The statistical basis of an unnatural alternation

Janet B. Pierrehumbert

The /k/-/s/ alternation in word pairs such as electric, electricity is not phonetically natural, and is learned by English speakers as a generalization over morphological relatives. Data gathered in an open-response experimental task show that it is productive before the suffix -ity. Lesser productivity is found for the same materials in a backformation task. Outcomes are analyzed as the result of a cognitive process of statistical inference. Abstract descriptions make crucial use of phonological variables. Cognitive preferences for certainty and for robust, redundant descriptions are argued to jointly determine the universe over which operational statistics are estimated.

1. Introduction

This paper presents an experimental study of the productivity of the /k/-/s/ alternation exhibited in derivational pairs such as *electric*, *electricity*. It is exemplified in numerous word pairs in English involving various suffixes, including *-ity*, *-ism*, and *-ist*. The Collins on-line English dictionary (distributed in 1990 through the ACL Data Collection Initiative) includes 108 clear examples of words ending in these suffixes in which a stem-final /k/ softens to /s/. The largest group, and the main topic of this paper, is words formed with *-ity*. The dictionary contains only twelve words with stems ending in /k/ which fail to soften before one of these three suffixes (e.g. *anarchy*, *anarchism*, *York*, *Yorkist*). All involve affixes other than *-ity*.

The productivity of the alternation is disputable. First, there are very few forms which would support extension of the alternation beyond an orthographic *-ic* followed by one of the triggering suffixes. For the suffixes just listed, the only common examples listed in the Collins are *Greek*/*Grecism*; *opaque*, *opacity*; *reciprocal*/*reciprocate*, *reciprocity*; and *pharmacology*, *pharmacist*. (A number of potentially relevant pairs, such as *caducous*, *caducity*; *cecum*, *cecity*; *paucal*, *paucity*; *raucous*, *raucity* would only be known to very erudite speakers.) Second, as Myers (1999) also notes, the

/k/-/s/ alternation as presently found in English is not natural (in the sense of Anderson 1981).

Velar Softening is not phonetically natural because the evident phonetic pressures on a /k/ in the target position would not produce /s/. If the suffix vowel is /ɪ/ (as transcribed in the dictionary) then coarticulation and lenition would yield an aspirated palatal approximant rather than the alveolar fricative /s/ (see Lavoie 2001). The alternative possibility for the vowel, /ə/, provides still weaker phonetic motivation for /s/. /s/ differs from the phonetically expected outcome by its maximal vocal fold abduction and its precise tongue shaping, which directs a jet of air against the teeth. These are active adjustments which cannot be characterized as accommodation to a following vowel. Thus, understanding the alternation of /k/ with /s/ requires recourse to some version of the concept of Structure Preservation in phonology (see Kiparsky 1985) which states that lexical alternations stay within a language's system of phonological categories. Since an aspirated palatal approximant is not a contrastive category in the English lexicon, it cannot be the outcome of a morphophonological rule, either. The reanalysis involved in lexicalizing the phonetically expected approximant as the lexically contrastive segment /s/ reveals the role of abstract cognitive factors, over and beyond phonetic ones.

Velar Softening is also unnatural because it is phonologically opaque. Though it originates historically in fronting and spirantization of the velar stop before a non-low front vowel, suffixes with such vowels on the surface do not in general trigger the softening of /k/ to /s/ in the synchronic phonology. /k/ never softens to /s/ before *-y*, as *smoke*, *smoky*. On the other hand, *-ize*, beginning with a low vowel, does trigger softening, because *-ize* formerly had a nonlow front vowel. In Chomsky and Halle (1968), this historical ordering is recapitulated in the extrinsic ordering of rules in the synchronic phonology. The phonological opacity created by such orderings is precisely one reason that the psychological validity of the Chomsky-Halle model became a matter of widespread dispute. The finding that the vowel shift is only partially productive (c.f. Jaeger 1984; McCawley 1986) also calls into question the productivity of the rule of Velar Softening, which is ordered before it.

Understanding productivity is important because it provides a crucial line of evidence about cognitive abstractions. The failure of an alternation to generalize suggests that no abstract generalization over the forms exhibiting the alternation has been formed. If the alternation is aggressively and reliably extended, even to forms which differ substantially from attested

ones, it follows that a very broad abstraction has been formed. For example, the reliable and aggressive extension of the regular English plural pattern indicates that it abstracts away from many properties of the word. If the situation lies somewhere in the middle, then the exact pattern of productivity can yield insights about the exact character of the abstraction that is formed.

Phonotactics is the area in which most research has been done on the availability of lexical patterns for use in novel forms. Numerous studies, reviewed in Pierrehumbert (2003), indicate that the type frequency (frequency in the lexicon) of a phonological pattern affects the likelihood and perceived well-formedness of novel words containing that pattern. This dependence is gradient; frequent sequences readily extended to new words, rare sequences are avoided, and moderately frequent sequences fall in between. For example, Hay, Pierrehumbert, and Beckman (2004) found that the perceived well-formedness of novel words containing nasal-obstruent clusters (such as /strɪnfi/ and /zæmpɹ/) was a gradient function of the frequency of the cluster. The frequency for a tautomorphemic cluster was estimated as its frequency in trochaic monomorphemic words with a lax front vowel in the CELEX monomorphemes. (see Baayen, Piepenbrock and Gulikens 1995, regarding CELEX; Hay, Pierrehumbert and Beckman 2004, regarding monomorphemes).

This choice of universe for estimating frequencies was opportunistic, and obscures a central issue in understanding the relation of lexical frequencies to pattern productivity. This issue is taken up with Figure 1. Figure 1 shows a partial lattice of heterosyllabic N.O clusters. The atoms on the bottom are individual heterosyllabic phoneme clusters. The nodes above the atoms are some of the various available natural classes of such clusters. As is well-known, natural classes can be formed using partial descriptions of phonological patterns. For example, the sequence /nt/ is an element of the set of clusters of /n/ followed by any stop; it is also an element of the set of clusters containing a homorganic nasal and stop. The cluster /np/ belongs to the former set but not the latter; the cluster /mp/ belongs to the latter set but not the former. The lattice is organized from specific (on the bottom) to general (at the top). Each node is labeled with the probability of the indicated descriptor with respect to the universe of N.O clusters, as estimated from counts in the CELEX monomorphemes. Clearly, the less specific the description, the more cases it encompasses and the larger the natural class it describes. Thus, the probabilities go up as we follow the lines up the lattice, but the exact way they go up depends on exactly what is lumped together in each class.

The topmost case, any nasal followed by any obstruent, is taken to define the universe for the probabilities which are indicated below each node. If the universe were larger, the probabilities shown on the figure would all be smaller, but their rankings would remain the same.

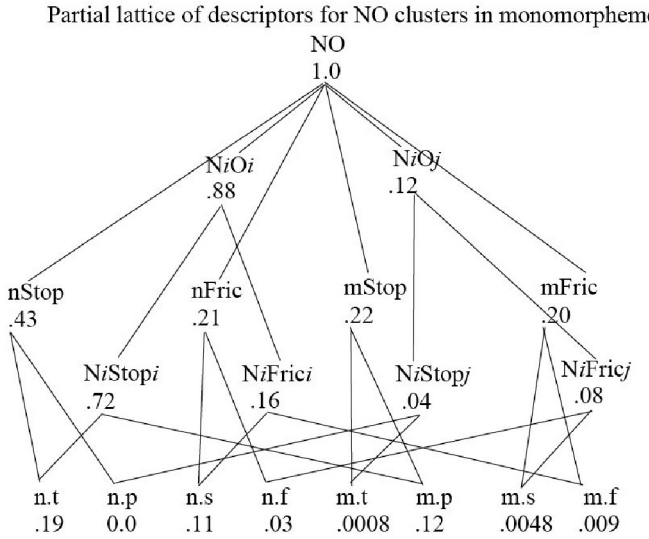

Figure 1. Partial lattice of probabilities for the universe of heterosyllabic nasal-obstruent clusters in monomorphemic words of English. Counts are established with respect to the Celex monomorphemes, as discussed in Hay et al. (2004).

