ASYMMETRIES IN THE REPRESENTATION OF CATEGORICAL PHONOTACTICS

GILLIAN GALLAGHER

New York University

A comparison of speakers' treatment of two categorically unattested phonotactic structures in Cochabamba Quechua reveals a stronger grammatical prohibition on roots with pairs of ejectives, *[k'ap'u], than on roots with a plain stop followed by an ejective, *[kap'u]. While the distribution of ejectives can be stated as a single restriction on ejectives preceded by stops (ejective or plain), *[−cont, −son][cg], speakers show evidence of having learned an additional constraint that penalizes cooccurring ejectives more harshly, *[cg][cg]. An inductive learning bias in favor of constraints with the algebraic structure of *[cg][cg] is hypothesized (Marcus 2001, Berent et al. 2002, Berent et al. 2012), allowing such constraints to be preferred by learners over constraints like *[−cont, −son][+cg], which penalize sequences of unrelated feature matrices.*

Keywords: obligatory contour principle, Quechua, learning, analytic bias, phonetic grounding

1. INTRODUCTION. Investigations of phonotactic and morphophonological knowledge have revealed that speakers do not generalize perfectly from the lexicon (Hayes et al. 2009, Becker et al. 2011, Becker et al. 2012, Kager & Pater 2012, Hayes & White 2013). Rather, some systematic patterns in the lexicon influence performance on nonceword tasks less than would be expected, while other patterns are generalized freely to new items. Typically, it is structurally simpler patterns that are generalized more readily than more complex or phonologically arbitrary patterns, a finding that is echoed in artificial grammar learning experiments (Moreton 2008, Moreton & Pater 2012, Pater & Moreton 2012) and suggests biases in phonological learning in favor of generalizations with certain structural properties.

This article presents evidence that a constraint with the algebraic structure *[αF][αF], that is, an OBLIGATORY CONTOUR PRINCIPLE (OCP) constraint (Leben 1973, Goldsmith 1976, McCarthy 1988, Suzuki 1998), is learned better than an otherwise comparable constraint without this structure. Algebraic or variable notation has been a part of phonological theory since the *Sound pattern of English* (Chomsky & Halle 1968) and was used in early work to represent assimilatory and dissimilatory rules. A substantial body of work has since established the necessity of representing abstract relationships between phonological atoms (segments or features), as variable notation allows. In an artificial grammar learning study, Marcus and colleagues (1999) find that infants generalize a phonotactic pattern based on repeated segments to novel segments, showing that they have learned the abstract identity relation. Looking at the existing phonotactics of a natural language, a series of studies shows that native Hebrew speakers have learned restrictions on the distribution of identical segments (Berent & Shimron 1997, Berent et al. 2001) and that they generalize these restrictions to novel segments (Berent et al.

2002). Marcus (2001) and Berent and colleagues (2012) then show that these empirical results can only be accounted for in a model that explicitly references variables and thus can learn algebraic dependencies (see also Gallagher 2013b).

The current study builds on previous work by suggesting that algebraic generalizations not only are possible grammatical representations, but are actually favored by learners. The case study is the distribution of ejective consonants in Cochabamba Quechua. In this language, ejectives are prohibited from appearing in root-medial position if the initial consonant in the root is either a plain stop, for example, *[kap'a], or another ejective, for example, *[k'ap'a]. While these two restrictions can be grouped under a single generalization, namely, that ejectives may not be preceded by stops of any kind, speakers' treatment of nonce words points to a stronger grammatical prohibition on pairs of cooccurring ejectives. This result suggests a privileged status for constraints with the algebraic structure of the OCP, here *[constricted glottis][constricted glottis] (henceforth *[cg][cg]), allowing such constraints to be learned even when they are not strictly necessary to account for the distributional patterns in a language.

The article is organized as follows. First, §2 lays the groundwork, presenting the descriptive facts of Cochabamba Quechua and introducing the issues faced by an inductive phonotactic learner. Section 3 then outlines the specific research questions addressed and motivates the experimental designs. The repetition task and two discrimination tasks are presented in §3 and §4, respectively. A summary of the findings (§6) and formalization of a grammatical analysis (§7) are then provided, followed by a conclusion (§8).

2. Inductive learning of phonological patterns. Recent work has argued that phonological patterns are learned via induction, based on a statistical analysis of the learning data combined with prior biases about the shape and content of phonological generalizations. Evidence that phonological patterns are learned via induction, as opposed to a universal set of constraints (Prince & Smolensky 2004 [1993]), comes from a range of studies showing that speakers have detailed knowledge of both the categorical and gradient sound patterns of their languages, some of which are highly irregular and phonologically arbitrary (Ernestus & Baayen 2003, Hayes et al. 2009, Gouskova & Becker 2013, Hayes & White 2013). Evidence that the process of induction is biased to favor certain patterns over others comes from mismatches between lexical statistics—that is, the learning data—and speaker behavior (Hayes et al. 2009, Becker et al. 2011, Becker et al. 2012, Kager & Pater 2012, Hayes & White 2013).

To take one example, Hayes and colleagues (2009) tested Hungarian speakers' generalization of vowel harmony to novel stems with the phonologically neutral front unrounded vowels [i i: e: ɛ]. Existing Hungarian stems with these vowels are lexically specified as to whether they take front or back allomorphs of suffixes (e.g. [tsi:m-nɛk] 'address, dative' but [hi:d-nɔk] 'bridge, dative'), though various phonological properties of a stem make selection of front/back suffixes more or less likely. Hayes and colleagues found that properties of the stem vowel, such as whether it was high or nonhigh, were equally strong predictors of suffix allomorphy in the lexicon and in nonce stems, while properties of the stem-final consonant, such as whether the stem ended in a sibilant, were stronger predictors of suffix allomorphy in the lexicon than in nonce stems. In this case, generalizations about the properties of the triggering stem vowel have been learned by Hungarian speakers better than arbitrary generalizations about the consonantal make-up of back- and front-selecting stems. Importantly, Hungarian speakers do show evidence of having learned the predictive properties of stem-final consonants, supporting learning through induction. Phonologically simple dependencies between

stem and suffix vowels, however, are learned more strongly, suggesting a bias in favor of inducing generalizations of this form.

Speakers' behavior on nonce forms reveals the generalizations they have learned, which constitute their phonological grammar. By comparing speakers' behavior to what would be expected given the statistical evidence available to them in their language, we gain insight into the principles that shape how grammars are learned. This section describes a categorical pattern in Cochabamba Quechua and evaluates the lexical evidence for the restriction in §2.1. How an inductive learner might represent the pattern is then discussed in §2.2.

2.1. COCHABAMBA QUECHUA: THE PATTERN. Cochabamba Quechua, a variety of South Bolivian Quechua, contrasts voiceless unaspirated (or plain) [p t tʃ k q], aspirated [pʰ tʰ tʃʰ kʰ qʰ], and ejective [p' t' tʃ' k' q'] stops and affricates. The aspirated and ejective series are subject to a range of restrictions, common to Quechua languages spoken in Bolivia and Southern Peru, as described previously in Parker & Weber 1996, Parker 1997, MacEachern 1999, and Gallagher 2011, 2013a. Ejectives and aspirates may only occur in the onset position of roots, which are overwhelmingly CV(C)CV. Within roots, there are multiple combinatorial restrictions on ejectives and aspirates. This article focuses on two of the restrictions on ejectives, illustrated in 1. Ejectives may occur root-initially (1a) and may also occur root-medially if the initial root consonant is a fricative or a sonorant (1b). Ejectives may not occur in pairs (1c), and an ejective may not occur medially in roots with an initial voiceless unaspirated stop (1d). The restriction on pairs of ejectives will be referred to as the COOCCURRENCE restriction, and the restriction on medial ejectives in plain stop-initial roots will be referred to as the ORDERING restriction.

(1) a. p'uɲu   'jug'   b. hap'i   'grab'
       t'impu   'boil'   rit'i   'snow'
       tʃ'uɲu   'freeze-dried potato'   satʃ'a   'tree'
       k'atʃa   'pretty'   ʎaŋk'a   'work'
       q'uwa   'ceremonial fire'   liq'i   'kind of hat'
    c. *k'atʃ'a   d. *katʃ'a
       *t'imp'u   *timp'u

Ejectives are further restricted from cooccurring with aspirates, in either order (*[t'impʰu], *[tʰimp'u]), and are also absent from roots that are vowel-initial (*[it'i]). Because of this restriction, vowel-initial roots are often analyzed as glottal stop-initial,*[ʔit'i], although glottal stop is not otherwise found in the language and is non-contrastive (Parker & Weber 1996, MacEachern 1999, Bennett 2013). Aspirates are subject to parallel restrictions: they may not cooccur in pairs (*[tʰimpʰu]) or occur root-medially if the initial consonant is a plain stop ([rakʰu] 'fat' but *[pakʰu]), and they also may not occur with the glottal fricative [h] (*[hakʰu]). While future experimental examination of the full range of restrictions will no doubt reveal interesting findings beyond those presented here, the current study is restricted to the cooccurrence and ordering restrictions on ejectives in order to make the project manageable, as well as to allow comparison with prior work.

Stops (ejective, aspirate, and plain) are found only in onset, or prevocalic, position in Quechua, and ejectives and aspirates are additionally found only in roots. There are no ejectives in suffixes (the language has no prefixes), meaning that the laryngeal structures that are prohibited in the root also do not arise at the word level. Plain stop-ejective sequences do arise across word boundaries with some frequency, and ejective-ejective sequences are also possible, though less common. Quechua is highly agglutinative, so roots in adjacent words within a phrase are often separated from one another by several suf-

fixes. Plain stop-ejective sequences may arise either when two roots are not separated by any suffixes, as in some noun phrases (e.g. [antʃa q'osɲi] 'a lot of trash'), or when a word with an ejective-initial root follows a word with a plain stop in the final suffix, as with the accusative marker -*ta* (e.g. [ajtʃata k'uturqa] 'she cut the meat'). Ejective-ejective sequences arise across word boundaries only when roots are not separated by suffixes (e.g. [misk'i t'anta] 'good bread').

There is strong lexical evidence for the systematic absence of roots with two ejectives or plain stop-ejective pairs, and the evidence is, moreover, comparable for the two restrictions. Both the ordering and cooccurrence restrictions are categorical, and both are dependencies between onset stops within roots. Roots that violate either restriction would be expected to occur a fair number of times if root consonants combined at chance. Given the 2,290 disyllabic roots in Laime Ajacopa's dictionary (2007), forty-four roots with two ejectives and forty-seven roots with an initial plain stop and a medial ejective are predicted.[1] Both ejectives and plain stops are relatively frequent in the language, providing good evidence that unattested onset combinations are systematically absent due to a phonotactic restriction, as opposed to accidentally absent due to a gap (cf. Gorman 2013). Both restrictions have comparable scope, in the sense of Albright & Hayes 2003, as both restrictions rule out twenty-five hypothetical onset combinations. The cooccurrence and ordering restrictions are thus matched in the support they receive from the lexicon; they are also matched in structural form in all ways except that algebraic structure is present for the cooccurrence restriction but not for the ordering restriction.

2.2. WHAT DOES THE GRAMMAR LOOK LIKE? Under one analysis, the grammar of Cochabamba Quechua contains a single phonotactic constraint, prohibiting ejectives from following stops, whether ejective or plain: *[−continuant, −sonorant][constricted glottis] (henceforth *[−cont, −son][cg]). An analysis in this spirit is presented in Mackenzie 2009, 2013. This is perhaps the simplest hypothesis, since it accounts for all of the data with just one constraint. This is also the analysis settled upon by an inductive phonotactic learner, the UCLA phonotactic learner (Hayes & Wilson 2008). Crucially, such an analysis predicts that both types of forms should be equally dispreferred by speakers.

The UCLA phonotactic learner searches through a space of possible constraints to build a grammar that rules out unattested or infrequent structures. The learner takes as input a list of attested forms in a language (the learning data) and featural representations that describe the segments of the language. The learner then searches through the constraints that can be constructed from these featural representations and selects constraints that penalize structures that are infrequent or unattested in the learning data. Constraints may penalize feature cooccurrence (e.g. *[−sonorant, +voice] would rule out voiced obstruents) or sequences of natural classes (e.g. *[−sonorant, −voice][−sonorant, +voice] would rule out voiceless-voiced obstruent clusters). The selected constraints are assigned weights according to the principle of maximum entropy (Jaynes 1983, Manning & Schütze 1999). The grammar (the constraints and their weights) then makes gradient predictions about the well-formedness of both attested and unattested forms. As in other weighted-constraint models, like HARMONIC GRAMMAR (Pater et al. 2007, Coetzee & Pater 2008, Pater 2009, McCarthy & Pater 2016), a given form is assigned a score by summing the weights of the constraint violations, and forms with higher scores are pre-

---

[1] There are 448 roots with initial plain stops, 419 roots with initial ejectives, and 238 roots with medial ejectives. The expected number of plain stop-ejective combinations is (448 * 238) / 2290 = 46.56, and the expected number of ejective-ejective combinations is (419 * 238) / 2290 = 43.55.

dicted to be less well formed. The UCLA phonotactic learner is a useful tool because it has very few preferences for the shape or content of constraints, and thus serves as a baseline for an unbiased inductive learner.[2]

When tasked with learning generalizations over pairs of Quechua onsets, the UCLA phonotactic learner was found to capture the restrictions on ejectives using a single, highly weighted constraint: *[−cont, −son][cg]. This is an intuitive result. Given that the goal of the model is to rule out classes of forms that are unattested, regardless of the featural content of the constraints required to do so, *[−cont, −son][cg] is a good choice. This constraint has the most support from the learning data and does the most work in accounting for the distribution of ejectives, since it rules out seventy-five unattested onset combinations. Any narrower restrictions that rule out only a subset of the unattested forms, such as *[cg][cg], have less support and are thus not selected by the learner.

