

Variation, multi-level selection and conflicts between phonological and morphological regularities

Andrew Wedel

University of Arizona

Draft April 2007

To appear in Juliette Blevins, (ed.), *Analogy in Grammar: Form and Acquisition*. Oxford University Press.

0. Introduction.

Every language system comprises many overlapping levels of organization, each with its own structures and patterns. Because these levels overlap, patterns at different levels can come into conflict. For example, phonological regularity may entail morphological irregularity, as when addition of an affix requires a change in a stem. Morphological regularity in turn can entail phonological irregularity, as when a stem fails to undergo an otherwise regular phonological change upon affixation in the service of maintaining paradigm uniformity. In previous work I have argued that similarity-biased variation can contribute to the entrenchment of regular patterns over many cycles of language use and transmission (Wedel 2007). A great deal of evidence indicates that lexical memory is richly detailed at a number of levels, rather than limited to storage of symbolic, contrastive features as proposed in many classical models (reviewed in Pierrehumbert 2003). Within a model incorporating this evidence for rich-memory, biases in production and perception toward previously experienced forms create a positive feedback loop promoting pattern entrenchment (Wedel 2006, 2007, reviewed in Pierrehumbert 2006). Given that a given system can potentially evolve toward many different meta-stable states, a task for anyone working within this evolutionary model of language pattern development is to understand what factors encourage or inhibit the transition from given pattern into another. Recent examples of work in this area can be found in Blevins and Blust (2003), Blevins (2004), Mielke (2004), Chitoran and Hualde (2007) and many

others. In this paper I argue that pattern conflict across distinct levels of organization can be understood in a feedback-driven model of change as an instance of multi-level selection, and that this can help us think productively about the role of category variance in promoting or inhibiting change throughout the language system.

In the following section I review the role of noise in creating similarity-biases in category processing, and then in section 2 I go over some of the kinds of language change in which similarity-biased error may plausibly play a role. In section 3 I review how variation introduced by error influences the development of patterns within a rich memory model, as well as the use of evolutionary theory to model this process. Section 4 discusses possible mechanisms for similarity biases in production and perception that can feed language change, and section 5 introduces multi-level selection as a potentially useful way to think about conflicts between different levels of generalization. Finally, section 6 presents an illustrative simulation of a multi-level selection at work in a model lexical system evolving under competing attractors formed by distinct phonological and morphological regularities, followed by brief concluding remarks.

1. Error and similarity-bias in categorization

Information processing is always errorful to some degree due to noise. The simplest error pattern arises in processing of individual bits of information in which there are just two possible states, e.g., 0 and 1. In this case, noise can only result in the transformation of one bit value into the other.

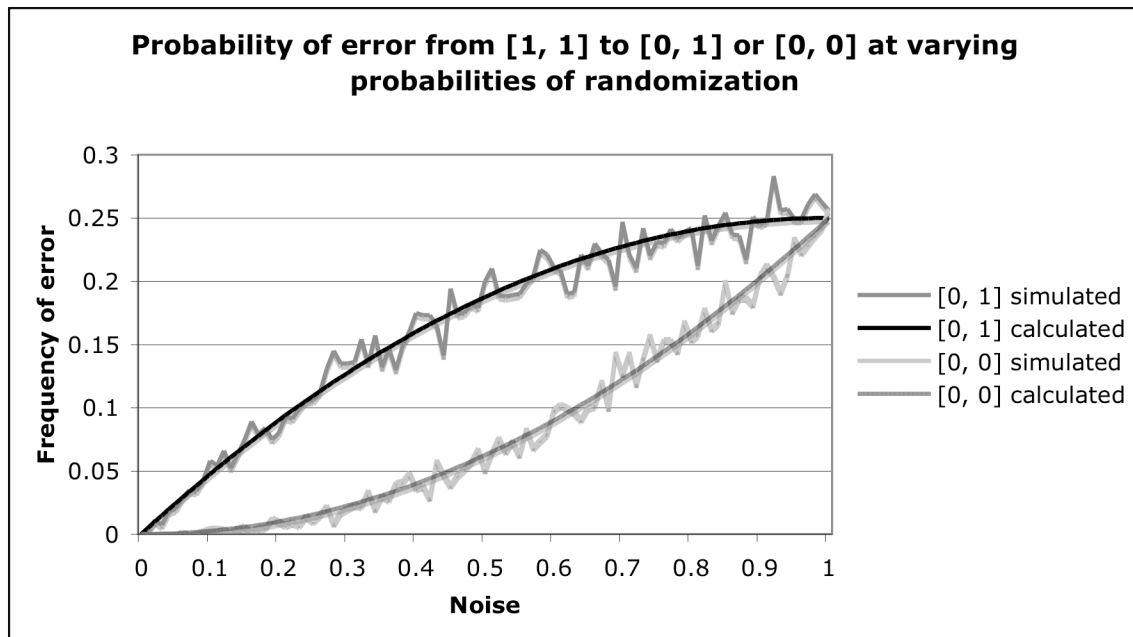
$$(1) \{ \dots, 1, 1, 1, \dots \} \rightarrow \text{noise} \rightarrow \{ \dots, 1, 0, 1, \dots \}$$

Much of the processing in language, on the other hand, involves processing compositional signals in which the unit of interest is above the level of an indivisible bit. For example, a word is composed of subsidiary units of information, such as segments. Successful transmission of a higher order category such as this requires that both the information source and target have access to a common lower-level information pattern

identifying the category, e.g., a segment sequence. In this case, there are two possible outcomes of noise in processing. If noise results in a pattern that is not successfully matched to any existing category, processing fails altogether at that level. However, noise can also result in a match to a different category, as when someone says *can't* and I understand *can*.

