Word-Edge Effects as Overphonologization of Phrase-Edge Effects

Jaye Padgett

1. Background

Phonetic realities shape phonological patterns. Phonological patterns (which for the purposes of this paper I take to be defined by their categorical nature) closely mirror independently existing phonetic ones (which are gradient, see Zsiga 2011 for recent discussion of this distinction). Hence the notion *phonologization*: the emergence of a phonological process based on phonetic underpinnings (Hyman 1976; 2013). However, there is no agreement among phonologists about precisely *how* phonetics shapes phonology and what this means for phonology itself.

This paper focuses on one facet of the question, the distinction between what I will call the *eliminative* versus *inclusive* views of phonology. To understand this distinction, consider first a joke.[1] A guy walks into a bar and asks "Why do people have opposable thumbs?" A formal phonologist who happens to be there answers, "Because we have genes that make it so". A phonetician nearby counters, "Because they help us grab bananas". Luckily there are functional phonologists in the room also, and they resolve the argument: "Because our genes want us to grab bananas".

In the actual world of phonology, we might ask, why does word-final obstruent devoicing happen in many languages? A formal phonologist working within Optimality Theory might answer with a constraint ranking like *D]$_{Wd}$ >> Ident(voice) (where 'D' is any voiced obstruent). A phonetician might respond that contrasts like [t] versus [d] are hard to discriminate in word-final position, and/or that [b,d,g] are hard to produce word-finally. A functional phonologist working within Optimality Theory might answer with something like this: Ident(voice)$_{Cues}$ >> *D >> Ident(voice), where "$_{Cues}$" is a stand-in here for the voicing cues we think matter to perceiving that contrast. These constraints and rankings protect an obstruent voicing contrast when the relevant obstruent occurs in a position providing robust cues to the voicing feature value; otherwise, as in word-final position, the obstruent must devoice (Steriade 1997). This last answer can evoke unease like that felt when hearing that genes want bananas. Does phonology need to encode or refer to aspects of its phonetic underpinnings?

Phonetic facts and constraints exist independently of phonology. Voicing in obstruents can be more difficult to identify word-finally even in languages that maintain the phonological contrast there (e.g., English, Benkí 2003). The aerodynamic challenges that voiced obstruents present (Ohala 1983) are also independent of phonology. Phonologization happens when gradient asymmetries caused by these purely phonetic factors lead to a categorical reanalysis of a sound pattern. If an analysis like *D]$_{Wd}$ >> Ident(voice) captures a phonological generalization as well as Ident(voice)$_{Cues}$ >> *D >> Ident(voice), how do we decide between the two analyses? How much phonetic knowledge is incorporated into the phonology by phonologization, and how can we know?

By *eliminative* phonology I mean the view that articulatory and perceptual biases shape phonology solely from the **outside** (e.g., Ohala 1974; Blevins 2004), leaving phonology 'substance-free' (Hale & Reiss 2000). By *inclusive* phonology I mean the view that articulatory and perceptual biases shape phonology from the **inside** (e.g., Flemming 1995 [2002]; Steriade 1997; Boersma 1998). Of course, these terms themselves attempt to impose a categorization on a gradient and multilayered spectrum of views. Nevertheless, there are large enough differences on this spectrum to cause controversy among researchers. Phonologists are eliminative to the extent that they minimize reference in phonological theory to phonetically-rooted distinctions (e.g., they employ a very limited set of phonological

[1] Adapted from a question I asked Donca Steriade about 20 years ago.

features), eschew gradient phenomena, or avoid constraining phonological theory to account for things that may be explained by phonetics (or acquisition, or language change, etc.). Inclusive phonologists go the other way. For example, Ident(voice)$_{Cues}$ encodes distinctions between positions based on perceptual robustness, and Ident(voice)$_{Cues}$ >> *D >> Ident(voice) encodes the fact that contrasts are lost in positions where they are poorly perceived.

Eliminative phonology has been claimed to have Occam's Razor on its side (Hale and Reiss 2000 are particularly explicit on this point). If phonetics (or acquisition, …) explains something, why recapitulate the explanation in phonology? There are several problems with this argument. First, it presupposes that the mechanisms that govern gradient (phonetic) and those that govern categorical (phonological) phenomena are completely disjoint; and to be fair this seems to be the view of many phonologists. Yet some researchers on the inclusive end of the spectrum seek to (partially) *unify* phonetics and phonology within a theoretical framework that outputs both categorical *and* gradient effects (e.g., Flemming 2001; Gafos 2006). If some of the very same mechanisms explain performance, sound change, etc., on the one hand, and phonological generalizations on the other, the argument from parsimony for eliminative phonology loses it force. Second, even if phonetics and phonology employ fundamentally different mechanisms, nothing obviously prevents language learners from consulting both when forming phonological hypotheses (Hayes & White 2013). Finally, biological systems are not obligated to respect Occam's Razor.

Despite years of argument, there are still eliminative and inclusive phonologists. The reason why seems clear: both kinds of theory predict that phonological patterns tend to mirror phonetic ones, so it is hard to adjudicate between them. Sweeping philosophical arguments don't tend to persuade, in part because it is not obviously an either/or question. A more fruitful approach seeks empirical territory that reveals specific ways in which phonology does or does not mirror phonetics. A recent example of this is the work of Moreton (2008, 2009), who coins the term 'underphonologization', meaning systematic failures of phonologization despite phonetic underpinnings (though see Yu 2011). In this paper I rather take up instances of what we might call *overphonologization*: phonologization *without* the phonetic underpinnings. The existence of 'unnatural' phonological rules has been discussed many times before in this context (Bach & Harms 1972; Anderson 1981; Hyman 2001; Blevins 2004). Here the focus is on a different kind of eliminative argument, taking up ideas in Myers and Padgett (to appear). Myers and Padgett argue that at least certain kinds of word edge segmental phonology have dubious phonetic underpinnings and are in fact examples of overphonologization. Their source is *domain generalization*: overgeneralization of effects from utterance edges, where they are phonetically motivated, to word edges, where they are not.

