2   Alveolars have been shown to behave uniquely in other contexts as well. In addition to the lexical effects described above, coronals (which include alveolars) are more frequently involved in exchange and substitution errors (Stemberger & Stoel-Gammon 1991), and are transparent to feature spreading in vowel harmony languages (Paradis & Prunet 1991).

3   These details of the recording methodology were constant for all experiments reported in this paper.

4   The vowel durations were identical in both series, even though one ended in a voiced consonant and the other in a voiceless consonant. Our primary concern was the perceived naturalness of the stimuli, and as they sounded natural to us and our subjects, we did not attempt to adjust vowel duration in any way. It is possible that this might have affected perception of the series in some manner. However, this is equally a potential factor in the present /ʃ/–/tʃ/ series (which showed a neighborhood effect), and in the /s/–/t/ series we present in Experiment 3, which did not.

5   Thus, the /d/ endpoint did have a complete /d/ burst, but the other members of the series did not. This is somewhat different than with the /d/–/t/ series, where all items had a burst. However, the items sounded natural, both to us and to our subjects. In addition, the second item in the series, which did not have a burst, was clearly unambiguous to the subjects (as seen in Figure 21.3).

6   Technically, TRACE has three levels: one for features, one for phonemes, and one for words. We are concerned here only with the latter two levels, however, and wish to make no claims regarding the existence of lower-level feature nodes.

22

Paradigm uniformity and the phonetics-phonology boundary

DONCA STERIADE

22.1 Paradigm uniformity

This is a study of paradigmatic relations and of their significance for the link between phonology and phonetics. A paradigm is a set of words sharing a morpheme, e.g. {*bomb, bomb-ing, bomb-ard, ...*}, or a set of phrases sharing a word, e.g. {*bomb, the bomb, ...*}. The main component of the analyses presented here is the preference for uniform paradigms, that is paradigms sharing contextually invariant morphemes. A Paradigm Uniformity (PU) condition is a statement of the type shown in (1), which promotes invariance of some sound property within a given paradigm:

(1)   All surface realizations of μ, where μ is the morpheme shared by the members of paradigm x, must have identical values for property P.

Examples of uniform and non-uniform paradigms appear in (2a) and (2b).

(2)   a.   {bɑm, bɑm-ɪŋ, bɑm-ɚ, bɑm-z}: uniform w.r.t. stem final C quality

b.   {bɑm, bɑmb-ɑrd, bɑmb-ɑrd-iɚ}: not uniform w.r.t. stem final C quality

Paradigm leveling represents the systematic generalization of one allomorph to positions where it is phonologically unjustified or unexpected, as a means of satisfying a PU condition. The degree of phonological invariance of shared morphemes depends on the paradigm's productivity and the transparency of the derivative's relation to its base (Bybee 1988): paradigms like (2a) are generated by productive word formation processes and involve less contextual variability than the unproductive paradigm (2b). We concentrate here on highly productive and compositional paradigms.

Paradigm leveling is a staple of the phonological literature. Different aspects of it have been studied under the names of *analogy* (Kuryłowicz 1949, Kiparsky 1978), *cyclic rules* (Chomsky & Halle 1968) *output-output correspondence* (Burzio 1994, Benua 1995, Flemming 1995, Kenstowicz 1996, Steriade 1995b, 1998). It is the Optimality Theoretic mode of analysis pursued in the last group of references that is assumed here. The term *paradigm uniformity* (along with *paradigm coherence* or *regularity*) was introduced into the generative tradition in Kiparsky's work.

22.1.2 *How much uniformity: phonetic vs. phonological?*

The critical question for this paper is just how much invariance counts as enough for the purpose of satisfying grammatical conditions like (1). The current assumption is that the measure of invariance is the *identity of phonological representations*. Thus we might assume that the stem final [m] of [bɑm] has the same phonological feature composition throughout the paradigm (2a): if so, the stem final consonant is invariant and the paradigm is uniform in its right edge. On this view, any differences in the phonetic implementation of this [m]—differences between a final-lengthened or voiceless [m] in [bɑm] and word-internal realizations of it—will not affect the uniformity of the paradigm, as they are not reflected in phonological representations. Therefore what can and cannot be part of a phonological feature matrix is highly relevant for our understanding of PU conditions. Conversely, the empirical study of PU conditions can shed light on where the boundary lies between phonetic and phonological features.

