

Sonority variation in Stochastic Optimality Theory: Implications for markedness hierarchies

Jennifer L. Smith and Elliott Moreton

Abstract. Constraints based on the sonority scale and other markedness hierarchies have been formalized in two different ways: as *scale-partition* constraint families, with a universally fixed ranking, and as *stringency* constraint families, where the constraints stand in subset relations and their ranking is not universal. We identify a new empirical domain where these two approaches make different predictions: within-speaker variation. Namely, in the framework of Stochastic Optimality Theory (Boersma and Hayes 2001), the scale-partition approach predicts *harmony reversals* in the presence of variation (i.e., sometimes a more marked form should be chosen in preference to a less marked form), while the stringency approach does not. We further demonstrate that the pattern of harmony reversals predicted by the scale-partition approach can be numerically quantified, based on the constraints' ranking values; this allows true harmony reversals to be efficiently distinguished from cases where an unrelated constraint causes apparent harmony-reversal behavior. Known cases of variation involving markedness hierarchies have not been described at the level of detail necessary to distinguish true harmony reversals from apparent ones, but we have identified empirical conditions that would strongly support the scale-partition approach to markedness hierarchies if they were to be uncovered in future work.

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 (Dell and Elmedlaoui 1985) and for low-sonority syllable onsets (Steriade 1982; Zec 1988). Sonority-related effects have been studied for many years (Sievers 1876/1893), and have influenced the study of markedness hierarchies in general (Aissen 1999).

In Optimality Theory (OT), markedness hierarchies have been formalized in one of two ways: as a *scale-partition constraint family* (Prince and 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 ordered in a universally fixed ranking, and as a *stringency constraint family* (Prince 1997, 1999; de Lacy 2002, 2004, 2006), in which constraints are formalized as increasingly larger 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 where the two approaches make distinct predictions, but where the implications for markedness hierarchies have remained largely unexplored: speaker-internal phonological variation. Specifically, we examine the role in variation patterns of *harmony reversals* – the selection of a less harmonic or less desirable form in preference to a more harmonic one, as seen schematically in the mapping of the form /plapna/ to [pV.la.pna]. Here, the potential onset cluster [pl] has been avoided through epenthesis, but the cluster [pn], which is more highly marked because of its smaller sonority distance, has surfaced faithfully.

We show that in the framework of Stochastic OT (Boersma and Hayes 2001), formalizing sonority effects (or other types of markedness hierarchies) as scale-partition constraints predicts the existence of harmony reversals under phonological variation, while the stringency approach makes no such prediction (sections 2–3). We confirm that patterns of variation involving multiple levels of the sonority scale exist, so the empirical scenario of interest is attested and theoretically relevant (section 4). However, it is difficult to test whether or not any particular case of sonority-related variation really does show harmony reversals, because interference from other constraints might lead to a pattern that coincidentally resembles a harmony-reversal pattern. 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 (section 5).

As our introductory discussion in this section illustrates, we take *sonority* to be a phonological property, to which formal phonological constraints can make reference, rather than a primarily phonetic property (although it may be ultimately based on or grounded in phonetic factors; see Parker 2002 for an extensive review). In addition, we view sonority in terms of a scale, to which, again, the phonological grammar can refer. However, the arguments we make in this chapter do not crucially depend on whether sonority is formalized as a single multivalued feature (de Lacy 2006), or whether instead the sonority scale is derived from combinations of values of binary major-class features such as [\pm vocalic, \pm approximant, \pm sonorant] (Clements 1990). Moreover, 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 will 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 constraints and predicted patterns of speaker-internal 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 and 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 sonority, place of articulation, animacy, or definiteness – 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 and Elmedlaoui 1985), stressed syllable peaks (Kenstowicz 1997), and moraic segments (Zec 1995; Gnanadesikan 2004) each show a preference for segments of the highest possible sonority level; that onsets show a preference for segments of the lowest possible sonority level (Steriade 1982); that onset clusters show a preference for rising and maximally dispersed sonority (Selkirk 1982; Clements 1990; Baertsch 1998; see also Parker, this volume); and that with respect to syllable contact, adjacent coda-onset sequences show a preference for falling and maximally dispersed sonority (Murray and Venneman 1983; Gouskova 2004).

