Version date 2011.7.1

Sonority variation in Stochastic OT: Implications for markedness hierarchies

Jennifer L. Smith and Elliott Moreton
University of North Carolina, Chapel Hill

1. Introduction

Sonority-related effects belong to a large class of phenomena—in phonology, and in linguistic theory more generally—that have been analyzed in terms of *markedness hierarchies.* A markedness hierarchy is a multi-step scale designed to model implicational universals, such as the cross-linguistic preferences for high-sonority syllable nuclei and for low-sonority syllable onsets.

In a constraint-based framework such as Optimality Theory, markedness hierarchies have been formalized in one of two ways: as a *scale-partition constraint family* (Prince & Smolensky 1993, 2004; see also much subsequent work on markedness scales in OT), in which there is one constraint per level of the hierarchy and the constraints in the family are universally ranked, and as a *stringency constraint family* (Prince 1997, 1999; de Lacy 2002, 2004, 2006), in which constraints are formalized as overlapping subsets of the hierarchy and can be freely ranked on a language-particular basis.

Previous comparison of the two formal approaches to markedness scales has focused on differences in their between-language typological predictions (see especially de Lacy 2004). This chapter identifies an additional empirical domain in which the two approaches make distinct predictions, but whose consequences for markedness hierarchies have remained largely unexplored: intra-speaker phonological variation. Specifically, we show that in the framework of Stochastic OT (Boersma & Hayes 2001), formalizing sonority constraints (or other types of markedness hierarchies) as scale-partition constraints predicts the existence of *harmony reversals*—the selection of a less harmonic or desirable form in preference to a more harmonic one—under phonological variation, while the stringency approach makes no such prediction (§§2–3). We confirm that patterns of variation involving multiple levels of the sonority scale are attested, so the empirical scenario of interest is attested and theoretically relevant (§4). However, it is difficult to test whether or not any particular case of sonority-related variation really does show harmony reversals or not, because it is always possible that interference from other constraints might also be at play. We present a new empirical method for testing whether a harmony reversal has actually been found, by showing that a true harmony reversal exhibits a particular mathematical relationship between the probabilities of ranking reversals among multiple constraints in the markedness-hierarchy constraint family (§5).

We take *sonority* to be a phonological property, to which formal phonological constraints can make reference, although it may be untimately based on or grounded in phonetic factors. Our work shows that whether the sonority scale is predicted to be universally consistent, or to have language-particular or variable aspects, crucially depends at least in part on the formalization of sonority constraints in a particular phonological framework; identical assumptions about the sonority scale itself lead to cross-linguistic consistency in the stringency approach but to cross-linguistic variation in the scale-partition approach when implemented under Stochastic OT.

While this chapter focuses on examples of sonority-related variation, the basic point about scale-partition versus stringency constaints and predicted patterns of intra-speaker variation is more widely relevant, with implications for any domain of linguistics where constraint families based on markedness hierarchies have been proposed.

2. Markedness hierarchies in Optimality Theory

The concept of the *markedness hierarchy* has proven useful in multiple domains of linguistics, in morphosyntax (Silverstein 1976; Keenan & Comrie 1977; Dixon 1979; Croft 1988; Aissen 1999, 2003; Lee 2006) as well as in phonology. In broad terms, a markedness hierarchy is a family of related linguistic features — such as the features encoding levels of animacy, definiteness, place of articulation, or sonority — that is structured in a cross-linguistically consistent hierarchy of implicational relationships, and plays a role in multiple linguistic patterns within and across languages.

A number of markedness hierarchies are based on the sonority scale.[1] For example, it has been proposed that syllable peaks in general (Dell & Elmedlaoui 1985), stressed syllable peaks (Kenstowicz 1994), and moraic segments (Zec 1995; Gnanadesikan 1995, 2004) each show a preference for segments of the highest possible sonority level; that onsets (Steriade 1982) show a preference for segments of the lowest possible sonority level; that onset clusters (Selkirk 1982; Baertsch 1988) show a preference for rising and maximally dispersed sonority; and that coda-onset sequences (Murray & Venneman 1983; Gouskova 2004) show a preference for falling and maximally dispersed sonority.

As a simplified illustration of a markedness hiearchy involving sonority, consider the preference for higher-sonority syllable peaks. Vowel height categories have the following sonority relationship (where '>' means 'is greater in sonority than').

(1) Sonority scale (partial)

high sonority *low sonority*
low vowels > mid vowels > high vowels

Since syllable peaks prefer higher-sonority segments, this basic sonority scale gives rise to a *harmony scale* (Prince & Smolensky 2004: 161), as in (2) (where '≻' means 'is more harmonic than,' i.e., 'is phonologically preferable to').

(2) Harmony scale

more preferred *less preferred*
peak/lowV ≻ peak/midV ≻ peak/highV

In Optimality Theory and related frameworks, the two main approaches to modeling such preference relationships by means of phonological constraints are the scale-partition approach (§2.1) and the stringency approach (§2.2).

---

[1] Indeed, sonority-related markedness hierarchies were the first to be formalized within OT (Prince & Smolensky 1993/2004).

2.1 The scale-partition approach to markedness hierarchies

Originally, a harmony scale such as that in (2) was mapped directly onto a family of phonological constraints according to a formal operation known as *constraint alignment* (Prince & Smolensky 2004: 161). By this operation, each phonological structure gives rise to a constraint that penalizes that structure (and only that structure). The *least*-preferred phonological configuration (here, a syllable peak with a high vowel) is associated with the *highest*-ranked constraint, i.e., incurs the most severe penalty.

(3) Scale-partition constraint family: *PEAK/X

*PEAK/HIGHV >> *PEAK/MIDV >> *PEAK/LOWV

The constraints in this family assign violations to candidates with different syllable peaks as follows.

(4) Violations assigned by *PEAK/X constraints

	*PEAK/HIGHV	*PEAK/MIDV	*PEAK/LOWV
a. [a]			*
b. [e]		*	
c. [i]	*		

As the tableau in (4) indicates, in order for the *PEAK/X constraints to select output candidates in accordance with the harmony scale in (2), they must always be ranked in the order shown in (3) in every language. This is because each constraint penalizes exactly one point on the harmony scale, rather than an interval or set of points along the scale. If these constraints could be ranked differently in different languages, then languages would vary as to which sonority levels were most preferred as nuclei. In the extreme case, taking the entire sonority scale into account, this predicts that some languages should prefer obstruent nuclei over low-vowel nuclei, a prediction that is not empirically supported. Therefore, Prince & Smolensky (2004: 162) explicitly propose that constraint families that are derived from markedess scales through constraint alignment in this way have a *universally fixed ranking* determined by the associated harmony scale (as in (2)).