This paper aims to show that paradigmatic uniformity is enforced through conditions that govern both phonological features and properties presently classified as phonetic detail, such as non-contrastive degrees in the duration of consonantal constrictions, non-contrastive details in the implementation of the voicing contrast, and degrees of gestural overlap. The form of the argument is as follows: one must show that some category has an identifying property of a phonetic category and then show that this category is being generalized through the effect of PU conditions, exactly like a phonological category. The result anticipated is that 'phonological' and 'phonetic' features are not being treated differently when it comes to enforcing morpheme invariance. There is a larger agenda behind this argument: the distinction between phonetic and phonological features is not conducive to progress and cannot be coherently enforced. It is unproductive because in order to understand phonological patterns one must be able to refer to details of their physical implementation, in perception and production (cf. Ohala 1995 and references therein; Lindblom 1990b and references therein, Flemming 1995, Jun 1995, Kirchner 1996, Steriade 1995b): separating phonological features from their phonetic realization creates the illusion of a

well-defined domain of facts that can be described and explained using an impoverished vocabulary of contrastive properties. The distinction is also unenforceable, because most contrasts are simultaneously implemented on several physical dimensions, most of which never function as independently contrastive: it is arbitrary in such cases to identify some one property as the contrast's flag bearer, the actual phonological feature, with the others as its predictable side-effects (cf. Fant's, Repp's and Port's contributions to Perkell & Klatt 1986). Phonologists may recognize this in discussing underspecification (Broe 1993, Stanley 1967, Steriade 1995a) but continue to assume the distinction between phonological and phonetic features. This study examines the effects of paradigm uniformity on the realization of phonetic detail in order to highlight the fact that grammatical structures and their physical implementation cannot be separately studied.

22.2 Phonetic vs. phonological features: the contrastivity test

I outline next what represents, in the current practice of phonologists, the criterion separating phonetic from phonological categories, since it is this distinction that I argue against. A candidate phonetic feature is the stop burst, the brief period of noise following the release of a closure. No phonological feature set includes [±burst], even though this is a perceptually important property and, as argued elsewhere, the class of segments that *can* produce a burst have distinct phonological behaviors (Steriade 1993). What membership criteria exclude the burst from the set of phonological features? One answer—implicit in the work of Jakobson and Halle (1962)—is that a phonological feature is an articulatory or auditory property that provides the sole basis of lexical contrast in at least some language. On this view, features are a subset of the physical properties of sounds: those that can function as independently contrastive in some language. The stop burst is excluded as a phonological feature because no language has phonemically distinct released and unreleased stops in any context.

A different way of defining the phonological feature set is suggested by Keating 1984; Ladefoged & Lindau 1986. For these writers, a phonological feature is a cover term for a class of lexical contrasts with identical phonological behavior and similar phonetic implementation. It is the abstract property distinguishing the phonological representations of contrasting sounds, rather than any of the physical correlates. On this definition too, [±burst] is not a feature because it does not correspond to a contrast.

Both conceptions of the phonological feature set—and others not mentioned here—are motivated by the belief that, however this set is defined, it must be a very small set. This is directly said by Jakobson & Halle (1962:483) "The supposed multiplicity of features proves to be largely illusory"; and by Keating (1984:289) in a passage criticizing proposals by Halle & Stevens (1973) on the

grounds that their proposals generate too many features: "Halle & Stevens [...] don't simply have the wrong features in these instances; *they will always have too many features* because they want to describe exactly how individual sounds are articulated. While we want phonological features to have some phonetic basis, we also want to distinguish possible contrasts from possible differences." The implication here is that the set of phonological features must be small because the set of contrasts employed in any given language is small. A mere difference between sounds is not linguistically significant if it is never contrastive.