In a system in which categories can overlap to varying degrees, that is, can share variable amounts of lower level information, noise will always favor mismatches between more similar, over less similar categories. As an example, consider four categories each comprising two bits of information: $\{[1, 1], [1, 0], [0, 1], [0, 0]\}$. At any level of noise below that producing complete randomization, the odds that $[1, 1]$ will be mismatched to $[0, 1]$ or $[1, 0]$ is always greater than the probability of matching to $[0, 0]$. As an illustration, Figure 1 shows the rate of matching $[1, 1]$ to $[0, 1]$, and to $[0, 0]$ respectively, at varying noise levels, where a noise level of 1 represents complete randomization of original information. Numerically predicted rates are shown as well as simulated rates averaged over 1000 trials at each noise level.

Figure 1. Pattern-matching error to similar versus less similar categories under noise.



Given that language involves the processing of compositional categories that vary in their similarity along various dimensions, noise-driven mismatch errors will always be biased toward similar categories. In previous work I have argued that similarity-biased, ‘analogical’ error can serve as a seed for phonological change and entrenchment of patterns within a rich-memory model of language production and perception (Wedel 2004, 2006, 2007). Here, I will explore some consequences of the hypothesis that a general similarity-biased error also contributes to analogical change at the morphological level. Because dimensions of phonological and morphological similarity can cut across one another, similarity biased error and variation should set up conflicts between these distinct kinds of regularity. My goal in the next sections is to show that considering analogical change of all kinds to be initiated by similarity-biased error has the potential to shed light on the outcomes of conflict between and among phonological and morphological regularities (Sturtevant 1947).

2. Pattern extension in phonology and morphology

Many sound changes are ‘unnatural’ in the sense that they do not appear to originate in common articulatory or perceptual tendencies. Some of these appear instead to originate in pattern extension. For example, in phonology, sound patterns can be extended from an original, ‘natural’ context into contexts in which the change is not clearly phonetically motivated (for examples, see Mielke 2004, pp 102-114, Blevins 2006).

In morphology, both *leveling* and *extension* changes can be considered instances of pattern extension (discussed in Hock 2003). In leveling, members *within* a paradigm become more alike in some way. For example, the historical stem-final [f~v] alternation in the singular-plural pair *dwarf* ~ *dwarves* has leveled for many speakers of American English to *dwarf* ~ *dwarfs*. Paradigmatic extension occurs when a change creates a relationship within one paradigm that is parallel in some way to a relationship holding in another. For example, the originally regular present-past paradigm of *dive* ~ *dived* has shifted for many speakers to *dive* ~ *dove*, presumably by extension on the model of the group containing *drive* ~ *drove*, *ride* ~ *rode*, etc.

None of these phonological or morphological patterns can be fully understood without making reference to the existing language system. Given that learners and adult users alike have some knowledge of the ambient linguistic system, there are two conceptually distinct pathways by which patterns in the existing system can influence change: (i) by influencing the range of variants presented by adults to learners as input, and (ii) by influencing the ways that learners organize this input as they bootstrap between input and their current system toward the adult system (Pierrehumbert 2003; cf. CHANCE and CHANGE in the framework of Evolutionary Phonology, Blevins 2004). In both cases, similarity-biases can accentuate asymmetries within the experience of an individual. Within a rich-memory model of language production and processing (e.g., Pierrehumbert 2001, Wedel 2004, 2007), this asymmetry in experience is recorded in a corresponding asymmetry in the language system at some level. What dimensions of similarity are most salient in a particular system is an empirical question, dependent on both relatively universal as well as system-specific details (see e.g., Albright (this volume) and Pierrehumbert (2006) for discussion of these issues).

3. Rich memory, feedback and evolution

There is abundant evidence that the mental lexicon stores much more information about lexical forms than would be in principle necessary to produce them accurately. In turn, there is evidence that new experiences continually contribute to this store of information, and that this information biases both subsequent perception (e.g., Johnson 1997, Guenther et al. 2004, Eisner and McQueen 2005) and production (Goldinger 2000, Harrington 2000). As a result, processing a particular instance of a form increases the probability that a similar form will be processed in the same way in the future, and that corresponding forms will be produced in a similar way in the future. This creates positive feedback that promotes the entrenchment of patterns over many cycles of production and perception in acquisition, and to some degree in adult usage as well (Wedel 2006, 2007, reviewed in Pierrehumbert 2006).

Within a model of language in which variation within and across categories can be stored in some form and reproduced, the system as a whole can change through

evolutionary processes (Wedel 2006). The most well-known mechanism for a reproducing population to evolve over time is through selection, in which some variant elements in the population reproduce more than others via some interaction within the system. As long as there is some mechanism for variation to arise and persist, selection can favor some variants over others in some way, thereby altering the distribution of characteristics within the population over time. Although some rich-memory models assume that the only content of categories is in the form of fully detailed exemplars (reviewed in Tenpenny 1995), there is evidence that behavior also proceeds through use of independent, more abstract generalizations about input data (e.g., Kuehne et al. 2000, Albright this volume). For the purposes of the argument here, provided that within-category variation can persist in the system at some level – whether at the level of exemplars or of generalizations about some form – selection among, and created by these variants will be able lead to change.