While a single phonotactic constraint is one possible analysis of Cochabamba Quechua, it is not the only one. The cooccurrence and ordering restrictions can both be represented independently: the cooccurrence restriction as an OCP constraint *[cg][cg] and the ordering restriction as *[−cont, −son, −cg][+cg] (note that binary laryngeal features are necessary to pick out plain stops to the exclusion of ejectives).[3] These constraints may be present in the grammar instead of, or in addition to, the broader restriction. One reason for thinking that at least the OCP constraint is present is the typological frequency of cooccurrence restrictions. As discussed in Gallagher 2013a, cooccurrence restrictions are found in a range of languages for a variety of features. Ordering restrictions, by contrast, are much less frequent. Coetzee (2014) documents a restriction on the ordering of voiced and voiceless stops in Afrikaans and shows that this restriction is only weakly reflected in the judgments of native speakers. Due to the frequency of cooccurrence restrictions, the OCP has been proposed as a universal organizing principle of grammars (Goldsmith 1976, Clements 1985, McCarthy 1986, 1988). As such, it may be expected that the cooccurrence restriction has a privileged status and is represented as an independent constraint in the grammar.

The current studies, described and motivated in the next section, show that forms violating the cooccurrence and ordering restrictions are not treated equivalently by Quechua speakers. Instead, forms that violate the cooccurrence restriction seem to be penalized more heavily, suggesting that an inductive learner needs to be equipped to learn the cooccurrence constraint *[cg][cg] in addition to the more general constraint *[−cont, −son][cg]. The broader implication of this finding is that the grammar is not learned simply by finding the most economical representation of licit and illicit structures. Rather, generalizations with certain structural form, like the algebraic structure of an OCP constraint, are preferentially employed in building a grammar.

3. Motivating the current studies. The current set of studies compares speaker behavior on nonce roots that violate the ordering and cooccurrence restrictions, with the goal of discovering whether these restrictions are represented differently in the grammar. Categorical phonotactic generalizations like these pose a challenge for experimental investigation, as the contribution of top-down grammatical knowledge must be teased apart from substantive difficulties with producing, perceiving, and judging non-

---

[2] The learning algorithm does implement a bias for constraints that (i) refer to fewer features and (ii) refer to fewer feature matrices. Other models of inductive phonotactic learning could also be explored (e.g. Adriaans & Kager 2010, Goldsmith & Riggle 2012), though these models are not equipped with biases of the type argued for in this article.

[3] The use of binary or privative laryngeal features does not affect the outcome of the UCLA phonotactic learner; under both scenarios, a single phonotactic constraint is learned.

native structures. Phonotactic distributions have been found to influence perception and production (Hallé et al. 1998, Dupoux et al. 1999, Berent et al. 2007, Hallé & Best 2007, Rose & King 2007, Coetzee 2008, Davidson 2010, Dupoux et al. 2011), as well as metalinguistic judgments (Scholes 1966, Frisch & Zawaydeh 2001, Albright 2009, Daland et al. 2011), though the source of these effects in top-down phonotactic knowledge as opposed to inexperience with articulatory plans and perceptual maps for unattested structures is often unclear (Davidson & Shaw 2012, Wilson et al. 2014).

In previous work (Gallagher 2013a), I found that both cooccurrence- and ordering-violating roots triggered more repetition errors than phonotactically legal controls, confirming that speakers have learned both restrictions. More interesting is the finding that speakers were more accurate on roots violating the ordering restriction than roots violating the cooccurrence restriction (more errors on targets like *[k'ap'a] than on *[kap'a]), suggesting that the cooccurrence restriction is stronger. A major question that arises from this finding, addressed in the present studies, is whether the difference in error rate on ordering- and cooccurrence-violating forms derives from phonetic or phonotactic differences between the two structures. Errors on the repetition task could be driven by high-level phonotactic knowledge, in which case the difference in error rate could reveal an asymmetry in how the restrictions are encoded in the grammar, or errors may arise in perception or production, in which case the difference in error rate could be due to substantive differences between ordering- and cooccurrence-violating forms.

Previous work has found some evidence for phonetic difficulties with ejective-ejective pairs. English speakers were found in Gallagher 2010a,b to be less accurate at discriminating forms with two ejectives from forms with one ejective (e.g. [k'ap'i] vs. [k'api]) than discriminating forms with one ejective from forms with no ejectives (e.g. [k'api] vs. [kapi]). While Quechua speakers are no doubt better than English speakers at discriminating ejectives from plain stops, Quechua speakers may also have difficulty, since they have no practice with this contrast. In production, Gallagher and Whang (2014) found that cross-word ejective-ejective sequences are subject to occasional de-ejectivization (e.g. /misk'i t'anta/ 'good bread' produced as [misk'i tanta]), suggesting that pairs of transvocalic ejectives may be articulatorily challenging as well. In contrast, the relevant studies have not been conducted with plain stop-ejective pairs. Given that the position of an ejective within a root in Quechua is predictable, however, perceiving or producing an ejective in an unattested position may be challenging.

The current work aimed to ascertain the role of perception, production, and top-down phonotactics by having speakers participate in both a repetition task and a discrimination task. The repetition task had two conditions. The 'local' condition was a near-replication of Gallagher 2013a, in which participants were asked to repeat disyllabic CVCV nonce words that contained a violation of the ordering restriction (e.g. *[kap'a]) or the cooccurrence restriction (e.g. *[k'ap'a]) or were phonotactically legal (e.g. [map'a]). The 'nonlocal' condition included trisyllabic CVCVCV nonce words with phonotactic violations in nonadjacent syllables, for example, *[kamip'a] (ordering), *[k'amip'a] (cooccurrence).

Repetition accuracy was compared between the local and nonlocal conditions, looking for improvement in the nonlocal condition. Both the ordering and cooccurrence restrictions are restrictions on combinations of consonants; ejectives and plain stops are not difficult to produce or perceive on their own; rather, it is particular combinations of ejectives and plain stops within a certain temporal domain that may be difficult. The perceptual and articulatory challenges of these unattested combinations should thus decrease the further apart the interacting consonants are. If perceptual and articulatory dif-

ficulty is at play in inaccurate repetitions, then an effect of distance sensitivity is expected. Improved performance on the nonlocal condition, however, will not necessarily point to a phonetic effect, as the phonological grammar may also be distance-sensitive. Indeed, previous work has frequently noted that cooccurrence restrictions on major place and laryngeal features are weaker for consonants that are not syllable-adjacent (Greenberg 1960, McCarthy 1994, Frisch et al. 2004, Rose & King 2007, Coetzee & Pater 2008). If improvement on the nonlocal condition is found, it may be phonological or phonetic in origin. If no improvement on the nonlocal condition is found, however, then phonetic difficulty is unlikely to be playing a dominant role.

Distance sensitivity is of further interest because of the small number of trisyllabic roots in Quechua. In Laime Ajacopa's (2007) dictionary, there are 118 trisyllabic roots, and only twelve of these have ejectives in the onset of the third syllable (e.g. [humint'a] 'tamale'). Thus, while longer roots are consistent with the phonotactic generalizations of the language, the evidence for the systematic absence of ejective-ejective or plain stop-ejective pairs in nonadjacent syllables is small.

The role of phonotactic knowledge can also be assessed by the types of errors on the repetition task. High-level phonotactic knowledge should only drive errors that are phonotactic repairs (e.g. repetition of target *[k'ap'i] as phonotactically legal [k'api]); nonrepairs (e.g. repetition of target *[k'ap'i] as phonotactically illegal *[kap'i]) cannot be attributed to the phonotactic grammar. While Gallagher 2013a looked at various error types, there was no explicit interpretation of the rate of repairs vs. nonrepairs, though both types of errors were found. Differences in the rate of repair to the two different phonotactic violations may indicate differences in the grammatical representation of these two restrictions.

In the discrimination task, stimuli from the repetition task were paired with all observed errors to diagnose if relative error rates could be due to misperception. For example, in the local condition of the repetition task, de-ejectivization is more common for forms that violate the cooccurrence restriction than for forms that violate the ordering restriction: *[k'ap'a] is repeated as [k'apa] more often than *[kap'a] is repeated as [kapa]. In the discrimination task, accuracy on discriminating pairs like *[k'ap'a] vs. [k'apa] and *[kap'a] vs. [kapa] was compared to assess whether the differing error rates on the repetition task could be due to different rates of misperception.

4. EXPERIMENT 1: REPETITION. In experiment 1, participants listened to nonsense words and were asked to repeat what they heard. Of interest are the rate and type of errors made on phonotactically illegal nonsense words that violate the cooccurrence and ordering restrictions, relative to controls. All materials for experiments 1 and 2 (presented in §5 below) are available at https://archive.nyu.edu/handle/2451/33774.

4.1. PARTICIPANTS. The participants were twenty native speakers of Quechua from Cochabamba Department, Bolivia, in their twenties and thirties. All participants were bilingual in Spanish, which was used as the metalanguage for the experiment. Data from one participant were excluded from the analysis because this participant produced target ejective consonants as aspirates.

4.2. STIMULI. The stimuli for the experiment were nonsense words, either disyllabic C_1VC_2V or trisyllabic $C_1VC_2VC_3V$. The disyllabic words made up the 'local' condition, and the trisyllabic words made up the 'nonlocal' condition; participants completed the local and nonlocal conditions in separate blocks.

Within each condition, there were three categories of target items, as well as some fillers. Stimuli were classified based on the features of the leftmost (C_1) and rightmost

(C_2 in the local condition, C_3 in the nonlocal condition) consonants. In the 'control' items, the rightmost consonant in the word was an ejective, and the preceding consonants were fricatives and/or sonorants, for example, [lap'a], [luwap'i]. In the 'ordering' items, the rightmost consonant in the word was an ejective, and the initial consonant was a plain stop, for example, [kup'a], [kuwip'a]. In the 'cooccurrence' items, the rightmost consonant and the initial consonant were both ejectives, for example, [p'it'a], [p'irut'u]. C_2 in the nonlocal condition was always a fricative or sonorant, so the only phonotactic violations in the stimuli were between the C_1 and C_3. Filler items had a plain stop in the rightmost position and could have an ejective, plain stop, fricative, or sonorant as the initial consonant, for example, [t'apu], [kusuta].

Stimuli were balanced between categories, as much as possible, for place of articulation of C_1 and C_2 and vowel combination. Stimuli in the control, ordering, and cooccurrence categories were also balanced for whether they had a neighbor with a plain stop in the rightmost position; for example, the phonotactically illegal stimulus in the cooccurrence category [k'it'a] has the legal neighbor [k'ita] 'wild', while the phonotactically legal neighbor for [k'up'i], [k'upi], is not an existing root. For the trisyllabic stimuli in the nonlocal condition, swapping a plain stop for an ejective never yielded an existing word, due to the very small number of trisyllabic roots. Instead, neighbors were assessed based on the last two syllables of the root; for example, [p'iɲuk'i] has the neighbor [ɲuk'i] 'lust'. Stress is fixed on the penultimate syllable in Cochabamba Quechua.

There were fifteen items in each category, for a total of sixty stimuli in each condition, shown in Tables 1 and 2.

	CONTROL	ORDERING	COOCCURRENCE	FILLER
	lap'a	tip'i	k'it'a	t'apu
	jup'a	kut'a	p'it'a	k'apu
	juk'u	pit'a	q'ap'i	k'uta
NEIGHBOR	juk'a	kip'u	k'ip'a	k'upi
	siq'a	qat'i	q'at'a	p'uka
	saq'u	tuk'i	p'ik'a	kupa
	muq'u	tuq'i	t'uq'i	tipu
				kipa
	lip'u	kup'a	t'ap'u	kapi
	nap'u	kap'a	k'ip'i	puki
	nat'u	tip'u	q'up'i	lipu
NO NEIGHBOR	nut'a	kip'a	k'ap'u	napu
	jap'i	kap'i	k'ut'a	natu
	lik'i	tip'a	k'up'i	japi
	luk'a	puk'i	p'uk'a	luka
	maq'u	taq'i	t'aq'u	

TABLE 1. Stimulus items for the local condition in the repetition experiment.

The stimuli were made by splicing syllables together from productions of a native Quechua speaker reading phonotactically legal nonce words. Since all stimulus items have a stop as the rightmost consonant, the stimuli were spliced together by making a cut during the silent closure of this consonant. For example, the stimulus [kup'a] was made by splicing the initial CV of [kupa] and the final CV of [jup'a], making the cut during the silent closure of the labial stop; the stimulus [punuk'i] was made by splicing the initial CVCV of [punuki] and the final CV of [liɲak'i], making the cut during the silent closure of the velar stop. All stimuli, regardless of category, were spliced in this fashion. The stimuli were normalized for amplitude, using the 'scale peak' function in Praat (Boersma & Weenink 2011), and ejectives were additionally normalized for a

	CONTROL	ORDERING	COOCCURRENCE	FILLER
NEIGHBOR	nusap'i	kuwip'a	q'usap'i	t'awipu
	jarup'a	tisap'i	p'irut'u	k'imapi
	juwat'u	karap'i	k'urit'a	p'uʎuka
	lumik'i	piɲuk'i	q'awit'a	p'aɲaqi
	siɲiq'a	tamuk'i	p'iɲuk'i	t'uwiqu
	musaq'i	paɲuk'i	t'imik'u	kamupa
	mawiq'i	tasaq'i		tisapu
NO NEIGHBOR	luwap'i	kamup'a	t'awip'u	kusuta
	limip'u	tisap'u	k'imap'i	punuki
	nasit'a	kusut'a	k'isit'a	kiʎupu
	juʎap'u	kurit'u	q'arip'i	yaʎuka
	jawit'i	kiʎup'u	k'iʎup'a	liɲaki
	jaʎuk'a	qawit'i	t'inak'i	nasita
	liɲak'i	punuk'i	p'uʎuk'a	luwapi
	sawiq'u	tuwiq'u	p'aɲaq'i	limipu
			p'aɲaq'u	

TABLE 2. Stimulus items for the nonlocal condition in the repetition experiment.

voice onset time (VOT) of 130 ms. VOT is a main cue, along with burst amplitude, for ejection in Quechua (Gallagher 2010a, 2011), and thus all ejectives in the stimuli were equally well cued along this dimension. The 130 ms value was chosen because it was on the long end of naturally occurring VOTs, resulting in stimuli with strong, clear cues to ejection. This normalization did result in obscuring natural differences in VOT by place of articulation.