What matters here more than the form of argument is its universally accepted conclusion: the phonological feature set is small and therefore some phonetic properties are not in it. To tell which phonetic properties to exclude, a test of contrastivity is being implicitly appealed to. A look at the recent literature indicates that this test is invoked every time the grammatical status of a phonetic property is questioned: stricture degree of secondary articulations (Sagey 1986); voicing in sonorants (Mester & Ito 1989); precedence between the phases of an affricate (Lombardi 1990); VOT-based categories (Selkirk 1992); vertical larynx movements (Kenstowicz 1994:40); [±consonantal] (Hume & Odden 1994); release-related properties (Rice 1994); segment-internal precedence relations (Scobbie 1995); and many others. In what follows, I assume that if a distinction is made between phonological and phonetic features then universally non-contrastive properties belong in the second class. The paper attempts a reductio ad absurdum of this belief by demonstrating the phonological relevance of categories classified as non-phonological.

22.3 Paradigm uniformity constraints

The central claim of this study is that grammatical conditions of the form in (1) apply equally to three types of sound properties: (a) those that are contrastive in the language under analysis; (b) those are contrastive in some language, though not the language analyzed; and (c) those not known to be independently contrastive in any language. The argument for rejecting a principled distinction between phonetic and phonological categories rests on the existence of paradigmatic effects involving type (c) properties. For type (a) and (b) properties, the effects of paradigmatic leveling are standardly attributed to cyclic rule application. Our view of the formal encoding of paradigmatic effects is different: whether or not rule-based analyses are generally defensible, the cyclic application account is unsuccessful because it describes only a fragment of a coherent class of related phenomena and does this incompletely.[1] We justify a mode of analysis that has appropriate coverage. The grammatical conditions thus motivated will promote paradigm uniformity for contrastive, non-contrastive and never contrastive properties alike. We then consider instances of leveling for type (c).

The key observation here is that forms belonging to the same lexical paradigm—base words and their derivatives—display phonological similarities that go beyond what may be expected from the fact that they share an underlying form. Consider a verb like *discipline* and its *-able* form *disciplinable*. In this paradigm, the accentual resemblance between base and derivative is unexpected: English morphemes do not allow strings longer than two stressless syllables, yet *disciplinable* contains four, [sə.plɪ.nə.bl̩]. An entire class of English affixes—Level 2 or unrestrictedly productive suffixes—are similar to *-able*: they maintain intact the distribution of stressed syllables in the bases they attach to, even when this entails some measure of deviance relative to constraints reflected in the underived vocabulary. The connection between productivity and phonological invariance has a clear rationale: productive formations give rise mostly to nonce formations, whose interpretation must be computed on-line. When the derivative maintains intact the phonology of the base, listeners can easily access the lexical entry to interpret the nonce form. If the derivative's stem is modified, this makes it harder to access the lexical entry. In particular, modifying the stress of the stem will slow down or impede access (Cutler 1979). What about Level 1 derivatives? If we assume that results of unproductive affixation are lexically listed, then the meaning of such derivatives can be looked up: for this reason the outcome of unproductive affixation is not subject to the same requirement of phonological compositionality.

At the same time, speakers are aware of the constraints defining a phonological 'possible word' and understand that upholding base invariance can lead to anomalous forms. For instance speakers asked to generate new forms in *-able* on antepenult stressed verbs (e.g. *díscipline, bénefit, jéttison, párody*) comment that the results (*dísciplinable, bénefitable, jéttisonable, párodiable*) are 'awkward', 'a mouthful', 'too long'. Nonetheless, they recognize that there are no viable alternatives: for the vast majority, well-stressed forms like *jètisónable* are out of question. In other words, speakers understand three essential points: first, that there is a conflict between phonological well-formedness—reflected in the underived words—and the requirements of base invariance. Second, that base invariance may carry higher priority: the stress of *dísciplinable* is due to it. And, third, speakers understand that words generated under this conflict are metrically imperfect: *dísciplinable* is not a very good word, but it is the best the system can generate.[2] All accounts must do justice to these points.