As a simplified illustration of a markedness hierarchy involving sonority, consider the preference for higher-sonority syllable peaks. Vowel height cate-

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

gories have the following sonority relationship² (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 corresponds to a *harmony scale* (Prince and Smolensky 1993/2004: section 8.1), 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 (section 2.1) and the stringency approach (section 2.2).

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 and Smolensky 1993/2004: section 8.1).³ 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

2. This fragment of the sonority scale is presented in order to facilitate a concrete discussion of the formalization of sonority-related constraints; for simplicity, no central (reduced) vowels and no consonants are considered here. See Kenstowicz (1996), Parker (2002), and de Lacy (2002, 2004, 2006) concerning the place of central vowels on the sonority scale; consonants have lower sonority than all vowels.

3. *Constraint alignment*, the operation which derives a family of universally ranked constraints from a markedness scale as discussed here, is distinct from *Alignment constraints* (McCarthy and Prince 1993a), which are phonological constraints requiring particular morphosyntactic and phonological constituents to align their edges in output representations.

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

(4) Violations assigned by *PEAK/X constraints⁴

	*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 be predicted to vary widely 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.⁵ Therefore, Prince and Smolensky (1993/2004: section 8.1) explicitly propose that constraint families that are derived from markedness 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 least-preferred end of the scale. In other words, 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

4. This is not a standard tableau, since no input is shown. This display simply shows how violation marks would be assigned to the given outputs by the constraints under discussion.

5. Sonority is not the only factor that influences which vowel is the least marked syllable peak in a given language; other constraints (such as those for featural markedness) are relevant as well. But it is clear that typological facts about default syllable nuclei do not support a constraint family such as that in (3) with free reranking.

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 scale-partition 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; lower-sonority peaks are harmonically bounded (with respect to the constraints in the stringency family). In other words, no single constraint of the *PEAK/≤X constraint family ever prefers [i] over [e] as a syllable peak, because any *PEAK/≤X constraint that assigns a violation to [e] *necessarily* assigns a violation to the lower-sonority [i] as well.

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 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. As we discuss in section 5, variation data described at the level of detail necessary to distinguish between the two approaches is not yet available. However, we identify empirical conditions that would provide strong support for the scale-partition approach if found.

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 1998b; Boersma and 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 value corresponds to a higher ranking, and so domination relations between constraints can be represented numerically.

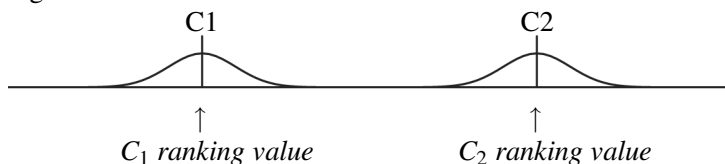
Stochastic OT models variation by adding an element of unpredictability to the ranking relationships between constraints. In the grammar of a particular language, each constraint has an intrinsic and consistent *ranking value*. However, every time an input is mapped to an output by the grammar, the ranking value for each constraint is perturbed by a “noise” component (stated informally, for each constraint in the grammar, a small amount is added to or subtracted from the ranking value), resulting in a value known as the *selection point* for the constraint in question. 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; each constraint’s selection point will be within three standard deviations (± 6.0 units) of the ranking value in more than 99% of all input-output mappings. Boersma and 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 section 3.2 below.

Because the value of the noise component varies for each constraint each time the grammar is invoked, so does the constraint's selection point. Since it is the *relative* ordering of two constraints that determines which constraint dominates the other, then when two constraints have ranking values that are very close together, the domination relation between them can actually vary, depending on whether their selection points on a given input-output mapping are higher or lower than their intrinsic ranking values.

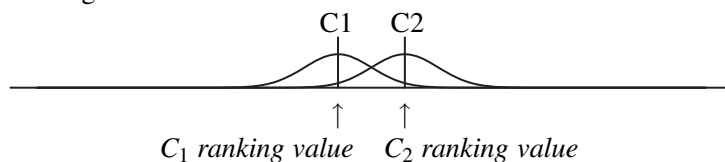
Specifically, if constraint C_1 has a ranking value that is higher than that of constraint C_2 by five standard deviations (here, 10.0 units) or more, then the relative ranking $C_1 \gg C_2$ is stable, because the probability of C_1 having a selection point that is lower than C_2 's selection point on a given input-output mapping is vanishingly small. This scenario is illustrated in (7), where ranking values and noise distribution curves are shown for two constraints whose ranking values are far apart.