2.2 The stringency approach to markedness hierarchies

An alternative to scale-partition constraint families for modeling markedness hierarchies is *stringency,* proposed by Prince (1997, 1999) and extensively developed by de Lacy (2002, 2004, 2006). In the stringency approach, for every point along the harmony scale, there is a constraint that assigns violations to that point and all points up to and including the less-preferred end of the scale. Thus, each constraint in the stringency family refers to the least-preferred structure on the harmony scale, and if the constraint refers to more than one point on the scale, all such points form a contiguous interval. For example, the harmony scale in (2) would give rise to the family of constraints in (5).

(5) Stringency constraint family: *Peak/≤X

*Peak/≤HighV	penalizes peaks associated with {HighV}
*Peak/≤MidV	penalizes peaks associated with {HighV, MidV}
*Peak/≤LowV	penalizes peaks associated with {HighV, MidV, LowV}

With this set of constraints, violations are assigned to candidates with different syllable peaks as in (6).

(6) Violations assigned by *Peak/≤X constraints

	*Peak/≤HighV	*Peak/≤MidV	*Peak/≤LowV
a. [a]			*
b. [e]		*	*
c. [i]	*	*	*

As discussed in detail by both Prince and de Lacy, stringency constraints differ from fixed-ranking constraints in that they do not require a universally fixed ranking in order to ensure that only grammars compatible with the harmony scale can be generated. Crucially, no ranking of the constraints in (6) can possibly produce a grammar in which a lower-sonority peak is allowed but a higher-sonority peak is not allowed, because any constraint that penalizes a higher-sonority peak will necessarily also penalize all peaks that are lower in sonority; higher-sonority peaks are harmonically bounded (with respect to the constraints in the stringency family).

Prince and de Lacy make an empirical case for the stringency approach to markedness scales on the basis of factorial typology. That is, they demonstrate that there are categorical phonological patterns occurring in natural-language phonologies that are predicted under a stringency approach, but not under a fixed-ranking approach. Crucial examples involve *scale conflation*, a pattern in which two or more points on the scale are treated by some phonological pattern as equally (un)desirable.

This chapter identifies a second domain in which the two approaches make distinct empirical predictions: in the formal learning of variation patterns involving the reranking of a single constraint with respect to multiple members of the constraints in a markedness scale. If we assume the scale-partition approach, then learning the variation pattern entails learning a pattern that produces *harmony reversals* — instances of variation in which a structure found lower on the harmony scale is actually chosen over a structure that is more harmonic. If we assume the stringency approach, then the variation pattern can be learned with a grammar that nevertheless continues to prohibit harmony reversals.

3. Markedness hierarchies, phonological variation, and harmony reversals

3.1 Phonological variation in Stochastic OT

A given constraint ranking produces one consistent output for each input. This means that a speaker (or a language community) showing variation between two or more output forms must in some way be making use of two or more distinct constraint rankings.

One influential approach to modeling linguistic variation in Optimality Theory is Stochastic OT (Boersma 1998; Boersma & Hayes 2001). In this framework, constraint rankings correspond to points on a number line. While the exact numerical values assigned to the constraints are arbitrary, a greater ranking value corresponds to a higher ranking, and so domination relations between constraints can be represented.

In the grammar of a particular language, each constraint has an intrinsic *ranking value*. However, this implementation of OT is stochastic in that every time an input is mapped to an output by the grammar, the ranking value for each constraint is perturbed by a noise component. The noise component is drawn from a normal distribution whose mean is the constraint's ranking value and whose standard deviation is some constant value (often set at 2.0 units by convention). Boersma & Hayes (2001: 50) propose that all constraints have the same standard deviation for their noise distribution, because the noise function is part of the grammar as a whole and not the property of an individual constraint. The proposal that the noise distribution is the same for all constraints has crucial consequences for variation involving markedness scales, as is discussed in detail in §3.2.

A constraint's intrinsic ranking value as modified by the noise component results in a numerical value known as the *selection point*. Because the value of the noise component varies for each constraint each time the grammar is invoked, so does the constraint's selection point. Crucially, if two constraints C_1, C_2 have ranking values that are close together, then the relative ordering of their selection points will vary, as shown in (7). Such a grammar does in essence make use of two different rankings because on some evaluations, $C_1 >> C_2$, as in (7)(b), but on others, $C_2 >> C_1$, as in (7)(c). (The probability of occurrence of $C_1 >> C_2$ versus $C_2 >> C_1$ depends on their ranking values, a point whose implications for markedness scales will be explored in §5.)

(7) Variable constraint ranking in Stochastic OT

(a) Ranking values and noise distributions for two constraints

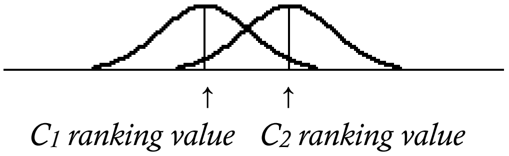

(b) Selection points on an evaluation where $C_1 >> C_2$

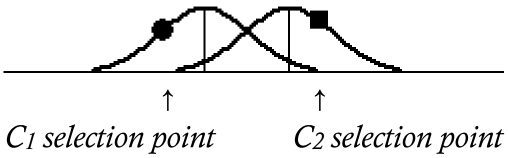

(c) Selection points on an evaluation where $C_2 >> C_1$

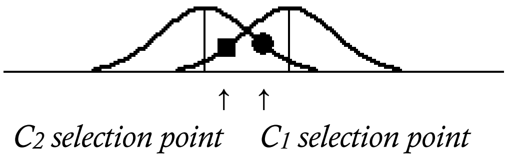

A grammar in which the relative ranking of two *conflicting* constraints can vary in this way is a grammar that produces discernable phonological variation, because the choice of which output wins on a given evaluation will depend on which of the two constraints happens to have a higher selection point for that evaluation.

3.2 Markedness hierarchies and harmony reversals under Stochastic OT

Under the assumption that the standard deviation of the noise distribution is the same for all constraints, the Stochastic OT model places restrictions on possible patterns of variation (a point discussed by Anttila 2007: 534 as well).

> ...it is worth noting that this model is quite restrictive: there are various cases of logically possible free rankings that it excludes. Thus, for example, it would be impossible to have a scheme in which A "strictly" outranks B (i.e., the opposite ranking is vanishingly rare), B "strictly" outranks C, and D is ranked freely with respect to both A and C. This scheme would require a much larger standard deviation for D than for the other constraints. (Boersma & Hayes 2001: 50)

In other words, given a ranked set of constraints *A* >> *B* >> *C*, and variation in the relative ranking between these constraints and a fourth constraint *D*, scenarios (8)(a-b) are possible, but (8)(c) is not.

(8) Variation scenarios for one constraint (*D*) versus ranked constraints (*A* >> *B* >> *C*)

(a) Possible: The relative ranking *A* >> *B* >> *C* does not vary
D varies with respect to at most two consecutive points on the scale

(b) Possible: The relative ranking *A* >> *B* >> *C* shows variation
D varies with respect to each of *A*, *B*, *C*

(c) Impossible: The relative ranking *A* >> *B* >> *C* does not vary
D varies with respect to each of *A*, *B*, *C*

As noted in §2.2, a major difference between the scale-partition and stringency approaches to markedness scales is whether or not the family of constraints requires a universal ranking in order to produce patterns in accordance with the associated harmony scale; the scale-partition approach does require such a universal ranking, but the stringency approach does not.