When incipient patterns conflict in systems including positive feedback, the most stable outcome is dominance of one pattern over the other. In previous work, I have argued that similarity-biased errors in production and perception may serve as an underlying cause of the development of regular phonological patterns in language through positive feedback, despite the ability of the language system to store and use otherwise predictable information (Wedel 2007). The development of consistent patterns in morphology has also been argued to arise through positive feedback over many cycles of errorful learning (e.g., Hare and Elman 1995). Here I will suggest that errors biased toward similarity assessed across distinct levels of categorization create multiple competing pathways for the resolution of conflicts between regularities, and that this conflict can be understood in terms of multi-level selection.

4. Similarity biases in production and perception

In any process that distinguishes between categories, the rate of error in element identification or manipulation due to noise will be greater between more similar categories relative to less similar categories. Processes in language use that provide opportunities for these kinds of similarity-biased errors include (i) motor entrenchment in

production, (ii) the magnet effect in perception, and (iii) the application of relational categories to compose related forms. Motor entrenchment is a general property of motor systems in which practiced motor routines bias future motor execution in some relation to similarity (Zanone and Kelso 1997). This sets up a positive feedback loop in which, *ceteris paribus*, less frequent production variants should be steadily deformed toward more frequent production variants over time (Bybee 2002, discussed in Wedel 2006, 2007). On the perceptual side, the perceptual magnet effect (Kuhl 1991, 1995) provides another potential source of positive feedback which can act to enhance the similarity of forms over time. The perceptual magnet effect refers to the finding that percepts tend to be biased systematically toward the centers of categories relative to the stimuli that gave rise to them. Within models of the perceptual magnet effect in which warping precedes categorization (Guenther and Gjaja 1996), this systematic warping should pull similar pronunciations closer together over time through feedback between perception and production (Wedel 2007).

Both motor gestures and linguistically relevant sound categories often have a relational internal structure, with the result that they cannot be fully characterized by a simple list of properties. Instead, these categories must include some higher-order relational information. Phonological examples include sound-categories with temporally ordered gestures such as diphthongs, affricates, and contour tones. The central importance for language of such ‘relational categories’ has been discussed at length by Dedre Gentner and colleagues in the context of semantics (Gentner and Kurtz 2005). Morphophonological patterns are also relational, in that they describe some mapping between forms (Bybee 1985). These patterns are often described in terms of rules, but they may be described as well in terms of relational categories, identified with, for example, the large number of possible patterns in the relationship between present and past forms of English verbs (Albright and Hayes 2002). Generalizations (whether expressed as rules or relational categories) play a role in production or identification of linguistic forms whenever some form is reconstructed from a related form. The parade example of this use is in the production of a novel form fitting a pattern. In this case, a large body of research indicates that the applicability of a generalization to a novel form is gradient and dependent on similarity to other forms that are covered under that

generalization (e.g., Long and Almor 2000, Albright 2002, Krott et al. 2002, Ernestus and Baayen 2003).

Production of previously learned, morphologically complex forms within a paradigm can proceed by direct retrieval from memory or through reconstruction from a base using an associated generalization (e.g., Baayen 1992, Alegre and Gordon 1999). Error in application of a generalization in this process can result in an extended or leveled output pattern depending on the source of the generalization (Hock 2003). Extension of a compositional pattern results in a leveled output, as when speakers of English occasionally produce the past tense of an irregular verb regularly. Extension of a pattern involving some stem-change appears as an extended output pattern, as when English speakers produce the past tense of a regular verb in a way that matches some irregular past-tense pattern (e.g., Bybee and Modor 1983). Extension of irregular patterns have been shown to be more likely to occur to bases that share phonological features with the set of forms exhibiting that irregular pattern (Bybee and Modor 1983, Long and Almor 2000, Albright and Hayes 2002).

There are a wide variety of generalizations that are potentially involved in production and perception of any linguistic form, from lower level phonotactic generalizations about feature groupings and segment sequences, to higher level relational, morphological generalizations about possible paradigmatic relationships. Because the sequences referred to by these generalizations can overlap to any degree, there is the possibility of conflict between distinct kinds of generalizations. The following section discusses the possible outcomes of this conflict in terms of competition between levels of selection.

5. Similarity biases and selection

In biological evolution, errors in the replication of a gene are thought to be random, at least with respect to the phenotype conferred by the gene. Selection on the basis of the interaction of a variant gene product with its environment influences the likelihood of reproduction of some unit containing the gene (such as a cell, a multi-cellular organism, or a kin-group). As a consequence, the production of variants and the filter on what

variants survive to reproduce are mechanistically distinct. On the other hand, within a model of language in which errors can be biased by similarity to other existing forms and patterns, variation in what is produced and what is perceived is non-random with regard to the ‘phenotype’ of the system (cf. CHANCE and CHOICE in the framework of Evolutionary Phonology (Blevins 2004)). In this regard, similarity-biased error acts in production as a selective filter acting on the pool of *potential* variants, influencing which variants actually emerge to become part of the exemplar set of the larger system. In perception, similarity-biased error acts as a selective filter by biasing identification and storage into categories.