- (7) Ranking values and noise distributions for two constraints whose ranking is stable

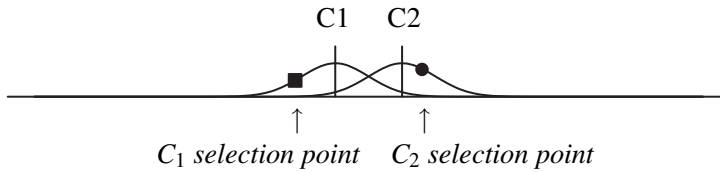


On the other hand, if two constraints C_1, C_2 have ranking values that are close together, then the relative ordering of their selection points will vary, as shown in (8). Such a grammar does in essence make use of two different rankings because on some evaluations, $C_1 \gg C_2$, as in (8)(b), but on others, $C_2 \gg C_1$, as in (8)(c). (The probability of occurrence of $C_1 \gg C_2$ versus $C_2 \gg C_1$ depends on their ranking values, a point whose implications for markedness scales will be explored in detail in section 5.)

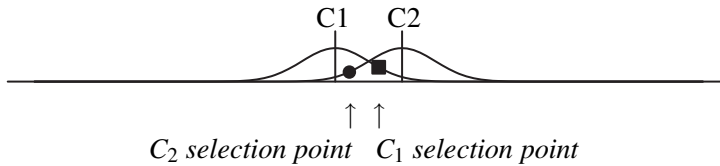
- (8) Variable constraint ranking in Stochastic OT
 (a) Ranking values and noise distributions for two constraints whose ranking is variable



(b) Example of selection points on an evaluation where $C_1 \gg C_2$



(c) Example of selection points on an evaluation where $C_2 \gg C_1$



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

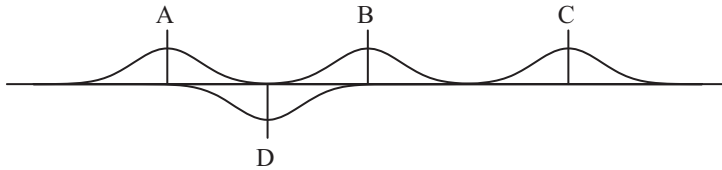
Under the assumption (Boersma and Hayes 2001) that the standard deviation of the noise distribution is the same for all constraints, the Stochastic OT model places restrictions on possible patterns of variation (a point discussed by Anttila 2007: 534 as well).

[...] 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 and Hayes 2001: 50)

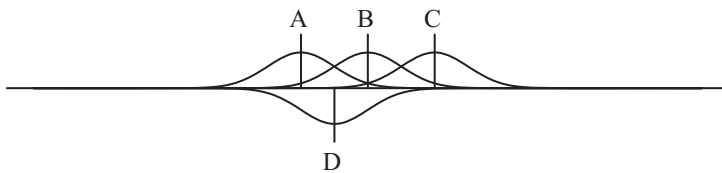
In other words, given a ranked set of constraints $A \gg B \gg C$, and variation in the relative ranking between these constraints and a fourth constraint D , scenarios (9)(a–b) are possible, but (9)(c) is not.

(9) Variation scenarios for one constraint (*D*) versus ranked constraints ($A \gg B \gg C$)

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



- (b) Possible:
- *D* varies with respect to each of *A*, *B*, *C*
 - The relative ranking $A \gg B \gg C$ also shows variation



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

As noted in section 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 Boersma and Hayes (2001) implementation of Stochastic OT, this difference leads to a difference in predicted phonological patterns. Assume a situation as in (9), 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 demonstrated in section 2.2 above. As a consequence, harmony reversals should never be observed, even in cases of variation as described here, because no constraint in a stringency family ever assigns a violation to a less-marked point on the scale without necessarily assigning a violation to all more-marked points at the same time.