Atoms at the bottom of the lattice represent specific nasal-obstruent sequences. For the sake of legibility, only eight atoms with a total P = 0.54 are shown. The remaining sequences, with a total P = 0.46, have been omitted. Superordinate nodes represent classes of nasal-obstruent sequences. Italic indices are used for convenience to indicate homorganicity or nonhomorganicity. (Actual phonological structures for homorganic consonants have feature sharing.) Capital N denotes any nasal. Probabilities of superordinate categories include frequencies of atoms which are not shown at the bottom, but which are properly described by the superordinate note. Superordinate nodes /ŋStop/ and /ŋFric/ are not shown.

The question raised by Figure 1 is: Of all the probabilities which may be defined using partial phonological descriptions of a pattern, which are relevant to productivity and perceived well-formedness? For example, if the

description N_iO_i (p = .88) versus N_iO_j (p = .12) were the relevant level of abstraction then the feature [+/– continuant] would have no importance for the evaluation and productivity of these clusters, and inhomorganic nasal-fricative clusters would seem every bit as bad as inhomorganic nasal-stop clusters. However, results in Hay et al. (2004) clearly show differential outcomes for nasal-stop and nasal-fricative clusters, indicating that this degree of generalization is too great. The next level down shows two alternative ways to break out the cases. One fixes the nasal consonant and generalizes over the following the obstruent. The other generalizes over place. It separates homorganic from inhomorganic clusters regardless of place, but it maintains information about the continuance of the obstruent. If the first alternative were the cognitively relevant description, then /np/ would be as acceptable as /nt/. This is false. The second alternative also groups /nt/ in a class with other clusters, namely /mp/ and /ŋk/ but not /np/. It is closer to the true state of affairs, since it captures both a strong effect of homorganicity on nasal-stop clusters and a weaker effect on nasal-fricative clusters. However, the lattice structure does not in itself say that one line of generalization is more relevant than the other. Although any phonologist would sensibly prefer one line of generalization to the other, there is no explicit formal account of what this "sensibleness" consists of. Still less is there an explanation of why subjects operated at a detailed level of description, rather than applying a simple overarching generalization about all N.O clusters.

The same issue arises in a different guise in dealing with morphophonological alternations. When such alternations are language particular, they must be learned from examples. There is by now abundant evidence that the productivity of an alternation depends on its type frequency (as well as on other factors). Alternations found in extremely few types, such as irregular conjugations for auxiliaries, are not productive no matter how frequently the irregular forms may be used. However, the universe of examples relevant for any given alternation, and the types of formal generalizations which are made over these examples, is not well understood.

In comparison to phonotactics, morphophonology provides both challenges and opportunities in addressing this issue. It is challenging because many morphophonological alternations are generalizations over word pairs rather than merely over words. For phonotactics, set theory provides a convenient hierarchy of abstraction over words, taking the shape of a lattice of partial descriptions as in Figure 1. For word pairs, in contrast, the proper formal toolkit is not as evident. Is it partial descriptions of the base which are relevant? Or partial descriptions of the complex form? Or relations of partial

descriptions of the base and the complex form? The research literature contains case studies which argue for all of these possibilities. As discussed in Myers (1999) and below, the /k/-/s/ alternation needs to be formalized with respect to word pairs; but this need does not in itself define the relevant universe of word pairs. For determining the pronunciation of a novel form *clemicity* given the base *clemic*, the pair *conic*, *conicity* is plainly relevant. But how about *Turk*, *Turkism* (involving a different affix)? Or *ferocious*, *ferocity* (for which no bare form of the base exists)? Or *morbid*, *morbidity* (illustrating preservation of a final stop before the same affix)?

In short, the expectation (based on results in phonotactics and in psychology) that implicit knowledge of morphophonology is stochastic does not in itself define what probabilities will be relevant. Probabilities can be estimated for any formal description that can be tabulated. Many of the "analytic biases" mentioned in Steriade's (2002) original commentary on this session can be viewed as claims about what statistics are available to the cognitive system. For example, in discussing Goldrick (2002), she suggests an analytic bias to the effect that voicing pairs should alternate alike. This is equivalent to the claim that statistics on formal descriptions which abstract over place, but not voicing, are highly available in the formation of phonological grammars.

The investigation of unnatural alternations provides special opportunities in understanding how people form abstract generalizations, because it sidesteps one of the most recalcitrant problems of phonology. This is the relationship of frequency to the phonetic foundations of phonological systems. Under the rubric of markedness theory, scholars have long observed that phonetic simplicity is related to frequency and to default phonological behavior. Quantitative phonetic models have now gone some distance towards elucidating the articulatory and psychoacoustic basis for more and less common segment types, and similar arguments can also be made about phonological sequences. For example, the tendency for languages to favor homorganicity between a nasal and a following obstruent is agreed to be founded in the tendency towards gestural overlap between successive consonants. In the light of such research, there is a risk of confusing correlation and cause when interpreting experimental findings such as Hay, Pierrehumbert and Beckman (2004). The high correlation ($r^2 = 0.65$) they report between lexical log frequency and perceived well-formedness could in principle arise from a concealed factor, namely markedness. Possibly, the phonetically simpler clusters are judged to be better because they are simpler and they are also used more often in words because they are simpler. This would lead to a

correlation between lexical frequency and perceived well-formedness, even in the absence of any ability to learn lexical frequencies.

Arguing against this viewpoint are reported findings of dissociations between universal markedness and stochastic generalizations within specific languages. The phoneme /t/ is rare in Arabic despite being unmarked (see counts in Frisch, Pierrehumbert and Broe 2004). Whalen (in press) observes that clicks are common in languages which have them, despite being marked. The learnability of language-specific sequential statistics is shown by comparative studies such as Cutler and Otake's (1998) findings on NO clusters in Japanese versus Dutch. Thus, empirically observed frequency effects are not in general reducible to markedness effects. However, present knowledge is very incomplete and this is an important issue for further research. Work on unnatural processes can make a contribution to this research by permitting an examination of frequency effects in a area where the phonetic foundation is poor. Since the /k/-/s/ alternation is neither pervasive in English nor ubiquitous across languages, any frequency effects which are observed can be presumptively attributed to the the learning from experiences with words of English.

2. Methods

The experiment uses a wugs paradigm, pioneered in Berko's (1958) experiments with children. In this paradigm, subjects are taught a novel stem and they use it as the base for a complex form. "Here is a wug. Look, now there are two of them. There are two ?????". This paradigm has been widely used to investigate the competition between regular and irregular inflectional forms; see, for example, Bybee and Pardo (1981) and Albright and Hayes (2003). Here, it is extended to derivational morphology, an extension also made in Zuraw's (2000) study of Tagalog morphology.

Most early studies of derivational morphology, such as the studies of the English Vowel Shift presented in Jaeger (1984), and McCawley (1986), use concept formation tasks or judgments of words presented in pairs, rather than the wugs paradigm. These tasks have the potential drawback of priming awareness of the regularity being studied through the very design of the stimulus materials. The wugs paradigm, with its open-response format, is more conservative. The materials for this study did not provide any examples of a /k/-/s/ alternation and the subjects were unaware that this alternation was being investigated.