In biological systems, genes exist within a Russian doll of nested units that are potentially the objects of selection, ranging from the gene itself, through the chromosome, the cell, the multi-cellular individual organism, the kin group and potentially beyond (Mayr 1997). Selection can potentially act at each of these levels, often mediated by distinct mechanisms and on different time scales¹. For example, selection at the level of the cell strongly favors cells that are unconstrained in their growth, which promotes the development of cancer within an individual’s lifetime. Selection at the level of the individual on the other hand strongly favors strong control over cell division. In concert, these two selective pressures lead both to selection for cancer within the population of cells within a single individual, and to selection against early development of cancer over a timescale of many lifetimes within the population of individuals.

Within a single level of selection, the net selection pressure deriving from multiple independent loci of selection can often be approximated as a simple sum. For example, if a trait increases fitness in some way to a given degree, but decreases it by to the same degree through an independent pathway, the net selection pressure on that trait may be near zero. In contrast, when selection pressures on a given trait operate at different levels of selection, say the cell versus the individual, these pressures can interact in a more complex way. Selection against a trait at one level can often proceed through

¹ The well-known phenomenon of kin-selection is a particular case of selection beyond the individual. In kin-selection, selection at the level of the kin-group favors the evolution of behavior detrimental to the self when it supports the greater reproductive success of a close relative.

creating a systemic change that makes selection *for* that trait at another level less efficient. One way to influence the efficiency of selection is through modifying the amount of variation present at a given level of selection; greater variance provides more opportunities for a fitter variant to be selected (e.g., Taylor et al. 2002).

Within the model presented here, the competition between selection for regularity at distinct levels of linguistic organization is similar to biological multi-level selection in that change at one level can influence the opportunities for change at another. Within the present model, a pattern serves as a self-reinforcing attractor by biasing variation/error toward itself. Because linguistic categories can overlap with or contain one another (as for example when a sound category is part of a sound-meaning category), a change that increases the regularity of a pattern at one categorial level can decrease it at another. A decrease in regularity of a pattern (i.e., an increase in variance) therefore has two interacting effects on further change: (i) as variance increases at that level, the range of future variation increases, potentiating change; (ii) as variance increases, similarity-driven selection pressure toward the mean is weakened. Both of these effects should independently potentiate a shift further away from regularity in the contents of a category through evolutionary change.

6. Illustrating multi-level, selection-driven pattern development by simulation.

In Wedel (2007) I illustrated the evolution of regular stress patterns through similarity-based positive feedback within a simulated lexicon over many cycles of production and perception. In these simulations, the only relationships encoded between lexical items were on the basis of shared segmental properties in temporal order. Segmental properties that were provided to the system included stress value, segmental category features and word-edge status. An example of a three syllable lexical entry is

[1, a, I] [-1, b] [1, c, F]

where square brackets enclose syllables. Each syllable is characterized by a stress value and one or more additional features: ‘1’ and ‘-1’ represent stress and stresslessness,

respectively, lower-case letters represent segmental features, and ‘I’ and ‘F’ correspond to ‘word-initial’ and ‘word-final’.

Lexical production in each round of the simulation proceeded by copying the information stored in the lexical entry into an output form with a low probability of error in the stress value, and then restoring it in the lexical entry, replacing the original. Directional change could intervene in this process through the action of two kinds of error-bias in output production, one external, and the other system-dependent. The external error bias was a constant, lexicon-independent bias favoring alternating stress, such that each word would eventually tend to exhibit alternating stress regardless of the initial state. The second kind of error consisted of a similarity-bias in which output stress values had a slight probability of deviating from the stored value toward the values of other forms, in relation to similarity and type-frequency. The simulation detected pattern trends within the lexicon by identifying every existing combination of features and stress values in the lexicon, and looking for robust generalizations. When an existing robust feature-set:stress-value generalization conflicted with the stored version of a word, the output based on that word had a greater than chance probability of shifting stress values to match the larger generalization. As a result, the system showed a strong tendency to create broad associations between stress values and features over many cycles of production and restorage.

Within the lexicon, there were many possible segmental features, but only two edge features (initial vs. final). Many words therefore failed to share any segmental features at all, while every word had both an initial and final edge feature associated with the initial and final syllables, respectively. As a consequence, the most robust generalization that the lexicon could possibly evolve was one in which a given stress value was consistently associated with the initial and/or final word edge, rather than to some other segmental feature(s). When both even- and odd-syllable words were included in the lexicon, the dominant pattern was the evolution of an alternating stress pattern consistently aligned *either* to the initial *or* the final syllable.

To illustrate multi-level selection within this model, I modified the simulation architecture to include two optional suffix syllables for a subset of the words in the lexicon, identified with the abstract features [y] and [z] respectively. A portion of a

sample lexicon is shown in Figure 3. The final syllable of every word contains a final-edge feature (F). The lexicon consists of 80 words. Half of the words in the lexicon do not have a related a suffixed form, illustrated in (3a). The other half, as illustrated in (3b), appear in addition in a suffixed form. An [F] feature appears on the final syllable in all forms².

Figure 3. Example of a statically regular lexicon

(a) Stem-only paradigms

[-1, a] [1, b, F]

...