This difference can be illustrated with a schematic example involving one sonority-based markedness scale, onset-sonority distance. (See section 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 (CV.CV or VC.CV, where *V* indicates an epenthetic vowel).⁶ 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 and 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; Clements 1990; Zec 2007). For example, consider the simplified consonant sonority scale in (10).

- (10) 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].⁷ 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.

6. Epenthesis is, of course, not the only possible way to avoid onset clusters, but it is used here in this schematic example for concreteness. Other types of faithfulness violations are discussed in section 4 below.

7. Exact numerical values for sonority distance will depend on the precise version of the sonority scale adopted. The scale shown in (10) is intended as a concrete illustration for use in the discussion, not a substantive claim about the exact structure of the sonority scale. Note also that this approach would model onset clusters with falling sonority as having a negative sonority distance, correctly predicting that they would be even more marked than a cluster with distance 0.

- (11) 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.

- (12) 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

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.

- (13) Ranking for a language with epenthesis into all potential onset clusters
 (a) With scale-partition constraints: DEP ranked below lowest constraint in scale⁸
 *DIST=0 >> *DIST=1 >> *DIST=2 >> *DIST=3 >> *DIST=4 >>

DEP

/pja/	*DIST=0	*DIST=1	*DIST=2	*DIST=3	*DIST=4	DEP
a. pja					*!w	L
→ b. pVja						*

- (b) With stringency constraints: DEP ranked below most stringent constraint
 *DIST≤4 >> **DEP** (the other *DIST ≤ n constraints can be ranked anywhere)

/pja/	*DIST≤4	DEP	*DIST≤2	*DIST≤0	*DIST≤3	*DIST≤1
a. pja	*!w	L				
→ b. pVja		*				

Conversely, in a language that allows all potential CC onset clusters to surface and never shows epenthesis, DEP is ranked above *all* sonority-distance

8. The tableau format used here and below is a “combination tableau” (McCarthy 2008), which indicates for each winner-loser pair the constraints that would favor the winner (W) and those that would favor the loser (L).

constraints, so that epenthesis is never chosen no matter how close the sonority distance in the onset cluster.

(14) 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

/pta/	DEP	*DIST=0	*DIST=1	*DIST=2	*DIST=3	*DIST=4
→ a. pta		*				
b. pVta	*!w	L				

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

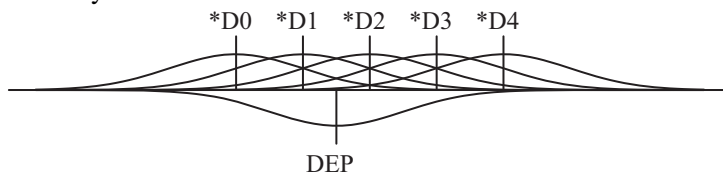
DEP >> { *DIST≤0, *DIST≤1, *DIST≤2, *DIST≤3, *DIST≤4 }
(the ranking among the *DIST ≤ n constraints is irrelevant)

/pta/	DEP	*DIST≤4	*DIST≤2	*DIST≤0	*DIST≤3	*DIST≤1
→ a. pta		* *	*	*	*	*
b. pVta	*!w	L L	L	L	L	L

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 (13) is invoked (epenthesis in even the best cluster), sometimes the grammar in (14) 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, as shown in (15)

(15) Variation between DEP and all members of the scale-partition *DIST = n family



Under standard Stochastic OT (Boersma and Hayes 2001), 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 (9) and the discussion thereof), leading to harmony reversals. A certain proportion of the time, clusters from lower on the sonority-distance harmony scale in (11) should actually be chosen in preference to clusters that are higher on the harmony scale, meaning that a form such as /plapna/ might surface as [pV.la.pna], with epenthesis into the more harmonic cluster /pl/ but not into the less harmonic cluster /pn/.

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 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.
- 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. For example, one approach would be to give any constraints that belong to a scale-partition family a noise distribution with a much smaller standard deviation – a “narrower curve” – than constraints that do not belong to such a family. A more radical modification would be a Stochastic OT model in which every constraint has a potentially distinct noise distribution, allowing some constraints to vary in their selection point much more widely than others; in such a model, what must be learned for each constraint in the course of language acquisition would be not only its relative ranking value as compared to other constraints, but the standard deviation of its noise distribution as well.