2. Word-final devoicing and domain generalization[2]
2.1. An inclusive phonological account and a problem

Russian word-final devoicing, as in (1), is a well known example of a word edge segmental rule. Such a rule occurs in many unrelated languages (Blevins 2006; Myers 2012).

(1)　/slʲed/　　slʲe**t**　　'track (nom.sg.)'　　cf. /slʲeˈda/　　slʲiˈ**da**　　'track (gen. sg.)'
　　　/knʲig/　　knʲi**k**　　'book (gen.pl.)'　　cf. /ˈknʲiga/　　ˈknʲi**ga**　　'book (nom. sg.)'

The typology of word-final devoicing reveals some clear asymmetries. First, there is an asymmetry in the alternation: final devoicing is common, but word-final obstruent *voicing* is highly under-attested (see discussion in Yu 2004). Second is a positional asymmetry: obstruent devoicing is under-attested in word-initial position compared to word-final position (Westbury & Keating 1986). Such asymmetries present a well known challenge to phonological theory. If a phonological formalism allows any statement to be applied to any position, phonology has no hope of explaining these asymmetries. For an eliminative phonologist, the response might be, "So what? These asymmetries are explained by phonetics". Assuming that an obstruent voicing contrast is hard to discriminate in word-final position, perhaps nudged by a variable tendency toward gradient devoicing for reasons of production and aerodynamics, learners or even adult listeners may miscategorize final intended /b,d,g/

---

as [p,t,k]. Given enough miscategorizations of this sort, speakers will end up with no [+voice] obstruents in word-final position in their phonological surface representations. They will infer a phonological generalization that only voiceless obstruents are possible in this position. The phonology infers this because, so to speak, this is what phonetics has handed it. If the world were to conspire to present speakers with the opposite generalization (word-final obstruent voicing), this would be inferred just as well. Phonology is equally equipped to infer either *[-son, +voice] or *[-son, -voice], in either position _]Wd or Wd[_ ; it doesn't care either way. But because phonetic asymmetries tend not to lead to miscategorizations of word-final /p,t,k/ as [b,d,g], or to miscategorizations in word-initial position, phonology simply never has the opportunity to infer such generalizations.

An inclusive phonologist instead builds a phonological account more directly from the known phonetic asymmetries. He may posit a constraint *D (*[-son, +voice]) but no constraint *T, or posit a fixed universal ranking *D >> *T, to explain the preferred alternant asymmetry. As for the positional asymmetry, he may assume a fixed ranking Ident(voice)$_{Cues}$ >> Ident(voice), as in the theory of licensing by cue (Steriade 1997, 2001). This theory assumes that contrasts are favored where the cues to the contrast are richly present. In the case of obstruent voicing, this is when an obstruent immediately precedes a sonorant (Lombardi 1995; Steriade 1997). The tableau in (2) exemplifies the analysis of Russian word-final devoicing in these terms (where the subscript 'PS' means 'pre-sonorant'). Given the inventory of constraints shown and a fixed ranking between the two faithfulness constraints, languages should favor voiceless but not voiced obstruents, and they should experience devoicing in word-final but not word-initial position.

(2)

Input: /dub/		Ident(voice)$_{PS}$	*D	Ident(voice)
a.	dub		**!	
b. ☞	dup		*	*
c.	tup	*!		**

This much is well rehearsed, but there is a problem for this inclusive phonological account that receives surprisingly little discussion: in many phrasal contexts where words actually appear, it doesn't work. Consider the contexts exemplified in (3), for underlying /dub/ 'oak tree', cf. [dub-a] 'oak tree (gen.)'. Word-final /b/ is realized as voiceless [p] even when the immediately following sound is voiced, as in (3)b-e.

(3) a. dup kustarnʲikovij   'coastal scrub oak'
    b. dup bʲelij   'white oak'
    c. dup mʲullʲera   'muller oak'
    d. dup lopastnoj   'California white oak'
    e. dup arʲizonskʲij   'Arizona white oak'

Most damaging is the fact that word-final devoicing obtains in Russian even when the immediately following sound is a sonorant. The account seen in (2) fails, as shown in (4). The correct output is (4)b, with final devoicing despite a cue-rich presonorant environment. (In normal speech there need be no phrasal boundary or pause between the relevant words.)

(4)

Input: /dub arʲizonskʲij/		Ident(voice)$_{PS}$	*D	Ident(voice)
a. ☞	dub arʲizonskʲij		***	
b. ☹	dup arʲizonskʲij	*!	**	*
c.	tup arʲizonskʲij	*!*	*	**

To avoid this outcome, Padgett (2012) defines the positional faithfulness constraint to apply only within a prosodic word:

(5)

Input: /dub arʲizonskʲij/	Ident(voice)$_{PS\text{-}Pwd}$	*D	Ident(voice)
a. [dub] [arʲizonskʲij]		***!	
b. ☞ [dup] [arʲizonskʲij]		**	*
c. [tup] [arʲizonskʲij]	*!	*	**

This works, since [arʲizonskʲij] lies outside the scope of the cue-based faithfulness constraint. However, the problem at this point is that the inclusive phonologist's appeal to phonetic underpinnings has become conveniently selective. If the phonetic underpinnings matter so much to explain typological asymmetries, then why do we get to make them irrelevant here?

It should be clear that this problem is not particular to Russian. Many languages have word-final devoicing, and in presumably all of them words with final devoicing can be followed without pause by words beginning with a sonorant. It is a general problem for the inclusive, cue-based approach.

2.2. Phonetic underpinnings of final devoicing

Myers and Padgett (to appear) question seeking the phonetic underpinnings of word-final devoicing in word-final position. They argue, rather, that those underpinnings lie in *utterance*-final position.

Consider first the bases for utterance-final devoicing in production. In non-speech breathing the vocal folds are open wide; as a speaker approaches an utterance-final pause, she begins to open her vocal folds in anticipation of this posture (Sweet 1877: 65; Lisker et al. 1969: 1545; Ohala 1983; Klatt & Klatt 1990; Shadle 1997: 42; Jessen 1998; Slifka 2006). In addition, the subglottal pressure needed to maintain vocal fold vibration declines as the utterance progresses (Westbury & Keating 1986: 156). Both of these facts contribute to a deterioration of voicing toward the end of an utterance, and in fact phonetic (gradient) utterance-final devoicing is observed in many unrelated languages (see references in Myers & Padgett to appear). There are no such articulatory pressures obtaining at the ends of *words*, putting aside, of course, those in utterance-final position. There are generally no pauses between words within an intonational phrase or utterance. And a decline in subglottal pressure throughout an utterance implies nothing about a word in utterance-initial or -medial position.