(b) Stem and Stem+Suffix paradigms

| <u>Stem</u> | <u>Stem + Suffix</u> |
|-------------------|----------------------------|
| [-1, c] [1, d, F] | ~ [1, c] [-1, d] [1, y, F] |

...

| | |
|-------------------|----------------------------|
| [-1, e] [1, f, F] | ~ [1, e] [-1, f] [1, z, F] |
|-------------------|----------------------------|

...

The lexicon in Figure 3 is ‘statically regular’ with regard to stress, because all entries show alternating stress aligned to the final syllable. It is ‘relationally irregular’ with regard to stress, because stems in bare and suffixed forms show opposite stress patterns.

The system retains the ability to detect robust static generalizations across all

² In order to focus the simulation on conflict between emergent phonological and morphological patterns, the development of stress associations at the final word-edge was encouraged by eliminating the initial-edge feature. Within over fifty independent trial simulations, the system always rapidly developed a stress pattern in stem-only paradigms in which stress was aligned to the final edge.

words in each cycle. In addition, the system has been equipped with the ability to identify the global stress pattern relationship between suffixed and unsuffixed forms and discover any robust associations between this relationship and any existing combination of features, using a parallel computational mechanism to that used for the discovery of static generalizations (described in Wedel 2004, 2007). This latter ability allows the emergence of relational generalizations of varying specificity. For example, a maximally specific generalization would match the stress pattern of a particular unsuffixed form to its related suffixed form, whereas a less specific generalization could emerge if a number of different unsuffixed forms shared the same stress pattern relationship with their suffixed forms. Although implemented in a computationally distinct way, this is conceptually parallel to the mechanism of relational generalization discovery in Albright and Hayes' Minimal Generalization Algorithm (2002)³. As before, the process of encoding an output corresponding to a stored form was subject to error biased toward existing patterns in the lexicon in proportion to similarity and type frequency. As in the single-level simulations in Wedel (2007), low level noise was also included, in the form of a very small probability of context-free error in correctly reproducing the stored stress value in any syllable.

The result is a system that has two distinct levels of system-dependent generalization that can influence error: static generalizations at the level of features, and relational generalizations between related words, where the targets of relational generalizations contain the targets of static generalizations. Pattern competition within and between these two levels of generalization resulted in three common classes of patterns. In one class of patterns, alternating stress developed with a given stress value consistently associated with the final edge of all words, with no reference to any semantic feature. This is the type of pattern that emerges in the absence of any possible relational generalization linking related words. This pattern represents full regularity with regard to phonological categories, and full irregularity within each morphologically related pair, as the stress pattern for the stem in the unsuffixed form is opposite that found in the

³ Previous work compared the mechanism of pattern discovery used here to several computational mechanisms including Analogical Language Modeling (Skousen 1989), and showed that they all produced qualitatively similar regular patterns (Wedel 2004).

corresponding suffixed form, as in Figure 3 above.

In a second common pattern, all two syllable forms in the lexicon had the same stress value associated with the final edge, while each suffixed form had the opposite stress value at its final edge, thereby preserving the stress pattern of the stem. This represents full morphological, or relational regularity with regard to stem stress pattern, and full phonological, or static irregularity with regard to final-edge stress alignment. The lexicon in Figure 4 below exhibits this pattern. A third pattern occasionally arose in which the paradigm associated with one suffix showed phonological regularity, and the other showed morphological regularity.

Figure 4. Example of a relationally regular lexicon.

(a) Stem-only paradigms

[-1, a] [1, b, F]

...

(b) Stem and Stem+Suffix paradigms

| <u>Stem</u> | | <u>Stem + Suffix</u> |
|-------------------|---|---------------------------|
| [-1, c] [1, d, F] | ~ | [-1, c] [1, d] [-1, y, F] |

...


| | | |
|-------------------|---|--------------------------|
| [1, e] [-1, f, F] | ~ | [1, e] [-1, f] [1, z, F] |
|-------------------|---|--------------------------|

...

In this model, both leveling and extension occur in the same way: through errorful application of a different relational generalization. If there is only one primary relational generalization that fits all the data, then the only available mechanism for change in stress pattern lies in the low level noise factor which provides a continual, small input of stress pattern variants into the system. This occasionally leads to the fortuitous emergence of a