In the framework of standard Stochastic OT, this difference leads to a difference in predicted phonological patterns. Assume a situation as in (8), in which *A*, *B*, and *C* are specifically a family of constraints associated with a markedness hierarchy. If there is another constraint *D* whose ranking is known to vary with respect to that of multiple members of the constraint family, then the ranking of the constraints within the family (as determined by their selection point values at the time of evaluation) *must also be variable.*

Crucially, if the constraints in the family are scale-partition constraints, then allowing them to vary will lead to cases of *harmony reversal,* in which a structure lower on the harmony scale is variably preferred to a structure higher on the scale. However, if the constraints in the family are stringency constraints, then they generate patterns consistent with the harmony scale under any ranking. As a consequence, harmony reversals should never be observed, even in cases of variation as described above.

This difference can be illustrated with a schematic example involving one sonority-based markedness scale, onset-sonority distance. (See §4 for language examples involving this and other sonority-based markedness scales.) Phonological patterns are sometimes attested in which there is variation in the production of a target syllable with an onset cluster (CCV): in some cases the cluster is produced intact (CCV), and in other cases the onset cluster is avoided through vowel epenthesis (C*v*.CV or *v*C.CV, where *v* indicates an epenthetic vowel). Crucially, the variation can be sensitive to the sonority profile of the onset cluster, so that a more harmonic cluster (such as obstruent+liquid) is produced with epenthesis less frequently than a less harmonic cluster (such as obstruent+obstruent). This pattern indicates that the ranking of the anti-epenthesis constraint DEP (McCarthy & Prince 1995) is varying with respect to multiple members of a sonority-based constraint family on onset clusters.

Sonority-based restrictions on onset clusters can be stated as a harmony scale, where a greater distance in sonority between the segments in a cluster is more strongly preferred (Selkirk 1982; Baertsch 1988).[2] For example, consider the simplified consonant sonority scale in (9).

(9) Sonority scale

[j] > [l] > [n] > [s] > [t]

On this scale, the cluster [tl] would have a sonority distance of 3, because [l] is three steps away from [t].[3] The cluster [nl] would have a distance of 1.

The cross-linguistic preference for larger sonority distance within an onset cluster can be modeled with the following harmony scale.

(10) Harmony scale for onset sonority distance

Dist=4 ≻ Dist=3 ≻ Dist=2 ≻ Dist=1 ≻ Dist=0

This harmony scale in turn corresponds to the following scale-partition and stringency constraint families respectively.

(11) Constraint families for onset sonority distance

(a) Scale-partition constraint family

*DIST=0 >> *DIST=1 >> *DIST=2 >> *DIST=3 >> *DIST=4

(b) Stringency constraint family

*DIST≤0, *DIST≤1, *DIST≤2, *DIST≤3, *DIST≤4

---

[2] Parker (this volume) argues that sonority dispersion, rather than sonority distance, is the relevant sonority-related criterion for onset clusters. If so, the same argument can be made here with slightly modified constraints; the important point is that it is a sonority-based harmony scale that is at stake.

[3] Exact numerical values for sonority distance will depend on the precise version of the sonority scale adopted. The scale shown in (9) is intended as a concrete illustration for use in the discussion, not a substantive claim about the exact structure of the sonority scale.

In a language that avoids all potential CC onset clusters through epenthesis, Dep is ranked *below* the sonority-distance constraint *against the least problematic cluster,* so that clusters are broken up no matter what their sonority distance.

(12) Ranking for a language with epenthesis into all potential onset clusters

(a) With scale-partition constraints: Dep ranked below lowest constraint in scale

*Dist=0 >> *Dist=1 >> *Dist=2 >> *Dist=3 >> *Dist=4 >> **Dep**

(b) With stringency constraints: Dep ranked below most stringent constraint

*Dist≤4 >> **Dep**
(the other *Dist≤*n* constraints can be ranked anywhere)

Conversely, in a language that allows all potential CC onset clusters to surface and never shows epenthesis, Dep is ranked above *all* sonority-distance constraints, so that epenthesis is never chosen no matter how close the sonority distance in the onset cluster.

(13) Ranking for a language where all potential onset clusters surface

(a) With scale-partition constraints: Dep ranked above entire scale

Dep >> *Dist=0 >> *Dist=1 >> *Dist=2 >> *Dist=3 >> *Dist=4

(b) With stringency constraints: Dep ranked above all stringency-family constraints

Dep >> { *Dist≤0, *Dist≤1, *Dist≤2, *Dist≤3, *Dist≤4 }
(the ranking among the *Dist≤*n* constraints is irrelevant)

Consequently, a language that shows variation between epenthesis and no epenthesis for target CCV forms of all sonority distances is one in which the ranking of Dep must vary with respect to the sonority constraints—sometimes the grammar in (12) is invoked (epenthesis in even the best cluster), sometimes the grammar in (13) is invoked (no epenthesis even in the worst cluster), and sometimes Dep takes an intermediate position (epenthesis in some clusters but not in others). Concretely, this means that under the stringency approach, the ranking of Dep must vary with respect to at least *Dist≤4 (and possibly with other *Dist≤*n* constraints as well, depending on the precise pattern of variation). And, crucially, under the scale-partition approach, the ranking of Dep must vary with respect to the entire scale-partition constraint family.

Under standard Stochastic OT, the scale-partition approach therefore predicts that the members of the scale-partition family *Dist=*n* must also be able to vary with respect to each other (see (8)), leading to harmony reversals. A certain proportion of the time, clusters from lower on the sonority-distance harmony scale in (10) should actually be chosen in preference to clusters that are higher on the harmony scale.

3.3 Implications for the phonological system

From the perspective of sonority, the empirical question is this: When there is phonological variation involving more than one member of the sonority scale, are patterns of harmony reversal ever observed?

If sonority-related harmony reversals are observed under phonological variation, then this supports a model that includes standard Stochastic OT and sonority constraints as scale-partition constraint families. If, on the other hand, sonority-related harmony reversals are never observed even under phonological variation, then this supports at least one of the following conclusions:

- Sonority-related constraints are instantiated as stringency families, not scale-partition families.
- Stochastic OT must be modified so that it is not necessary for the standard deviation of the noise function for all constraints to be the same.
- Stochastic OT must be modified so that the addition of the noise component to each constraint's ranking value never causes constraints in a fixed-ranking family to alter their family-internal ordering. One implementation of this modification would be to add the same exact noise value to each member of a fixed-ranking constraint family on every evaluation, so that the relative ranking of the whole family might vary with respect to other constraints but the ranking distance between members of the family would never vary.

The remainder of this chapter first presents language examples confirming that sonority-related variation is observed in phonological patterns (§4). Then, a new empirical heuristic for distinguishing a true harmony reversal from the interference of an additional constraint in an otherwise harmony-scale-consistent pattern is presented in §5. Conclusions and implications are considered in §6.