As for perception, some works in the cue-based framework discussing the paucity or weakness of cues in word-final position really seem to be talking about words in utterance-final position (e.g., Padgett 1995; Steriade 1997). For example, Padgett (1995) ties place and laryngeal neutralization in word-final position in some languages to the absence of a consonantal release in that position (where 'release' implies offset formant transitions and/or consonantal burst). But word-final releases in this sense are reliably absent only in utterance-final position (in some languages); there will be a burst and/or formant transitions whenever a vowel- or sonorant-initial word follows within a phrase. Turning to experimental work, Myers (2012) finds that English listeners more frequently confuse pairs like *leaf* versus *leave* when the stimuli come from utterance-final position compared to utterance-medial position. As for direct comparisons of the perception of word-initial versus word-final obstruent voicing, these do not seem to be numerous. Some studies demonstrate poorer identification or discrimination (involving voicing or other features) in word-final position, but employ words in citation form or drawn from utterance-final position in a carrier phrase (e.g., Miller & Nicely 1955; Wang & Bilger 1973). These studies therefore do not tell us whether word-final contrasts are difficult outside of utterance-final position. Redford and Diehl (1999) *do* address this question, though not for voicing, and conclude that even in identical segmental environments within an utterance English listeners do not perceive word-final consonants as well as word-initial ones. Redford and Diehl also find acoustic differences between the word-initial and word-final consonants that likely underlie the perceptual difference; these presumably reflect production differences. Benkí (2003) finds a perceptual difference for the English voicing contrast specifically, for CVC stimuli heard in the carrier phrase 'You will write ___ please". It is unclear whether this carrier phrase was pronounced with a pause after the target word, but in any case it does not allow us to infer what happens in pre-sonorant position, where devoicing would seem to be least motivated. In addition, the error rates involving voicing reported by Benkí (2003) are under 10% even under higher noise levels and differ little by position in the word.

More experimental work is needed, but our current state of knowledge does not provide strong support for the claim that specifically *word*-final devoicing is motivated in production or perception.

2.3. Domain generalization

If the phonetic underpinnings of final obstruent devoicing obtain at the ends of utterances (before pause) and not at the ends of words, than there is a puzzling mismatch between the domain of the phonetic effects and that of phonological word-final devoicing. Some researchers have already suggested a solution to this mismatch problem. Assuming that final devoicing is utterance-final at an initial stage, this means that speakers do encounter *words* with final devoicing – those words produced utterance-finally. Perhaps word-final devoicing happens by a process of generalization (Ewert 1943: 75; Wackernagel 1957: 301-309; Vennemann 1974; Hyman 1978b; Westbury & Keating 1986: 161; Hock 1991: 239; Blevins 2006). Myers and Padgett (to appear) call this 'domain generalization', and discuss other common word-final phonological effects that are good candidates for domain generalization. For example, in many languages with tone, a high tone is disallowed in word-final position (Hyman 1978a: 265; Myers 1999: 216). A likely phonetic precursor to such a rule is utterance-final F0 lowering (Liberman & Pierrehumbert 1984; Pierrehumbert & Beckman 1988; Herman 1996); there is no obvious phonetic motivation for F0 lowering at the ends of *words*, since words can be utterance-initial or -medial. Some languages with a vowel length contrast also categorically disallow long vowels word-finally. Myers and Hansen (2007) argue that such a rule, as with final obstruent devoicing, has its precursor in utterance-final phonetic devoicing. As we have seen, there is no clear phonetic *word*-level correlate of final devoicing.

Under this view, word-final phonological rules like final obstruent devoicing come about in the stages shown in (6). These stages would have occurred many hundreds of years ago in Russian, but here we use contemporary Russian forms for convenience. The first innovation is utterance-final devoicing, which is phonetically motivated (Stage 2). A word like /sad/ 'garden' (bolded) will undergo devoicing in utterance final position but not otherwise. At a later stage (Stage 3), speakers have generalized final devoicing to all positions.

(6) Hypothesized stages leading to word-final obstruent devoicing

Stage 1	/**sad** vixod^j it v drugoj **sad**/	No devoicing
	[**sad** vixod^j it v drugoj **sad**]	
	'The garden lets out onto another garden'	
Stage 2	/**sad** vixod^j it v drugoj **sad**/	**Utterance**-final devoicing
	[**sad** vixod^j it v drugoj **sat**]	Phonetically motivated
Stage 3	/**sad** vixod^j it v drugoj **sad**/	**Word**-final devoicing
	[**sat** vixod^j it v drugoj **sat**]	Domain generalization

Polish dialects taken together show these stages underway: some have only utterance-final devoicing, while others have innovated word-final devoicing (Jassem & Richter 1989: 317). Similarly, some Mayan languages have word-final devoicing of sonorant consonants, while in others it occurs only utterance-finally (Bennett 2014 and references therein). In Reykjavík Icelandic, [l] is always devoiced utterance-finally, but devoices more variably in utterance-medial positions, with word-final devoicing within phonological phrases occurring only occasionally (Dehé 2014). This is a possible intermediate stage between utterance-final devoicing and invariant word-final devoicing.