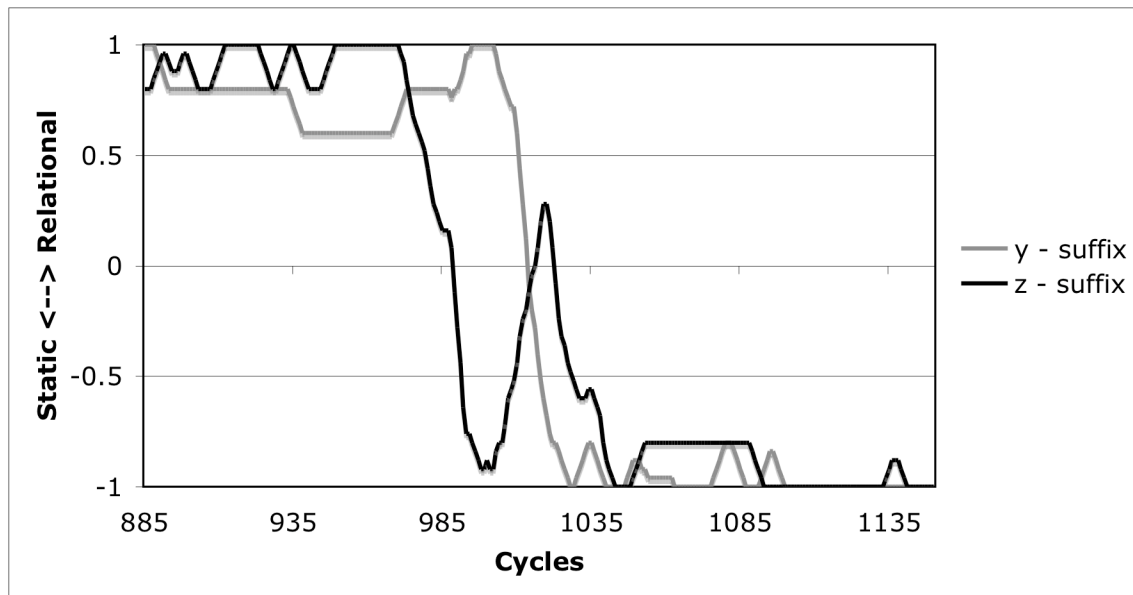
different generalization, which can then spread through similarity based error. As expected, it first spreads through the most similar subset of words within the lexicon, after which it may spread further. This is illustrated in the graph in Figure 5 below, where a value of '1' represents full relational regularity in the stress of stems in related suffixed and unsuffixed forms, and '-1' represents full static identity in the stress patterns of all words in the lexicon with respect to the final word edge. The simulation is seeded with a lexicon exhibiting full relational regularity in both the 'y' and 'z' suffix paradigms, like that shown above in Figure 4. This pattern remains stable for 1000 cycles despite the steady introduction of low-level variation in stress patterns by noise.

Figure 5. Competition between static and relational regularities



Shortly after the z-suffix paradigm switches to a pattern in which all forms have the same stress with regard to the final word edge, the y-suffix forms are able to follow suit. This change is potentiated because the z-suffix paradigm presents a similar group of words governed by a distinct generalization which can itself be errorfully applied to members of the y-suffix paradigm. In other words, as soon as a new generalization emerges, its misapplication provides a new pathway of change. A close-up of this transition is shown in Figure 6 below.

Figure 6. Competition between static and relational regularities: cycle 885-1150.

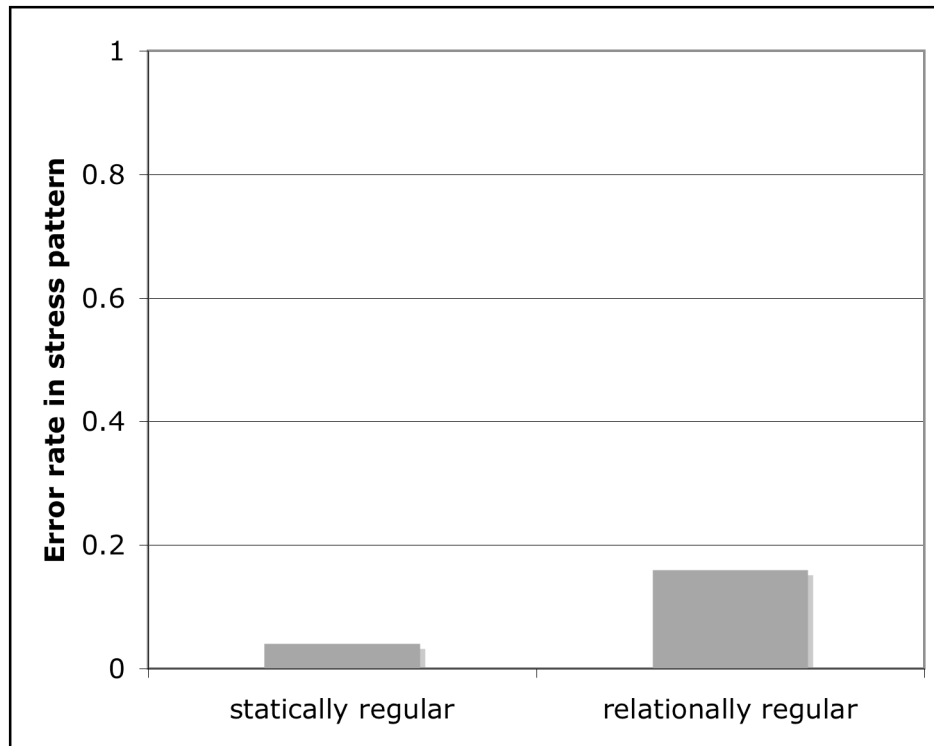


The dependence of a change between a static and relational stress pattern on the pre-existence of a target pattern in the lexicon holds in both the ‘extension’ and ‘leveling’ directions. In 10 simulations in which the seed lexicon was statically regular throughout (that is, where stress was aligned to the final edge of all words), it took an average of ~1400 cycles for a change to a relational pattern within one of the suffix paradigms to emerge. Likewise, in 10 simulations seeded with a lexicon exhibiting relational regularity in all suffixed-unsuffixed word pairs (that is, where the stem in each pair had the same stress), the average time to a change to static regularity was near 1700 cycles. In contrast, if the seed lexicon started with static regularity in one suffix paradigm, and relational regularity in the other, in 10 simulations it took on average less than 300 cycles for a further change to occur in one of the paradigms. The rate of change is greater when there are multiple existing patterns in the lexicon because each pattern represents a template for analogical extension.