4. Sonority-related phonological variation: Examples

The preceding discussion has shown that the empirical predictions of scale-partition and stringency constraint families are different under Stochastic OT in cases of phonological variation involving a markedness hierarchy. This section reviews a selection of case studies demonstrating that phonological variation involving multiple points on a sonority-related harmony scale does indeed occur, and therefore that the theoretical points raised in this chapter have empirical relevance.

Anttila (1997) presents an analysis of genitive plural allomorphy in Finnish according to which a markedness-hierarchy family of constraints preferring higher sonority for stressed vowels shows variation in ranking with respect to other constraints on syllable weight, stress, and sonority.

Berent et al. (2006) and Berent et al. (2009) report indirect evidence for phonological variation between epenthesis and different onset-cluster types. When listeners were exposed to clusters that were illegal in their native language, they sometimes perceived the target clusters accurately, and other times as though they had been separated by an epenthetic vowel. Such 'perceptual epenthesis' occurred at a higher rate for clusters with a less desirable sonority profile, but variability was shown at several sonority levels.

A role for sonority-related constraints has been found in first-language (L1) acquisition; for example, in determining which consonant in a cluster is retained when the cluster is reduced to a singleton (Pater & Barlow 2003; Gnanadesikan 1995, 2004). Sonority-related variation in the L1 acquisition of Dutch is described by Jongstra (2003ab); see §5 below for discussion.

Cases of sonority-related variation have been reported in studies of second-language (L2) phonological acquisition as well. For example, Petrič (2001) studied the pronunciation of German word-final clusters by 48 children, aged 11 to 13, who were learning German in school in Slovenia. For clusters consisting of a liquid, nasal, or fricative followed by a nasal, fricative, or stop, the pattern in the aggregated data is that the production error rate increased as the sonority distance falls (Petrič 2001, Table 10), although stop-fricative and stop-stop clusters were an exception to this pattern, being easier than expected.

There are a number of L2 studies investigating cases in which learners' productions show variation between the target-language realization of an onset cluster and some non-target form, generally involving epenthesis, in which the frequency of a target CC production decreases as the sonority profile of that cluster becomes more marked. For example, Cardoso (2008) presents results from a study of Brazilian Portuguese speakers learning English that examined the production of target [st], [sn], and [sl] clusters in word-initial position, and whether these clusters were produced accurately or with vowel epenthesis ([is.C]). A GoldVarb analysis indicated that [sn] and [sl] were both produced more accurately than [st] in Cardoso's learner corpus. A similar study is presented in Boudaoud & Cardoso (2009), examining the production of target [st], [sn], and [sl] clusters in the L2 English of Farsi speakers. This time, the GoldVarb results showed greater accuracy for [sl] as compared to both [sn] and [st]. Both cases are consistent with the generalization that clusters with a higher sonority distance are produced accurately more often than clusters with a lower sonority distance, indicating ranking variation between constraints in the onset-sonority distance family and the anti-epenthesis constraint Dep.

Carlisle (2006) examines the L2 English productions of [sl], [sn], and [st] clusters by Spanish-speaking learners; realizations varied between target [CC] productions and forms with epenthesis ([es.C]). When the results of all speakers were pooled, success at target [sl] was highest, followed by [sn] and then [st], once again in accordance with decreasing sonority distance in the cluster.

One additional example may be found in Broselow & Finer (1991), who present results from a study of English onset-cluster production by Japanese and Korean speakers which they interpret as showing better accuracy for clusters with a larger sonority distance. However, it is possible that other aspects of segmental markednes (such as the marginal status and non-existence of [f] in Japanese and Korean respectively) might also be a factor in their findings (in particular, error rates for most clusters involved epenthesis, [CCV] —> [C*V*CV], but errors for [fC] clusters largely concerned a featural change from [f] to [p]).

As this example from Broselow & Finer (1991) illustrates, a generally sonority-based phonological pattern can sometimes include aspects that do not follow directly from the predictions of the sonority scale. In some cases, these sonority-exception patterns really are caused by interactions with other phonological constraints or processes, as is likely true in the case of [f] in Broselow & Finer's results. Along similar lines, Pater & Barlow (2003) present an analysis in which the outcome of cluster simplification in the phonology of children learning English is generally driven by a sonority-related markedness scale, specifically, by constraints based on a harmony scale that relates onset consonants to low sonority. Some of the children's productions appear to go against a sonority-based pattern, but Pater & Barlow (2003) account for these not as cases of actual harmony reversal (requiring a reranking among the sonority constraints), but as cases where other, unrelated constraints such as *FRICATIVE interact with the sonority-based constraint family in particular ways. Likewise, Bouaoud & Cardoso (2009) consider whether their findings on cluster production in Farsi speakers' L2 English (in which [sl] clusters are more accurately produced than either [sn] or [st]) are best explained with sonority constraints, or instead with reference to the [±continuant] values of segments in clusters.

Consequently, although §3 has shown that the scale-partition approach predicts harmony reversals under variation, and the stringency approach does not, it is not a trivial problem to determine whether or not harmony reversals are actually observed. It is essential that we find a way to distinguish true cases of harmony reversal from cases of harmony scales merely interacting with additional constraints. §5 uses the properties of constraints and their noise distributions under Stochastic OT to propose a method for making this distinction.

5. Deriving empirical predictions

Stringency hierarchies exclude all possibility of harmony reversal under within- or between-speaker variation, whereas scale-partition hierarchies do not (§3, above). The stringency hypothesis would thus at first glance seem to have the virtue of easy falsifiability, since a single case of markedness reversal would refute it. However, the effects of a stringency hierarchy can be interfered with by constraints outside it in ways that could produce the appearance of a harmony reversal. We are thus faced with the problem of distinguishing actual counterexamples from spurious ones. This section of the paper describes a class of situations in which the effects of a scale-partition hierarchy can be recognized unambiguously, in the form of a transparent relationship between the frequencies of the observed variants.

5.1 Example: Cluster simplification

For a concrete example of a harmony reversal involving sonority, we consider the simplification of onset clusters from two consonants (C_1C_2) to one (C_1 or C_2) by first-language learners. The process is illustrated by data from Gita, a two-year-old American-English learner studied by Gnanadesikan (1995, 2004). Gita regularly reduced target biconsonantal onsets to the less sonorous of the two consonants; e.g., *blue* [bu], *sky* [ɡaɪ], *snow* [soʊ]. Pater and Barlow (2003) propose that the simplification is driven by highly-ranked *COMPLEXONSET, while the output consonant is chosen by a markedness scale that penalizes sonorous segments in onsets:

(14) **x*-ONS: Give one violation mark for every segment in sonority class *x* that is in an onset.

*GLIDE-ONS >> *LIQUID-ONS >> *NASAL-ONS >> *FRICATIVE-ONS

When only one consonant can surface in the onset, the *x-Ons hierarchy favors the retention of the less-sonorous one. This is shown in Tableau (15) (after Pater & Barlow 2003, Tableau 7).