The remainder of this chapter first presents language examples confirming that sonority-related variation is observed in phonological patterns (section 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 section 5. Conclusions and implications are considered in section 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 exist, and therefore that the theoretical points raised in this chapter have empirical relevance. For reasons discussed below and in section 5, these case studies are not described in the literature in enough detail for us to determine whether harmony reversals are actually observed. The goal of section 4 is specifically to confirm that the right types of phonological variation occur such that it is possible, in principle, to look for harmony reversals among them.

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.⁹

Berent et al. (2007) and Berent et al. (2009) report indirect evidence for phonological variation between epenthesis and the faithful realization of 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 and Barlow 2003; Gnanadesikan 2004). Sonority-related *variation* in the L1 acquisition of Dutch is described by Jongstra (2003a, 2003b); see section 5 below for discussion.

9. We strongly suspect that additional cases of sonority-related variation occur in adult L1 phonology as well, but examples discussed in detail are difficult to find.

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 consonant clusters by 48 children, aged 11 to 13, who were learning German as a second language in school in Slovenia. For German-legal 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 increases as the sonority distance falls (Petrič 2001, Table 8, summarized in our Table 1 below), although stop-fricative and stop-stop clusters were an exception to this pattern, being easier than expected.

Table 1. Proportion of target clusters successfully produced by L2 German learners in the study of Petrič (2001, Table 8)

C1	C2		
	Nasal	Fricative	Stop
Liquid	0.62	0.94	0.96
Nasal	–	0.84	0.97
Fricative	–	0.78	0.92
Stop	–	0.92	1.00

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 variable-rules analysis in Goldvarb¹⁰ indicated that accurate CC clusters were more likely for [sn] and [sl] than for [st] in Cardoso's learner corpus. A similar study is presented in Boudaoud and 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 CC 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 be-

10. Goldvarb is multivariate analysis software for performing a variable-rules analysis based on the methodology first developed in Cedergren and Sankoff (1974). The program outputs factor weights, which indicate the degree to which each linguistic or sociolinguistic factor in the analysis (here, onset cluster type) predicts the appearance of a variable pattern (here, accurate [CC] realization).

tween constraints in the onset-sonority distance family and the anti-epenthesis constraint DEP.

Carlisle (2006) examines the L2 English productions of initial [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 and 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 markedness (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, $CCV \rightarrow CVCV$, but errors for [fC] clusters largely concerned a featural change from [f] to [p]).

As this example from Broselow and 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 and Finer's results. Along similar lines, Pater and Barlow (2003) present an analysis in which the outcome of cluster simplification in the phonology of children acquiring L1 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, as when a word with a target [sn] onset cluster is realized as [n], rather than [s], even though [n] has higher sonority. However, Pater and 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 and 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 [\pm continuant] values of segments in clusters, such that a cluster whose members are both [+continuant] ([sl]) is preferred over clusters with a mismatch in [\pm continuant] values ([sn], [st]).

Consequently, although section 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 harmony reversals are actually observed. In order to compare these two competing approaches to markedness scales, 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. Section 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 (section 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. For example, the hierarchy in (12) predicts that no language should favor epenthesis into a stop-liquid onset (distance = 3) more than epenthesis into a fricative-liquid onset (distance = 2). English seems to counterexemplify this prediction by permitting syllables to begin with /sl/ but not with /tl/. The apparent markedness reversal is really due to a constraint outside the sonority-distance hierarchy; several such plausible constraints are discussed by Bradley (2006).

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 that is highly unlikely to arise by a chance conspiracy of unrelated constraints.