4.3. PROCEDURE. The experiment was conducted in a hotel room in Cochabamba, Bolivia. All participants completed both the local and nonlocal conditions, in separate blocks. The order in which participants completed the two conditions was balanced. Within each condition, the stimuli were presented in a unique, random order to each participant. On each trial, a single stimulus was repeated twice, 400 ms apart, and then participants were asked to repeat what they heard as precisely as possible. The experiment was run using the Psyscope X software, running on a Macbook Air. Participants listened to stimuli via Audio Technica headphones, and their responses were recorded with a Marantz PMD560 digital recorder, at a sampling rate of 44 kHz, and an Audio Technica 831b lapel microphone. Participants were told that the words had no meaning in Quechua or in any other language, though they contained sounds that were used in existing Quechua words. The instructions were given in Spanish, both in written form on the computer screen and orally by the experimenter. Participants were allowed to practice with as many trials as they wanted (about five trials for most people) until they were comfortable with the task. Practice trials were randomly selected stimulus items drawn from the full set of experimental items. Items selected for presentation in practice were then represented to participants once they began the experiment. Responses to practice trials were not analyzed. Participants had as long as they needed to respond; once they had repeated the stimulus item, they pressed the space bar on the keyboard to move on to the next trial. There was a break in between the two conditions. The entire experiment took ten to fifteen minutes, depending on the pace of the individual participant.

4.4. ANALYSIS. Responses were transcribed and coded for accuracy and type of error, if any. All responses were transcribed by two people, and by a third in cases of disagreement. Thirty-three responses (1.4% of the data) were removed from the analysis entirely, sixteen from the local condition and seventeen from the nonlocal condition,

because (i) the participant paused or was otherwise disfluent, (ii) a target stop was produced as an aspirate, or (iii) the transcribers could not agree on a transcription.

The overwhelming majority of errors involved productions of a target ejective as a plain stop; for example, a participant heard [k'it'a] and produced [k'ita]. To determine whether a given stop was ejective or plain, the spectrogram and waveform were examined for the characteristic large burst and long VOT of a Quechua ejective.

In addition to mismatches between ejectives and plain stops, there were two other, less common, types of errors: change in a target vowel, for example, [luka] produced as [lika], and change in place of articulation of a target stop, for example, [k'it'a] produced as [k'ip'a]. These errors were ignored in the analysis; that is, responses with these errors were coded as 'correct'.

Various types of errors were found among inaccurate responses, both errors that repaired phonotactic violations and errors that did not. The repair and nonrepair errors on ordering and cooccurrence category targets are shown in Table 3, with examples from the local and nonlocal conditions. The only attested error on control items was de-ejectivization of C_2, for example, [lap'a] repeated as [lapa].

CATEGORY	ERROR	EXAMPLE		REPAIR?
		TARGET	REPEATED AS	
ORDERING	C_2 de-ejectivization	kup'i kuwip'a	kupi kuwipa	yes
	movement	kup'i kuwip'a	k'upi k'uwipa	yes
	doubling	kup'i kuwip'a	k'up'i k'uwip'a	no
COOCCURRENCE	C_2 de-ejectivization	k'up'i k'imap'i	k'upi k'imapi	yes
	C_1 de-ejectivization	k'up'i k'imap'i	kup'i kimap'i	no

TABLE 3. Repair and nonrepair errors on ordering and cooccurrence stimuli in the repetition experiment.

In addition to the errors in Table 3, there was one instance of a local cooccurrence target, [p'uk'a], repeated as [puka], with de-ejectivization of both stops. This item was coded as a repair and included in the analyses below as such, though this error type is not discussed further.

4.5. RESULTS.

ACCURACY. Accuracy on the control, ordering, and cooccurrence categories in both the local and nonlocal conditions is shown in Figure 1.

Participants were more accurate at repeating the phonotactically legal nonce words in the control category than the phonotactically illegal nonce words in the ordering and cooccurrence categories. To confirm the effect of grammaticality—the phonotactically legal control category as compared to the illegal ordering and cooccurrence categories— a MIXED LOGIT MODEL (MLM) was fit using the lmer() function in the lme4 package (Bates & Maechler 2010) in R (R Core Team 2015). All predictors were centered, and a maximal random-effects structure was used with random slopes for all predictors by subject and random intercepts by subject and stimulus item. There was a significant effect of grammaticality on accuracy ($\beta = 6.67$, $SE = 1.08$, $z = 6.17$, $p < 0.0001$). Control responses were not considered in any further analyses.

Turning to the ordering and cooccurrence categories, accuracy on the cooccurrence category was higher than the ordering category in the local condition, but this pattern

reversed in the nonlocal condition, where accuracy on the ordering category was higher than on the cooccurrence category. Accuracy on the ordering condition was higher in the nonlocal condition than in the local condition, while accuracy on the cooccurrence condition was relatively constant across conditions. To assess differences between the ordering and cooccurrence categories, and the effect of locality condition on these two categories, a second MLM was fit with predictors of category (ordering trials were coded as +0.5, cooccurrence as −0.5) and condition (local −0.5, nonlocal +0.5), as well as their interaction.[4] All predictors were centered, and a maximal random-effects structure was used with random slopes for all predictors by subject and random intercepts by subject and stimulus item. The model, shown in Table 4, showed a significant interaction between condition and category. The significant interaction was a crossover interaction: the direction of the effect of category reversed between the local and nonlocal conditions. A series of planned comparisons showed that accuracy on the ordering category is lower than on the cooccurrence category in the local condition ($\beta = -1.07$, $SE = 0.47$, $z = -2.29$, $p < 0.03$), but the pattern reversed in the nonlocal condition, where accuracy is higher on the ordering category than the cooccurrence category, though this difference was only marginally significant ($\beta = 0.66$, $SE = 0.37$, $z = 1.78$, $p = 0.08$). Locality condition had a significant effect on accuracy for the ordering category ($\beta = 1.48$, $SE = 0.32$, $z = 4.70$, $p < 0.0001$), but not for the cooccurrence category ($\beta = -0.28$, $SE = 0.51$, $z = -0.54$, $p = 0.59$).


FIGURE 1. Accuracy in the repetition task, by stimulus category and locality condition. Error bars indicate *SE*.

	β	*SE*	*z*	*p*
(intercept)	1.16	0.36	3.25	< 0.002
condition	0.66	0.32	2.05	< 0.05
category	−0.16	0.31	−0.50	0.62
condition : category	1.74	0.60	2.92	< 0.01

TABLE 4. Results of mixed logit model with binomial dependent variable of accuracy and predictors of locality condition and stimulus category.

ERRORS. The proportions of repair and nonrepair errors are shown in Figure 2. Repair errors for the cooccurrence restriction are C_2 de-ejectivization (e.g. *[k'ap'a] repeated as [k'apa]), and nonrepair errors are C_1 de-ejectivization (e.g. *[k'ap'a] repeated as

---

[4] A model with a single predictor testing for whether a stimulus having a neighbor had an effect on accuracy revealed no significant effect.

*[kap'a]). Repair errors for the ordering restriction were either C_2 de-ejectivization (e.g. *[kap'a] repeated as [kapa]) or movement (e.g. *[kap'a] repeated as [k'apa]), and nonrepair errors were doubling (e.g. *[kap'a] repeated as *[k'ap'a]). Differences in the rate of repair types on the ordering restriction are analyzed further below.


FIGURE 2. Proportion of total responses that were repairs (light gray) and nonrepairs (dark gray), for ordering and cooccurrence stimuli.

The rate of repair was analyzed for effects of condition and category. All errors were coded as repair/nonrepair, and a binomial model was fit with predictors of category (ordering trials were coded as +0.5, cooccurrence as −0.5) and condition (local −0.5, nonlocal +0.5), as well as their interaction. Both predictors were centered, and a maximal random-effects structure was used with random slopes for each predictor by subject and random intercepts by subject and item. The model, shown in Table 5, showed a significant effect of category, due to the greater proportion of repairs in the cooccurrence category than the ordering category. In the cooccurrence category, repairs accounted for 94% of errors in the local condition and 80% in the nonlocal, while in the ordering category repairs were just 52% of errors in the local condition and 50% in the nonlocal. Neither condition nor the condition-by-category interaction was significant.

	β	SE	z	p
(intercept)	1.03	0.66	1.56	0.12
condition	−0.58	0.67	−0.87	0.39
category	−3.68	0.78	−4.75	< 0.0001
condition : category	2.00	1.27	1.57	0.17

TABLE 5. Results of mixed logit model with binomial dependent variable of repair (for errorful trials only) and predictors of locality condition and stimulus category.

The rates of different error types in the ordering condition, shown in Figure 3, were further analyzed to determine if the greater overall accuracy for the ordering category in the nonlocal condition was due to a uniform decrease in all types of errors, or if only certain error types were affected by locality.

Three separate models were fit to assess the effect of locality on the three different errors types in the ordering category. For each model, participants that never made the relevant error were removed from analysis. The model for C_2 de-ejectivization included

FIGURE 3. Responses to stimuli violating the ordering restriction, by error type.

ten participants, the model for movement included eleven participants, and the model for doubling included seventeen participants. In each model, all utterances were coded for whether they exhibited the relevant error; for example, for the C_2 de-ejectivization model, all utterances that showed C_2 de-ejectivization were coded as '1', and all other utterances, whether they were correct or showed some other error, were coded as '0'. Each model included a single centered predictor of condition, comparing the local and nonlocal conditions and random slopes by subject and random intercepts by subject and item. There was no effect of locality for C_2 de-ejectivization ($\beta = -0.46$, $SE = 0.70$, $z = -0.66$, $p = 0.51$). Both movement and doubling were less common in the nonlocal condition than in the local condition (movement: $\beta = -1.40$, $SE = 0.68$, $z = -2.08$, $p < 0.04$; doubling: $\beta = -1.07$, $SE = 0.27$, $z = -3.97$, $p < 0.0001$); adjusting for multiple comparisons, this difference is significant only for the doubling errors.

4.6. DISCUSSION. The repetition task found a strong effect of phonotactic legality, with accurate repetition much more likely for phonotactically legal control roots than phonotactically illegal roots in the ordering and cooccurrence categories, consistent with previous work showing that speakers are sensitive to phonotactic distributions in experimental settings (Frisch & Zawaydeh 2001, Rose & King 2007, Gallagher 2013a, 2014, Coetzee 2014).

Of main interest are the two clear differences between the ordering and cooccurrence categories. First, accuracy on the ordering category is affected by locality, while accuracy on the cooccurrence category is not. Second, errors on the cooccurrence category are overwhelmingly repairs, while errors on the ordering category are evenly split between repairs and nonrepairs.[5]

---

[5] To check whether the prevalence of doubling errors on ordering-category forms was a task effect, due to interference from the cooccurrence-category forms, a follow-up experiment was run. The follow-up presented the forms from the filler, control, and ordering categories from both the local and nonlocal conditions and was completed by twelve participants. While the overall error rate was lower in the follow-up than in the original experiment, errors were still evenly split between repairs and nonrepairs (10% repairs vs. 13% nonrepairs for local items, 2% repairs vs. 2.5% nonrepairs for nonlocal items). These results show that the presence of cooccurrence-violating forms has an effect on the task: more errors are triggered when the phonotactic structure of the stimuli is more varied. The distribution of errors, however, does not change. Neither the high rate of doubling errors nor the effect of locality can be attributed to interference from cooccurrence-category targets.

The differences between the ordering and cooccurrence categories suggest that these restrictions are represented quite differently by speakers. Accuracy on the ordering category is sensitive to distance, and errors are not consistently phonotactic repairs, suggesting that this restriction is partially synchronically rooted in perceptual or articulatory factors. In contrast, the results for the cooccurrence category suggest that this restriction is represented as a phonotactic constraint in the traditional sense, as errors on cooccurrence-category roots are unaffected by locality and are overwhelmingly phonotactic repairs.

It may seem that the effect of locality on the ordering category could also be consistent with a higher-level phonotactic constraint. Many restrictions on consonant combinations in other languages show effects of locality, with consonants in adjacent syllables more restricted than those in nonadjacent syllables (Greenberg 1960, McCarthy 1994, Frisch et al. 2004, Rose & King 2007, Coetzee & Pater 2008). At first blush, then, Quechua may be showing a similar pattern: the ordering restriction is a phonotactic constraint that applies more strongly between consonants in adjacent positions than nonadjacent positions. The lack of a locality effect for the cooccurrence category makes this explanation unsatisfactory, however. Cochabamba Quechua speakers have no direct evidence that phonotactic restrictions weaken with distance, since there are so few trisyllabic roots, and if distance sensitivity were the default interpretation of a phonotactic restriction (cf. Zymet 2014), we would expect to see comparable effects for both restrictions. The presence of a locality effect for only the ordering restriction suggests that this effect is not due to the phonotactic grammar.

The local condition of the current experiment is largely a replication of Gallagher 2013a. The structure of the stimuli in the control, ordering, and cooccurrence categories is the same between the two experiments, though the results differ somewhat. Gallagher 2013a found a significant difference in accuracy between the ordering and cooccurrence categories, with accuracy higher on ordering than cooccurrence stimuli, while the current experiment found the opposite effect: accuracy is higher on the ordering category than the cooccurrence category in the local condition. The most likely explanation for this discrepancy is that the number of participants in the current experiment is more than twice that in Gallagher 2013a: nineteen vs. eight. While the details of the methods, analysis, and results in Gallagher 2013a do differ from those here, the main interpretation of the results is the same between the two experiments: there are substantial, meaningful differences in how ordering- and cooccurrence-violating nonce roots are treated by Cochabamba Quechua speakers. The content of these differences is explored in more detail in the current work, through a greater number of participants as well as the comparison between the local and nonlocal conditions and the discrimination experiments, to which we now turn.