(15) Retention of the less-sonorous onset in onset-cluster simplification

		*Glide-Ons	*Liquid-Ons	*Nasal-Ons	*Fricative-Ons
sky	/skaɪ/				
	[saɪ]				*!W
>	[gaɪ]				
smell	*/smɛl/*				
>	[sɛl]				*L
	[mɛl]			*!W	

Gnanadesikan (2004) does not describe variation in Gita's choice of reduction output, but the scale-partition hypothesis predicts that it is possible for a learner to show such variation. We can imagine a Gita-like grammar in which the *x-Ons constraints are ranked close enough to each other to be observed exchanging places, so that, e.g., *smell* surfaces sometimes as [sɛl] and sometimes as [mɛl], depending on whether *Fricative-Ons is sampled above or below *Nasal-Ons.

The pattern of simplification to the less sonorous segment is common across children and is well attested in Dutch as well as English learners (see Jongstra 2003a, Ch. 2, for a review). In a picture-naming study of 45 typically-developing Dutch-learning children around two years of age, Jongstra (2003ab) found that when word-initial two-consonant clusters were reduced to a single consonant, most children followed consistent reduction patterns, but there were also several cases of within-child variation (Jongstra 2003a, §4.2.2.3). Clusters of the form plosive+/l/, plosive+/r/, fricative+/r/, and fricative+plosive were reduced to the less-sonorous member by most children. Clusters of the form fricative+/l/, fricative+nasal, and /sx/ were more variable within and/or across children; e.g., /sm/ is consistently produced as [s] by five children, [m] by three, and variably as [s] or [m] by four (as in the hypothetical *smell* example above). This variation would not be predicted by an alternative grammar model in which the *x-Ons constraints were replaced by a stringency hierarchy—e.g., *Glide-Ons, *(Glide or Liquid)-Ons, *(Glide or Liquid or Nasal)-Ons, etc.—because the less-sonorous output harmonically bounds the more-sonorous one (§2.2, above).

However, that does not allow us to reject outright the hypothesis that the sonority effects are governed by a stringency hierarchy. It is clear that factors other than sonority can be involved in the choice of output. For example, many children consistently or variably reduce /sm/ and /sn/ to the more-sonorous [m] and [n]. A stringency constraint family cannot be ranked to produce those outputs, but neither can a scale-partition constraint family, since the ranking value of *Nasal-Ons cannot be smaller than that of *Fricative-Ons. Some other constraint from outside the hierarchy must be responsible, e.g., a context-free *Fricative constraint (Pater &

Barlow 2003). Such a constraint could produce the harmony-reversal effect with either kind of sonority-constraint hierarchy.

If even outright reversals of harmony cannot distinguish between the two hypotheses, what can? The next two subsections of this paper describe a characteristic signature left by interactions among scale-partition constraints in the form of a particular relationship among the frequencies of the different variants.

5.2 Ranking distance and domination probability

Consider a system of three constraints, C_0, C_1, and C_2, in a Stochastic OT grammar. Ultimately, we will want to interpret them as markedness constraints from the same harmony scale, but since the logic applies equally well to any three constraints, we will start out speaking as generally as possible. For simplicity's sake, we renumber the ranking scale so that the noise distribution has a standard deviation of 1, and we let μ_i be the ranking value of C_i. Let p_{ij} = Pr ($C_i >> C_j$), the probability that C_i will be seen to dominate C_j on any particular optimization, and suppose our language data allow us to estimate p_{10} and p_{20}. Since these probabilities tell us how far C_1 and C_2 are ranked from C_0, they also tell us how far they are from each other, and so p_{12} is predictable from them. The next bit of this paper derives a simple approximate method for making that prediction.

It is intuitively clear that if p_{10} and p_{20} are similar, then C_1 and C_2 must be ranked near each other, and so p_{12} must be about 0.5. If, on the other hand, p_{10} is much larger than p_{20}, then C_1 must be ranked well above C_2, and so p_{12} should be close to 1. This intuition can be made more precise quantitatively. The difficult step is converting back and forth between ranking values and variation frequencies. For any given selection point x, the probability that C_i is observed in an interval of width dx around x is $\phi(x-\mu_i)\mathrm{d}x$, where $\phi(x)$ is the standard normal probability-density function. The probability that C_j is observed below x is $\Phi(x-\mu_j)$, where $\Phi(x)$ is the standard normal cumulative distribution function. The probability of both events happening at once is $\Phi(x-\mu_j)\phi(x-\mu_i)\mathrm{d}x$. Summing this up for each possible selection point x, we get the following equation.

(16) $$p_{ij} = \int_{-\infty}^{\infty} \Phi(x-\mu_j)\phi(x-\mu_i)\mathrm{d}x$$

Although this equation can be solved numerically for any *specific* values of p_{ij} or of μ_i–μ_j (e.g., via a simulation using the Gradual Learning Algorithm), it is opaque to intuition and provides no help in thinking about *general* relationships between constraint ranking probabilities.

To find a more convenient approximation, we start by restricting our attention to variant frequencies whose magnitudes are typical for linguistic data, say, between 1% and 99%. Next, we convert the observed frequencies from probabilities to log-odds, where log-odds(p) = ln(p/(1–p)). As Figure 1 shows, it turns out that log-odds(p_{ij}) is approximately linear in μ_i–μ_j over that range.