Two related experiments are reported. In both, the stimuli were two sentence paragraphs. The first sentence introduced a target word. The subject's task was to supply the morphologically related word missing in the second sentence. For one group of subjects, the first sentence introduced a base adjective, and the task was to create an abstract noun. The instructions mentioned the variety of means in English for turning adjectives into nouns, such as affixing *-ity* (as in *virgin bride*/*virginity*) and affixing *-ness* (as in *bright*/*brightness*). Subjects were told to make a noun in any way they wished, and to respond as soon as an idea occurred to them. For the second group, the format of the materials was reversed, and the subject's job was to backform the base adjective from an abstract noun. Subjects were young adults recruited through Northwestern University and Ohio State University. Some of the subjects were members of a subject pool comprised of students in introductory linguistics courses, and others were paid $8. There were 10 subjects in the noun formation task and 7 in the backformation task. No subject did both tasks.

For both groups, the instructions and the materials were presented entirely orally. There were 64 items: 16 baseline items (extant words with an established nominal form ending in *-ness* or *-ity*); 16 fillers (existing and novel words that present some uncertainty between *-ity* and *-ness*), and 32 novel target words. None of the example items, baseline items, or fillers involved /k/, /s/, or an alternation between /k/ and /s/. The target word in each stimulus was the last or next-to-last word in the sentence. The same set of bases figured in both the noun formation and the backformation task. A full listing of baseline, target, and filler words can be found in the appendix, including IPA transcriptions for nonwords used in filler and target items.

Of the 16 baseline items in the experiment, eight were words for which there is an established noun in *-ness* and no established noun in *-ity*, such as (1).

(1) When Anna discovered a new doughnut shop, she was very happy. For her, a warm doughnut means ?????.
(ANSWER: happiness, *happyity)

Eight were words for which there was an established form in *-ity*, with the *-ness* form, if any, having an inappropriate meaning.

(2) Bob's short-term bonds were among his most liquid assets. After he got arrested, he was able to post bail because of his high ?????.
(ANSWER: liquidity, *liquidness)

The sixteen filler items were evenly divided between real words and non-words. For both, *-ness* and *-ity* forms were possible, as indicated by pilot testing.

(3) My brother has always been very frugal. Reusing aluminum foil is just one symptom of his ?????.
(ANSWER: frugalness OR frugality)

(4) Anthropologists working in Manuka found all the hallmarks of a caustive society. In fact, it became a textbook example of ?????.
(ANSWER: caustiveness OR caustivity).

There were four different types of target items, differing in their prototypicality as hosts for a /k/-/s/ alternation. Eight items were Latinate pseudowords, ending in the phonetic form of the suffix *-ic*, /ɪk/). In the pronunciation used by the experimenter, this is a front schwa, as in the well-known minimal pair *roses* /ɹoʊzɪz/, *Rosa's* /ɹoʊzəz/. All Latinate pseudowords were polysyllabic, and some suggested existing words through their prefixes or stems.

(5) Halley's comet is a very interponic comet. Its orbital period varies because of its ?????.

The second set of eight target items, the semi-Latinate set, had a main stress on the initial syllable and a secondary stress on the last syllable. These items also ended in /k/. The (unreduced) vowel in the last syllable was /ɛ/, /æ/ or /a/. Thus, the last syllable clearly differed from *-ic*.

(6) Before Pierre stood an electrifyingly hovac sculpture. In his entire career as curator, he had never before seen such a perfect example of ????? .

The third set, the non-Latinate bases, were monosyllabic pseudowords (in some cases with a prefix *over-* or *under-*).

(7) Inside, the light was so dim it was entirely mork. We couldn't read the instructions in the ?????.

The stimuli also included a fourth set of target items, non-Latinate pseudowords ending in /s/. These are not eligible for velar softening, but are included because they are needed in the backformation task to ensure a balance between /...knɛs/ and /...snɛs/ forms.

In the backformation condition, the two sentence paragraphs were reworked so that the first sentence introduced a complex noun, and the second sentence had a missing adjective. The semantic content of each paragraph was minimally modified so as to maintain the contexts for the forms. For example, the reversed version of (5) is (8) and the reversed version of [6] is (9).

(8) The period of Halle's comet varies because of its interponicity. It is a very ????? comet.

(9) In Pierre's entire career as a curator, he had never before seen such a perfect example of hovacity. It was an electrifyingly ????? sculpture.

Stimuli were block randomized in eight blocks of eight, and read aloud to each subject individually. Subjects repeated the target adjective out loud during a pause after the first sentence. Their pronunciation was corrected if necessary. Intersubstitutions of /æ/ and /ɛ/ in words such as *bowdec* and *hovac* were, however, accepted, as some speakers have apparently merged these vowels. The responses were transcribed as they occurred. The entire session was recorded, and the recordings were used to resolve the few uncertainties in transcription.

To score the noun formation data, the subset of responses in which subjects selected the affix *-ity* after a stem ending in /k/ was extracted. The frequency with which /k/ is softened to /s/ is computed on this subset of "hits". The size of this subset differed considerably across subjects due to the open response format. The hits for the backformation task are the responses in which all and only the affix *-ity* was removed from a noun ending in /sɪti/. The frequency with which the bare stem was produced with final /k/ (as opposed to /s/ or some other consonant) was computed on this set.

Subjects in both groups understood the instructions and generally succeeded at the task. Debriefing after the noun formation task revealed that only one subject was able to guess before the end of the experiment that the /k/-/s/ alternation was being investigated. In the backformation task, no subject guessed what was being scored in the experiment.

3. Results

In the noun formation condition, subjects produced *-ity* and *-ness* responses about equally often. 80% of baseline items predicted to have *-ity* did indeed have it. 82% of baseline items predicted to have *-ness* had *-ness*. 53% of fillers were produced with *-ity* and 47% with *-ness*.

The ten subjects in the noun formation task produced a total of 71 hits. Of the 10 subjects, eight produced examples of velar softening, with the number of examples ranging from 4 to 13 per subject. These results indicate that velar softening is productive for most educated adults. Results by target type are shown in Table 1. For Latinate hits, softening applied nearly 100% of the time. It was somewhat less reliable for the semi-Latinate hits, but the sample size is not big enough to be confident of a difference. In the few cases in which *-ity* was attached to a non-Latinate target, softening never applied. Though the number of such cases was small, the lack of softening is so readily confirmed by native speaker judgments that I will view it as a fact which needs to be explained. Thus, the main effects which require explanation are the high productivity of softening for the semi-Latinate stems, a group for which there is no critical mass of extant forms, and the lack of softening in the non-Latinate stems (given that softening was observed in the Latinate and semi-Latinate stems).

Table 1. Outcomes in the noun formation task.

Target Type	Hits	Softening Before *-ity*	% Softening
Latinate targets	30	28	93
Semi-Latinate targets	36	30	83
Non-Latinate targets	5	0	0

Subjects were generally successful on the backformation task. There were only two errors on the baseline items (*inanity* → *inate* and *profanity* → *profound*). Over all items, subjects produced a form with a bare stem 86% of the time. (Other responses involved either addition of a suffix, stem truncation, or lexical intrusions.) A total of 68 hits reflect an implicit choice to preserve the surface /s/ or to undo velar softening to yield /k/. Results for these forms are shown in Table 2. In interpreting this table, recall that the materials did not include non-Latinate targets ending in *-ity*. The non-Latinate targets all involved /k/ or /s/ before *-ness*.

Table 2. Outcomes in the backformation task.

Target type	Hits	Backformations to /k/	% /k/ Responses	% /s/ Responses
Latinate targets	32	6	18	82
Semi-Latinate targets	36	5	13	87

The fact that some reversals of /s/ to /k/ occur is evidence of the psychological reality of the /k/-/s/ alternation. However, the rate of back-formation of /k/ from /s/ is much lower than the rate of velar softening in the noun formation task. Two subjects out of seven were responsible for all cases of backformation of /s/ to /k/. These backformed at rates of 33% and 75%, respectively. (Both produced examples of /k/ for both Latinate and semi-Latinate bases). The finding that backformation to /k/ is less frequent and reliable than softening of /k/ to /s/ requires explanation.