Thus for a proper understanding of paradigmatic effects in phonology one must recognize the notion of constraint conflict, the central element of Optimality Theory (Prince & Smolensky 1993). There is conflict between the preference for stress invariance and the preference against long stressless strings: for certain forms, both preferences cannot be satisfied. One can formulate these preferences in a number of ways but the right picture emerges only if their formulation

reflects this conflict. We provide this below, employing the notion of correspondence between strings developed by McCarthy & Prince (1995).

Correspondence constraints evaluate the extent of similarity between two linguistic expressions. The typical purpose of this evaluation is to determine whether the two expressions resemble each other sufficiently to be considered realizations of the same linguistic sign. These constraints can be read as asking the question: Assuming that two surface strings are, at the lexical level, one and the same unit, are they identical with respect to some specified phonological property P? If the answer is yes, the constraint in question is said to be satisfied; otherwise, it is violated. The strings thus compared are said to be *correspondent* strings: this means that they are, by hypothesis, variant realizations of the same string in the mental lexicon. Whether this is the right hypothesis or not in any given case is determined by the overall constraint system, not by any single correspondence constraint, as we see below.

The property P named in specific correspondence constraints may be segmental identity, featural identity, or identity with respect to some aspect of prosodic structure. Some constraints determine whether every specified element (say every segment or syllable) in one string has a (not-necessarily identical) counterpart in the correspondent string. Other constraints determine whether these pairs of corresponding elements are identical in detail, i.e. with respect to specified features. One should emphasize that the system of correspondence constraints for any language must be set up so that it allows the recognition that two different strings may in fact stand in correspondence, i.e. be lexically related. Thus the English system must accept the hypothesis that the suffixal strings [d] (as in [læg-d]) and [t] (as in [læk-t]) are one and the same lexical element, despite the voicing difference, while rejecting the hypothesis that [mæd] and [mæt] are lexically related. This result is obtained by stating that the requirement of voicing agreement in obstruent clusters outranks the requirement of correspondence for voicing. Therefore, the hypothesis that two expressions are lexically related does not hinge on the answer to each and every correspondence constraint: two strings Σ1 and Σ2 may differ with respect to voicing and still emerge as correspondents, if this satisfies higher ranked constraints.

There are two variable elements in every correspondence constraint. One is the property P for which the constraint mandates identity between correspondent strings. The other is the lexical relation holding between the pairs of forms considered: these may involve an underlying string and its surface counterpart; or a pair of related surface forms. The cases of interest to us here are pairs consisting of the surface realization of a morpheme in isolation (e.g. *díscipline*) and its realization when affixed (e.g. *disciplinable*). Since this aspect of the correspondence constraint is constant throughout, we will state it in the first constraint ((3) below) and omit it subsequently.

(3) PARADIGM UNIFORMITY (STRESS), abbreviated PU (STRESS)
Let F be a form exhaustively analyzable into the constituents A(F), an affix, and S(F), a stem. If a realization of S(F) occurs as a free word W, then, for every syllable σ in S(F), if σ has a correspondent σ′ in W then σ has the same stress category (stressed or stressless) as σ′.

PU (STRESS) states that pairs of correspondent syllables are identical for stress: therefore PU (STRESS) is satisfied by the pair {*díscipline, disciplin(-able)*} and violated by pairs like {*díscipline, discíplin(-able)*}. The constraint that conflicts with PU (STRESS) in the forms discussed earlier is *LAPSE:

(4) *LAPSE
Strings of stressless syllables longer than two do not occur within one word.

The tableau in (5) records the conflict between PU (STRESS) and *LAPSE in the realization of *disciplinable*. We assume that the former outranks the latter: PU (STRESS) >> *LAPSE. The candidate ['dɪsəplɪnəbl̩] is marked below as violating *LAPSE twice because it contains two distinct sequences of three stressless syllables: [səplɪnə] and [plɪnəbl̩]. This method of evaluation allows us to recognize distinct degrees of violations of this constraint.