Do speakers actually generalize as suggested for Stage 3? Domain generalization is motivated by considerations such as the comparative lack of phonetic underpinnings outside of utterance-final position, but could there be more direct evidence for it? Myers and Padgett (to appear) present results of two artificial language learning experiments that provide such evidence. Following Wilson (2003,

2006) and Finley and Badecker (2011), Myers and Padgett use a *poverty of stimulus* method. Here I briefly describe one of the experiments and its results.[3]

Forty three English-speaking undergraduate student participants were assigned to one of two learning groups. One group implicitly learned a generalization of utterance-final obstruent devoicing, the other learned utterance-final obstruent *voicing*. In both the learning and testing phases, participants heard nonce words of the form ˈCVC(V(C(V))), where V = any of [i, e, a, o, u] and C = any of [p, t, s, z, m, n]. Codas occurred only word-finally and were always one of [s, z, m, n]. Both [s] and [z] occurred in any onset position, but in the learning phase only [s] occurred in coda position for the final devoicing group, and only [z] occurred there for the final voicing group. Participants heard target words in two sentence contexts, ˈ*santa* ___ (target word in utterance-final position) and ˈ*santa* ___ ˈ*mizupu* (utterance-medial position). In the learning phase, participants encountered words with [s] or [z] in coda position only in utterance-final position. They encountered target words with [n] or [m] as coda in both utterance-final and utterance-medial position. Thus participants were familiarized with forms in either position, but they had no information about word-final obstruent voicing in utterance-medial position. It is in this sense that the experiment uses a poverty of the stimulus design. During the testing phase, all participants encountered words with any of [s, z, m, n] in coda position in both utterance-medial and -final position. The table in (7) summarizes the kinds of stimuli that participants encountered depending on which learning group they were in, the phase of the experiment, and the position of the stimulus in the utterance. The crucial obstruent-final stimulus types are bolded.

(7) Stimuli encountered by learning group, utterance position, and experimental phase

Experiment phase	Learning group	Utterance-final (ˈ*santa* ___)	Utterance-medial (ˈ*santa* ___ ˈ*mizupu*)
Learning phase	Natural	**pis**, pum, pamu	pum, pamu
	Unnatural	**puz**, pum, pamu	pum, pamu
Testing phase	Natural and Unnatural	**nos**, **naz**, nom, nusu	**nos**, **naz**, nom, nusu

Participants were told that they would be listening to sentences that would teach them new words. In the learning phase, they were instructed to listen carefully to each sentence and repeat it aloud. In the testing phase, they were instructed to listen to each sentence and judge whether the sentence belonged to the language they had learned. Also in the testing phase, participants only heard sentences with target words that they had not heard in the learning phase, and participants of both learning groups heard the same sentences.

As can be seen in (8), judging by the frequency with which they accepted forms, participants in both learning groups showed evidence of learning the generalization that was implicit in their respective learning phase: utterance-final devoicing or voicing.[4] (The differences seen within each learning group are significant.) The figure in (9) shows the more interesting results for utterance-*medial* position. Though they had no evidence about it in the learning phase, participants in the final devoicing group extended the generalization to this position.[5] In other words, the experimental participants performed domain generalization. Similar results are reported for a second experiment testing the effect on a broader range of consonants ([p, t, k] vs. [b, d, g]).

---

[3] For space reasons, some questions about the methods and analysis cannot be addressed here. See Myers and Padgett.

[4] Figures (8) and (9) reproduced from Myers and Padgett (to appear).

[5] The difference seen for the final voicing group is not significant. Though this difference between the two learning groups is interesting, it is not directly relevant to the point here. See Myers and Padgett (to appear) for discussion.

(8) Proportion of 'accept' responses by word-final voicing category and learning group
     Utterance-**final** position
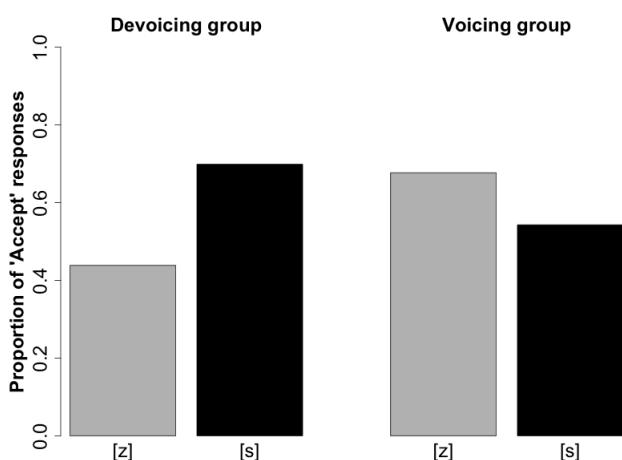
(9) Proportion of 'accept' responses by word-final voicing category and learning group
     Utterance-**medial** position
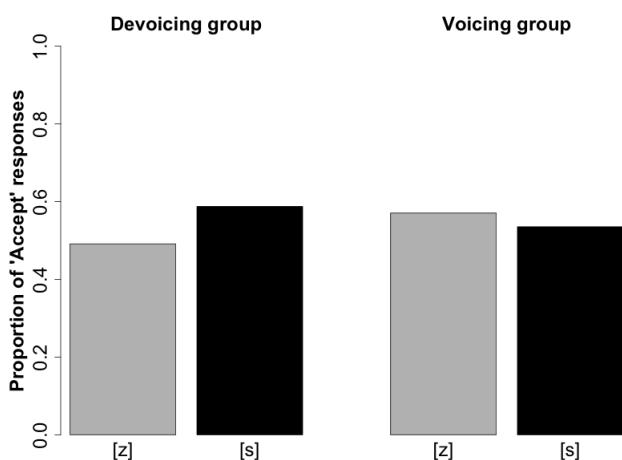
2.4. Why domain generalization?

Why does domain generalization happen? Recall the problem for an inclusive, cue-based account of word-final devoicing seen in (4), repeated here:

(10)

Input: /dub arʲizonskʲij/	Ident(voice)$_{PS}$	*D	Ident(voice)
a. ☞ dub arʲizonskʲij		***	
b. ☹ dup arʲizonskʲij	*!	**	*
c. tup arʲizonskʲij	*!*	*	**

An account directly tied to phonetic underpinnings wrongly predicts no devoicing here, since the consonant in question is intervocalic, the best possible position to realize an obstruent voicing contrast according to cue-based theory.

Steriade (1997) is aware of this issue, and suggests (pp. 55-8) a solution via 'paradigm uniformity' (roughly speaking). The idea is that /dub/ devoices in [dup arʲizonskʲij] because it should be identical to [dup] in utterance-final position, or in citation form, a particular case of utterance-final position. The tableau in (11) shows a possible implementation of this idea in Optimality Theory, making use of output-output faithfulness constraints (Benua 1995; Kenstowicz 1996). Candidate (11)a is no longer optimal under this analysis, because it violates the high-ranking Ident(voice)$_{OO}$ constraint requiring that /dub/ have an output identical to utterance-final [dup].