4. Discussion

The experiment showed that the /k/-/s/ alternation is productively applied in an open-response task. However, it not completely productive; it fails to apply to non-Latinate stems and for backformation, there is large variability across subjects. My goal will be to explain this exact pattern of productivity as a reflex of statistical learning over patterns in the lexicon. In exploring this issue, I will make several simplifying assumptions. One is that the relevant probabilities can be approximated over word pairs involving the exact affix in the experiments, *-ity*. This assumption is made because the relevance of exceptions to velar softening involving other affixes is unclear; notably, Ohala's (1974) experimental study of *-ism* reports a much lower (30%) productivity level for Velar Softening, possibly as a consequence of the rather many exceptions involving this suffix. A second assumption is that any given word pair either is, or is not, in the universe over which a probability is defined. I consider only in passing models in which the importance of a word pair is weighted on a scale by its similarity to the current target.

A straightforward extension of previous experimental studies of phonotactics would seek to identify stochastic constraints on word form to which the products of *-ity* affixation must conform. This extension would be in the spirit of Optimality Theory, as well as of many prior studies in prosodic morphology, in seeking to explain morphophonological alternations as the

result of constraints that are generally true of the language. However, it is not the case that the consonant preceding word-final /ɪti/ (or the morpheme *-ity*), is usually /s/. In the Collins on-line dictionary, only 25% of words ending in /ɪti/ end in /sɪti/. /l/ is more common than /s/, although even /l/ does not achieve a majority of the forms. Simple frequency matching on the surface forms would predict that subjects would tend to substitute /l/ for /k/ (or for any other consonant!) but only at a rate of about 25%. Surface statistics do not explain the extremely high rate of substitution observed for the Latinate and semi-Latinate stems ending in /k/, or the failure of other consonants to be affected in the same way. Similarly, the responses by two subjects in the back-formation task also appear to reflect implicit knowledge of specific morphological relationships. Without knowledge of such relationships, there would be no reason to backform /s/ to /k/, since s/ is both more faithful to the stimulus and more common in word-final position.

As a result, knowledge of the alternation must be a generalization over morphologically related word pairs. This conclusion echoes the treatment of velar softening as a derivational rule in Chomsky and Halle (1968). In Optimality Theory, constraints generalizing over word pairs have been used since McCarthy and Prince (1995) proposed them in order to overcome the limitations of constraints over word forms in explaining the behavior of reduplication. Generalizations over word pairs also figure in the non-OT literature on computational morphology, notably Skousen (1989), Daelemans et al. (1999), Ernestus and Baayen (2002), Baayen (2003), and Albright and Hayes (2003). These works all share the assumption that variable outcomes in morphophonology are related to conditional probabilities defined on word pairs.

The acknowledgment that the alternation is learned as a generalization over word pairs goes far towards explaining the amount of variation observed across individuals. Individuals differ both in the size and the contents of their vocabularies. To know a relevant example of an alternation, they must know both words in the pair. Furthermore, they must view them as related to each other. Not everyone infers a decomposition of *Mediterranean* on the basis of words such as *medium* and *terrestrial*. The difficulties of assessing such implicit semantic relationships mean that most computational studies, including the present one, rely on phonological matching in large dictionaries and on morphological analyses by linguists. They probably overestimate the pool of relevant word pairs known to the subject pool.

In the following discussion, I will be particularly concerned with the claim, advanced in the analogical models of Skousen (1989) and Baayen (2003),

that the productivity of morphophonological alternations in new forms is determined by a general statistical inference. Unlike Chomsky and Halle (1968), these models draw no fundamental distinction between morphological derivation and back-formation. They therefore make precise predictions in both directions. The pronunciation of an unknown form is inferred from a related known form in the light of a universe of known word pairs exemplifying the same relationship. The models set up analogies in which the unknown variable in the analogy may be either a base or a derived form: they pose equally questions such as *conic* : *conicity* :: *clemic* : *?* and questions such as *conicity* : *conic* :: *clemicity* : *?* . An analogical approach does predict that an alternation may display different rates of productivity for morphological derivation and back-formation, because the probability of A given B, and the probability of B given A are mathematically distinct and often have different values. Exact predictions about the outcome probabilities depend on the exact assumptions about the universe of generalization over which the probabilities are estimated. Thus, the key issue is how the universe of generalization is established, and why the cognitive system takes some universes of generalization to be the operational ones, as opposed to others which are equally available from a mathematical point of view.

The following discussion is premised on three hypotheses about the operational level of generalization, which bring together some of the threads of the literature on morphological processing by people and by machines.

HYPOTHESIS 1: All other things being equal, the cognitive system prefers generalizations which yield more certainty about the outcome to those which yield less certainty.

This claim, a plain language statement of the information-theoretic proposals of Daelemans et al. (1999), means that descriptions of the data which are associated with extreme probabilities are more relevant than ones which characterize the outcome as a random choice. The extreme probabilities of 1.0 and 0.0 provide complete certainty; an outcome with probability 1.0 is the only one possible, and one with probability 0.0 is absolutely impossible. For a two-way choice, a probability of 0.5 represents complete uncertainty, providing no information either way.

Typically, hypothesis 1 will tend to favor generalizations based on small sets of words over generalizations based on bigger sets, as smaller sets tend to be more homogeneous (to exhibit more uniform outcomes) and bigger sets tend to be more heterogeneous (to exhibit more diverse outcomes). But

this is not always the case, as the study of learning of English verb morphology by Derwing and Skousen (1994) indicates. They successfully apply Skousen's (1989) AML approach (Analogical Modeling of Language), which anticipates the conclusions of Daelemans et al. (1999) by providing an algorithm for automatically growing smaller analogical sets to bigger ones exactly when the homogeneity of the universe is not compromised. Albright and Hayes (2003) also adopt information-theoretic weighting of generalizations over examples.

HYPOTHESIS 2: All other things being equal, the cognitive system prefers generalizations based on larger sets of examples to those based on smaller sets.

This preference is justified because increasing the sample size increases the reliability of the estimate of the probability of a pattern. In a morphological analyzer directed towards automatic part-of-speech tagging, Mikheev (1997) brings together premises 1 and 2 by assigning rule scores based on the lower edge of the 90% confidence interval for the probability associated with the rule. For rules which positively specify the nature of the outcome, this number increases both with the estimated probability and with the size of the sample from which the probability is estimated.

A further claim made in Mikheev (1997) is:

HYPOTHESIS 3: All other things being equal, longer phonological descriptors are preferable to shorter ones.

This claim is directly at odds with the assumption of classical generative grammar that the structural descriptions in rules should be as simple and general as possible. Notably, the Velar Softening rule in Chomsky and Halle (1968) is maximally general, targeting all biphonemic sequences with certain distinctive features. It states that in derived environments,

(10) [–anterior, –continuant, <–voice >] → [–back, <+anterior>]
/ ___ [–back, –low, –cons]

(10) groups together the /k/-/s/ alternation and the alternation of /g/ with /dʒ/. It neglects potential conditioning by the specific suffix involved, and by the etymological class, length, or structure of the stem. It thus predicts softening in neologisms such as *taskism*. It presupposes that all surface ex-

amples of the /s/ variant before low or back vowels (such as *focus*, *foci*) involve vowel shifting from underlying nonlow front vowels. In the interests of unifying the analyses of pairs such as *opaque*, *opacity*; *critic*, *criticize*; *analogue*, *analogy*, Chomsky and Halle thus pursued an aggressive program of generalization and abstraction and developed a complex theory of rule interaction.

Mikheev's hypothesis receives support from recent experimental work in phonetics and psycholinguistics. Numerous results indicate that cognitive representations are more redundant than was imagined in the early days of generative grammar; see discussion in Baayen (2003), Broe and Pierrehumbert (2000), and Bybee (2001). These results bear on the present data in suggesting that phonological characteristics which are true of all word pairs exhibiting an alternation would be maintained in the general template for that alternation even if they are redundant with respect to the universe over which the generalization has been learned.