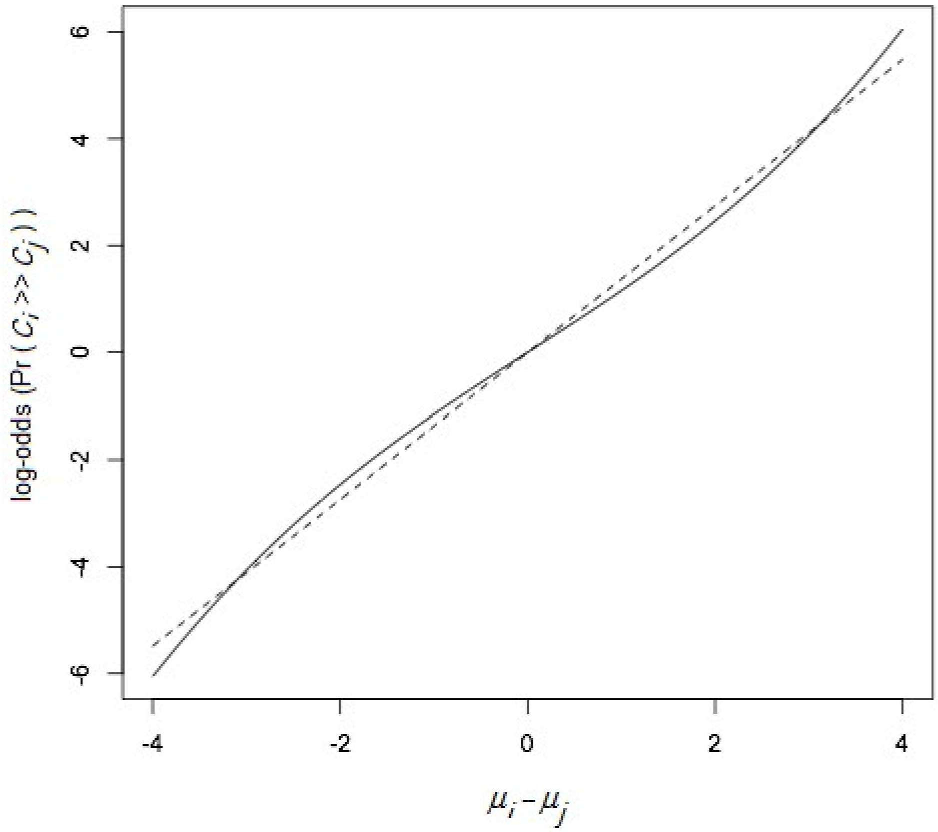

Figure 1. There is an approximately linear relationship between μ_i–μ_j and the log-odds of the probability that C_i will be observed to dominate C_j, as long as the log-odds is between about –6 and 6 (corresponding to a probability between about 1% and 99%). The dashed line, an ordinary-least-squares regression line, has slope s=1.371.

Thus, μ_i–μ_j, the difference in ranking values, is approximately a constant factor s times the log-odds of the probability of observing C_i >> C_j:

(17) $$\text{log-odds}(p_{ij}) \approx s(\mu_i - \mu_j)$$

That is, if variation probabilities are expressed as log-odds, they can be treated as distances between constraints, as if we had simply rescaled the ranking continuum using a different length unit. Consequently,

(18) $$\begin{aligned}\text{log-odds}(p_{12}) &\approx s(\mu_i - \mu_j)\\ &\approx s(\mu_1 - \mu_0) - s(\mu_2 - \mu_0)\\ &\approx \text{log-odds}(p_{10}) - \text{log-odds}(p_{20})\end{aligned}$$

In other words, because of the near-linear relationship between log-odds and ranking distance, we can (approximately) predict log-odds directly from log-odds without going through ranking values at all.

To show how much accuracy is lost in the approximation, we calculated p_{12} as p_{10} and p_{20} jointly ranged over 0.025, 0.05, 0.10, 0.15, 0.25, 0.5, 0.75, 0.85, 0.9, 0.95, and 0.975, excluding combinations for which log-odds(p_{10}) and log-odds(p_{20}) differed by more than 4. The exact and approximate p_{12} are plotted against each other in Figure 2. The largest difference in absolute terms was 0.061.

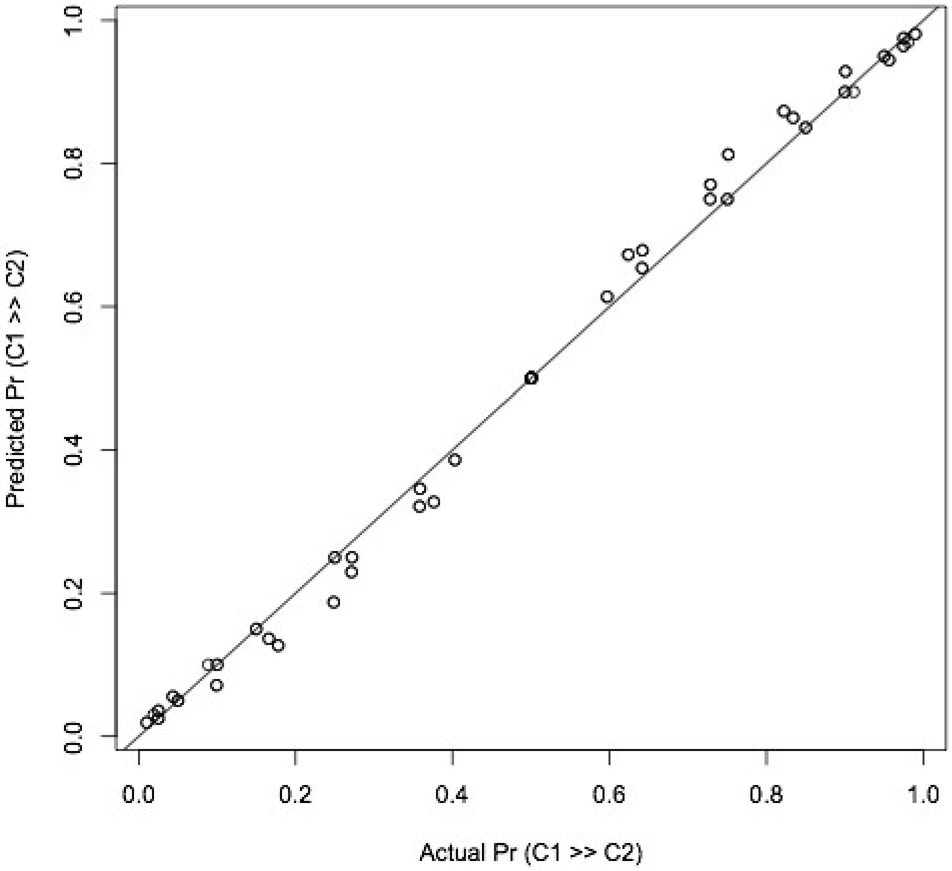

Figure 2. Approximate vs. exact p_{12} as calculated at multiple levels of p_{10} and p_{20}.

5.3. Variant frequency in scale-partition hierarchies

The foregoing is general and abstract, applying to any three constraints whatsoever as long as their relative ranking probabilities p_{ij} can be unambiguously inferred from the data. Conveniently, when the three constraints belong to a scale-partition hierarchy, there are circumstances in which the p_{ij} are not just inferrable from, but actually equal to, the variation probabilities.

A concrete example can be constructed from Pater and Barlow (2003)'s analysis of the data from Gnanadesikan (1995, 2004), shown above in Tableau (15). Since the *x-Ons constraints do not overlap with each other (there are no entailment relations between them), *Fricative-Ons and *Nasal-Ons are relevant for the [sɛl]/[mɛl] decision. Hence the probability that fricative-nasal clusters are reduced to the fricative rather than the nasal is exactly equal to Pr(*Nasal-Ons >> *Fricative-Ons), and likewise for any other pair of sonority classes. The non-overlapping property of scale-partition constraints thus means that the domination probabilities p_{ij} can be read directly off the variation probabilities in the data. Therefore, by the argument made in §5.2, the variation probabilities should stand in predictable relations to each other, e.g.,

(19) log-odds (Pr (*sk* ⇢ *s*)) ≈ log-odds (Pr (*sl* ⇢ *s*) – log-odds (Pr (*kl* ⇢ *k*))

If the empirical variation probabilities are *not* so related, then one of the hypotheses must be wrong: Either the relevant sonority constraints do not form a scale-partition hierarchy (they overlap with each other, perhaps as in a stringency hierarchy, so that the variation probabilities are not equal to the domination probabilities), or other constraints are also involved in the choice between sonority classes. On the other hand, if the predicted relationship does hold, it would be strong evidence in favor of a scale-partition hierarchy.