(11)

Input: /dub arʲizonskʲij/ Base: [dup]$_{Utterance-final}$	Ident(voice)$_{OO}$	Ident(voice)$_{PS}$	*D	Ident(voice)
a. dub arʲizonskʲij	*!		***	
b. ☞ dup arʲizonskʲij		*	**	*
c. tup arʲizonskʲij	*!	**	*	**

This sort of influence of utterance-final forms on their utterance-medial realizations may play a role in domain generalization, but it cannot be the whole story. Speakers of languages like Russian devoice words even with no experience of those same words' utterance-final or citation forms. (A Russian attempting to speak a foreign language, for example, will apply final devoicing.) Speakers must infer a bona-fide generalization that word-final obstruents be devoiced. How does this happen?

First, generalizations obviously must be *over something*, and this 'something' must be stored representations (Pierrehumbert 2003; Edwards et al. 2004). Contemporary Russian learners, for example, infer a generalization of word-final devoicing over many stored representations like [sat] 'garden', [trup] 'pipes (gen.)', and [luk] 'meadow', as informally depicted in (12).

(12)   Generalization over stored representations

[sada ~ sat]   [truba ~ trup]   [luga ~ luk]   etc.

final obstruent devoicing

In this context, it is plausible and relevant to assume that our inventory of stored *utterances* is small compared to our inventory of stored *words*. This must be so because we encounter more words than utterances, since utterances consist of words. In addition, words may be stored more successfully than utterances, since they are shorter and they are encountered more often. If generalizations are formed over stored representations, then it follows that it will be easier to infer generalizations over words than over utterances.

To see the relevance to domain generalization, consider again the hypothesized stage ('Stage 2' in (6) above) of utterance-final devoicing that precedes the innovation of word-final devoicing. Since a word like /sad/ 'garden' can occur in any utterance such as (13)a-c (and likewise for /trub/ and /lug/ etc.), learners of such a language will encounter and store many instances of /sad/, but they will encounter and store few instances of the precise utterances that contain /sad/. If generalizations are inferred over stored representations, then domain generalization occurs simply because generalizations at the level of the *word* are the ones most robustly supported.

(13)   Encounters of words vs. utterances (hypothesized utterance-final devoicing)

a. 'The garden lets out onto   b. 'She looked at the garden'   c. 'The garden is located across
   another garden'   the street'

/**sad** vixod^jit v drugoj **sad**/   /ona smotr^jela na **sad**/   /**sad** naxod^jits^ja tʃer^jez ul^jitsu/
[**sad** vixod^jit v drugoj **sat**]   [ona smotr^jela na **sat**]   [**sad** naxod^jits^ja tʃer^jez ul^jitsu]

[…sad…sat…sat…sad…]
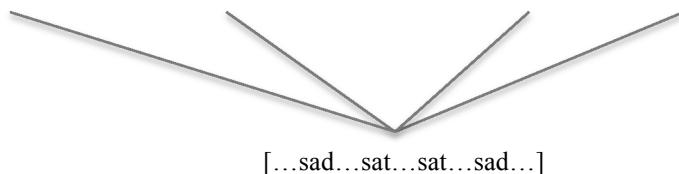
The main impediment for this understanding of domain generalization is the following: though it is less robustly supported in terms of stored representations, utterance-final devoicing is *more* robustly supported than word-final devoicing by a different criterion. Utterance-final devoicing is exceptionless at this hypothesized stage of the language; in contrast, devoiced forms like [sat] co-mingle in the learner's experience with non-devoiced forms like [sad] – the realizations that are not utterance-final. This raises two related questions: why does generalization happen at all, and why is it in favor of the voiceless variant?

There are reasons to believe that utterance-final word forms (including those uttered in isolation) are better perceived and stored by young language learners than corresponding word forms encountered in other positions. Infants both produce and perceive words at utterance edges better than they do utterance-medially, and they learn to segment words from phrases at utterance edges sooner than utterance-medially. Utterance-final variants of words may therefore figure more prominently in early lexicons, explaining why learners are more likely to infer generalizations based on them, leading (in this case) to a generalization of word-final devoicing. (See discussion and references in Myers & Padgett to appear.)

3. Implications: inclusive vs. eliminative phonology

Domain generalization has implications for our understanding of how 'inclusive' vs. 'eliminative' phonology is. The hypothesized stages leading to word-final devoicing, shown in (6) above, are repeated below. Since it has clear phonetic underpinnings, phonologization of *utterance*-final devoicing (Stage 2) is consistent with either the eliminative or the inclusive view.

(14)   Hypothesized stages leading to word-final obstruent devoicing

Stage 1	/**sad** vixod^jit v drugoj **sad**/ [**sad** vixod^jit v drugoj **sad**] 'The garden lets out onto another garden'	No devoicing
Stage 2	/**sad** vixod^jit v drugoj **sad**/ [**sad** vixod^jit v drugoj **sat**]	**Utterance**-final devoicing Phonetically motivated
Stage 3	/**sad** vixod^jit v drugoj **sad**/ [**sat** vixod^jit v drugoj **sat**]	**Word**-final devoicing Domain generalization

The same cannot be said of *word*-final devoicing. In fact, at least in contexts such as [dup ar^jizonsk^jij], where a vowel or sonorant consonant follows the word-final consonant without pause, final devoicing happens *in spite of* the phonetic facts on the ground. Domain generalization is in this sense a kind of *overphonologization*: phonologization without the phonetic underpinnings. As we saw in section 2.3, other word-edge effects are also plausibly understood as resulting from domain

generalization. The conclusion that this leads to is that word-final devoicing, and possibly many other word-edge phonological generalizations, represent an unusually clear and pervasive area of eliminative (or 'substance-free') phonological generalizations.