Given these hypotheses, the high rate of velar softening for Latinate and semi-Latinate stems can be readily explained. Extant Latinate stems provide a core of multisyllabic stems ending in /ɪk/, for which velar softening occurs at P=1.0 before the affix -*ity*. By hypothesis 1, this is a highly relevant generalization. Expanding the set to include all multisyllabic stems ending in /k/ (e.g to also include the few pairs such as *opaque*, *opacity*, *reciprocate*, *reciprocity*) increases the sample size without compromising the reliability of the generalization, since the probability of softening is still 1.0. However, dropping the shared prosodic features, such that the relevant universe is simply words ending in /k/ with a counterpart in -*ity*, simplifies the description without expanding the sample size. The sample size remains the same because there are no monosyllabic words ending in /k/ with a derivative in -*ity*. By hypothesis 3, then, the generalization stops short of this simplification. Hence, its structural description is not met in novel forms such as *bleckity*, and it fails to apply. A general implication of Mikheev's claim is that morphological alternations would generalize to forms which lie in the cracks between existing forms, so to speak, but not to forms which break entirely new ground.

This analysis does not explain why productivity of velar softening was higher for Latinate than for semi-Latinate forms. If the small observed difference proves replicable, it will provide an example of a prototypicality effect. The Latinate words are extremely similar to the core of the distribution for the alternation, the many forms involving stems in /ɪk/. The semi-Latinate forms are less similar. A trading relationship between similarity and frequency is

widely attested, and as discussed by Dell (2000) is a hallmark of connectionist models. However, the same trading relationship is also characteristic of stochastic analogical models. In all of these approaches, the behavior of prototypical examples is observed only part of the time for nonprototypical examples, because the nonprototypical examples can also be captured by the pattern of a competing generalization. For the present case, the competing generalization is either a generalization over multiple affixes (a complexity which exceeds the scope of this paper), or a default, namely that of maintaining the base form consonant without modification. This default is one to which we will return in connection with the backformation data.

Turning now to the backformation data, a generalization at the same level relates words of four or more syllables ending in /sɪti/ to words of two or more syllables ending in any consonant. This universe includes everything in the universe for the noun formation task, plus pairings such as *diverse*, *diversity* and *porous*, *porosity*. The probability of /k/ in the stem given /s/ in the derived form is 0.42 over this universe. The minority of /k/s is due to the greater number of examples such as *porosity*, *immensity*, in which the surface /s/ really does relate to a base /s/. The probability of 0.42 represents considerable uncertainty, very unattractive according to hypothesis 1. A second problem is that the observed rate of backformation to /k/ was much lower.

Some light may be shed on this low rate by considering more carefully the competing outcome, in which the consonant in the backformation is identical to the corresponding consonant in the suffixed from. The universe just discussed did not include pairs such as *liquid*, *liquidity* or *concave*, *concavity*, because these do not end in /sɪti/. At a superficial level of description, introducing such pairs into the universe would lead to extremely heterogenous outcomes, since all the various consonants except /s/ that appear in *-ity* words are faithfully copied from the stem. However, such pairs may be characterized in a homogeneous way by introducing variables. All examples except for the /k/-/s/ pairs involve identity between the stem-final consonant and the consonant appearing before *-ity*. Just as the alpha notation of Chomsky and Halle (1968) can enforce matches between structural description and output, indexing can also enforce points of identity in the reversed direction. The pairs *liquidity*, *liquid* and *concavity*, *concave* are then both examples of the abstract pairing (C_iɪti, C_i). The need for such variables is extensively discussed in Marcus (2001) on the basis of other results on phonological learning and productivity.

Figure 2 shows the observed probability of maintaining the same C in the stem, for complex words ending in *-ity*, as a function of the universe

of lexical items used to make the estimate. The leftmost point on the graph corresponds to the universe just discussed, in which the probability of maintaining the same C is 0.58 (that is, 1.0 minus the 0.42 probability of /k/). Successively bigger universes are various supersets of this universe. As shown, the probability increases monotonically as the description of the universe is expanded to include more and more cases. Given that the probability of maintaining the same consonant is already above 0.5 for the smallest set, enlarging the set steadily increases certainty about the outcome.

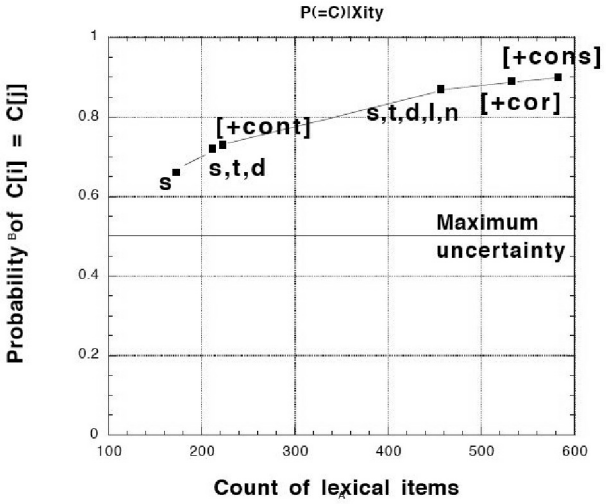

Figure 2. Rate of faithful pairings between the final phoneme preceding the suffix *-ity* and the final phoneme of the stem when this stem occurs as a independent word. The universe over which the rate is computed varies along the x-axis. 1) Complex words ending in /sɪti/, however spelled. 2) Complex words ending in voiceless coronal obstruents (/s/, /t/,/ʃ/,/θ/) plus /ɪti/. Points 1) and 2) appear superimposed because the sets differ by only one word pair. 3) Complex words ending in any coronal obstruent plus /ɪti/. 4) Complex words ending in any fricative (whether coronal or not) plus /ɪti/. 5) Complex words ending in any coronal consonant (including stops, fricatives, and sonorants) plus /ɪti/. 6) Complex words ending in any true consonant (excluding glides) plus /ɪti/. 7) All complex words in which stripping /ɪti/ yields an identifiable stem.

Insofar as enlarging the universe of comparison includes more and more examples of nonalternating (faithful) consonant pairings, this line of reason-

ing about backformation yields a bias towards a representation of the base form that is faithfully reflected in the complex form. We may compare it to the OT principle of lexicon optimization proposed in Prince and Smolensky (1993), according to which languages select underlying representations which are maximally harmonic with the surface representations. In the present case, the /s/ candidate is faithful, whereas the /k/ candidate is not. Given that the /s/ candidate is otherwise unproblematic, it would always win over the /k/ candidate. The /k/ candidate could only win if the surface representation of a base form with /k/ were also available to the learner, which was not the situation in the backformation experiment.

Unamended, OT lexical optimization predicts a single outcome, but actually the outcomes are variable. For five individuals, the results are consistent with the suggestion that the maximally inclusive universe of Figure 2 is relevant. This universe predicts 11% /k/s per individual, or 1.7 examples. This is surely within the statistical error of zero, especially if one allows for a low vocabulary level or for the confidence interval calculation proposed by Mikheev (1997). The two individuals who did backform to /k/ produced five examples (out of 15 hits) and six examples (out of eight hits), respectively. The two subjects differed in that one removed only the *-ity*, whereas the second often removed additional material. For *nodacity*, she responded *nodal* and for *runomicity*, *runate*, leaving only eight responses with just the *-ity* removed. Thus, one way of understanding these two subjects is that both worked with a comparatively narrow universe, but one had more active competition from other morphological interpretations of the input forms. Although this explanation is not complete, the statistical inference model at least provides ways by which vocabulary level, morphological awareness, and individual decision-making traits can manifest themselves in variable outcomes.

