, Mona. 1984. Phonetic differences in glottalic consonants. *Journal of Phonetics* 12: 147-155.
- MacEachern, Margaret. 1999. *Laryngeal co-occurrence restrictions*. New York: Garland.
- Mackenzie, Sara. 2009. Contrast and similarity in consonant harmony processes. Doctoral dissertation, University of Toronto.
- McCarthy, John and Alan Prince. 1993. Generalized alignment. *Yearbook of Morphology*: 79-153.
- McCarthy, John. 1986. OCP effects: gemination and anti-gemination. *Linguistic Inquiry* 17:207-63.
- McCarthy, John. 1988. Feature geometry and dependency: A review. *Phonetica* 43: 84-108.
- McCarthy, John. 2003. OT constraints are categorical. *Phonology* 20:75-138.

- Meinhof, C. 1932. *Introduction to the phonology of the Bantu languages*. Berlin:Dietrick Reimer/Ernst Vohsen.
- Mester, Armin and Junko Itô. 1986. The phonology of voicing in Japanese. *Linguistic Inquiry* 17:49-73.
- Moreton, Elliott. 2002. Structural constraints in the perception of English stop-sonorant clusters. *Cognition* 84:55-71.
- Moreton, Elliott. 2002. Structural constraints in the perception of English stop-sonorant clusters. *Cognition* 84:55-71.
- Moreton, Elliott. 2008. Analytic bias and phonological typology. *Phonology* 25:83-127.
- Munson, Benjamin. 2001. Phonological pattern frequency and speech production in adults and children. *Journal of Speech, Language and Hearing Research* 44(4):778–792.
- Ohala, John J. 1993. The phonetics of sound change. In Charles Jones, ed. *Historical Linguistics: Problems and Perspectives*, 237-278. London: Longman.
- Ohala, John. J. 1981. The listener as a source of sound change. In C. S. Masek, R. A. Hendrick and M. F. Miller (eds.), *Papers from the Parasession on Language and Behavior*, 178-203. Chicago Linguistics Society: Chicago.
- Parker, Steve and David Weber. 1996. Glottalized and aspirated stops in Cuzco Quechua. *International Journal of American Linguistics* 62:70-85.
- Parker, Steve. 1997. An OT account of laryngealization in Cuzco Quechua. *Working Papers of the Summer Institute of Linguistics, University of North Dakota Session* 41:1-11.
- Parker, Steve. 2007. Un análisis acústico del quechua del Cusco. *GIALens* 3.
- Pierrehumbert, Janet and David Talkin. 1991. Lenition of /h/ and glottal stop. In G. Docherty and D. Ladd, eds. *Papers in Laboratory Phonology II*. Cambridge: Cambridge University Press.
- Pierrehumbert, Janet. 1993. Dissimilarity in Arabic verbal roots. In A. Schafer (ed.), *NELS 23:Proceedings of the North East Linguistic Society*, 367-381. GLSA: Amherst.
- Pitt, Mark. 1998. Phonological processes and the perception of phonotactically illegal consonant clusters. *Perception & Psychophysics* 60:941-951.
- Prince, Alan and Paul Smolensky. 2004 [1993]. *Optimality Theory: Constraint Interaction in Generative Grammar*. Maldin, MA: Blackwell.
- Rose, Sharon and Lisa King. 2007. Speech error elicitation and co-occurrence restrictions in two Ethiopian Semitic languages. *Language and Speech* 50:451-504.
- Saffran, Jenny R. & Erik D. Thiessen. 2003. Pattern induction by infant language learners. *Developmental Psychology* 39. 484–494.
- van den Berg, Rene & Ode Sidu. 1996. *Muna-English Dictionary*. Leiden:KITLV Press.
- Vitevitch, Michael and Paul Luce. 1999. Probabilistic phonotactics and neighborhood activation in spoken word recognition. *Journal of Memory and Language* 40:374-408.
- Vitevitch, Michael and Paul Luce. 2005. Increases in phonotactic probability facilitate spoken nonword repetition. *Journal of Memory and Language* 52(2):193–204.
- Vitevitch, Michael, Paul Luce, Jan Charles-Luce, and David Kemmerer. 1997. Phonotactics and syllable stress: Implications for the processing of spoken nonsense words. *Language and Speech* 40:47–62.
- Walter, Mary Ann. 2007. *Repetition avoidance in human language*. Doctoral dissertation, Massachusetts Institute of Technology.
- Werker, Janet and Richard Tees. 1984. Cross-language speech perception: Evidence for perceptual reorganization during the first year of life. *Infant Behavior and Development* 7:49-63.