To put it differently, this reasoning vindicates the use of ']$_{Wd}$' over cues in at least some phonological rules. However, the context in which this conclusion is reached is very different from, and better than, that in which ']$_{Wd}$' would have been invoked decades ago. We are now in a better position to understand the typological asymmetries discussed earlier involving word-final devoicing. Voiceless alternants are favored, and in final position, because of the phonetic conditions associated with utterance-final devoicing. And it is domain generalization that provides a bridge from utterance-final devoicing, which is phonetically motivated, to word-final devoicing, which is not.

Taking up discussion in the introduction, the question of inclusiveness vs. elimination is not plausibly an either/or question. If the question is to what degree phonology behaves independently of phonetics, we can seek answers in specific mismatches between phonetic underpinnings and phonological generalizations. The general conclusion here is not that phonology is eliminative or substance-free. Domain generalization is consistent with the view that phonology incorporates phonetic bias but that such bias can be overcome (see Hayes & White 2013 and references therein). Domain generalization reveals one way in which phonology operates independently of phonetic biases.

There are other kinds of overphonologization – instances of 'unnatural' phonology – that bear on how eliminative vs. inclusive phonology is. Cases of apparent cyclicity, level-ordering in lexical phonology, or paradigm uniformity all represent well known examples of overphonologization. For example, a phonological rule of Russian (whose synchronic status is questionable) backed /e/ to [o] before non-palatalized consonants (which were velarized). This led to paradigms like the one shown in (15), in which historical /e/ became [o] in many present tense verb forms. The rule should not have affected the first /e/ in the second plural [idⁱetⁱe] 'you (pl.) go', because the following consonant was palatalized. However, the vowel of this form is [o] in all verbs of this type, a fact that is attributable to paradigm uniformity. (There are also output-output or cyclic exceptions like [pⁱosⁱik] 'dog (dim.)', derived from [pⁱos] 'dog'.) This is overphonologization, because /e/-backing is not phonetically motivated when /e/ is surrounded by palatalized consonants.

(15)   Paradigmatic exceptions to /e/ to [o] backing (forms of the verb 'go')

	Sg.	Pl.	
1ˢᵗ	idu	idⁱ**om**	
2ⁿᵈ	idⁱ**oş**	idⁱ**ot**ⁱe	*idⁱetⁱe
3ʳᵈ	idⁱ**ot**	idut	

However, domain generalization is interestingly different, the reification of a word edge independently of any effect of morphologically related forms.

Many have also pointed to allegedly unnatural or 'crazy' phonological rules as evidence that phonology is not 'natural' (e.g., Bach & Harms 1972; Anderson 1981; Hyman 2001; Blevins 2004). The problem that dogs this kind of evidence is that many such cases can be (wisely or not) written off as lexicalized, morphologized, or unproductive, and therefore not relevant to phonology at all. The case from word-final devoicing cannot be dismissed in this way. It is eminently attested, and often applies *automatically* (having no lexical or morphological exceptions).

4. Word-initial domain-generalization?

The reasoning behind domain generalization is not particular to the ends of domains. What about word-initial phenomena? Could at least some of these be generalizations from utterance-initial position? Here I briefly suggest possible avenues worth exploring.

Word-initial processes seem to be less common than word-final ones. For example, Fabian (2014) counted 80 word-initial vs. 198 word-final processes in the P-Base database of sound patterns (Mielke 2014).[6] Assuming that domain generalization is a factor here, this may be because utterance-initial

---

[6] Counting processes involving a single sound (e.g., no metatheses), including deletions but not insertions. Also excluding processes that were restricted morphologically or by lexical class.

position, in contrast to utterance-final position, is a position of phonetic *strength* (Keating et al. 1998; Cho & Keating 2001). This makes neutralizations less likely in initial position.[7] We might therefore look for processes that can plausibly be understood as contrast *enhancement*.

Consider word-initial aspiration in English, which, unlike word-internal aspiration, obtains even when the following vowel is unstressed, as in [pʰəˈtʰeɪɾoʊ] 'potato' and [kʰəˈlæmɪɾi] 'calamity'. Kingston and Diehl (1994) suggest a possible phonetic motivation for *utterance*-initial (post-pausal) aspiration. In this position, voicing of voiced obstruents like /b,d,g/ is difficult, causing these sounds to partially or completely devoice. Aspiration of the voiceless stops /p,t,k/ can be seen as a means of enhancing the contrast with (partially) devoiced /b,d,g/ in this position. However, as with word-final devoicing, this phonetic motivation does not reliably obtain at the *word* level. As Kingston and Diehl note, voicing of /b,d,g/ in intervocalic position, as in [ɹæbɪd] 'rabid' is not a problem; hence the comparative *lack* of aspiration in [ɹæpɪd] 'rapid', where it is not needed for contrast enhancement. Even in phrasal intervocalic contexts like 'a **b**aloney sandwich', full voicing of /b,d,g/ obtains. For an overly-inclusive account of word-initial aspiration, then, the presence of word-initial aspiration of 'potato' even in such contexts, e.g., 'a **p**otato omelette', is hard to explain. If word-initial aspiration in English is the result of domain generalization from utterance-initial position, then its apparent lack of phonetic underpinnings is no longer a puzzle.

Given the status of initial position as one of strength and contrast preservation, we might also look for cases of domain generalization in the form of rule *underapplication*. A candidate example, discussed in Myers and Padgett (to appear), involves suppression of lenition word-initially in Judeo-Spanish (Hualde 2013). Hualde posits two stages in the historical evolution of intervocalic lenitions as commonly found in Romance languages. In the first stage, lenition is across the board, affecting even word-initial consonants in the appropriate phrasal context. Thus in modern Spanish we have post-pausal [boka] 'mouth', but lenition in, e.g., [la βoka] 'the mouth'. However, in a second stage, affecting Judeo-Spanish, stops are resurrected in word-initial position even in such phrasal contexts: [la βoka] > [la boka]. This second stage is plausibly overphonologization of an utterance-initial requirement to word-initial position by domain generalization. This is because strengthening is best motivated in utterance-initial position, since greater phonetic strengthening is associated with the onsets of higher prosodic domains across languages (Keating et al. 1998; Cho & Keating 2001). Of course, the alternative is to assume that strengthening occurs word-initially because word-initial position is also strong, if not as strong as utterance-initial position. The question for this alternative account is why in that case word-initial consonants *always* undergo lenition in the earliest stages of lenition, as Hualde argues they do.