(5)   PU (STRESS) >> *LAPSE

W = *díscipline* ['dɪsəplɪn]	PU (STRESS)	*LAPSE
☞ i. *disciplinable* ['dɪsəplɪnəbl̩]		**
ii. *disciplínable* [ˌdɪsə'plɪnəbl̩]	* (plɪn ≠ 'plɪn)	

The conflict between PU (STRESS) and *LAPSE revealed in (5) is reflected in the fact that both *-able* derivatives of *díscipline* considered in (5) violate one of these two constraints.[3] The prevalent form in actual use, *disciplinable*, is better than the alternative, but it is imperfect nonetheless: it violates *LAPSE. Our approach differentiates three classes of *-able* forms: forms which violate *LAPSE twice (*disciplinable, jéttisonable, párodiable, bénefitable*), forms which violate it once (*éditable, estáblishable, devélopable*) and, finally, forms which do not violate *LAPSE at all (*represéntable, redeplóyable, disinféctable*). All three classes are in potential use, in contrast with stress-shifted *jettisónable, *develópable, etc. which satisfy *LAPSE but violate PU (STRESS). But some of the useable forms are better than others, in proportion to the extent they violate *LAPSE. We are in the process of confirming these claims through a study of speakers' well-formedness judgments for nonce *-able* words. Results so far, from 15 speakers, confirm that when one controls for length and familiarity subjects show a systematic preference for words which do not violate *LAPSE at all; in addition, there is a preference for words that minimize *LAPSE violations,

e.g. for *estáblishable* as against *dísciplinable*. See also Steriade 1998 for grammatical evidence on this.

This is, in outline, the analysis of stress uniformity effects based on the idea of conflict between PU (STRESS) and metrical well-formedness constraints like *LAPSE. There are no satisfactory alternatives to it. Suppose, for instance, that we describe the stress differences between non-derived and productively derived words of English by imposing domain limitations on some metrical constraints. One can restate *LAPSE so that its proper domain of application are roots and stems derived by unproductive, Level 1, affixes. We can state that sequences of stressless syllables longer than two are impermissible inside roots or in stems derived at Level 1. Under this analysis, no need arises to recognize a *LAPSE violation in *dísciplinable*, etc. since the reformulated *LAPSE constraint will not apply to any Verb-*able* combination. If there is no *LAPSE violation in *dísciplinable*, there is no need to find a higher ranked constraint like PU (STRESS) to compel *LAPSE violations. Thus the domain restriction is a descriptive substitute for the PU constraint. But it is not a good substitute, because it fails to reflect precisely the intuition that forms like *dísciplinable* are metrically imperfect and that they are being used only for lack of a better alternative.

An equally unsatisfactory analysis consists of restricting the stress rules so that they are inapplicable to words created through Level 2 affixation. This is the Lexical Phonology approach and it is open to the same objection as above. In any rule-based approach, well-formedness is a function of the correctness of rule application in the derivational history of the form evaluated. By this standard, *dísciplinable* is as well formed as, say, *redeplóyable*, as neither gives evidence of an illegitimate rule application. But these forms, as we have seen, are not equivalent in well-formedness. The right account of Level 2 phonology must distinguish better formed *redeplóyable* from awkward *dísciplinable*, while at the same time identifying the factor that induces accentual invariance in both. This factor is PU (STRESS).

We have seen so far that phonological constraints like *LAPSE, whose effects are categorical in the core vocabulary of underived forms, may fail to shift stress in productively affixed words. This is not because constraints like *LAPSE 'fail to apply' to affixed words: they do. Speakers' intuitions of relative well-formedness reflect this directly. Rather, stress invariance must be attributed to a constraint insuring the surface similarity between base forms and their counterparts in affixed words. The constraint is PU (STRESS) and its functional rationale is, we speculate, facilitation of lexical access.