- Wilson, Colin. 2003. Analytic bias in artificial phonology learning: consonant harmony vs. random alternation. Handout from presentation at the Workshop on Markedness and the Lexicon, MIT.
- Wright, Richard, Sharon Hargus and Katharine Davis. 2002. On the categorization of ejectives: data from Witsuwit'en. *Journal of the International Phonetic Association* 32:43-77.
- Zoll, Cheryl. 1998. Positional asymmetries and licensing. In John McCarthy (ed.), *Optimality Theory in Phonology: A reader*, 365-378. Oxford:Blackwell.
- Zsiga, Lisa. 2003. Articulatory timing in a second language: Evidence from Russian and English. *Studies in Second Language Acquisition* 25: 399-432.

Appendix A – acoustic properties of the medial ejectives in the stimuli

Acoustic measurements of the medial ejectives in each stimulus item were taken in Praat:

Burst amplitude - measured in arbitrary units directly off the waveform in Praat by subtracting the minimum amplitude from the maximum.

Intensity - the difference between maximum and minimum intensity in dB for the period from the beginning of the burst until the onset of the following vowel, measured as the beginning of the first glottal pulse.

Voice quality - When available, pitch in the first 30 ms of the vowel was taken. In cases of extreme glottalization, pitch is often not resolved in Praat; in these tokens, the duration of glottalization was measured. For each stimulus item, only one measure of voice quality is given, and the cell for the other measure is shaded.

stimulus	category	burst (arbitrary units)	intensity difference (db)	pitch (Hz)	glottalization (ms)
haʎp'a	real	0.4	11	133	
hutʃuj	real	2	14		20
ʎank'ay	real	1.01	25	155	
ʎant'a	real	0.87	22		20
ʎimp'ij	real	0.71	17	138	
ʎuntʃu	real	2	17	109	
matʃi	real	2	22	171	
map'a	real	0.48	14	181	
mast'a	real	0.71	17		26
mat'i	real	0.42	15	162	
misk'i	real	2	28		57
ruk'ana	real	0.91	21		33
satʃa	real	1.58	22	94	

suʌk'a	real	1.07	27	171	
wirp'a	real	0.4	11	133	
hutʃ'a	gap	2	25	174	
hup'i	gap	0.51	15	196	
limp'u	gap	0.31	5	214	
lip'u	gap	0.31	5	213	
luk'a	gap	1.06	14		70
mant'u	gap	0.37	15		77
nump'i	gap	0.71	12		22
rank'i	gap	2	27	149	
rik'u	gap	0.58	17	192	
sanq'e	gap	2	17		17
saq'o	gap	2	21		31
yak'i	gap	2	25		66
yank'u	gap	0.58	18	186	
yap'i	gap	0.31	18		56
yump'a	gap	0.48	22	182	
tʃʌnq'o	co-occurrence	2	20	157	
tʃʌp'a	co-occurrence	0.48	18	180	
kʌnt'i	co-occurrence	0.42	15	174	
kʌp'u	co-occurrence	0.31	13		56
kʌt'u	co-occurrence	0.47	25	244	
pʌnk'u	co-occurrence	0.58	19	187	
pʌntʃ'a	co-occurrence	2	25	175	
pʌnt'i	co-occurrence	0.42	15	172	
qʌmp'a	co-occurrence	0.48	18	180	
qʌntʃ'a	co-occurrence	2	25	174	
qʌp'u	co-occurrence	0.31	7	198	
tʌp'u	co-occurrence	0.31	16	213	
tʌik'aj	co-occurrence	1.01	25	159	
tʌink'aj	co-occurrence	1.01	25	159	
tʌuk'i	co-occurrence	2	32		61
tʃink'aj	ordering	1.01	25	155	
tʃuk'i	ordering	2	28		57
kap'ij	ordering	0.71	14		26
kat'u	ordering	0.37	15		77
kitʃ'i	ordering	2	23	172	
kuntʃ'a	ordering	2	25	174	
kunt'a	ordering	0.87	20		21

pank'u	ordering	0.58	19	189	
pitʃ'u	ordering	2	14		20
pint'a	ordering	0.87	20		23
pit'u	ordering	0.37	15		77
puk'ay	ordering	1.01	24	154	
tik'aj	ordering	1.01	24	154	
tip'a	ordering	0.48	18	180	
tump'ij	ordering	0.71	18	180	

Appendix B – acoustic properties of medial stops

ejectives – burst amplitude (arbitrary units)

	avg (SD)	<i>n</i>
real	1.39 (0.06)	101
gap	1.18 (0.07)	85
ordering	1.26 (0.11)	46
co-occurrence	1.17 (0.12)	33

ejectives - intensity difference (dB)

	avg (SD)	<i>n</i>
real	14.25 (0.62)	99
gap	12.87 (0.72)	78
ordering	13.09 (1.03)	45
co-occurrence	12.55 (1.18)	34

Target ejectives produced as plain stops

	avg (SD)	<i>n</i>
VOT (ms)	36 (23)	93
burst (arbitrary units)	.47 (.46)	91
intensity difference (dB)	4.84 (3.88)	92
pitch (Hz)	195.5 (66.5)	86
glottalization (ms)	3 (11.5)	88