We have not succeeded in finding any published data sets which conform to the scale-partition variation predictions. Only a few have the relevant quantitative data in any case (e.g., Tropf 1987, Ohala 1999, Hansen 2001, Jongstra 2003a, b), so not too much should be made of this failure yet. To illustrate how the predictions are tested, we applied the model sketched above to a subset of the Jongstra (2003ab) data. We focus here on the three clusters [sk sn kn], and on the ten children who reduced each of those clusters (C_1C_2) to a single consonant at least eight times in the sample (that being the author's criterion of frequent attestation). Table 1 shows the rate of reduction to C_1, or to a consonant in the same sonority class, as a proportion of all reductions to a single consonant.

	Target cluster		
Child	sk	sn	kn
3	0.08	1.00	0.86
4	0.00	0.97	0.69
5	0.00	0.96	0.95
6	0.00	0.85	1.00
13	0.00	0.75	0.91
14	0.13	0.94	0.86
15	0.05	1.00	1.00
23	0.00	0.92	1.00
28	0.08	1.00	1.00
34	0.00	1.00	1.00
MEAN	0.034	0.939	0.927

Table 1. Reductions of [sk sn kn] to a segment in the same sonority class as one of the onset consonants, showing the proportion where the class of the output consonant was the same as that of the first target consonant rather than the second (Jongstra 2003a, Table 5.2b). Each proportion is based on at least eight observations.

All ten of these children preferentially reduce [sk] to a stop and [sn] to a fricative; i.e., they choose the least-sonorous segment, just as Gita did. Since stops are preferred to fricatives, and fricatives to nasals, we expect [kn] should be reduced to a stop, as indeed it is. If the choice is determined entirely by a scale-partition constraint family like the **x*-ONS constraints, then we would expect the preference for stops over nasals to be even greater than that for stops over fricatives or that for fricatives over nasals; indeed, expressed as log-odds, it should be approximately equal to their sum.

However, this is not the case. Child 14, for example, prefers stops over fricatives 87% of the time (a log-odds of 1.90) and fricatives over nasals 94% of the time (2.75). The **x*-ONS hypothesis predicts that he or she should prefer stops over nasals 99% of the time (4.65 = 1.90 + 2.75), but the observed rate is only 86% (1.82), which is numerically *less* than the rate of preferring stops to fricatives or fricatives to nasals. If we assume that all of the children have the same constraint ranking values, and combine their data (by averaging with equal weight), the same pattern occurs. Stops are preferred to fricatives for [sk] 96.6% of the time (log-odds of 3.35), and fricatives to nasals for [sn] 93.9% of the time (2.73), which predicts that stops should be preferred to nasals for [kn] 99.7% of the time (6.08 = 3.35 + 2.73). The actual rate is 92.7% (2.54), less than either of the other two preferences. There is thus no way to assign ranking values to the constraints in Tableau (15) that will match the observed frequencies.[4]

This section has identified a clear empirical signature of a scale-partition hierarchy, which crucially depends on the lack of overlap between the constraints which are in variation. Deviation from the predicted relationship indicates overlap, either between constraints in the hierarchy itself, or between constraints inside and outside the hierarchy. Cases which conform to the relationship may be rare (since there are many outside constraints that could interfere), but would strongly support the scale-partition hypothesis if found.

6. Conclusions and implications

In this chapter, we have shown that variation provides a new empirical domain for comparing the two competing approaches to markedness hierarchies in a constraint-based model. The scale-partition approach predicts harmony reversals, while the stringency approach does not. Further, we have shown that the pattern of harmony reversals predicted by the scale-partition approach can be empirically distinguished from superficially similar patterns caused by interactions of markedness-hierarchy constraints with other, unrelated constraints.

The ideal data set for distinguishing the two hypotheses would describe a sonority-sensitve process involving at least three distinct sonority classes. It would provide individual-level data, so that within-speaker variation could be separated from between-speaker variation. If an acquisition study, it would break the data down further by recording session, to distinguish variation from change over time. Finally, it would have enough tokens in each cell to allow reasonably precise probability estimates. All this could prove a tall order, since pinning down even a single variant frequency to a 95% confidence interval of ±0.1 could take up to 100 observations (Tortora 1978), and there would need to be at least three cells per speaker. Data sets of this size may become more common as technology improves.

---

[4] We caution against making too much of this particular example, since the number of data points per child does not allow meaningful comparision between such close proportions, and there are other complications such as changes in children's productions over the five-month course of the study.

It is worth noting that the general predictions we have made about the differences between scale-partition constraint families and stringency constraint families are independent of whether a constraint-based phonological framework is implemented as Optimality Theory (Prince & Smolensky 1993, 2004), in which higher-ranked constraints strictly dominate lower-ranked constraints, or as Harmonic Grammar (Legendre et al. 1990; Smolensky and Legendre 2006), in which constraints are weighted rather than strictly ranked and the effects of violations of different constraints are additive. The same predictions are made under HG as under OT because even in HG the scale-partition constraints will not show additive effects, as their violation profiles are completely independent of one another. As for the stringency constraints, they will show additive effects under HG, and this would likely affect their overall position with respect to the entire constraint hierarchy in a given language, but it does not change the fact that stringency constraints rule out harmony reversals altogether.

The results of this chapter have implications beyond sonority, and in fact beyond phonology. The use of markedess hierarchies, and of constraint families based on harmony scales, is a technique that has been applied in morphosyntax as well. Moreover, analyses involving just the crucial scenario we have identified here, where there is variation in the ranking of some constraint with respect to multiple members of a harmony-scale constraint family, have been proposed by, for example, Aissen (2003) and Lee (2006). However, the implications for harmony reversals have not generally been explored, beyond a brief remark by Dingare (2001: 8) acknowledging that Stochastic OT might allow for the selection points for constraints in a markedness hierarchy to end up in reverse order from their usual harmony scale. Thus, the predictions we identify and questions we raise may be fruitfully pursued both within and beyond phonology.

References

Aissen, Judith. 1999. Markedness and subject choice in Optimality Theory. *Natural Language and Linguistic Theory* 17: 673–711.

Aissen, Judith. 2003. Differential object marking: Iconicity vs. economy. *Natural Language and Linguistic Theory* 21: 435–483.

Anttila, Arto. 1997. Deriving variation from grammar. In Frans Hinskens, Roeland van Hout, & Leo Wetzels (eds.), *Variation, change, and phonological theory,* 35–68. Amsterdam: John Benjamins.

Baertsch, Karen. 1998. Onset sonority distance constraints through local conjunction. In M. Catherine Gruber, Derrick Higgins, Kenneth S. Olson, and Tamra Wysocki (eds.), *CLS 34, Part 2: Papers from the Panels,* 1–15. Chicago: CLS.