5. Experiment 2: discrimination. The discrimination task has two goals. First, it serves to further diagnose the status of the ordering and cooccurrence restrictions for Cochabamba Quechua speakers, as phonotactic restrictions have been shown to influence perception (Werker & Tees 1984, Berent & Shimron 1997, Hallé et al. 1998, Pitt 1998, Dupoux et al. 1999, Moreton 2002, Berent et al. 2007, Hallé & Best 2007, Coetzee 2008). If the ordering and cooccurrence restrictions are represented as abstract phonotactic restrictions, then a top-down effect of the phonotactic grammar on perception is expected, with phonotactically illegal forms being highly confusable with phonotactically legal forms. Second, it serves to assess the perceptual confusability of targets and their various repairs from the repetition task, allowing an assessment of the contribution of misperception to repetition errors. While phonotactic restrictions have an impact on perception, not all phonotactically illegal forms are misperceived equally

often (Davidson 2006, 2010, Kabak & Idsardi 2007, Davidson & Shaw 2012). Consequently, some of the errors on the repetition task may be more likely to arise from misperception than others.

Two discrimination tasks were run with different sets of participants. Experiment 2a compared participants' accuracy at discriminating ordering- and cooccurrence-category targets relative to repairs with C_2 de-ejectivization (e.g. [kup'i]-[kupi] and [k'up'i]-[k'upi]), compared to controls (e.g. [lap'a]-[lapa]). Experiment 2b looks at the other attested mismatches between target and repetition: doubling/C_1 de-ejectivization (e.g. [kup'i]-[k'up'i]) and movement (e.g. [kup'i]-[k'upi]).

5.1. EXPERIMENT 2A: C_2 DE-EJECTIVIZATION.

PARTICIPANTS. Twenty-two native speakers of Cochabamba Quechua participated in the experiment. The participants were sixteen students at the Quechua Indigenous University of Bolivia (Universidad Indígena Boliviana Quechua) in Chimore, a town in the Chapare region of Cochabamba Department, Bolivia. An additional six participants were recruited in the city of Cochabamba. All participants were in their twenties and thirties and were bilingual in Spanish. None of the participants in the discrimination task had completed the repetition task. The data from one participant, who gave the same response to every trial but one, were removed from analysis.

STIMULI. The stimuli were items from the control, cooccurrence, and ordering categories from both the local and nonlocal conditions in the repetition study, paired with corresponding items with a plain stop in the rightmost position. The stimulus pairs all contrasted an ejective and a plain stop in the final consonantal position. In the control category, the contrast was between two phonotactically legal structures, whereas in the cooccurrence and ordering categories, the contrast was between a phonotactically illegal structure and a legal one. The contrasting pairs are shown for the local condition in Table 6 and the nonlocal condition in Table 7.[6]

CONTROL	ORDERING	COOCCURRENCE
lap'a – lapa	tip'i – tipi	k'it'a – k'ita
jup'a – jupa	kut'a – kuta	p'it'a – p'ita
juk'u – juku	pit'a – pita	q'ap'i – q'api
juk'a – juka	kip'u – kipu	k'ip'a – k'ipa
siq'a – siqa	qat'i – qati	q'at'a – q'ata
saq'u – saqu	tuk'i – tuki	p'ik'a – p'ika
muq'u – muqu	tuq'i – tuqi	t'uq'i – t'uqi
lip'u – lipu	kup'a – kupa	t'ap'u – t'apu
nap'u – napu	kap'a – kapa	k'ip'i – k'ipi
nat'u – natu	tip'u – tipu	q'up'i – q'upi
nut'a – nuta	kip'a – kipa	k'ap'u – k'apu
jap'i – japi	kap'i – kapi	k'ut'a – k'uta
lik'i – liki	tip'a – tipa	k'up'i – k'upi
luk'a – luka	puk'i – puki	p'uk'a – p'uka
maq'u – maqu	taq'i – taqi	t'aq'u – t'aqu

TABLE 6. Stimulus items for the local condition in experiment 2a.

[6] Because some of the phonotactically illegal disyllabic stimuli from the repetition experiment had existing lexical neighbors, some of the stimulus pairs here contrast not just a legal and illegal form, but an existing legal form and a nonexistent illegal form. Post-hoc comparison revealed no differences in accuracy between trials that included an existing item and those that did not. While perhaps surprising, this effect is likely due to the high frequency of nonce words in the task focusing participants' attention away from the lexicon. Another possible source of the (non)effect is that roots in Cochabamba Quechua rarely appear in unsuffixed form and are perhaps less immediately recognizable in isolation.

CONTROL	ORDERING	COOCCURRENCE
nusap'i – nusapi	kuwip'a – kuwipa	q'usap'i – q'usapi
jarup'a – jarupa	tisap'i – tisapi	p'irut'u – p'irutu
juwat'u – juwatu	karap'i – karapi	k'urit'a – k'urita
lumik'i – lumiki	piɲuk'i – piɲuki	q'awit'a – q'awita
siɲiq'a – siɲiqa	tamuk'i – tamuki	p'iɲuk'i – p'iɲuki
musaq'i – musaqi	paɲuk'i – paɲuki	t'imik'u – t'imiku
mawiq'i – mawiqi	tasaq'i – tasaqi	t'awip'u – t'awipu
luwap'i – luwapi	kamup'a – kamupa	k'imap'i – k'imapi
limip'u – limipu	tisap'u – tisapu	k'isit'a – k'isita
nasit'a – nasita	kusut'a – kusuta	q'arip'i – q'aripi
juʎap'u – juʎapu	kurit'u – kuritu	k'iʎup'a – k'iʎupa
jawit'i – jawiti	kiʎup'u – kiʎupu	t'inak'i – t'inaki
jaʎuk'a – jaʎuka	qawit'i – qawiti	p'uʎuk'a – p'uʎuka
liɲak'i – liɲaki	punuk'i – punuki	p'aɲaq'i – p'aɲaqi
sawiq'u – sawiqu	tuwiq'u – tuwiqu	t'uwiq'u – t'uwiqu

TABLE 7. Stimulus items for the nonlocal condition in experiment 2a.

Each stimulus pair was used to create four ABX trials, fully balancing which stimulus appeared as A, B, and X. For example, the stimulus pair [lap'a]-[lapa] appeared in the four trials *lap'a-lapa-lapa*, *lap'a-lapa-lap'a*, *lapa-lap'a-lapa*, and *lapa-lap'a-lap'a*. The X stimuli were repetitions of the same A or B token; they were not distinct tokens of the same stimulus.

PROCEDURE. The experiment was conducted in a quiet room. As in the repetition study, the local and nonlocal conditions were presented as separate experiments, with a break in between. All participants completed both conditions, with the order in which they completed the experiments balanced. Each participant heard all of the stimuli in a different random order. The three stimulus items in each trial were presented with a 300 ms interstimulus interval. After the third stimulus was heard, a green line of asterisks appeared on the screen to indicate the response period. During the response period, participants had two seconds to respond; if they did not respond within this period, they automatically went on to the next trial.

Participants were told that the words they would hear had no meaning in Quechua or any other language, but would contain sounds familiar from existing Quechua words. They were told that the first two words they heard would always be different from one another, and that the third would match either the first or the second. If the third item matched the first item, they were instructed to press the '1' key, and if the third item matched the second, they were instructed to press the '0' key (marked with a sticker saying '2'). The instructions were presented in Quechua orthography on the screen, but as the majority of participants are not comfortable reading in Quechua, the instructions were given orally by either the author or a native Quechua-speaking research assistant. The experiment took about ten minutes to complete, and participants were compensated the equivalent of about $7.

RESULTS. For each participant, category (control, ordering, or cooccurrence), and condition (local or nonlocal), a d-prime score was calculated (Macmillan & Creelman 2005, Boley & Lester 2009).[7] The results, given in Figure 4, show that perception of an

[7] D-prime is calculated by taking the *z*-score of the proportion of hits and false alarms, and then subtracting false alarms from hits: $d' = z(H) - z(F)$. In the ABX task, a 'hit' was defined as a correct 'A' response and a 'false alarm' as an incorrect 'A' response. In other words, accuracy on ABA trials was relativized to errors on ABB trials. Proportions with the values of 0 and 1 were converted to 0.01 and 0.99, respectively, yielding a maximum d-prime score of 4.5 (Macmillan & Creelman 2005). The use of d-prime means that the statistical

ejective-plain contrast in medial position was most accurate in the control condition, where both an ejective and a plain stop are phonotactically legal, followed by the ordering category and then the cooccurrence category. The pattern was the same in both the local and nonlocal conditions. To test for an effect of grammaticality, a LINEAR MIXED MODEL (LMM) was fit with d-prime as the dependent variable and a single predictor comparing the control category (coded as +0.5) and the ordering and cooccurrence categories (coded as −0.25). The model had random intercepts and slopes by subject. There is a significant effect of grammaticality on d-prime ($\beta = 0.98$, $SE = 0.22$, $t = 4.53$), as indicated by a t-value of greater than ±2 (Gelman & Hill 2006).


FIGURE 4. D-prime by category and condition for medial ejective contrasts in experiment 2a.
Error bars indicate *SE*.

To assess differences between the ordering and cooccurrence categories, and the effect of locality condition on these two categories, a second LMM was fit with predictors of category (ordering trials were coded as +0.5, cooccurrence as −0.5) and condition (local −0.5, nonlocal +0.5), as well as their interaction. The model had random intercepts and slopes by subject. D-prime on the ordering category was significantly higher than on the cooccurrence category. There was no effect of locality condition, nor was the interaction between category and condition significant. The model is given in Table 8.

	β	*SE*	*t*
(intercept)	2.17	0.22	9.81
category	0.82	0.19	4.24
condition	0.22	0.17	1.24
category : condition	−0.05	0.26	−0.18

TABLE 8. Results of linear mixed model with dependent variable of d-prime and predictors of category (ordering or cooccurrence) and locality condition (local or nonlocal) for experiment 2a.

analyses include only a single value for each category and condition combination for each participant. While each participant responded to 240 trials, these were reduced to six d-prime scores.

DISCUSSION. Experiment 2a found that a contrast in ejection in C_2 was most perceptible in control pairs (e.g. [lap'a]-[lapa]), where both an ejective and a plain stop are phonotactically legal. A medial ejective contrast was less perceptible in ordering-category pairs (e.g. [tip'i]-[tipi]) and was less perceptible still in cooccurrence-category pairs (e.g. [k'ip'a]-[k'ipa]). While a medial ejective is phonotactically illegal in both the cooccurrence- and ordering-category pairs, the contrast was still perceived better in ordering-category pairs than in cooccurrence-category pairs. These perception results are consistent with the rates of C_2 de-ejectivization in the repetition task: the more confusable a contrast in ejection in C_2 in the discrimination task, the more frequent C_2 de-ejectivization errors in the repetition task. In the repetition task, C_2 de-ejectivization errors were most common for the cooccurrence category (local: 29.75%, nonlocal: 20.6%), followed by the ordering category (local: 6.8%, nonlocal: 5%), and were least common in the control category (local: 0.7%, nonlocal: 2.8%). The discrimination task did not reveal any effect of locality on the perceptibility of C_2 contrasts, a finding that is also consistent with the repetition task. Rates of C_2 de-ejectivization in the repetition task did not differ significantly between the local and nonlocal conditions.

The results of the discrimination task are consistent not only with the repetition task, but also with the discrimination results for English speakers reported in Gallagher 2010a,b. English speakers, like Quechua speakers, find pairs contrasting a single ejective with no ejectives (e.g. [kip'a]-[kipa]) easier to discriminate than pairs contrasting two ejectives and one ejective (e.g. [k'ip'a]-[k'ipa]), though for English speakers there is presumably no phonotactic restriction underlying this effect. Quechua speakers' difficulty with the cooccurrence category, then, may derive both from the presence of a phonotactic restriction in the language and from a language-independent perceptual or processing difficulty with this particular contrast.

5.2. EXPERIMENT 2B: MOVEMENT AND DOUBLING.

PARTICIPANTS. Twenty native speakers of Cochabamba Quechua participated in experiment 2b, all recruited in the city of Cochabamba, Bolivia. All participants were in their twenties and thirties, from the Cochabamba area, and were students in the applied linguistics program at the Universidad Mayor de San Simón in Cochabamba. All participants were bilingual in Spanish, and none had participated in experiment 1 or experiment 2a.

STIMULI. The stimuli for experiment 2b compared forms that differed in the position of ejection or number of ejectives. In the 'movement' category, ordering-category items from the repetition experiment were paired with items with an ejective in C_1, for example, [kip'u]-[k'ipu], which is the change made in movement errors in the repetition study. In the 'doubling/C_1 de-ejectivization' category, ordering category items from the repetition study were paired with forms with two ejectives, for example, [kip'u]-[k'ip'u]. A change between these two forms is found in the repetition task in C_1 de-ejectivization errors on the cooccurrence category (target *[k'ip'u] repeated as *[kip'u]) and doubling errors on the ordering category (target *[kip'u] repeated as *[k'ip'u]). Stimulus pairs in the movement category contrast a phonotactically legal and an illegal form, while pairs in the doubling/C_1 de-ejectivization category contrast two phonotactically illegal forms. The same control items from experiment 2a were also included, contrasting an ejective and a plain stop in C_2 in forms where they are both phonotactically legal. The contrasting pairs are shown for the local condition in Table 9 and the nonlocal condition in Table 10.

Each stimulus pair was used to create four ABX trials, fully balancing which stimulus appeared as A, B, and X, as for experiment 2a.