A second difference between the OT accounts and the present one is that the pressure towards faithfulness which can be read into Figure 2 depends critically on monotonicity of the graph and on the statistics of the most narrowly described universe. Faithfulness is not relevant in all situations, but only in those in which it resolves uncertainty. The outcomes for the noun formation task provide a case in point. Although this task was an exact counterpart to the backformation task, the productivity of the alternation proved to be entirely different; there was at best a sporadic penchant for a faithful outcome. This difference can be explained by noting that in the forward direction, the /k/-/s/ alternation is statistically reliable; expanding the universe of description would only weaken a certain inference. For the backformation task, in contrast, expanding the universe increases certainty.

Given this line of argument, it is important to determine why the analogy set based on words in *-icity* /ˈɪsɪti/ did not appear to be active or relevant in the backformation experiment. This descriptor is more specific than those used to calculate the figure, and the set it describes has a baseform ending in /k/ with P = 1.0. This generalization reliably covers both words ending in the suffix *-ic* and words ending in /ɪk/ which are not synchronically decomposable, as in (11).

(11) eccentric/eccentricity (*eccenter)
plastic/plasticity (*plast)
public/publicity (*puble)
rustic/rusticity (*rust)
toxic/toxicity (*tox)

If this set were operational, the experimental results should have displayed a strong distinction between the Latinate targets (which would backform to /k/) and the others (for which /s/ would be the more probable outcome).

The failure of words in *-icity* to support a reliable pattern of backformation casts doubt on the form of hypothesis 3. Even if moderately long descriptions are preferable to short ones, very long descriptions might not be preferable to long ones. Consider the joint effect of hypotheses 2 and 3. Extremely detailed descriptions tend to be statistically unstable, since they pertain to so few cases that their statistics are not robust across individual differences in vocabulary; this point is developed in more depth in Pierrehumbert (2001). The joint pressure towards detailed descriptions and large sample sizes could mean that the best entry level for generalizations over phonological patterns is moderately detailed – more detailed than the simplest descriptions that Mikheev dismisses, but still more broadly applicable than the worst cases considered by Pierrehumbert. These entry-level descriptions are then further generalized when the generalization increases certainty about the outcome, as discussed above.

5. Conclusion

In conclusion, the /k/-/s/ alternation was found to be highly productive in noun formation, and some evidence of its psychological reality is also found in backformation. An approach based on statistical inference over word pairs enjoys considerable success in explaining the outcomes. A key assumption

is that the universe of comparison grows more general to provide a critical mass of examples and to reduce uncertainty in the predicted outcome. A second key to success is the idea that word pairings can involve variables, so that pairings in which the same consonant is preserved no matter what its character can act together in influencing the outcome.

Explaining why /k/ is preserved in forms such as /blɛk/ – /blɛkɪti/ requires the assumption that in abstracting over a universe of examples, the cognitive system prefers to maintain somewhat rich and redundant descriptions. Abstractions are not simplified beyond what is required by differences amongst the examples. At the same time, the preservation of /s/ in some backformations such as /mtɹpanɪsɪti/ – /mtɹpanəs/ indicates some limits on how fine-grained and detailed abstractions can be. The behavior of these cases suggests that arbitrary phonological descriptors are not actually relevant to forming the universe for the morphophonological inference. Instead, there may be a privileged degree of granularity in analysis, reminiscent of the basic level of categorization which privileges the concept DOG over DALMATION or ANIMAL. Initial generalizations made at this level can then be adjusted upward or downward to achieve more certainty about the outcomes.

Acknowledgments

I'm very grateful to the Spring 2001 students in Laboratory Phonology II who participated in the initial development of the stimuli for this experiment as a laboratory exercise: Martin Armstrong, Jungwon Choy, Erin McMahon, Jessica Peterson, Rebecca Schwarzlose, and Xanthe Wigfall.

Appendix

(1) Baseline Items

NESS-WORDS: calm, calmness; correct, correctness; happy, happiness; icy, iciness; kind, kindness; light, lightness; lonely, loneliness; sweet, sweetness.

ITY-WORDS: abnormal, abnormality; captive, captivity; inane, inanity; liquid, liquidity; profane, profanity; valid, validity; virile, virility.

(2A) Word Fillers

USED WITH -NESS IN BACKFORMATION TASK: arid, frugal, human, odd.

USED WITH -ITY IN BACKFORMATION TASK: arboreal, binomial, liberal, real.

[2B] Nonword fillers

USED WITH -NESS IN BACKFORMATION TASK

Spelling	*IPA*	*Spelling*	*IPA*
clipid	klˈɪpɪd	clipidness	klˈɪpɪdnəs
demarte	dəmˈaɹt	demarteness	dəmˈaɹtnəs
flader	flˈedɹ	fladerness	flˈedɹnəs
mastive	mˈæstɪv	mastiveness	mˈæstɪvnəs

USED WITH -ITY IN BACKFORMATION TASK

Spelling	*IPA*	*Spelling*	*IPA*
bordal	bˈɔɹdl	bordality	bˌɔɹdˈælɪti
caustive	kˈɔstɪv	caustivity	kˌɔstˈɪvɪti
justical	dʒˈʌstɪkl	justicality	dʒˌʌstɪkˈælɪti
tromucal	trˈamjukl	tromucality	trˌamjukˈælɪti

[3A] “Latinate” targets ending in /ɪk/

Spelling	*IPA*	*Spelling*	*IPA*
clemic	klˈɛmɪk	clemicity	klˌɛmˈɪsɪti
criotic	kɹˌajˈatɪk	crioticity	kɹˌajətˈɪsɪti
extric	ˈɛkstɹɪk	extricity	ˌɛkstɹˈɪsɪti
hytronic	hˌajtɹˈanɪk	hytronicity	hˌajtɹənˈɪsɪti
interponic	ˌɪntɹpˈanɪk	interponicity	ˌɪntɹpanˈɪsɪti
malatonic	mˌælətˈanɪk	malatonicity	mˌælətənˈɪsɪti
phynomic	fˌajnˈomɪk	phynomicity	fˌajnəmˈɪsɪti
runomic	ɹˌunˈamɪk	runomicity	ɹˌunəmˈɪsɪti

[3B] “Semi-Latinate” targets ending in a secondary stressed syllable not construable as *-ic*.

Spelling	*IPA*	*Spelling*	*IPA*
bowdec	bˈodˌɛk	bowdecity	bˌodˈɛsɪti
hovac	hˈovˌæk	hovacity	hˌovˈæsɪti
nodac	nˈodˌæk	nodacity	nˌodˈæsɪti
pavoc	pˈævˌak	pavocity	pˌævˈasɪti
solvoc	sˈalvˌak	solvocity	sˌolvˈasɪti
stanorac	stˈænəɹˌæk	stanoracity	stˌænəɹˈæsɪti
strenoc	stɹˈɛnˌak	strenocity	stɹˌɛnˈasɪti
trylec	tɹˈajlˌɛk	trylecity	tɹˌajlˈɛsɪti

[3C] “Non-Latinate” targets ending in /k/

Spelling	*IPA*	*Spelling*	*IPA*
bleck	blˈɛk	bleckness	blˈɛknəs
mork	mˈɔɹk	morkness	mˈɔɹknəs
over-glique	ˌovɹglˈik	over-gliqueness	ˌovɹglˈiknəs

shruk	ʃɹˈʌk	shrukness	ʃɹˈʌknəs
snilk	snˈɪlk	snilkness	snˈɪlknəs
toque	tˈuk	toqueness	tˈuknəs
twake	twˈek	twakeness	twˈeknəs
under-grack	ˌʌndɹgɹˈæk	under-grackness	ˌʌndɹgɹˈæknəs

[3D] "Non-Latinate" targets ending in /s/

Spelling	*IPA*	*Spelling*	*IPA*
blarse	blˈaɹs	blarseness	blˈaɹsnəs
deploose	dəplˈus	deplooseness	dəplˈusnəs
dwess	dwˈɛs	dwessness	dwˈɛsnəs
druss	dɹˈʌs	drussness	dɹˈʌsnəs
jace	dʒˈes	jaceness	dʒˈesnəs
melse	mˈɛls	melseness	mˈɛlsnəs
queece	kwˈis	queeceness	kwˈisnəs
under-dass	ˌʌndɹdˈæs	under-dassness	ˌʌndɹdˈæsnəs

References

Albright, Adam and Bruce P. Hayes
2003 Rules vs. analogy in English past tenses: A computational/experimental study. *Cognition* 90: 119–161.