The potentiating effect of multiple patterns can also be seen in the rate of error in output stress among the stem-only paradigms within the lexicon. When the stress pattern is statically regular, all stress patterns are edge-aligned across the entire lexicon. Under this condition, random noise is the only source of error in stress-output in stem-only

paradigms. When the stress pattern is relationally regular, however, some lexical entries have the opposite stress pattern as that in stem-only paradigms. In this case, there is an additional pattern in the lexicon to provide a pathway for variation in stress-output beyond random noise. Within the simulation, this can be seen by comparing the number of stem-only paradigm outputs with a variant stress patterns in a statically regular, versus relationally regular lexicon. Figure 7 below shows the error rate in stress within stem-only paradigm outputs over 10 independent runs of 1000 cycles each in the context of either a statically, or relationally regular lexicon. When the stress pattern is consistently aligned to the final word-edge over the entire lexicon (i.e. is statically regular), the error rate in stress in stem-only paradigm outputs is .04. However, when the stress pattern is instead aligned to a stem-edge within stem ~ stem+suffix paradigms (i.e., is relationally regular), the average error rate in stem-only paradigm outputs goes up to .16.

Figure 7. The error rate in production of the dominant stress pattern in stem-only paradigm outputs is lowest when the entire lexicon is statically regular. The error rate goes up in these paradigms when other stem ~ stem+suffix paradigms are relationally regular.



This higher error rate in the relationally regular lexicon comes about through the existence of an additional pattern in the lexicon, which provides an additional pathway to a change in stress. The resulting increased variance in stress patterns within stem-only paradigms has two related effects: (i) the dominant stress pattern of stem-only paradigms is *less* stable, and therefore more likely to change over time, and (ii), the dominant stress pattern of stem ~ stem+suffix paradigms is *more* stable, because any similarity-bias promoting static regularity is weakened. This is conceptually parallel to cases in biological evolution in which selection at one level acts by modulating variance at another (e.g., Taylor et al. 2002).

7. Summary and conclusions

The statistics of error in pattern matching ensure that similar patterns will substitute for each other more often than less similar patterns in production and perception. In any model of language production and perception in which intra-categorical variants can coexist and compete within the system, positive feedback promotes the evolution of

regular patterns. Under the assumption that both static and relational generalizations are manipulated during language production and perception, error between similar generalizations should produce a wide range of similarity effects at different levels, from phonotactics to morphology (see e.g., Burzio 2002).

When similarity at separate levels of organization cannot be simultaneously maximized, similarity-biased error and feedback promotes the entrenchment of one pattern at the expense of the other. The potentiation of change by similarity-biased error allows this snowball-effect to proceed in opposite directions at different levels: as variance decreases at one level, further change to solidify the spreading pattern is potentiated by feedback; at the same time, the more variance increases at the other level, the less similarity-bias brakes further change. When we view similarity-bias as a form of selection on the range of possible variants that enter the linguistic system over time (Wedel 2006), the interaction between overlapping levels of organization in the lexicon can be understood as a form of multi-level selection. As an illustration, I presented results from a simple simulation showing that in a system in which errorful pattern extension is a primary pathway of change, competition between a static regularity and a relational regularity resulted in the rapid stabilization of one over the other, in part by modulating variance at distinct levels of organization. Further, if a new pattern establishes itself in a subset of the lexicon, the existence of this new generalization potentiates development of a similar pattern in other, similar words.

References

- Albright, Adam and Bruce Hayes (2002). Modeling English Past Tense Intuitions with Minimal Generalization. In Michael Maxwell, (ed.) *Proceedings of the Sixth Meeting of the ACL Special Interest Group in Computational Phonology*. Philadelphia, July 2002. ACL.
- Alegre, Maria and Peter Gordon (1999) Frequency effects and the representational status of regular inflections. *Journal of Memory and Language* 40: 41-61.