CONTROL	MOVEMENT	DOUBLING/ C_1 DE-EJECTIVIZATION
lap'a – lapa	tip'i – t'ipi	tip'i – t'ip'i
jup'a – jupa	kut'a – k'uta	kut'a – k'ut'a
juk'u – juku	pit'a – p'ita	pit'a – p'it'a
juk'a – juka	kip'u – k'ipu	kip'u – k'ip'u
siq'a – siqa	qat'i – q'ati	qat'i – q'at'i
saq'u – saqu	tuk'i – t'uki	tuk'i – t'uk'i
muq'u – muqu	tuq'i – t'uqi	tuq'i – t'uq'i
lip'u – lipu	kup'a – k'upa	kup'a – k'up'a
nap'u – napu	kap'a – k'apa	kap'a – k'ap'a
nat'u – natu	tip'u – t'ipu	tip'u – t'ip'u
nut'a – nuta	kip'a – k'ipa	kip'a – k'ip'a
jap'i – japi	kap'i – k'api	kap'i – k'ap'i
lik'i – liki	tip'a – t'ipa	tip'a – t'ip'a
luk'a – luka	puk'i – p'uki	puk'i – p'uk'i
maq'u – maqu	taq'i – t'aqi	taq'i – t'aq'i

TABLE 9. Stimulus pairs for the local condition in experiment 2b.

CONTROL	MOVEMENT	DOUBLING/ C_1 DE-EJECTIVIZATION
nusap'i – nusapi	kuwip'a – k'uwipa	kuwip'a – k'uwip'a
jarup'a – jarupa	tisap'i – t'isapi	tisap'i – t'isap'i
juwat'u – juwatu	karap'i – k'arapi	karap'i – k'arap'i
lumik'i – lumiki	piɲuk'i – p'iɲuki	piɲuk'i – p'iɲuk'i
siɲiq'a – siɲiqa	tamuk'i – t'amuki	tamuk'i – t'amuk'i
musaq'i – musaqi	paɲuk'i – p'aɲuki	paɲuk'i – p'aɲuk'i
mawiq'i – mawiqi	tasaq'i – t'asaqi	tasaq'i – t'asaq'i
luwap'i – luwapi	kamup'a – k'amupa	kamup'a – k'amup'a
limip'u – limipu	tisap'u – t'isapu	tisap'u – t'isap'u
nasit'a – nasita	kusut'a – k'usuta	kusut'a – k'usut'a
juʎap'u – juʎapu	kurit'u – k'uritu	kurit'u – k'urit'u
jawit'i – jawiti	kiʎup'u – k'iʎupu	kiʎup'u – k'iʎup'u
jaʎuk'a – jaʎuka	qawit'i – q'awiti	qawit'i – q'awit'i
liɲak'i – liɲaki	punuk'i – p'unuki	punuk'i – p'unuk'i
sawiq'u – sawiqu	tuwiq'u – t'uwiqu	tuwiq'u – t'uwiq'u

TABLE 10. Stimulus pairs for the nonlocal condition in experiment 2b.

PROCEDURE. The procedure was the same as for experiment 2a.

RESULTS. For each participant, category (control, movement, or doubling/C_1 de-ejectivization), and condition (local or nonlocal), a d-prime score was calculated. The results, given in Figure 5, show that contrasts between phonotactically illegal ordering-category and cooccurrence-category forms in the doubling/C_1 de-ejectivization category are most confusable. Contrasts in the position of ejection in the movement category, which contrasts phonotactically legal and illegal forms, are somewhat more perceptible. D-prime is highest in the control category, where both forms are phonotactically legal. To confirm the effect of grammaticality, an LMM was fit with d-prime as the dependent variable and a single predictor comparing the control category (coded as +0.5) with the movement and doubling/C_1 de-ejectivization categories (coded as −0.25). The model had random intercepts and slopes by subject. There is a significant effect of grammaticality on d-prime ($\beta = 6.67$, $SE = 0.20$, $t = 6.33$).

To assess differences between the movement and doubling/C_1 de-ejectivization categories, and the effect of locality condition on these two categories, a second LMM was fit with predictors of category (movement trials were coded as +0.5, doubling/C_1 de-ejectivization as −0.5) and condition (local −0.5, nonlocal +0.5), as well as their inter-

FIGURE 5. D-prime by category and condition in experiment 2b. Error bars indicate *SE*.

action. The model had random intercepts and slopes by subject. D-prime on the movement category was significantly higher than on the doubling/C_1 de-ejectivization category. There was no effect of locality condition, nor was the interaction between category and condition significant. The model is given in Table 11.

	β	*SE*	*t*
(intercept)	1.77	0.19	9.26
category	−0.35	0.07	−5.11
condition	0.29	0.18	1.56
category : condition	0.14	0.10	1.40

TABLE 11. Results of linear mixed model with dependent variable of d-prime and predictors of category (movement or doubling/C_1 de-ejectivization) and locality condition (local or nonlocal) for experiment 2b.

DISCUSSION. Experiment 2b found a significant effect of grammaticality, as well as a significant difference in the perceptibility of contrasts in the position of ejection (movement category) and the number of ejectives (doubling/C_1 de-ejectivization category). The low rate of discrimination on the doubling/C_1 de-ejectivization category likely arises because this category compares two unattested structures, while the movement category compares an unattested structure to an attested one. There was no effect of locality.

These results are only partially consistent with the results of the repetition task. The main result of interest in the repetition task was that both movement and doubling errors were less frequent in the nonlocal condition than the local condition, resulting in an increase in accuracy on roots violating the ordering restriction. Experiment 2b suggests that this decrease in movement and doubling errors does not have a perceptual source. Movement errors involve switching the position of ejection, and this error was more common in the local condition of the repetition experiment than the nonlocal condition, but a contrast in position of ejection was not more confusable in the local condition of the discrimination task than the nonlocal condition; for example, target *[kap'a] was repeated as [k'apa] more often than target *[kamup'a] was repeated as [k'amupa], but the pair [kap'a]-[k'apa] was not more confusable than the pair [kamup'a]-[k'amupa]. Similarly, doubling errors were more common in the local condition of the repetition experiment than the nonlocal condition, but a contrast in the number of ejectives was equally perceptible across the two conditions in the discrimination task; for example, target

*[kap'a] was repeated as *[k'ap'a] more often than target *[kamup'a] was repeated as *[k'amup'a], but the pair [kap'a]-[k'ap'a] was not more confusable than the pair [kamup'a]-[k'amup'a].

The significant difference in discriminability of the movement category and the doubling/C_1 de-ejectivization category is also not wholly consistent with the repetition task. While movement errors are less frequent than doubling errors in the nonlocal condition of the repetition task (movement: 4%, doubling: 10.8%), movement and doubling errors were equally common in the local condition of the repetition task (movement: 17.14%, doubling: 17.86%). In the discrimination task, discriminability in the movement category is much higher than in the doubling/C_1 de-ejectivization category in both the local and nonlocal conditions.

Finally, the high rate of confusability on the doubling/C_1 de-ejectivization category in the perception task is consistent with the relatively high rate of doubling errors on the ordering category in the repetition task (local: 17.86%, nonlocal: 10.8%), but not with the relatively low rate of C_1 de-ejectivization errors on the cooccurrence category (local: 2.2%, nonlocal: 6.4%). This mismatch again suggests that errors on the repetition task are not solely due to misperception.

6. INTERIM SUMMARY. The repetition task revealed two substantial differences between the ordering and cooccurrence restrictions. First, accuracy on repeating forms that violate the ordering restriction is higher in the nonlocal condition than the local condition, due to the decrease in movement and doubling errors, while accuracy on forms violating the cooccurrence restriction is the same across locality conditions. Second, errors on cooccurrence-violating forms are almost entirely phonotactic repairs, while errors on ordering-violating forms are evenly split between repairs and nonrepairs. The regular repairs and insensitivity to distance are consistent with the errors on cooccurrence-violating forms being driven by a high-level phonotactic constraint. In contrast, the prevalence of nonrepair errors and the sensitivity to distance for ordering-violating forms suggests a strong role for forces outside the phonotactic grammar.

The discrimination tasks revealed that while some errors on the repetition task could be due to misperception, the major results of the repetition task cannot be attributed to misperception alone. Specifically, neither the effect of locality on ordering-violating forms nor the difference in rate of repairs on cooccurrence- and ordering-violating forms is predicted from the discrimination tasks. For the cooccurrence restriction, the discrimination and repetition results are comparable, further supporting a high-level phonotactic restriction that has a top-down influence on performance on these tasks. The mismatch between the discrimination and repetition tasks for forms violating the ordering restriction, however, suggests a substantial role for articulatory factors and perception-production mapping in the treatment of these forms on the repetition task.

In the following section, I present an analysis of these results that relies on two main claims. First, the ordering and cooccurrence restrictions are enforced by the grammar via two distinct constraints: one general constraint that rules out all stop-ejective sequences, and an additional, higher-weighted constraint that specifically penalizes ejective-ejective pairs. These constraints are also reflected in the results of the discrimination tasks. Second, errors on the repetition task are generated by the interaction of the phonotactic grammar with (i) task-specific faithfulness constraints and (ii) speech errors resulting from blending competing speech plans.

The structure of the analysis is as follows. The phonotactic grammar is responsible for driving repairs, which are more common for forms violating the cooccurrence re-

striction than for forms violating the ordering restriction, suggesting a stronger phonotactic dispreference for cooccurrence-violating forms. Additionally, the phonotactic grammar is unaffected by locality: pairs of ejectives and plain stop-ejective sequences are penalized equally, regardless of the amount of material intervening between the interacting consonants.

Task-specific faithfulness interacts with the phonotactic grammar, overriding the grammar in the case of accurate repetition and choosing between possible repairs in the case of errorful repetition. The effect of locality for movement errors on ordering-category forms is driven by faithfulness constraints that penalize reassociation of a feature over longer distances more than reassociation over shorter distances. Doubling errors are then derived from movement errors as speech errors that blend the speech plan for a movement error with the speech plan for an accurate repetition; for example, [k'apa] competes with [kap'a] to yield [k'ap'a].

7. ANALYSIS.

7.1. THE GRAMMAR. At the core of the proposed analysis are two phonotactic constraints that overlap in the forms they penalize: a broader constraint against stops followed by ejectives and a more specific constraint penalizing pairs of ejectives. These two constraints are formalized in 2.

> (2) a. *[cg][cg]$_{RT}$: Assign one violation mark for each pair of [constricted glottis] specifications within a root.
>   b. *[−cont, −son][cg]$_{RT}$: Assign one violation mark for each [−continuant, −sonorant] specification that is followed by a [constricted glottis] specification within a root.

The overlapping formulation of these markedness constraints means that violations of the cooccurrence restriction, which is enforced by both constraints, will always be penalized by the grammar more than violations of the ordering restriction, which is enforced by only one constraint. This is consistent with the greater rate of repair found for cooccurrence-violating roots in the repetition task. The markedness constraints in 2 reference the domain of the root, and thus they penalize local and nonlocal violations of either the cooccurrence or ordering restriction equally. This allows for a straightforward account of the lack of a locality effect for cooccurrence-violating forms, which were repaired equally often in the local and nonlocal conditions of the repetition experiment. The effect of locality on forms violating the ordering restriction is related to the choice of repair, which is regulated by task-specific faithfulness constraints presented below.

REPAIR ERRORS AND THE PHONOTACTIC GRAMMAR. For forms violating the ordering restriction, both the rate and choice of repair on the repetition task were variable. To account for the choice of repair, faithfulness constraints are needed to determine how a given form will be modified to conform to the phonotactic constraints in 2. To account for the variability in rate of repair, constraints must be weighted, implemented here in a MAXIMUM ENTROPY grammar (Goldwater & Johnson 2003, Hayes & Wilson 2008). Maximum entropy grammars, like NOISY HARMONIC GRAMMAR (Coetzee 2009, Coetzee & Pater 2011) and STOCHASTIC OPTIMALITY THEORY (Boersma 1997, Boersma & Hayes 2001), are probabilistic and can model variable outputs.

Two faithfulness constraints, given in 3, are active: one penalizes de-ejectivization, and the other penalizes reassociation of ejection from one segment to another. Both constraints evaluate autosegmental representations. The constraint against reassociation, NOFLOP, is distance-sensitive, penalizing reassociation over longer distances more than reassociation over shorter distances. For the present analysis, the relevant distance

d for this constraint is a vowel. In the disyllabic forms in the local condition, ejection re-associates over a single vowel ([kip'u] → [k'ipu]), violating NoFLOP once, and in the trisyllabic forms in the nonlocal condition, ejection reassociates over two vowels ([kamip'u] → [k'amipu]), violating NoFLOP twice. The distance sensitivity of NoFLOP may be rooted in articulatory pressures, such that smaller gestural reorganizations are penalized less than larger reorganizations. This would be consistent with the effect of distance being found in the repetition task, but not in the perception task.

    (3) a. *MAX[cg]: Assign one violation mark for any [constricted glottis] specification present in the input that is not present in the output.

         b. NoFLOP[cg]: Given a [constricted glottis] specification that is uniquely associated to a segment S_1 in the input and uniquely associated to a distinct segment S_2 in the output, assign one violation mark for every unit of distance *d* that separates S_1 and S_2.

The relative weights of these constraints needed to generate faithful and repair mappings on the repetition task are shown in Table 12, where *w*C indicates the weight *w* of a constraint C.