Anderson, Stephen R.
1981 Why phonology isn't natural. *Linguistic Inquiry* 12: 493–539.

Baayen, R. Harald
2003 Probabilistic approaches to morphology. In: Rens Bod, Jennifer Hay and Stefanie Jannedy (eds.), *Probabilistic Linguistics*, 229–288. Cambridge, MA: MIT Press.

Baayen, R. Harald, R. Piepenbrock and L. Gulikens
1995 *The CELEX lexical database (release 2) CD-ROM*. Philadelphia, PA: Linguistic Data Consortium.

Berko, Jean
1958 The child's learning of English morphology. *Word* 14: 150–177.

Broe, Michael and Janet Pierrehumbert
2000 Introduction. In: Michael Broe and Janet Pierrehumbert (eds.), *Papers in Laboratory Phonology V: Acquisition and the Lexicon*, 1–8. Cambridge: Cambridge University Press.

Bybee, Joan
2001 *Usage-Based Phonology*. Cambridge: Cambridge University Press.

Bybee, Joan and Elly Pardo
1981 Morphological and lexical conditioning of rules: Experimental evidence from Spanish. *Linguistics* 19: 937–968.
Chomsky, Noam and Morris Halle
1968 *The Sound Pattern of English*. New York: Harper and Row.
Cutler, Anne and T. Otake
1998 Assimilation of place in Japanese and Dutch. In: *Proceedings of the Fifth International Conference on Spoken Language Processing*, 1751–1754. Sydney.
Daelemans, Walter, Jakub Zavel, Ko van der Sloot and Antal van den Bosch
1999 TiMBL: Tilburg memory-based learner version 2.0 reference guide. http://citeseer.ist.psu.edu/article/daelemans99timbl.html.
Dell, Gary S.
2000 Commentary: Counting, connectionism, and lexical representation. In: Michael Broe and Janet Pierrehumbert (eds.), *Papers in Laboratory Phonology V: Acquisition and the Lexicon*. Cambridge: Cambridge University Press.
Derwing, Bruce and Royal Skousen
1994 Productivity and the English past tense: Testing Skousen's analogy model. In: S. D. Lima, R. L. Corrigan and G. K. Iverson (eds.), *The Reality of Linguistic Rules*, 193–218. Amsterdam: John Benjamins.
Ernestus, M. and R. Harald Baayen
2002 The functionality of incomplete neutralization in Dutch: The case of past-tense formation. Paper presented at: *Eight Conference on Laboratory Phonology*. New Haven, CT.
Frisch, S., Janet Pierrehumbert and Michael Broe
2004 Similarity avoidance and the OCP. *Natural Language and Linguistic Theory* 22: 179–228.
Goldrick, M.
2002 Probabilistic phonotactic constraints and sub-segmental representations: Evidence from speech errors. Paper presented at: *Eight Conference on Laboratory Phonology*. New Haven, CT.
Hay, Jennifer, Janet Pierrehumbert and Mary E. Beckman
2004 Speech perception, well-formedness, and the statistics of the lexicon. In: John Local, Richard Ogden and Rosalind Temple (eds.), *Papers in Laboratory Phonology VI*, 58–74. Cambridge: Cambridge University Press.
Jaeger, Jeri J.
1984 Assessing the psychological reality of the Vowel Shift Rule. *Journal of Psycholinguistic Research* 13: 13–36.

Kiparsky, Paul
1985 Some consequences of lexical phonology. *Phonology Yearbook* 2: 85–138.
Lavoie, Lisa M.
2001 *Consonant Strength: Phonological Patterns and Phonetic Manifestations.* New York: Routledge.
Marcus, Gary
2001 *The Algebraic Mind: Integrating Connectionism and Cognitive Science (Learning, Development and Conceptual Change)*. Cambridge, MA: MIT Press.
McCarthy, John and Alan Prince
1995 Faithfulness and reduplicative identity. ROA-60, Rutgers Optimality Archive, http://roa.rutgers.edu.
McCawley, J. D.
1986 Today the world, tomorrow phonology. *Phonology Yearbook* 3: 27–45.
Mikheev, A.
1997 Part-of-speech guessing rules: Learning and evaluation. *Computational Linguistics* 23: 405–423.
Myers, James
1999 Lexical phonology and the lexicon. ROA-330-0699, Rutgers Optimality Archive, http://roa.rutgers.edu.
Ohala, John J.
1974 Experimental historical phonology. In: J. M. Anderson and C. Jones (eds.), *Historical Linguistics II: Theory and Description in Phonology*, 353–389. Amsterdam: North Holland.
Pierrehumbert, Janet
2001 Why phonological constraints are so coarse-grained. In: James M. McQueen and Anne Cutler (eds.), *Spoken Word Access Processes (SWAP)*. Psychology Press.
2003 Probabilistic Phonology: Discrimination and Robustness. In: Rens Bod, Jennifer Hay and Stefanie Jannedy (eds.), *Probabilistic Linguistics*. Cambridge, MA: MIT Press.
Prince, Alan and Paul Smolensky
1993 Optimality theory: Constraint interaction in generative grammar. *Technical Report TR-2, Rutgers University Cognitive Science Center*.
Skousen, Royal
1989 *Analogical Modelling of Language*. Dordrecht: Kluwer Academic Publishers.
Steriade, Donca
2002 Commentary. Paper presented at: *Eight Conference on Laboratory Phonology*. New Haven, CT.

Whalen, Douglas H.
in press How the study of endangered languages will revolutionize linguistics. In: Piet van Sterkenburg (ed.), *Linguistics Today*. Amsterdam: John Benjamins.
Zuraw, Kie
2000 Patterned exceptions in phonology. Ph.D. dissertation, UCLA.

/k/-/s/ alternation as presently found in English is not natural (in the sense of Anderson 1981).

Velar Softening is not phonetically natural because the evident phonetic pressures on a /k/ in the target position would not produce /s/. If the suffix vowel is /t/ (as transcribed in the dictionary) then coarticulation and lenition would yield an aspirated palatal approximant rather than the alveolar fricative /s/ (see Lavoie 2001). The alternative possibility for the vowel, /ə/, provides still weaker phonetic motivation for /s/. /s/ differs from the phonetically expected outcome by its maximal vocal fold abduction and its precise tongue shaping, which directs a jet of air against the teeth. These are active adjustments which cannot be characterized as accommodation to a following vowel. Thus, understanding the alternation of /k/ with /s/ requires recourse to some version of the concept of Structure Preservation in phonology (see Kiparsky 1985) which states that lexical alternations stay within a language's system of phonological categories. Since an aspirated palatal approximant is not a contrastive category in the English lexicon, it cannot be the outcome of a morphophonological rule, either. The reanalysis involved in lexicalizing the phonetically expected approximant as the lexically contrastive segment /s/ reveals the role of abstract cognitive factors, over and beyond phonetic ones.