- Baayen, R. H. (1992). Quantitative Aspects of Morphological Productivity. In Booij, Geert. and van Marle, Jan. (eds.), *Yearbook of Morphology 1992*. Kluwer Academic Publishers. 181–208.
- Blevins, Juliette (2004). *Evolutionary Phonology: The emergence of sound patterns*. Cambridge: Cambridge University Press.
- Blevins, Juliette. (2006). New perspectives on English sound patterns: 'Natural' and 'Unnatural' in Evolutionary Phonology. *Journal of English Linguistics* 34:6-25.
- Blevins, Juliette, and Robert Blust (2003). Towards a theory of drift: Patterns of Austronesian syncope. Ms. University of California, Berkeley.
- Blevins, Juliette and Andrew Garrett (1998). The origins of consonant-vowel metathesis. *Language* 74: 508-56.
- Burzio, Luigi (2005). Sources of Paradigm Uniformity. In Laura J. Downing, T. Alan Hall, Renate Raffelsiefen, (eds.), *Paradigms in Phonological Theory*, pp. 65-106. Oxford: Oxford University Press.
- Bybee, Joan (1985). *Morphology: A study of the relation between meaning and form*. Amsterdam: J. Benjamins.
- Bybee, Joan., and Carol L Modor. (1983). Morphological Classes as Natural Categories, *Language* 59: 251-270.
- Chitoran, Ioana and Jose Hualde (2007). On the origin and evolution of the contrast between diphthongs and hiatus sequences in Romance. *Phonology*
- Eisner, Frank and James McQueen (2005). The specificity of perceptual learning in speech processing. *Perception and Psychophysics* 67: 224-238.
- Ernestus, Mirjam and R. Harald Baayen (2003). Predicting the unpredictable: Interpreting neutralized segments in Dutch. *Language* 79, 5-38.
- Goldinger, Stephen D. (2000). The role of perceptual episodes in lexical processing. In Cutler, A., McQueen, J. M, and Zondervan, R. *Proceedings of SWAP (Spoken Word Access Processes)*, Nijmegen: Max-Planck-Institute for Psycholinguistics.
- Guenther, Frank H., and Gjaja, Marin. N. (1996). The perceptual magnet effect as an emergent property of neural map formation. *Journal of the Acoustical Society of America* 100, 1111-1121.

- Guenther, Frank H., Alfonso Nieto-Castanon, Satrajit S. Ghosh, and Jason A. Tourville (2004). Representation of sound categories in auditory cortical maps. *Journal of Speech, Language, and Hearing Research* 47, 46-57.
- Hare, Mary. and Jeffrey L. Elman (1995). Learning and morphological change. *Cognition* 56, 61-98.
- Harrington, Jonathan, Sallyanne Palethorpe and Catherine I. Watson. (2000). Does the Queen speak the Queen's English? *Nature* 408, 927-928.
- Hock, Hans H. (1991). *Principles of Historical Linguistics* (2nd ed.). The Hague: Mouton de Gruyter.
- Hock, Hans H. (2003). Analogical Change. In: Brian Josephs and Richard Janda, (eds.) *The Handbook of Historical Linguistics*. Blackwell 441-480.
- Johnson, Keith (1997). Speech perception without speaker normalization. In Keith Johnson and John W. Mullennix (eds.) *Talker Variability in Speech Processing*. San Diego: Academic Press. 145-166.
- Krott, Andrea., Robert Schreuder. and R. Harald Baayen (2002). Linking elements in Dutch noun noun compounds: constituent families as analogical predictors for response latencies. *Brain and Language*, 81 (1-3), 723-735.
- Kuehne, Sven. E., Kenneth D. Forbus, Dedre Gentner and Brynn Quinn (2000). SEQL - Category learning as progressive abstraction using structure mapping. *Proceedings of the Twenty-second Annual Conference of the Cognitive Science Society*, 770-775.
- Kuhl, Patricia K. (1991) Human adults and human infants show a 'perceptual magnet effect' for the prototypes of speech categories, monkeys do not. *Perception and Psychophysics*. 50: 93-107.
- Kuhl, Patricia. K. (1995). Mechanisms of developmental change in speech and Language. In Kjell Elenius and Peter Branderud (eds.), *Proceedings of the XIIIth International Congress of Phonetic Sciences*, Vol. 2. Stockholm: Stockholm University, 132-139.
- Long, Christopher and Amit Almor (2000). Irregularization: The interaction of item frequency and phonological interference in regular past tense production. *Proceedings of the twenty-second annual conference of the cognitive science society*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Mayr, Ernst. (1997). The objects of selection. *Proceedings of the National Academy of*

Sciences 94: 2091-2094.

Mielke, Jeff (2004). The Emergence of Distinctive Features. PhD Dissertation, Ohio State University.

Pierrehumbert, Janet (2001). Exemplar dynamics: Word frequency, lenition, and contrast. In Joan Bybee and Paul Hopper (eds.) *Frequency effects and the emergence of linguistic structure*. John Benjamins, Amsterdam. 137-157.

Pierrehumbert, Janet (2003). Phonetic Diversity, Statistical Learning, and Acquisition of Phonology. *Language and Speech* 46: 115-154.

Pierrehumbert, Janet (2006). The new toolkit. *Journal of Phonetics* 34: 516-530.

Sturtevant, Edgar H. (1947). *An Introduction to Linguistic Science*. New Haven: Yale University Press.

Taylor, Douglas, Clifford Zeyl and Erin Cooke. (2002) Conflicting levels of selection in the accumulation of mitochondrial defects in *Saccharomyces cerevisiae*. *Proceedings of the National Academy of Sciences* 99: 3690-3694..

Tenpenny, Patricia. L. 1995. Abstractionist versus episodic theories of repetition priming and word identification. *Psychonomic Bulletin and Review* 2: 339-363.

Wedel, Andrew (2004). *Self-organization and categorical behavior in phonology*. Santa Cruz CA: University of California Santa Cruz Ph.D Dissertation.

Wedel, Andrew. (2006). Exemplar models, evolution and language change. *The Linguistic Review* 23: 247-274.

Wedel, Andrew (2007). Feedback and regularity in the lexicon. *Phonology*

Zanone, Pier G., and Kelso, J. A. Scott (1997). The coordination dynamics of learning and transfer, Collective and component levels. *Journal of Experimental Psychology, Human Perception and Performance* 23, 1454–1480.