TARGET	OUTPUT	WEIGHTING
COOCCURRENCE		
local:	k'ap'a	$w\text{MAX[cg]} > w\text{*[cg][cg]}_{RT} + w\text{*[−cont, −son][cg]}_{RT}$
k'ap'a		
	k'apa	$w\text{*[cg][cg]}_{RT} > w\text{MAX[cg]}$
		or $w\text{*[−cont, −son][cg]}_{RT} > w\text{MAX[cg]}$
nonlocal:	k'amip'a	$w\text{MAX[cg]} > w\text{*[cg][cg]}_{RT} + w\text{*[−cont, −son][cg]}_{RT}$
k'amip'a		
	k'amipa	$w\text{*[cg][cg]}_{RT} > w\text{MAX[cg]}$
		or $w\text{*[−cont, −son][cg]}_{RT} > w\text{MAX[cg]}$
ORDERING		
local:	kap'a	$w\text{NoFLOP[cg]}, w\text{MAX[cg]} > w\text{*[−cont, −son][cg]}_{RT}$
kap'a		
	k'apa	$w\text{*[−cont, −son][cg]}_{RT}, w\text{MAX[cg]} > w\text{NoFLOP[cg]}$
	kapa	$w\text{*[−cont, −son][cg]}_{RT}, w\text{NoFLOP[cg]} > w\text{MAX[cg]}$
nonlocal:	kamip'a	$2w\text{NoFLOP[cg]}, w\text{MAX[cg]} > w\text{*[−cont, −son][cg]}_{RT}$
kamip'a		
	k'amipa	$w\text{*[−cont, −son][cg]}_{RT}, w\text{MAX[cg]} > 2w\text{NoFLOP[cg]}$
	kamipa	$w\text{*[−cont, −son][cg]}_{RT}, 2w\text{NoFLOP[cg]} > w\text{MAX[cg]}$

TABLE 12. Relative weightings of the constraints in the grammar to generate faithful and repair outputs in the repetition task.

    In the repetition task, faithful mappings of targets that violate either restriction are common, ranging from 55% to 81% of responses. This accuracy rate is high from the perspective of phonotactic well-formedness and is hypothesized to be a task effect. Participants are aware that the repetition task involves nonsense words that may not sound perfectly natural, but that their task is to focus on the details of what they hear and be as precise as possible in their repetition. The frequency of faithful mappings is hypothesized to reflect an increased importance given to faithfulness constraints in the context of this task, as compared to the task-independent phonotactic grammar; for an implementation of how nongrammatical factors may perturb faithfulness constraints, see

Coetzee & Kawahara 2013. In general, the phonotactic grammar is assumed to weight markedness constraints higher than faithfulness constraints, as in 4, favoring repairs.

(4) w*[cg][cg]$_{RT}$, w*[−cont, −son][cg]$_{RT}$ > wMax[cg], 2wNoFlop[cg]

The greater weight of faithfulness in the repetition task allows phonotactic violations to surface, in turn revealing the relative penalties for violating the cooccurrence and ordering restrictions. Because the cooccurrence restriction violates two constraints while the ordering restriction violates just one, a greater increase in faithfulness is necessary to allow cooccurrence-restriction violations to surface faithfully. In more specific terms, for cooccurrence-restriction violations to surface faithfully, wMax[cg] must be greater than the sum of w*[cg][cg]$_{RT}$ and w*[−cont, −son][cg]$_{RT}$. For an ordering-restriction violation to surface faithfully, however, either wMax[cg] or wNoFlop[cg] must be greater than w*[−cont, −son][cg]$_{RT}$ alone. Task-specific increased faithfulness then results in fewer repairs to ordering-restriction violations than to cooccurrence-restriction violations.

Nonrepair doubling errors and gestural blending. A main challenge for formalizing an analysis of repetition errors is to explain why doubling errors (e.g. target *[kap'a] repeated as *[k'ap'a]) are possible. The similar patterning of doubling and movement errors suggests a connection between these two error types. Both error types are sensitive to distance, being more frequent in the local condition than in the nonlocal condition, and these error types occur at similar rates (21% and 16% in the local condition, 11% and 5% in the nonlocal condition).

The hypothesis here is that doubling errors are speech errors caused by the blending of two competing articulatory plans. Given a target that violates the ordering restriction, the grammar often favors realignment of ejection to the initial stop, but task-specific faithfulness favors a faithful repetition. A speech plan with realignment of ejection, [k'apa], is then in competition with the faithful speech plan, [kap'a]. While on many trials only one of these two speech plans is executed, doubling errors arise when these two plans are blended.

If doubling errors are the result of gestural blending between a faithful repetition of an ordering-violating form and a movement repair, then blending of the other attested repair, C_2 de-ejectivization, with faithful repetitions of both cooccurrence- and ordering-violating forms is also expected. In the case of C_2 de-ejectivization, however, the blending of the repair with the faithful repetition is not distinct from the faithful repetition itself. Blending of a target violating the ordering restriction, *[kap'a], with a C_2 de-ejectivization repair, [kapa], would yield [kap'a], which would be coded as an accurate repetition. Similarly, blending of a target violating the cooccurrence restriction, *[k'ap'a], with a C_2 de-ejectivization repair, [k'apa], would yield *[k'ap'a], also coded as an accurate repetition. Blending of targets and repairs is thus consistent with the observed repetition patterns, though it is directly observable only in the case of a target violating the ordering restriction and a movement repair.

Support for this line of analysis comes from previous work that has observed blending errors in production tasks (Frisch & Wright 2002, Pouplier & Goldstein 2005, Goldstein et al. 2007, Pouplier 2008). Examining the acoustic properties of tongue-twister data involving English [s] and [z], Frisch and Wright (2002) found evidence of errors intermediate between voiced and voiceless. Such tokens may have a percept of an error, but maintain some acoustic properties of the intended segment; for example, some [s] → [z] errors have less voicing than a faithful [z]. Pouplier & Goldstein 2005, Goldstein et al. 2007, and Pouplier 2008 have found articulatory evidence that some speech errors arise from gestural blending that results in a phonotactically illegal structure, as is seen in doubling errors. In some cases, the intrusive gesture may be reduced in magnitude.

Such errors particularly arise in tasks that involve participants repeating similar yet nonidentical structures. For example, when asked to repeat the phrase *top cop*, where the two words differ only in the place of articulation of the onset consonant, articulatory errors are observed where the alveolar and velar closure gestures are produced simultaneously. The fixed prosodic structure of target items in the repetition task, as well as the frequency of ejectives, may provide the structural similarity that makes blending particularly likely.

If doubling errors on the repetition task result from the blending of competing speech plans, then the resulting form should bear some articulatory or acoustic artifacts of this blending. In particular, reduction of one or both ejectives may be expected. The current data are insufficient to conclusively determine whether such reduction occurs, because doubling errors are not evenly distributed over participants, and place of articulation is not evenly distributed over stimulus items. Some exploratory observations are possible, though, as follows.

The three participants with the highest rates of doubling errors in the local condition are participant 1 (twelve doubling errors), participant 6 (seven doubling errors), and participant 7 (ten doubling errors), and the most frequent place of articulation for a medial stop in an ordering-violating form was labial. For these three participants, then, the VOT of [p'] in C_2 of a doubling error was measured and compared to the VOT for [p'] in C_2 of control forms and accurate repetitions of cooccurrence forms.[8] Table 13 gives the average VOT and standard deviation for C_2 [p'] for each participant. While the standard deviations are all quite large, average VOT is indeed longer in control and cooccurrence forms than in doubling errors.

	C_2 [p']		
	CONTROL	DOUBLING	COOCCURRENCE
PARTICIPANT 1	139 (14)	119 (28)	161 (20)
	$n = 5$	$n = 10$	$n = 8$
PARTICIPANT 6	102 (15)	73 (28)	106 (44)
	$n = 5$	$n = 5$	$n = 7$
PARTICIPANT 7	102 (24)	89 (21)	81 (61)
	$n = 5$	$n = 6$	$n = 5$

TABLE 13. Average VOT (and *SD*) of C_2 [p'] for accurate repetitions of control-category targets and doubling errors on ordering-category targets for participants 1, 6, and 7.

The measurements reported here are based on a small amount of data from a small number of speakers, but they are not inconsistent with some instances of ejectives in doubling errors being reduced. To thoroughly explore whether doubling errors are in fact due to the blending of competing speech plans, additional studies that aim to elicit doubling errors more frequently than in the current task, and in stimuli that are more balanced for place of articulation, are necessary.

THE FULL ANALYSIS OF THE REPETITION TASK. The analysis of the repetition results has two parts: the grammar first generates repair errors, and then some portion of these repair errors is subject to blending with the faithful speech plan. Since movement and doubling errors are observed at similar rates, the simplifying assumption made here is that blending and repairs apply with equal frequency. Figures 6–9 give a schematic representation of the analysis. The frequencies at which repairs and blending apply are idealized, and the output frequencies predicted from this idealization are compared to the observed frequencies in each figure. The MaxEnt grammar tool was used to verify that the two markedness (*[cg][cg]$_{RT}$ and *[−cont, −son][cg]$_{RT}$) and two faithfulness con-

---

[8] The VOT of C_1 was not examined because there were no control forms with an ejective in C_1.

straints (MAX[cg] and NoFLOP[cg]) can be weighted to produce the idealized output frequencies.[9] In contrast, a grammar with only a single markedness constraint (*[−cont, −son][cg]$_{RT}$) fails to reproduce the idealized output frequencies, showing that the distribution of repairs on the repetition task cannot be due to differences in faithfulness constraints alone. The grammar with one markedness constraint overpredicts accuracy on cooccurrence targets and underpredicts accuracy on ordering targets. This grammar also underpredicts C_2 de-ejectivization rates for cooccurrence targets and overpredicts C_2 de-ejectivization rates for ordering targets.

In Figure 6, a cooccurrence target from the local condition is considered. The grammar favors a phonotactic repair 30% of the time, generating C_2 de-ejectivization errors. An additional 30% of the time, a speech plan for a C_2 de-ejectivization repair blends with the faithful speech plan, resulting in an accurate repetition. Task-specific faithfulness overrides the grammatical preference for repair 40% of the time, also resulting in an accurate repetition. Together, accurate repetitions arising from the grammar and from blending account for a predicted 70% of total responses, compared to an observed rate of 68%. Repairs are predicted and observed 30% of the time. In the nonlocal condition, the analysis is the same, as shown in Figure 7.

target:	GRAMMAR		IDEALIZED
k'ap'a	accurate k'ap'a		40%
	C_2 de-ejectivization k'apa		30%
	BLENDING		
	k'ap'a + k'apa = k'ap'a		30%

output:		PREDICTED	OBSERVED
accurate	k'ap'a	70%	68%
C_2 de-ejectivization	k'apa	30%	30%
C_1 de-ejectivization	kap'a	0%	2%

FIGURE 6. Analysis of a local cooccurrence-restriction violation.

target:	GRAMMAR		IDEALIZED
k'amip'a	accurate k'amip'a		40%
	C_2 de-ejectivization k'amipa		30%
	BLENDING		
	k'amip'a + k'amipa = k'amip'a		30%

output:		PREDICTED	OBSERVED
accurate	k'amip'a	70%	70%
C_2 de-ejectivization	k'amipa	30%	24%
C_1 de-ejectivization	kamip'a	0%	6%

FIGURE 7. Analysis of a nonlocal cooccurrence-restriction violation.

[9] The MaxEnt grammar tool, developed by Colin Wilson, Bruce Hayes, and Ben George, is available at http://www.linguistics.ucla.edu/people/hayes/MaxentGrammarTool/. The model matches the trained frequencies, using the following weights: *[cg][cg]$_{RT}$ = 2.28, *[−cont, −son][cg]$_{RT}$ = 0.64, NoFLOP[cg] = 1.12, MAX[cg] = 2.51.

A target violating the ordering restriction is shown in Figure 8. Repairs are generated by the grammar at a rate of 5% for C_2 de-ejectivization and 17.5% for movement. An additional 5% of the time, a C_2 de-ejectivization repair is blended with the speech plan for a faithful repetition, resulting in an accurate repetition. Movement repairs are blended with the faithful speech plan 17.5% of the time, resulting in doubling errors at the same rate as movement errors. Task-specific faithfulness favors an accurate repetition 55% of the time. These idealized frequencies result in predicted response rates that are a close match to observed rates. The same predicted rate of movement and doubling errors, 17.5%, overestimates the rate of movement, 16%, and underestimates the rate of doubling, 21%.

target:	GRAMMAR		IDEALIZED
kap'a	accurate	kap'a	55.0%
	C_2 de-ejectivization	kapa	5.0%
	movement	k'apa	17.5%
	BLENDING		
	kap'a + kapa = kap'a		5.0%
	kap'a + k'apa = k'ap'a		17.5%

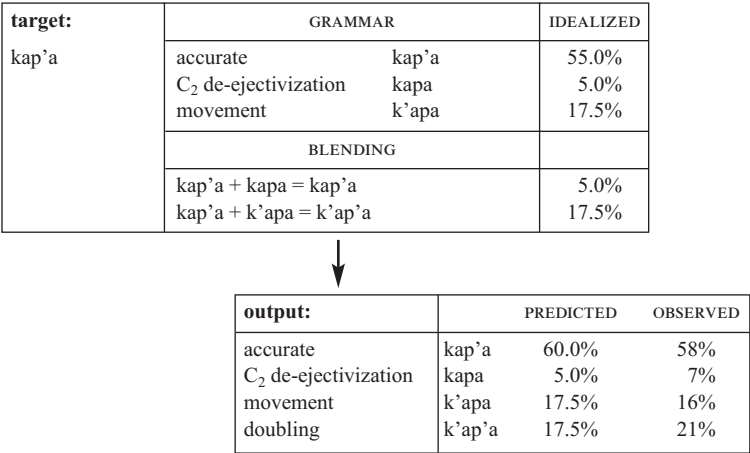| output:		PREDICTED	OBSERVED
accurate	kap'a	60.0%	58%
C_2 de-ejectivization	kapa	5.0%	7%
movement	k'apa	17.5%	16%
doubling	k'ap'a	17.5%	21%

FIGURE 8. Analysis of a local ordering-restriction violation.

Figure 9 shows a target violating the ordering restriction from the nonlocal condition, where errors are overall less frequent than in the local condition. The grammar generates C_2 de-ejectivization repairs at a rate of 5%, as in the local condition, but generates movement repairs only 7.5% of the time. These repairs are subject to blending with the faithful speech plan an additional 5% and 7.5% of the time, and task-specific faithfulness favors an accurate repetition 75% of the time. The predicted response rates are a good match to the observed rates, with the exception again that the frequency of movement errors is somewhat overestimated and the rate of doubling errors is underestimated.