5. Conclusion

The idea of domain generalization is motivated by an apparent mismatch between the domain of phonetically motivated final devoicing – the utterance – and that of phonological final devoicing in Russian and many languages – the word. It is further supported by the experiments reported in Myers and Padgett (to appear). As noted above and by Myers and Padgett, there are other word edge effects that may be best understood via domain generalization as well. The general question raised by this work is to what extent *any* word edge effect might arise from domain generalization, or put differently, to what extent any word edge effects are phonetically motivated as such. If these ideas are on the right track, they bear in a particularly robust way on the debate about eliminative vs. inclusive phonology: common word edge effects in at least some cases represent overphonologization.

References

Anderson, Stephen R. (1981). Why phonology isn't "natural". Linguistic inquiry 12, 493-539.
Bach, Emmon & Robert T. Harms (1972). How do languages get crazy rules? In Stockwell, R. & R. Macaulay (eds.) Linguistic change and generative theory. Bloomington: Indiana University Press, 1-21.

---

[7] Unfortunately Fabian does not report how many of the counted processes are neutralizing. As a separate point, a systematic exception to the relationship assumed here between phonetic strength and neutralization involves processes related to prosodic prominence like stress or tone, and sonority, e.g., a requirement that stressed vowels be long (Smith 2005).

Benkí, José (2003). Analysis of English nonsense syllable recognition in noise. Phonetica 60, 129-157.

Bennett, Ryan (2014). Mayan phonology. Ms., Yale University.

Benua, Laura (1995). Identity effects in morphological truncation. In Beckman, J., S. Urbanczyk & L. Walsh (eds.) Papers in Optimality Theory. Amherst, MA: GLSA, 77-136.

Blevins, Juliette (2004). Evolutionary phonology. Cambridge: Cambridge University Press.

Blevins, Juliette (2006). A theoretical synopsis of Evolutionary Phonology. Theoretical linguistics 32.2, 117-166.

Boersma, Paul (1998). Functional phonology: formalizing the interactions between articulatory and perceptual drives. The Hague: Holland Academic Graphics.

Cho, Taehong & Patricia Keating (2001). Articulatory and acoustic studies on domain-initial strengthening in Korean. Journal of phonetics 29.2, 155-190.

Dehé, Nicole (2014). Final devoicing of /l/ in Reykjavík Icelandic. In Campbell, N., D. Gibbon & D. Hirst (eds.) Proceedings of the Speech Prosody 2014 Conference. Poster presented at the 2014 Speech Prosody Conference, Trinity College Dublin, 757-761.

Edwards, Jan, Mary Beckman & Benjamin Munson (2004). The interaction between vocabulary size and phonotactic probability effects on children's production accuracy and fluency in nonword repetition. Journal of speech, language, and hearing research 47, 421-436.

Ewert, Alfred (1943). The French language. London: Faber.

Fabian, Peter (2014). Initial domain generalization: an experimental approach to a typological asymmetry in phonological alternations. Ms., UC Santa Cruz.

Flemming, Edward (1995 [2002]). Auditory representations in phonology. Doctoral dissertation, UCLA [published in 2002 by Routledge, New York].

Flemming, Edward (2001). Scalar and categorical phenomena in a unified model of phonetics and phonology. Phonology 18.1, 7-44.

Gafos, Adamantios (2006). Dynamics in grammar: comment on Ladd and Ernestus and Baayen. In Goldstein, L., D. H. Whalen & C. T. Best (eds.) Laboratory phonology 8. Berlin: Mouton de Gruyter, 51-79.

Hale, Mark & Charles Reiss (2000). 'Substance abuse' and 'dysfunctionalism': Current trends in phonology. Linguistic inquiry 31, 157-169.

Hayes, Bruce & James White (2013). Phonological naturalness and phonotactic learning. Linguistic inquiry 44.1, 45-75.

Herman, Rebecca (1996). Final lowering in Kipare. Phonology 13, 171-196.

Hock, Hans Heinrich (1991). Principles of historical linguistics. New York: Mouton de Gruyter.

Hualde, José Ignacio (2013). Intervocalic lenition and word-boundary effects: evidence from Judeo-Spanish. Diachronica 30.2, 223-266.

Hyman, Larry (1978a). Historical tonology. In Fromkin, V. (ed.) Tone: A Linguistic Survey. New York: Academic Press, 257-269.

Hyman, Larry (1978b). Word demarcation. In Greenberg, J. (ed.) Universals of human language, volume 2: phonology. Stanford: Stanford University Press, 443-470.

Hyman, Larry (2001). The limits of phonetic determinism: *NC revisted. In Hume, E. & K. Johnson (eds.) The role of speech perception in phonology. New York: Academic Press, 141-185.

Hyman, Larry (2013). Enlarging the scope of phonologization. In Yu, A. (ed.) Origins of sound change: approaches to phonologization. Oxford: Oxford University Press, 3-28.

Hyman, Larry M. (1976). Phonologization. In Juilland, A. (ed.) Linguistic Studies Offered to Joseph Greenberg on the Occasion of his Sixtieth Birthday. Saratoga, CA: Anma Libri, 407-419.

Jassem, Wiktor & Lutoslawa Richter (1989). Neutralization of voicing in Polish obstruents. Journal of phonetics 17, 317-325.

Jessen, Michael (1998). Phonetics and phonology of tense and lax obstruents in German. Amsterdam: John Benjamins.

Keating, Patricia, Taehong Cho, Cécile Fougeron & Chai-Shune Hsu (1998). Domain-initial articulatory strengthening in four languages. In Local, J., R. Ogden & R. Temple (eds.) Phonetic interpretation: papers in laboratory phonology vi. Cambridge: Cambridge University Press, 145-163.

Kenstowicz, Michael (1996). Base identity and uniform exponence: alternatives to cyclicity. In Durand, J. & B. Laks (eds.) Current trends in phonology: models and methods. Salford, Manchester: European studies research institute, 363-393.