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 (2004). Gita regularly reduced target biconsonantal onsets to the less sonorous of the two consonants; e.g., *blue* [bu], *sky* [kaɪ], *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:

- (16) **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-ONS¹¹

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 (17) (after Pater and Barlow 2003, Tableau 7).¹²

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

	*GLIDE-ONS	*LIQUID-ONS	*NASAL-ONS	*FRICATIVE-ONS
<i>sky</i> /skaɪ/				
a. [saɪ]				*! _W
→ b. [kaɪ]				
<i>smell</i> /smɛl/				
→ a. [sɛl]				*
b. [mɛl]			*! _W	<i>L</i>

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 (though we know of no concrete examples) 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 2003b, Ch. 2, for a review). In a picture-naming study of 45 typically-developing Dutch-learning children around two years of age, Jongstra (2003a, 2003b) 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 2003b: section 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

11. Pater and Barlow (2003) assume an approach to markedness scales that does not include a constraint against the least marked member of the scale; here, stop onsets.

12. Deletion in (17) is motivated by the ranking *COMPLEXONSET ≫ MAX (Pater and Barlow 2003). These two constraints are not shown.

across children; e.g., /sm/ is produced consistently as [s] by five children, consistently as [m] by three, and variably as [s] or [m] by four (as in the hypothetical *smell* example above). This pattern of variation is consistent with the predictions of a scale-partition constraint family whose members are ranked close enough together to change places, as described 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 (section 2.2, above, and (18) below).¹³

- (18) A failed alternative using a stringency hierarchy; the output is the same no matter how the ranking of the constraints varies

	*GLIDE-ONS	*(GLIDE OR LIQUID)-ONS	*(GLIDE, LIQUID, OR NASAL)-ONS	*(GLIDE, LIQUID NASAL, OR FRICATIVE)-ONS
<i>sky</i> /skar/				
a. [sar]				*!w
→ b. [kar]				
<i>smell</i> /smɛl/				
→ a. [sɛl]				*
b. [mɛl]			*!w	*

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 usually reduce /sm/ and /sn/ to the more-sonorous [m] and [n] (Pater and Barlow 2003). 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 re-

13. Following most first-language acquisition work in Optimality Theory, we adopt the hypothesis that adult and child grammars differ only in constraint ranking, and our conclusions are only valid under that assumption. It has been clear from the start that constraints must to some extent be learned (see, e.g. the discussion of Lardil FREE-V in Prince and Smolensky 1993/2004: section 7.1), and hence that the constraint set may differ between adults and children, or even between children (Pater and Werle 2003). It may be possible to construct alternative explanations based on maturation or constraint learning for any of the phenomena discussed in this paper.

sponsible, e.g., a context-free *FRICATIVE constraint (Pater and Barlow 2003).¹⁴ Such a constraint could produce the apparent harmony-reversal effect with either kind of sonority-constraint hierarchy. An apparent harmony reversal therefore cannot be taken as evidence against the stringency hypothesis unless we are quite certain that no interfering constraint, known or unknown, exists outside the hierarchy.

If even outright reversals of harmony cannot necessarily distinguish the 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: If the odds of preferring *A* to *C* are o_{AC} to 1, the odds of preferring *B* to *C* are o_{BC} to 1, and the odds of preferring *A* to *B* are o_{AB} to 1, then $\log o_{AB}$ should be approximately equal to $\log o_{AC} - \log o_{BC}$. This relationship is very specific and unlikely to arise from a chance arrangement of other constraints, even unknown ones; if observed, it is good evidence for a scale-partition hierarchy. (Failure to observe it may be due to interference from other constraints and so does not have an unambiguous interpretation.)

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 $p_{ij} = \Pr(C_i \gg C_j)$, 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 $\varphi(x - \mu_i)dx$, where $\varphi(x)$ is the standard

14. For justification of *FRICATIVE on typological and developmental grounds, see Pater and Barlow (2003).

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 $\Phi(x - \mu_j)\phi(x - \mu_i)dx$. Summing this up for each possible selection point x , we get the following equation.