To summarize, there were three key results in the repetition task that the formal analysis aimed to account for. The first was that forms violating the cooccurrence restriction are repaired more often than forms violating the ordering restriction, accounted for here through a formulation of markedness constraints such that the cooccurrence restriction is penalized by two constraints ($*[cg][cg]_{RT}$ and $*[-cont, -son][cg]_{RT}$) and the ordering restriction by only one ($*[-cont, -son][cg]_{RT}$). The second result was that accuracy on ordering-violating forms was sensitive to distance, while accuracy on cooccurrence-violating forms was not. The analysis of this effect is that both markedness constraints are stated over the domain of the root, meaning that local and nonlocal violations are penalized equally. The favored repair for violations of the ordering restriction is reassociation of ejection, and this reassociation is disfavored by distance-sensitive faithfulness constraints. Reassociation over longer distances is penalized more than reassociation over shorter distances, resulting in fewer movement errors in the nonlocal condition. Finally, the analysis needed to account for the nonrepair doubling

target:	GRAMMAR		IDEALIZED
kamip'a	accurate	kamip'a	75.0%
	C_2 de-ejectivization	kamipa	5.0%
	movement	k'amipa	7.5%
	BLENDING		
	kamip'a + kamipa = kamip'a		5.0%
	kamip'a + k'amipa = k'amip'a		7.5%

output:		PREDICTED	OBSERVED
accurate	kamip'a	80.0%	81%
C_2 de-ejectivization	kamipa	5.0%	5%
movement	k'amipa	7.5%	4%
doubling	k'amip'a	7.5%	10%

FIGURE 9. Analysis of a nonlocal ordering-restriction violation.

errors on forms violating the ordering restriction. These errors are analyzed as the blending of speech plans, a task effect that arises from the high faithfulness required on the task interacting with the strong categorical phonotactics of the language. While blending of faithful and repair forms likely happens throughout the experiment, it results in a unique, observable output only in doubling errors, which blend a form violating the ordering restriction with a movement repair.

ACCOUNTING FOR THE DISCRIMINATION RESULTS. The main difference between the discrimination and the repetition tasks is that the distance-based effects are found only in the repetition task. Crucially, in the proposed analysis, these distance-based effects derive from the task-specific ranking of faithfulness constraints, not from the relative weight of markedness constraints that constitute the task-independent grammar.

The discrimination results are consistent with forms being misperceived in proportion to the penalty they are assigned by the phonotactic grammar. Misperceived forms are then less likely to be distinguished from a phonotactically legal, expected form. On control trials (e.g. [hap'u] vs. [hapu]), where both forms are phonotactically legal, discrimination is very high. Discrimination is lower on trials that include a stimulus violating the ordering restriction; a form like [kap'a] is somewhat likely to be misperceived, and as a consequence is less distinguishable from phonotactically expected forms like [kapa] and [k'apa]. Discrimination is even lower on trials that compare a form violating the cooccurrence restriction with a C_2 de-ejectivization repair (e.g. [k'ap'a] vs. [k'apa]), consistent with the greater penalty for cooccurrence-violating forms than for ordering-violating forms. Finally, the weakest discriminability is found on trials that compare the two phonotactically illegal types of forms, for example, [k'ap'a] vs. [kap'a], both of which are penalized by the phonotactic grammar and are expected to be misperceived.

Given the discrimination results, it is likely that some of the errors on the repetition task arise from misperception, which in turn arises in proportion to the magnitude of phonotactic violation. Most importantly, the discrimination results are consistent with the central claim of the proposed analysis, which is that the cooccurrence restriction is more strongly enforced by the phonotactic grammar than the ordering restriction.

7.2. IMPLICATIONS OF THE ANALYSIS FOR THE STRUCTURE OF THE PHONOTACTIC GRAMMAR. The proposed phonotactic grammar contains a general constraint that rules

out both plain stop-ejective and ejective-ejective combinations, *[−cont, −son][cg], and also a more specific constraint that penalizes only ejective-ejective pairs, *[cg][cg]. The existence of this more specific cooccurrence constraint is nontrivial from a learning perspective. As shown in §2.2, the ordering and cooccurrence restrictions can both be accounted for with the single general phonotactic constraint on stop-ejective combinations, so the more specific constraint enforcing the cooccurrence restriction is not necessary to capture the phonotactic distribution of ejectives. It is thus unclear how a learner would discover the more specific cooccurrence constraint.

To arrive at the proposed phonotactic grammar, an inductive learner must be equipped with a bias to discover the cooccurrence constraint. The bias in favor of this constraint is likely a form of structural bias (Moreton 2008, Moreton & Pater 2012, Pater & Moreton 2012). Cooccurrence constraints are very common crosslinguistically, and are attested for a wide range of features, providing further support for a bias in favor of this type of constraint. An inductive learning model like the UCLA phonotactic learner could be modified to both preferentially select and use (i.e. assign high weights to) constraints that penalize a sequence of identical feature matrices over those that penalize sequences of nonidentical feature matrices.

While the current results speak directly only to a bias in favor of an OCP constraint on ejectives, this result may reveal a broader learning preference for generalizations that are stated over algebraic relations. Much of the work on algebraic relations has focused on generalizations over identical or nonidentical classes of segments, which cannot be picked out with standard natural classes (Berent & Shimron 1997, Marcus et al. 1999, Berent et al. 2001, Marcus 2001, Berent et al. 2002, Berent et al. 2012). OCP restrictions also fall under the scope of algebraic restrictions, as do many assimilatory restrictions (Frisch & Zawaydeh 2001, Rose & King 2007). Crucially, an inductive learner without variable notation is capable of capturing dissimilatory and assimilatory restrictions and of accounting for how such restrictions may generalize. A restriction like *[cg][cg] rules out all pairs of ejectives, regardless of place of articulation, and thus is expected to apply to nonce forms with novel ejective segments as well as native segments. This property makes algebraic restrictions on features fundamentally different from restrictions on the distribution of identical or nonidentical segments, classes that cannot be picked out without variable notation (Marcus 2001, Colavin et al. 2010, Berent et al. 2012). As shown in Berent et al. 2002, restrictions on the distribution of identical segments in Hebrew do generalize to novel segments, necessitating the algebraic representation of these restrictions. The role of variable notation in the current case, then, is in distinguishing dissimilatory and assimilatory generalizations from restrictions on sequences of unrelated natural classes. Enriching a baseline inductive learner, like the UCLA phonotactic learner, with both variable notation and a bias in favor of generalizations that use this notation may achieve a better fit with speaker knowledge.

A bias in favor of constraints with algebraic structure makes predictions about the relative well-formedness of other unattested structures in Cochabamba Quechua. As summarized in §2.1, aspirates are also subject to ordering and cooccurrence restrictions, and the transparent prediction is that the cooccurrence restriction on aspirates should be stronger than the ordering restriction. A more interesting question is how forms with unattested combinations of one ejective and one aspirate are treated, for example, *[k'apʰi] or *[kʰap'i]. These forms would fall under the scope of general phonotactic restrictions on stops followed by ejectives or aspirates (*[−cont, −son][cg] and *[−cont, −son][sg]), but would not fall under the scope of more specific algebraic cooccurrence restrictions on pairs of ejectives (*[cg][cg]) or pairs of aspirates (*[sg][sg]). If forms

with one ejective and one aspirate are treated comparably to forms with plain stop-ejective and plain stop-aspirate sequences, then the hypothesis that the algebraic structure of cooccurrence restrictions is special would be supported. If forms with one ejective and one aspirate are treated like forms violating the cooccurrence restriction, however, it is less clear what conclusions can be drawn. One possibility is that ejectives and aspirates are represented with a single acoustic feature [long VOT], which groups these segments together as a class, as argued for in Gallagher 2011. With this feature, the restrictions on pairs of ejectives, pairs of aspirates, and pairs of one ejective and one aspirate can all be grouped under a single OCP restriction: *[long VOT][long VOT]. Another possibility, suggested by a referee, is that restrictions that can be stated on a feature-geometric laryngeal tier (Clements 1985, McCarthy 1988) have a special status.

8. Conclusion. The experiments in this article show strong asymmetries in the representation of two categorical phonotactic restrictions, despite the parallel evidence for these restrictions in the lexicon. The distribution of ejectives in Cochabamba Quechua roots can be neatly accounted for with a single grammatical restriction that prohibits ejectives preceded by stops of any kind in the root. Despite the adequacy of this single restriction in accounting for the data, speakers of Cochabamba Quechua show substantial differences in how they treat forms with ejective-ejective and plain stop-ejective pairs, supporting a stronger prohibition on ejective-ejective pairs.

To account for the behavioral asymmetries, the phonotactic grammar must contain both the general constraint on all stop-ejective pairs, *[−cont, −son][cg], and a more specific constraint that penalizes only cooccurring ejectives, *[cg][cg]. Learning such a grammar requires a bias in favor of the cooccurrence constraint, as it is not strictly necessary to account for the data. It is hypothesized that the cooccurrence constraint benefits from its structural form, and that the bias that allows learners to find and use this constraint is a form of analytic bias (Moreton 2008, Pater & Moreton 2012); OCP constraints reference identical feature matrices, as opposed to a sequence of unrelated feature matrices, and as such are structurally distinguishable from other types of constraints.

In addition to the experimental results here, a bias in favor of OCP constraints is consistent with the crosslinguistic frequency of this type of restriction. Dissimilatory cooccurrence restrictions on a wide range of features are common typologically, and a learning bias in favor of phonotactic generalizations with the form of the OCP may partially underlie the typological frequency of this pattern. While the current results support an inductive learning bias in favor of an OCP restriction on ejectives, both articulatory and perceptual precursors have also been identified and argued for (Ohala 1981, 1993, Gallagher 2010a,b, Gallagher & Whang 2014). One striking aspect of the current results is that Cochabamba Quechua speakers do not show the effects of locality that would be predicted if articulatory or perceptual difficulty were directly encoded in their synchronic representation of this generalization. Thus, while phonetic considerations may influence learning in addition to inductive biases, the phonotactic grammar may also contain constraints that abstract over phonetic difficulty.

REFERENCES

Adriaans, Frans, and René Kager. 2010. Adding generalization to statistical learning: The induction of phonotactics from continuous speech. *Journal of Memory and Language* 62.311–31. DOI: 10.1016/j.jml.2009.11.007.

Albright, Adam. 2009. Feature-based generalisation as a source of gradient acceptability. *Phonology* 26.9–41. DOI: 10.1017/S0952675709001705.

ALBRIGHT, ADAM, and BRUCE HAYES. 2003. Rules vs. analogy in English past tenses: A computational/experimental study. *Cognition* 90.119–61. DOI: 10.1016/S0010-0277 (03)00146-X.

BATES, DOUG, and MARTIN MAECHLER. 2010. lme4: Linear mixed effects models using S4 classes. Online: https://cran.r-project.org/web/packages/lme4.

BECKER, MICHAEL; NIHAN KETREZ; and ANDREW NEVINS. 2011. The surfeit of the stimulus: Analytic biases filter lexical statistics in Turkish laryngeal alternations. *Language* 87.84–125. DOI: 10.1353/lan.2011.0016.

BECKER, MICHAEL; ANDREW NEVINS; and JONATHAN LEVINE. 2012. Asymmetries in generalizing alternations to and from initial syllables. *Language* 88.231–68. DOI: 10.1353 /lan.2012.0049.

BENNETT, WILLIAM. 2013. *Dissimilation, consonant harmony, and surface correspondence*. New Brunswick, NJ: Rutgers University dissertation.

BERENT, IRIS; DANIEL EVERETT; and JOSEPH SHIMRON. 2001. Do phonological representations specify variables? Evidence from the obligatory contour principle. *Cognitive Psychology* 42.1–60. DOI: 10.1006/cogp.2000.0742.

BERENT, IRIS; GARY MARCUS; JOSEPH SHIMRON; and ADAMANTIOS GAFOS. 2002. The scope of linguistic generalizations: Evidence from Hebrew word formation. *Cognition* 83.113–39. DOI: 10.1016/S0010-0277(01)00167-6.

BERENT, IRIS, and JOSEPH SHIMRON. 1997. The representation of Hebrew words: Evidence from the obligatory contour principle. *Cognition* 64.39–72. DOI: 10.1016/S0010-0277 (97)00016-4.

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. DOI: 10.1016/j.cognition.2006.05.015.

BERENT, IRIS; COLIN WILSON; GARY MARCUS; and DOUG BEMIS. 2012. On the role of variables in phonology: Remarks on Hayes and Wilson 2008. *Linguistic Inquiry* 43.97–119. DOI: 10.1162/LING_a_00075.

BOERSMA, PAUL. 1997. How we learn variation, optionality, and probability. *University of Amsterdam Institute of Phonetic Sciences Proceedings* 21.43–58.

BOERSMA, PAUL, and BRUCE HAYES. 2001. Empirical tests of the gradual learning algorithm. *Linguistic Inquiry* 32.45–86. DOI: 10.1162/002438901554586.

BOERSMA, PAUL, and DAVID WEENINK. 2011. Praat: Doing phonetics by computer. Version 5.2.43. Online: http://www.praat.org.

BOLEY, JON, and MICHAEL LESTER. 2009. Statistical analysis of ABX results using signal detection theory. *Audio Engineering Society Convention Papers* 127:7826. Online: http://www.aes.org/e-lib/browse.cfm?elib=15022.

CHOMSKY, NOAM, and MORRIS HALLE. 1968. *The sound pattern of English*. Cambridge, MA: MIT Press.

CLEMENTS, G. NICK. 1985. The geometry of phonological features. *Phonology Yearbook* 2.225–52. Online: http://www.jstor.org/stable/4419958.

COETZEE, ANDRIES. 2008. Grammaticality and ungrammaticality in phonology. *Language* 84.218–57. DOI: 10.1353/lan.0.0000.