Berent, Iris, Donca Steriade, Tracy Lennertz, & Vered Vaknin. 2007. What we know about what we have never heard: Evidence from perceptual illusions. *Cognition* 104: 591-630.

Berent, Iris, Tracy Lennertz, Paul Smolensky, & Vered Vaknin-Nusbaum. 2009. Listeners' knowledge of phonological universals: Evidence from nasal clusters. *Phonology* 26: 75–108.

Boersma, Paul. 1998. *Functional Phonology: Formalizing the interactions between articulatory and perceptual drives.* The Hague: Holland Academic Graphics.

Boersma, Paul, & Bruce Hayes. 2001. Empirical tests of the Gradual Learning Algorithm. *Linguistic Inquiry* 32: 45-86.

Boudaoud, Malek, & Walcir Cardoso. 2009. Vocalic [e] epenthesis and variation in Farsi-English interlanguage speech. *Concordia Working Papers in Applied Linguistics* 2: 1–34.

Broselow, Ellen, & Daniel Finer. 1991. Parameter setting in second language phonology and syntax. *Second Language Research* 7: 35–59.

Cardoso, Walcir. 2008. The development of sC onset clusters in interlanguage: Markedness vs. frequency effects. In Roumyana Slabakova, Jason Rothman, Paula Kempchinsky, & Elena Gavruseva (eds.), *Proceedings of the 9th Generative Approaches to Second Language Acquisition Conference (GASLA 2007)*, 15–29. Somerville, MA: Cascadilla Proceedings Project.

Carlisle, Robert. 2006. The sonority cycle and the acquisition of complex onsets. In Barbara O. Baptista & Michael Watkins (eds.), *English with a Latin beat: Studies in Portuguese/Spanish-English interphonology,* 105–138. Amsterdam: John Benjamins.

Croft, William. 1988. Agreement vs. case marking and direct objects. In Michael Barlow & Charles A. Ferguson (eds.), *Agreement in natural language: Approaches, theories, descriptions,* 159–179. Stanford, CA: CSLI.

de Lacy, Paul. 2002. *The formal expression of markedness.* PhD dissertation, University of Massachusetts, Amherst.

de Lacy, Paul. 2004. Markedness conflation in Optimality Theory. *Phonology* 21: 145–199.

de Lacy, Paul. 2006. *Markedness: Reduction and preservation in phonology.* Cambridge: Cambridge University Press.

Dell, François, and Mohamed Elmedlaoui. 1985. Syllabic consonants and syllabification in Imdlawn Tashlhiyt Berber. *Journal of African Languages and Linguistics* 7:105–130.

Dingare, Shipra. 2001. *The effect of feature hierarchies on frequencies of passivization in English.* MA thesis, Stanford University.

Dixon, R. M. W. 1979. Ergativity. *Language* 55: 59–138.

Gnanadesikan, Amalia. 1995. Markedness and faithfulness constraints in child phonology. Ms., University of Massachusetts, Amherst. ROA #(???).

Gnanadesikan, Amalia. 2004. Markedness and faithfulness constraints in child phonology. In René Kager, Joe Pater, & Wim Zonneveld (eds), Constraints in phonological acquisition, 73–108. Cambridge: Cambridge University Press.

Gouskova, Maria. 2004. Relational hierarchies in Optimality Theory: The case of syllable contact. *Phonology* 21: 201–250.

Hansen, Jette G. 2001. Linguistic constraints on the acquisition English syllable codas by native speakers of Mandarin Chinese. Applied Linguistics 22: 338–365.

Jongstra, Wenckje. 2003a. *Variation in reduction strategies in Dutch word-initial consonant clusters.* PhD dissertation, University of Toronto.

Jongstra, Wenckje. 2003b. Variable and stable clusters: Variation in the realisation of consonant clusters. The Canadian Journal of Linguistics 48: 265–288

Keenan, Edward, & Bernard Comrie. 1977. Noun phrase accessibility and universal grammar. *Linguistic Inquiry* 8: 63–99.

Kenstowicz, Michael. 1994. Sonority-driven stress. Ms., MIT. ROA #33.

Lee, Hanjung. 2006. Parallel optimization in case systems: Evidence from case ellipsis in Korean. *Journal of East Asian Linguistics* 15: 69–96.

Legendre, Géraldine, Yoshiro Miyata, & Paul Smolensky. 1990. Harmonic Grammar—a formal multilevel connectionist theory of linguistic wellformedness: An application. In *Proceedings of the Twelfth Annual conference of the Cognitive Science Society,* 884–891. Hillsdale, NJ: Lawrence Erlbaum.

McCarthy, John, & Alan Prince. 1995. Faithfulness and reduplicative identity. In Jill N. Beckman, Laura Walsh Dickey, & Suzanne Urbanczyk (eds.), *Papers in Optimality Theory.* University of Massachusetts Occasional Papers 18, 250–384. Amherst, MA: GLSA.

Murray, Robert W., and Theo Venneman. 1983. Sound change and syllable structure in Germanic phonology. *Language* 59: 514–528.

Ohala, Diane. 1999. The influence of sonority on children's cluster reductions. *Journal of Communication Disorders* 32: 397–442.

Parker, Steve. This volume. Sonority distance vs. sonority dispersion: A typological survey.

Pater, Joe, & Jessica Barlow. 2003. Constraint conflict in cluster reduction. *Journal of Child Language* 30: 487–526.

Petrič, Teodor. 2001. Acquisition of marked consonant clusters in German as a foreign language. *Poznań Studies in Contemporary Linguistics* 37: 157 – 186.

Prince, Alan, & Paul Smolensky. 1993. *Optimality Theory: Constraint interaction in generative grammar.* Ms., Rutgers University and University of Colorado, Boulder.

Prince, Alan, & Paul Smolensky. 2004. *Optimality Theory: Constraint interaction in generative grammar.* Oxford: Blackwell.

Selkirk, Elisabeth O. 1982. The syllable. In Harry van der Hulst & Norval Smith (eds.), *The Structure of Phonological Representations, vol. II,* 337–383. Dordrecht: Foris.

Silverstein, Michael. 1976. Hierarchy of features and ergativity. In R. M. W. Dixon (ed.), *Grammatical Categories in Australian Languages,* 112–171. Canberra: Australian Institute of Aboriginal Studies.

Smolensky, Paul, & Géraldine Legendre. 2006. *The harmonic mind: From neural computation to optimality-theoretic grammar.* Cambridge, MA: MIT Press.

Steriade, Donca. 1982. *Greek prosodies and the nature of syllabification.* PhD dissertation, MIT.

Tortora, Robert D. 1978. A note on sample size estimation for multinomial proportions. *American Statistician* 32: 100–102.

Tropf, Herbert S. 1987. Sonority as a variability factor in second language phonology. In Allan R. James & Jonathan Leather (eds.), *Sound patterns in second language acquisition,* 173–191. Dordrecht: Foris.

Zec, Draga. 1995. Sonority constraints on syllable structure. *Phonology* 12:85-129.