Velar Softening is also unnatural because it is phonologically opaque. Though it originates historically in fronting and spirantization of the velar stop before a non-low front vowel, suffixes with such vowels on the surface do not in general trigger the softening of /k/ to /s/ in the synchronic phonology. /k/ never softens to /s/ before *-y*, as *smoke, smoky*. On the other hand, *-ize*, beginning with a low vowel, does trigger softening, because *-ize* formerly had a nonlow front vowel. In Chomsky and Halle (1968), this historical ordering is recapitulated in the extrinsic ordering of rules in the synchronic phonology. The phonological opacity created by such orderings is precisely one reason that the psychological validity of the Chomsky-Halle model became a matter of widespread dispute. The finding that the vowel shift is only partially productive (c.f. Jaeger 1984; McCawley 1986) also calls into question the productivity of the rule of Velar Softening, which is ordered before it.

Understanding productivity is important because it provides a crucial line of evidence about cognitive abstractions. The failure of an alternation to generalize suggests that no abstract generalization over the forms exhibiting the alternation has been formed. If the alternation is aggressively and reliably extended, even to forms which differ substantially from attested

(Guion (1998) also discusses the role of perceptual structuring for the typologically related change /k/ → /tʃ/.

ones, it follows that a very broad abstraction has been formed. For example, the reliable and aggressive extension of the regular English plural pattern indicates that it abstracts away from many properties of the word. If the situation lies somewhere in the middle, then the exact pattern of productivity can yield insights about the exact character of the abstraction that is formed.

Phonotactics is the area in which most research has been done on the availability of lexical patterns for use in novel forms. Numerous studies, reviewed in Pierrehumbert (2003), indicate that the type frequency (frequency in the lexicon) of a phonological pattern affects the likelihood and perceived well-formedness of novel words containing that pattern. This dependence is gradient; frequent sequences readily extended to new words, rare sequences are avoided, and moderately frequent sequences fall in between. For example, Hay, Pierrehumbert, and Beckman (2004) found that the perceived well-formedness of novel words containing nasal-obstruent clusters (such as /strɪnfi/ and /zæmpɪ/) was a gradient function of the frequency of the cluster. The frequency for a tautomorphemic cluster was estimated as its frequency in trochaic monomorphemic words with a lax front vowel in the CELEX monomorphemes. (see Baayen, Piepenbrock and Gulikens 1995, regarding CELEX; Hay, Pierrehumbert and Beckman 2004, regarding monomorphemes).

This choice of universe for estimating frequencies was opportunistic, and obscures a central issue in understanding the relation of lexical frequencies to pattern productivity. This issue is taken up with Figure 1. Figure 1 shows a partial lattice of heterosyllabic N.O clusters. The atoms on the bottom are individual heterosyllabic phoneme clusters. The nodes above the atoms are some of the various available natural classes of such clusters. As is well-known, natural classes can be formed using partial descriptions of phonological patterns. For example, the sequence /nt/ is an element of the set of clusters of /n/ followed by any stop; it is also an element of the set of clusters containing a homorganic nasal and stop. The cluster /np/ belongs to the former set but not the latter; the cluster /mp/ belongs to the latter set but not the former. The lattice is organized from specific (on the bottom) to general (at the top). Each node is labeled with the probability of the indicated descriptor with respect to the universe of N.O clusters, as estimated from counts in the CELEX monomorphemes. Clearly, the less specific the description, the more cases it encompasses and the larger the natural class it describes. Thus, the probabilities go up as we follow the lines up the lattice, but the exact way they go up depends on exactly what is lumped together in each class.

result of constraints that are generally true of the language. However, it is not the case that the consonant preceding word-final /rti/ (or the morpheme *-ity*), is usually /s/. In the Collins on-line dictionary, only 25% of words ending in /rti/ end in /srti/. /l/ is more common than /s/, although even /l/ does not achieve a majority of the forms. Simple frequency matching on the surface forms would predict that subjects would tend to substitute /l/ for /k/ (or for any other consonant!) but only at a rate of about 25%. Surface statistics do not explain the extremely high rate of substitution observed for the Latinate and semi-Latinate stems ending in /k/, or the failure of other consonants to be affected in the same way. Similarly, the responses by two subjects in the back-formation task also appear to reflect implicit knowledge of specific morphological relationships. Without knowledge of such relationships, there would be no reason to backform /s/ to /k/, since /s/ is both more faithful to the stimulus and more common in word-final position.

As a result, knowledge of the alternation must be a generalization over morphologically related word pairs. This conclusion echoes the treatment of velar softening as a derivational rule in Chomsky and Halle (1968). In Optimality Theory, constraints generalizing over word pairs have been used since McCarthy and Prince (1995) proposed them in order to overcome the limitations of constraints over word forms in explaining the behavior of reduplication. Generalizations over word pairs also figure in the non-OT literature on computational morphology, notably Skousen (1989), Daelemans et al. (1999), Ernestus and Baayen (2002), Baayen (2003), and Albright and Hayes (2003). These works all share the assumption that variable outcomes in morphophonology are related to conditional probabilities defined on word pairs.

The acknowledgment that the alternation is learned as a generalization over word pairs goes far towards explaining the amount of variation observed across individuals. Individuals differ both in the size and the contents of their vocabularies. To know a relevant example of an alternation, they must know both words in the pair. Furthermore, they must view them as related to each other. Not everyone infers a decomposition of *Mediterranean* on the basis of words such as *medium* and *terrestrial*. The difficulties of assessing such implicit semantic relationships mean that most computational studies, including the present one, rely on phonological matching in large dictionaries and on morphological analyses by linguists. They probably overestimate the pool of relevant word pairs known to the subject pool.

In the following discussion, I will be particularly concerned with the claim, advanced in the analogical models of Skousen (1989) and Baayen (2003),

Bybee, Joan and Elly Pardo
1981 Morphological and lexical conditioning of rules: Experimental evidence from Spanish. *Linguistics* 19: 937–968.

Chomsky, Noam and Morris Halle
1968 *The Sound Pattern of English*. New York: Harper and Row.

Cutler, Anne and T. Otake
1998 Assimilation of place in Japanese and Dutch. In: *Proceedings of the Fifth International Conference on Spoken Language Processing*, 1751–1754. Sydney.

Daelemans, Walter, Jakub Zavel, Ko van der Sloot and Antal van den Bosch
1999 TiMBL: Tilburg memory-based learner version 2.0 reference guide. http://citeseer.ist.psu.edu/article/daelemans99timbl.html.

Dell, Gary S.
2000 Commentary: Counting, connectionism, and lexical representation. In: Michael Broe and Janet Pierrehumbert (eds.), *Papers in Laboratory Phonology V: Acquisition and the Lexicon*. Cambridge: Cambridge University Press.

Derwing, Bruce and Royal Skousen
1994 Productivity and the English past tense: Testing Skousen's analogy model. In: S. D. Lima, R. L. Corrigan and G. K. Iverson (eds.), *The Reality of Linguistic Rules*, 193–218. Amsterdam: John Benjamins.

Ernestus, M. and R. Harald Baayen
2002 The functionality of incomplete neutralization in Dutch: The case of past-tense formation. Paper presented at: *Eight Conference on Laboratory Phonology*. New Haven, CT.

Frisch, S., Janet Pierrehumbert and Michael Broe
2004 Similarity avoidance and the OCP. *Natural Language and Linguistic Theory* 22: 179–228.

Goldrick, M.
2002 Probabilistic phonotactic constraints and sub-segmental representations: Evidence from speech errors. Paper presented at: *Eight Conference on Laboratory Phonology*. New Haven, CT.

Hay, Jennifer, Janet Pierrehumbert and Mary E. Beckman
2004 Speech perception, well-formedness, and the statistics of the lexicon. In: John Local, Richard Ogden and Rosalind Temple (eds.), *Papers in Laboratory Phonology VI*, 58–74. Cambridge: Cambridge University Press.

Jaeger, Jeri J.
1984 Assessing the psychological reality of the Vowel Shift Rule. *Journal of Psycholinguistic Research* 13: 13–36.

Guion, S. (1998) The role of perception in the sound change of velar palatalization. Phonetica 55, 18–52.