COETZEE, ANDRIES. 2009. An integrated grammatical/non-grammatical model of phonological variation. *Current issues in linguistic interfaces, vol. 2*, ed. by Young-Se Kang et al., 267–94. Seoul: Hankookmunhwasa.

COETZEE, ANDRIES. 2014. Grammatical change through lexical accumulation: Voicing cooccurrence restrictions in Afrikaans. *Language* 90.693–721. DOI: 10.1353/lan.2014 .0066.

COETZEE, ANDRIES, and SHIGETO KAWAHARA. 2013. Frequency biases in phonological variation. *Natural Language and Linguistic Theory* 31.47–89. DOI: 10.1007/s11049- 012-9179-z.

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. DOI: 10.1007/s11049-008-9039-z.

COETZEE, ANDRIES, and JOE PATER. 2011. The place of variation in phonological theory. *The handbook of phonological theory*, 2nd edn., ed. by John Goldsmith, Jason Riggle, and Alan Yu, 401–34. Oxford: Blackwell.

Colavin, Rebecca; Roger Levy; and Sharon Rose. 2010. Modeling OCP-Place in Amharic with the maximum entropy phonotactic learner. *Chicago Linguistic Society* 46.2.27–41.

Daland, Robert; Bruce Hayes; James White; Marc Garellek; Andrea Davis; and Ingrid Norrman. 2011. Explaining sonority projection effects. *Phonology* 28.197–234. DOI: 10.1017/S0952675711000145.

Davidson, Lisa. 2006. Phonology, phonetics, or frequency: Influences on the production of non-native sequences. *Journal of Phonetics* 34.104–37. DOI: 10.1016/j.wocn.2005.03 .004.

Davidson, Lisa. 2010. Phonetic bases of similarities in cross-language production: Evidence from English and Catalan. *Journal of Phonetics* 38.272–88. DOI: 10.1016/j .wocn.2010.01.001.

Davidson, Lisa, and Jason Shaw. 2012. Sources of illusion in consonant cluster perception. *Journal of Phonetics* 40.234–48. DOI: 10.1016/j.wocn.2011.11.005.

Dupoux, Emmanuel; Yuki Hirose; Kazuhiko Kakehi; Christophe Pallier; and Jacques Mehler. 1999. Epenthetic vowels in Japanese: A perceptual illusion? *Journal of Experimental Psychology: Human Perception and Performance* 25.1568–78. DOI: 10.1037/0096-1523.25.6.1568.

Dupoux, Emmanuel; Erika Parlato; Sonia Frota; Yuki Hirose; and Sharon Peperkamp. 2011. Where do illusory vowels come from? *Journal of Memory and Language* 64.199–210. DOI: 10.1016/j.jml.2010.12.004.

Ernestus, Mirjam, and Harald Baayen. 2003. Predicting the unpredictable: Interpreting neutralized segments in Dutch. *Language* 79.5–38. DOI: 10.1353/lan.2003.0076.

Frisch, Stefan; Janet Pierrehumbert; and Michael Broe. 2004. Similarity avoidance and the OCP. *Natural Language and Linguistic Theory* 22.179–228. DOI: 10.1023/B :NALA.0000005557.78535.3c.

Frisch, Stefan, and Richard Wright. 2002. The phonetics of phonological speech errors: An acoustic analysis of slips of the tongue. *Journal of Phonetics* 30.139–62. DOI: 10.1006/jpho.2002.0176.

Frisch, Stefan, and Bushra Zawaydeh. 2001. The psychological reality of OCP-Place in Arabic. *Language* 77.91–106. Online: DOI: 10.1353/lan.2001.0014.

Gallagher, Gillian. 2010a. *The perceptual basis of long-distance laryngeal restrictions*. Cambridge, MA: MIT dissertation.

Gallagher, Gillian. 2010b. Perceptual distinctness and long-distance laryngeal restrictions. *Phonology* 27.435–80. DOI: 10.1017/S0952675710000217.

Gallagher, Gillian. 2011. Acoustic and articulatory features in phonology—The case for [long VOT]. *The Linguistic Review* 28.281–313. DOI: 10.1515/tlir.2011.008.

Gallagher, Gillian. 2013a. Speaker awareness of non-local ejective phonotactics in Cochabamba Quechua. *Natural Language and Linguistic Theory* 31.1067–99. DOI: 10 .1007/s11049-013-9200-1.

Gallagher, Gillian. 2013b. Learning the identity effect as an artificial language: Bias and generalisation. *Phonology* 30.253–95. DOI: 10.1017/S0952675713000134.

Gallagher, Gillian. 2014. An identity bias in phonotactics: Evidence from Cochabamba Quechua. *Laboratory Phonology* 5.3.337–78. DOI: 10.1515/lp-2014-0012.

Gallagher, Gillian, and James Whang. 2014. An acoustic study of trans-vocalic ejective pairs in Cochabamba Quechua. *Journal of the International Phonetic Association* 44.2.133–54. DOI: 10.1017/S0025100314000048.

Gelman, Andrew, and Jennifer Hill. 2006. *Data analysis using regression and multilevel/hierarchical models*. Cambridge: Cambridge University Press.

Goldsmith, John. 1976. *Autosegmental phonology*. Cambridge, MA: MIT dissertation.

Goldsmith, John, and Jason Riggle. 2012. Information theoretic approaches to phonological structure: The case of Finnish vowel harmony. *Natural Language and Linguistic Theory* 30.859–96. DOI: 10.1007/s11049-012-9169-1.

Goldstein, Louis; Marianne Pouplier; Larissa Chen; Elliot Saltzman; and Dani Byrd. 2007. Dynamic action units slip in speech production errors. *Cognition* 103. 386–412. DOI: 10.1016/j.cognition.2006.05.010.

Goldwater, Sharon, and Mark Johnson. 2003. Learning OT constraint rankings using a maximum entropy model. *Proceedings of the Workshop on Variation within Optimality*

Theory, Stockholm University. Online: http://homepages.inf.ed.ac.uk/sgwater/papers /OTvar03.pdf.

GORMAN, KYLE. 2013. *Generative phonotactics*. Philadelphia: University of Pennsylvania dissertation.

GOUSKOVA, MARIA, and MICHAEL BECKER. 2013. Nonce words show that Russian yer alternations are governed by the grammar. *Natural Language and Linguistic Theory* 31. 735–65. DOI: 10.1007/s11049-013-9197-5.

GREENBERG, JOSEPH. 1960. The patterning of root morphemes in Semitic. *Word* 6.162–81. DOI: 10.1080/00437956.1950.11659378.

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. DOI: 10.1121/1.2534656.

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. DOI: 10.1037/0096-1523.24.2.592.

HAYES, BRUCE, and JAMES WHITE. 2013. Phonological naturalness and phonotactic learning. *Linguistic Inquiry* 44.45–75. DOI: 10.1162/LING_a_00119.

HAYES, BRUCE, and COLIN WILSON. 2008. A maximum entropy model of phonotactics and phonotactic learning. *Linguistic Inquiry* 39.379–440. DOI: 10.1162/ling.2008.39.3 .379.

HAYES, BRUCE; KIE ZURAW; PÉTER SIPTÁR; and ZSUZSA LONDE. 2009. Natural and unnatural constraints in Hungarian vowel harmony. *Language* 85.822–63. DOI: 10.1353/lan .0.0169.

JAYNES, EDWIN T. 1983. *Papers on probability, statistics, and statistical physics*, ed. by R. D. Rosenkrantz. Dordrecht: Reidel.

KABAK, BARIŞ, and WILLIAM IDSARDI. 2007. Perceptual distortions in the adaptation of English consonant clusters: Syllable structure or consonantal contact constraints? *Language and Speech* 50.23–52. DOI: 10.1177/00238309070500010201.

KAGER, RENÉ, and JOE PATER. 2012. Phonotactics as phonology: Knowledge of a complex restriction in Dutch. *Phonology* 29.81–111. DOI: 10.1017/S0952675712000048.

LAIME AJACOPA, TEOFILO. 2007. *Diccionario Bilingüe, Iskay simipi yuyayk'ancha: Quechua–Castellano, Castellano–Quechua* [Bilingual dictionary: Quechua–Spanish, Spanish–Quechua]. La Paz.

LEBEN, WILLIAM. 1973. *Suprasegmental phonology*. Cambridge, MA: MIT dissertation.

MACEACHERN, MARGARET. 1999. *Laryngeal co-occurrence restrictions*. New York: Garland.

MACKENZIE, SARA. 2009. *Contrast and similarity in consonant harmony processes*. Toronto: University of Toronto dissertation.

MACKENZIE, SARA. 2013. Laryngeal co-occurrence restrictions in Aymara: Contrastive representations and constraint interaction. *Phonology* 30.297–345. DOI: 10.1017/S0952675 713000146.

MACMILLAN, NEIL, and DOUGLAS CREELMAN. 2005. *Detection theory: A user's guide*. 2nd edn. Mahwah, NJ: Lawrence Erlbaum.

MANNING, CHRISTOPHER, and HINRICH SCHÜTZE. 1999. *Foundations of statistical natural language processing*. Cambridge, MA: MIT Press.

MARCUS, GARY. 2001. *The algebraic mind: Integrating connectionism and cognitive science*. Cambridge, MA: MIT Press.

MARCUS, GARY; S. VIJAYAN; S. BANDI RAO; and P. M. VISHTON. 1999. Rule learning by seven-month-old infants. *Science* 283.77–80. DOI: 10.1126/science.283.5398.77.

MCCARTHY, JOHN. 1986. OCP effects: Gemination and anti-gemination. *Linguistic Inquiry* 17.207–63. Online: http://www.jstor.org/stable/4178486.

MCCARTHY, JOHN. 1988. Feature geometry and dependency: A review. *Phonetica* 45.84–108. DOI: 10.1159/000261820.

MCCARTHY, JOHN. 1994. The phonetics and phonology of Semitic pharyngeals. *Phonological structure and phonetic form: Papers in laboratory phonology 3*, ed. by Patricia Keating, 191–233. Cambridge: Cambridge University Press.

MCCARTHY, JOHN, and JOE PATER (eds.) 2016. *Harmonic grammar and harmonic serialism*. London: Equinox.

MORETON, ELLIOTT. 2002. Structural constraints in the perception of English stop-sonorant clusters. *Cognition* 84.55–71. DOI: 10.1016/S0010-0277(02)00014-8.

MORETON, ELLIOTT. 2008. Analytic bias and phonological typology. *Phonology* 25.83–127. DOI: 10.1017/S0952675708001413.

MORETON, ELLIOTT, and JOE PATER. 2012. Structure and substance in artificial-phonology learning. Part 1: Structure. Part II: Substance. *Language and Linguistics Compass* 6. 686–718. DOI: 10.1002/lnc3.363. DOI: 10.1002/lnc3.366.

OHALA, JOHN J. 1981. The listener as a source of sound change. *Chicago Linguistic Society (Parasession on language and behavior)* 17.2.178–203.

OHALA, JOHN J. 1993. The phonetics of sound change. *Historical linguistics: Problems and perspectives*, ed. by Charles Jones, 237–78. London: Longman.

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. Online: http://arts-sciences.und.edu/summer-institute-of-linguistics/work-papers /_files/docs/1997-parker.pdf.

PARKER, STEVE, and DAVID WEBER. 1996. Glottalized and aspirated stops in Cuzco Quechua. *International Journal of American Linguistics* 62.70–85. DOI: 10.1086/466276.

PATER, JOE. 2009. Weighted constraints in generative linguistics. *Cognitive Science* 33.999–1035. DOI: 10.1111/j.1551-6709.2009.01047.x.

PATER, JOE; RAJESH BHATT; and CHRIS POTTS. 2007. Linguistic optimization. Amherst: University of Massachusetts, Amherst, MS.

PATER, JOE and ELLIOTT MORETON. 2012. Structurally biased phonology: Complexity in learning and typology. *The EFL Journal* 3.2. Online: http://www.openhumanities alliance.org/journals/eflj/article/view/63.

PITT, MARK. 1998. Phonological processes and the perception of phonotactically illegal consonant clusters. *Perception & Psychophysics* 60.941–51. DOI: 10.3758/BF03211 930.

POUPLIER, MARIANNE. 2008. The role of a coda consonant as error trigger in repetition tasks. *Journal of Phonetics* 36.114–40. DOI: 10.1016/j.wocn.2007.01.002.

POUPLIER, MARIANNE, and LOUIS GOLDSTEIN. 2005. Asymmetries in the perception of speech production errors. *Journal of Phonetics* 33.47–75. DOI: 10.1016/j.wocn.2004 .04.001.

PRINCE, ALAN, and PAUL SMOLENSKY. 2004 [1993]. *Optimality theory: Constraint interaction in generative grammar*. Malden, MA: Blackwell.

R CORE TEAM. 2015. R: A language and environment for statistical computing. Vienna: R Foundation for Statistical Computing. Online: http://www.R-project.org.

ROSE, SHARON, and LISA KING. 2007. Speech error elicitation and co-occurrence restrictions in two Ethiopian Semitic languages. *Language and Speech* 50.451–504. DOI: 10.1177 /00238309070500040101.

SCHOLES, ROBERT. 1966. *Phonotactic grammaticality*. The Hague: Mouton.

SUZUKI, KEIICHIRO. 1998. *A typological investigation of dissimilation*. Tucson: University of Arizona dissertation.

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. DOI: 10.1016/S0163-6383(84)80022-3.

WILSON, COLIN; LISA DAVIDSON; and SEAN MARTIN. 2014. Effects of acoustic–phonetic detail on cross-language speech production. *Journal of Memory and Language* 77.1–24. DOI: 10.1016/j.jml.2014.08.001.

ZYMET, JESSE. 2014. Distance-based decay in long-distance phonological processes. Paper presented at the West Coast Conference on Formal Linguistics 32.

[ggillian@nyu.edu]