$$(19) \quad p_{ij} = \int_{-\infty}^{\infty} \Phi(x - \mu_i)\phi(x - \mu_j)dx$$

Although this equation can be solved numerically for any *specific* values of p_{ij} or of $\mu_i - \mu_j$ (e.g., via a simulation using the Gradual Learning Algorithm of Boersma 1998b, Boersma and Hayes 2001), it is opaque to intuition (i.e., it does not suggest an interpretation in terms of linguistically-relevant concepts) 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 between 1% and 99%. Next, we convert the observed frequencies from probabilities to log-odds, where $\text{log-odds}(p) = \ln(p/(1-p))$. As Figure 1 shows, it turns out that $\text{log-odds}(p_{ij})$ is approximately linear in $\mu_i - \mu_j$ over the range from 1% to 99%. Thus, $\mu_i - \mu_j$, the difference in ranking values, is approximately a constant factor s times the log-odds of the probability of observing $C_i \gg C_j$:

$$(20) \quad \text{log-odds}(p_{ij}) \approx s(\mu_i - \mu_j)$$

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,

$$(21) \quad \begin{aligned} \text{log-odds}(p_{12}) &\approx s(\mu_1 - \mu_2) \\ &\approx s(\mu_1 - \mu) - s(\mu_2 - \mu) \\ &\approx \text{log-odds}(p_1) - \text{log-odds}(p_{20}) \end{aligned}$$

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.¹⁵

15. If the normally-distributed noise in Stochastic OT is replaced by logistically-distributed noise, then the log-odds is *exactly* proportional to the ranking distance, and there is no need to approximate. The normal and logistic distributions are very similar, making the two different versions of Stochastic OT difficult to distinguish empirically (Evanini 2007). In the absence of compelling empirical evidence favoring

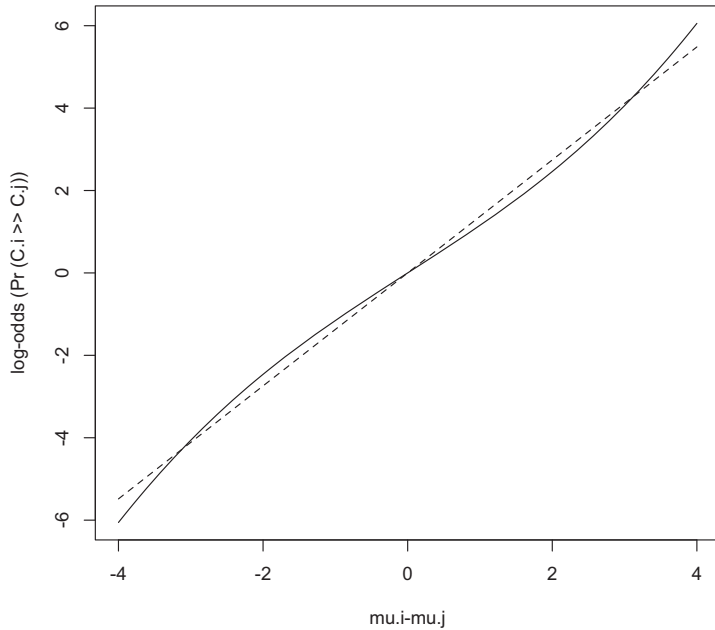


Figure 1. There is an approximately linear relationship between $\mu_i - \mu_j$ and the log-odds of the probability that C_i will be observed to dominate C_j , as long as the log-odds is between -6 and 6 (corresponding to a probability between about 1% and 99%). The dashed line, an ordinary-least-squares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

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 $\log\text{-odds}(p_{10})$ and $\log\text{-odds}(p_{20})$ 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, which occurred when p_{10} and p_{20} were 0.025 and 0.100 (in either order).

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

either distribution, there is no reason for linguists not to prefer the mathematically more tractable option.