A large class of the phenomena thought to motivate cyclic rule application fall into the category of unexpected accentual similarity between a base and its derivatives. We suggested that the better account will rely on constraints requiring similarity between paradigmatically related surface forms. An equally large class of processes described earlier through cyclic rules involves segmental

similarity between a stressed syllable and its stressless correspondent in related forms. For instance, in Levantine Arabic (Kager 1998), a vowel in the derivative corresponding to a stressed vowel in the base cannot delete, even though other vowels, in comparable contexts, do. The pair *fíhim* 'he understood' and its inflected derivative *fihímna* 'he understood us' illustrates this. *Fihímna* is expected to surface as **fhímna*, since in general stressless high vowels delete in open syllables. But this particular *i*, at the initial of *fihímna*, is the correspondent of the stressed *i* of its base, *fíhim*: if it had deleted, the base stressed vowel would have no counterpart in the derivative at all. A similar situation is observed in Catalan (Mascaró 1976), where the base–derivative relation inhibits a process of glide formation. In general high vowels become glides after vowels, hence /franku-italjá/ 'Franco-Italian' realized as [frankujta'lja]. But the correspondents of base stressed vowels do not undergo glide formation: the derivative of [ru'inǝ] 'ruin' is [ru.i.'nos], not [ruj.'nos]. Here too, correspondent vowels differ in stress, but an element of the base stressed vowel is maintained in the derivative and signals in this way its lexical connection to that stressed syllable.

With such cases in mind, we suggest that PU (STRESS) should characterize not only stress identity between syllables but also the use of individual stress correlates (such as duration, pitch accents, vowel quality) to flag the stress profile of the lexical item to be accessed. To implement this, we modify PU (STRESS) so that it promotes identity between a stressed syllable in one form and corresponding strings in paradigmatically related forms, whether or not such strings constitute a syllable. The new class of constraints evaluates the similarity between corresponding strings with respect to specific stress correlates such as duration. An instance of this class appears below.

(6)   PU (STRESS: DURATION) : If two strings, Σ and Σ', stand in correspondence and if Σ is a stressed syllable, then Σ and Σ' are durationally equivalent.

'Durationally equivalent' means that corresponding strings have the same range of durational values, with the same distribution. The requirement of durational equivalence admits multiple degrees of satisfaction/violation: thus in the two pairs of Levantine Arabic forms {*fíhim*–*fihímna*} vs. {*fíhim*–**fhímna*} the correspondent strings are closer to being durationally equivalent in the first pair {*fí*-, *fi*-} than in the second {*fí*-, *f*-}. Similarly for the Catalan pairs {*ru.í.na*–*ru.i.nós*} vs. {*ru.í.na*–*ruj.nós*}: neither the pair {-*í*-, -*i*-} nor the pair {-*í*-, -*j*-} may count as fully equivalent durationally, but the first pair is closer to equivalence than the second and thus better satisfies PU (STRESS: DURATION). Further evidence supporting PU (STRESS: DURATION) is provided by Kenstowicz & Abdul-Karim (1980): the correspondent of a base-stressed vowel is exempt from a general vowel shortening process. The same notion of durational equivalence will play a further role in what follows.

To summarize, the approach suggested here consists of decomposing the notion of accentual correspondence into multiple components, each of which represents the equivalence between two strings with respect to an individual stress correlate. (6) provides just one of these constraints, the one for which some empirical evidence is being presented. Full accentual correspondence between corresponding strings is evaluated through the entire set of constraints of the form PU (STRESS: <STRESS CORRELATE>).

22.4 Tapping and paradigm leveling: the Withgott effect

We now turn to the type of paradigmatic leveling that generalizes phonetic detail, a phenomenon I refer to as *phonetic analogy*. The case discussed first is that of American English tapping. The difference between [ɾ] and [t]/[d] is a function of closure duration (Zue & Laferriere 1979, Banner-Inouye 1995): the tap's mean duration is 26 ms, as against 75 ms for *d* and 129 ms for *t*. The extra-short duration of [ɾ] is a candidate for a never-contrastive property, one that cannot correspond to a phonological feature in the contrast-based Jakobsonian feature theory outlined earlier.