Kingston, John & Randy L. Diehl (1994). Phonetic knowledge. Language 70.3, 419-454.

Klatt, D. H. & L. Klatt (1990). Analysis, synthesis, and perception of voice quality variations among female and male talkers. Journal of the Acoustical Society of America 87, 820-857.

Liberman, Mark & Janet Pierrehumbert (1984). Intonational invariance under changes in pitch range and length. In Aronoff, M. & R. T. Oehrle (eds.) Language sound structure. Cambridge, MA: MIT Press, 157-233.

Lisker, Leigh, Arthur Abramson, F. Cooper & M. Schvey (1969). Transillumination of the larynx in running speech. Journal of the Acoustical Society of America 45, 1544-1546.

Lombardi, Linda (1995). Laryngeal neutralization and syllable wellformedness. Natural language and linguistic theory 13.1, 39-74.

Mielke, Jeff (2014). A database of sound patterns in 500+ languages. http://aix1.uottawa.ca/~jmielke/pbase/index.html.

Miller, George A. & Patricia E. Nicely (1955). An analysis of perceptual confusions among some English consonants. Journal of Acoustical Society of America 27, 338-352.

Moreton, Elliott (2008). Analytic bias and phonological typology. Phonology 25, 83-127.

Moreton, Elliott (2009). Underphonologization and modularity bias. In Parker, S. (ed.) Phonological argumentation: essays on evidence and motivation. London: Equinox, 79-101.

Myers, Scott (1999). Tone association and f_0 timing in Chichewa. Studies in African linguistics 28, 215-239.

Myers, Scott (2012). Final devoicing: production and perception studies. In Borowsky, T., S. Kawahara, T. Shinya & M. Sugahara (eds.) Prosody matters: essays in honor of Elisabeth Selkirk. London: Equinox Publishing.

Myers, Scott & B. Hansen (2007). The origin of vowel length neutralization in final position: evidence from Finnish speakers. Natural language and linguistic theory 25, 157-193.

Myers, Scott & Jaye Padgett (to appear). Domain generalization in artificial language learning. Phonology.

Ohala, John (1974). Phonetic explanation in phonology.

Ohala, John J. (1983). The origin of sound patterns in vocal tract constraints. In MacNeilage, P. (ed.) The Production of Speech. New York: Springer-Verlag, 189-216.

Padgett, Jaye (1995). Partial class behavior and nasal place assimilation. Proceedings of the South Western Optimality Theory Workshop 1995. Tuscon, AZ: The University of Arizona Coyote Papers, 145-183.

Padgett, Jaye (2012). The role of prosody in Russian voicing. In Borowsky, T., S. Kawahara, T. Shinya & M. Sugahara (eds.) Prosody matters: essays in honor of Lisa Selkirk. London: Equinox, 181-207.

Pierrehumbert, Janet (2003). Probabilistic phonology: discrimination and robustness. In Bod, R., J. Hay & S. Jannedy (eds.) Probabilistic linguistics. Cambridge, MA: MIT Press, 177-228.

Pierrehumbert, Janet & Mary Beckman (1988). Japanese Tone Structure. Cambridge, MA: MIT Press.

Redford, M. & Randy L. Diehl (1999). The relative perceptual distinctiveness of initial and final consonants in CVC syllables. Journal of Acoustical Society of America 106, 1555-1565.

Shadle, C. (1997). The aerodynamics of speech. In Hardcastle, W. & J. Laver (eds.) The handbook of phonetic sciences. Oxford: Blackwell, 33-64.

Simmons, J. P., L. D. Nelson & U. Simonsohn (2011). False-positive psychology: undisclosed flexibility in data collection and analysis allows presenting anything as significant. Psychol Sci 22, 1359-1366.

Slifka, J. (2006). Some physiological correlates to regular and irregular phonation at the end of an utterance. Journal of voice 20, 171-186.

Smith, Jennifer L. (2005). Phonological augmentation in prominent positions. New York: Routledge.

Steriade, Donca (1997). Phonetics in phonology: the case of laryngeal neutralization. Ms., UCLA.

Steriade, Donca (2001). Directional asymmetries in place assimilation: a perceptual account. In Hume, E. & K. Johnson (eds.) The role of speech perception in phonology. New York: Academic Press, 219-250.

Sweet, Henry (1877). Handbook of phonetics. Oxford: Oxford University Press.

Vennemann, Theo (1974). Words and syllables in Natural Generative Phonology. In Bruck, A., R. Fox & M. La Galy (eds.) Papers from the Parasession on Natural Phonology. Chicago: Chicago Linguistic Society, 346-374.

Wackernagel, J. (1957). Altindische grammatik. Band I: lautlehre (reprinted from the 1896 edition). Göttingen: Vandenhoeck and Ruprecht.

Wang, M. D. & R. C. Bilger (1973). Consonant confusions in noise: a study of perceptual features. Journal of the Acoustical Society of America 54, 1248-1266.

Westbury, John R. & Patricia A. Keating (1986). On the naturalness of stop consonant voicing. Journal of linguistics 22, 145-166.

Wilson, Colin (2003). Experimental investigation of phonological naturalness. In Garding, G. & M. Tsujimura (eds.) Proceedings of WCCFL 22. Somerville, MA: Cascadilla Press, 101-114.

Wilson, Colin (2006). Learning phonology with substantive bias: an experimental and computational study of velar palatalization. Cognitive science 30.5, 945-982.

Yu, Alan (2011). On measuring phonetic precursor robustness: a response to Moreton 2008. Phonology 28, 491-518.

Yu, Alan C. L. (2004). Explaining final obstruent voicing in Lezgian: phonetics and history. Language 80.1, 72-97.

Zsiga, Elizabeth (2011). External Sandhi in a Second Language: The Phonetics and Phonology of Obstruent Nasalization in Korean-Accented English. Language 87, 289-345.

Proceedings of the 32nd West Coast Conference on Formal Linguistics

edited by Ulrike Steindl, Thomas Borer, Huilin Fang, Alfredo García Pardo, Peter Guekguezian, Brian Hsu, Charlie O'Hara, and Iris Chuoying Ouyang

Cascadilla Proceedings Project   Somerville, MA   2015