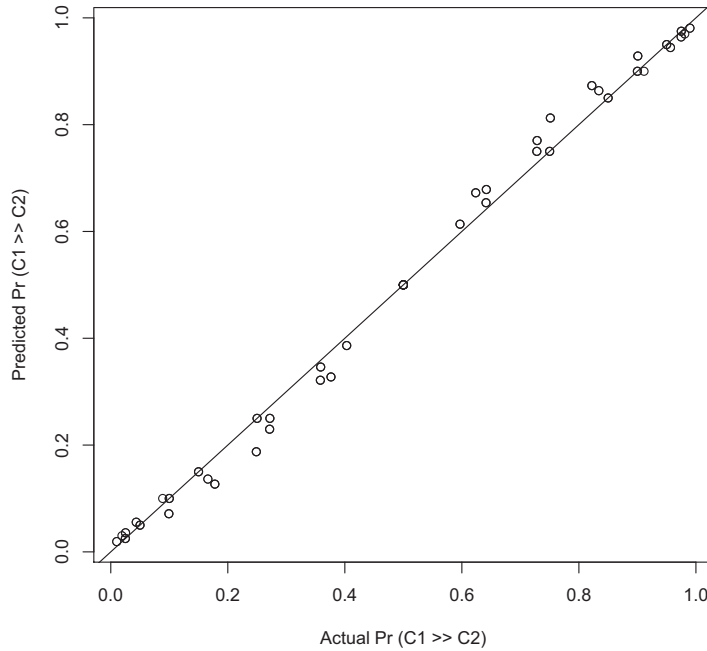


Figure 2. Approximate vs. exact p_{12} as calculated at multiple levels of p_1 and p_{20}

scale-partition hierarchy, there are circumstances in which the p_{ij} are not just inferable from, but actually equal to, the variation probabilities.

A concrete example can be constructed from Pater and Barlow's (2003) analysis of the data from Gnanadesikan (2004), shown above in (17). Since the sets of forms assigned violations by each $*x$ -ONS constraint do not intersect (there are no entailed violations between constraints in this family), it is only $*\text{FRICATIVE-ONS}$ and $*\text{NASAL-ONS}$ that are relevant for the $[\text{s}\epsilon\text{l}]/[\text{m}\epsilon\text{l}]$ decision. Hence the probability that fricative-nasal clusters are reduced to the fricative rather than the nasal is exactly equal to $\text{Pr}(*\text{NASAL-ONS} \gg * \text{FRICATIVE-ONS})$, and likewise for any other pair of the sonority classes listed in (16). The non-intersecting 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 section 5.2, the variation probabilities should stand in predictable relations to each other, e.g.,

$$(22) \quad \log\text{-odds}(\text{Pr}(sk \rightarrow s)) \approx \log\text{-odds}(\text{Pr}(sl \rightarrow s)) - \log\text{-odds}(\text{Pr}(kl \rightarrow k))$$

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 (the sets of forms to which they assign violations intersect, 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., Tropic 1987; Ohala 1999; Hansen 2001; Jongstra 2003a, 2003b), 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 (2003a, 2003b) data. We focus here on the three Dutch-legal 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 2 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.

Table 2. 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 2003b, Table 5.2b). Each proportion is based on at least eight observations

Child	Target cluster		
	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

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. By (21), there is no way to assign ranking values to the constraints in (17) that will match the observed frequencies: In order to model large preferences for stops over nasals and nasals over fricatives, the anti-stop and anti-nasal constraints have to be far apart, but in order to model the smaller preference for stops over nasals, the same two constraints have to be close together.¹⁶

This section has identified a clear empirical signature of a scale-partition hierarchy, which crucially depends on the lack of intersecting sets of violations between the constraints which are in variation. Deviation from the predicted relationship indicates such intersection, 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.

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

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-sensitive process involving at least three distinct sonority classes (three, because (21) presupposes three distinct constraints). 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,¹⁷ and we would need to determine at least three such frequencies per speaker. Data sets of this size may become more common as technology improves.

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 and Smolensky 1993/2004), in which higher-ranked constraints strictly dominate lower-ranked constraints, or as Harmonic Grammar (Legendre, Miyata and Smolensky 1990b; 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

17. In order to estimate p_{10} to a precision of ± 0.10 with 95% confidence, we need a sample of size at least $(1.96)^2 p_{10}(1 - p_{10}) / (0.1^2) = 384 p_{10}(1 - p_{10})$ (Tortora 1978, Eqn. 3.1). The value of p_{10} is unknown, but in the worst case it could be 0.5, requiring 96 observations. The number will be more favorable the smaller p_{10} really is; e.g., for $p_{10} = 0.05$, we need only 18 observations.

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 markedness 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.

Abbreviations

C	consonant
Dist	(sonority) distance
HG	Harmonic Grammar
L1	first language
L2	second language
OT	Optimality Theory
Pr	probability
V	vowel

Acknowledgments

Many thanks to Paul de Lacy, Jaye Padgett, and Steve Parker for comments and suggestions.