We reach this conclusion in two ways. First, the tap-related contrasts surveyed by Banner-Inouye (1995) involve either alveolar rhotics—distinguishable by [+sonorant, +continuant] from the corresponding stops—or voiced alveolar stops—distinguished by [+voice] from *t* and by [–son] from *r*. Banner-Inouye's instances of tap/trill contrasts involve, in a restricted feature set, differences of either syllable weight or point of articulation. The tap [ɾ] is never in clear contrast, in the same system, with a homorganic voiced stop *and* a homorganic rhotic of identical moraic count. This is one reason, for a feature theory that adheres to the contrastivity test, to reject any expansion of the feature set meant to accommodate [ɾ]. The second reason is that any closure-duration feature that distinguishes [ɾ] from [d] will be restricted to coronal obstruents: no comparable contrasts exist between short and extra-short labials or short and extra-short velar stops. The conclusion then is that a feature such as [extra-short closure] will not pass the contrastivity test.

This is reflected in the fact that most phonological discussions of English flapping avoid mentioning by name the feature that distinguishes [ɾ] from [t]/[d]. It is not satisfactory to use [+sonorant] for this purpose. The context that induces tapping is one where *all* oral constrictions are shortened (Browman & Goldstein 1992a) to a greater or lesser degree: but the difference between sonorants and obstruents is not a function of gesture duration. Therefore identifying taps on the basis of their durational category—the feature [extra-short closure]—allows the tapping effect to emerge from a general statement of closure shortening, whereas labeling them as [+sonorant] results in a description that ignores any connection between the tap and the shortening context that produces

---

it. I assume then that at least one of the features of the tap is [extra-short closure] a reference to a duration category distinct from that of both voiced and voiceless stops (cf. also Williamson 1977 and Banner-Inouye 1995).

Despite being non-contrastive, the feature [extra-short closure] plays a role in English phonology: this is shown by the paradigmatic extension of the unflapped stop [t] in contexts where taps are normally expected. The extension is systematic, language specific and serves the purpose of generating uniform paradigms, in exactly the same way as the extension of the phonological properties discussed earlier. If the transfer of stress analyzed above belongs in the phonology, then so does the paradigmatic extension of the [extra-short closure] feature. We consider now the data leading to this conclusion.

Withgott (1983) notes the near-contrast between [t] and [ɾ] in the accentually parallel forms *militaristic* [ˌmɪlətəˈɹɪstɪk] and *capitalistic* [ˌkʰæpərəˈlɪstɪk]. In both words, -*ta*- is realized with a stressless [ə] nucleus, but the onset of [ə] is—or can be—unflapped in *militaristic*, whereas it must be flapped in *capitalistic*, [tə] vs. [ɾə]. The essential fact is that unflapped [t] in [ˌmɪlətəˈɹɪstɪk] corresponds to unflapped and stressed [tʰ] in the corresponding base *military* [ˈmɪləˌtʰæɹi], while the flap in *capitalistic* corresponds to the flap in *capital* [ˈkʰæpərəl]. The principle at work is Paradigm Uniformity: the paradigm of {*military, militar(-istic)*} becomes less variable phonologically if the stop *t* is generalized to the stressless syllable. Further, non-flapping maintains to a greater extent the durational equivalence between [tə] in [ˌmɪlətəˈɹɪstɪk] and [tʰæ] in [ˈmɪləˌtʰæɹi]. Withgott's observation thus fits into the large class of cases in which correspondence between stressed and stressless syllables is signaled by maintaining a partial durational equivalence.

To verify this, one must first establish that the Withgott effect is systematic. This has been done on the basis of the list in (7), checked with 12 speakers of American English.

(7)   a.   Bases: *positive, primitive, relative, negative voluntary*
           Derivatives: *positivistic, primitivistic, relativistic, negativistic, voluntaristic*

       b.   Bases: *rotary, fatal, fetish, totem, notary*
           Derivatives: *rotaristic, fatalistic, fetishistic, totemistic, notaristic*

We used -*istic* adjectives because this formation is productive and compositional, and thus most likely to display stem invariance effects. Most -*istic* words in (7) are nonce formations. The base forms were selected according to the following criterion: we anticipated that speakers will differ individually in stress in stressing -*tive* forms and the penult in *voluntary*. Under stress, the *t* in -*tive* would not flap; without stress, we expect [ɾɪv]. This expectation was borne out and subjects did differ on this point. Thus there appears to be no established

norm on whether to tap or not in the *-tive* forms of (7a). In the absence of a clear norm in the pronunciation of the bases, the prediction is that tapping in the derivative will occur subject to the effects of Paradigm Uniformity: if the base contains [ɾɨv], the derivative will too, whereas if the base contains [tɨv], its derivative will be more likely to contain a stop. When PU is moot, in monomorphemic -*V̀CVtV*- strings, stressless *t* is generally tapped: *mèri[ɾ]ocrátic, hèma[ɾə]génesis, pèri[ɾə]nítis, hèma[ɾə]cýstic.*[4] Therefore the interaction of phonological principles alone, without the contribution of the PU factor, favors a tap in this context: any stops observed in the *-tivistic* forms should therefore be counted as effects of paradigmatic leveling.

The 12 speakers were asked to read at a normal rate the randomized list of base forms in (7). After a pause, they were asked to read a randomized list of corresponding derivatives. All were naive as to the purpose of the experiment. Most of the forms presented are nonce formations (e.g. *voluntaristic, rotaristic, totemistic, primitivistic*) yet the speakers produced the forms without hesitation or false starts.

The list of forms in (7a) was interspersed with those in (7b) to detect the possibility that the pronunciation of one *-tive* word might influence that of the following items on the list. This did not happen: the (7b) bases contained almost invariably taps, regardless of what words preceded. Similarly, the *-tivistic* forms (7a) were interspersed with the *-istic* forms listed in (7b). An added reason to include these items on the list was to verify that speakers would not produce artificially untapped pronunciations: the *-istic* forms in (7b) should contain taps, both because of the segmental context where the *t* occurs (*V́_V*) and because their bases contain taps. Indeed, we observed no artificial stops: all base forms in (7b) where a tap was predicted did in fact contain a tap.

The results of the survey show tapping variation for most of the bases in which *t* is separated by one syllable from stress. The observed ratios of tappers to non-tappers were 1/5 for *positive*, 7/5 for *primitive* and *relative*, 2/1 for *negative*. One speaker failed to tap in *fetish* but, aside from this, all directly posttonic *t*'s were tapped. On the other hand, we observed virtually no variation on the issue of base–derivative correspondence: of the 12 speakers, 11 had [t] correspond to [t] and [ɾ] correspond to [ɾ] in every one of the relevant base–derivative pairs. There were 6 word pairs in which a stop was possible in the unaffixed base while a tap or a reduced stop would be expected in the derivative, absent the PU factor, hence 72 pairs of forms that could in principle have shown a disparity between the stop quality of the base and the tapped quality of *t* in the derivative. In fact, however, only one speaker produced a tap in *primitive* and then a stop in the corresponding form *primitivitistic*. With this exception, the forms produced showed complete correspondence between the tap or stop quality of the base and its counterpart in the derivative.

The tapping variability we encountered contrasts with the strict correspondence observed in the quality of base-derivative consonant pairs. This suggests that a productive correspondence principle—rather than rote learning of lexical properties—ensures the complete identity in tap/stop quality between the *-tive* and *-tivistic* forms. This point was confirmed by the observation that one speaker had an atypical stop pronunciation for the *t* in both *fetish* and the corresponding item *fetishistic*. These observations confirm that the medial stop in the *military/militaristic* pair must be attributed to a paradigm leveling effect. An abbreviated analysis of the phonological pattern is shown in (9). I focus first on the constraints that predict the occurrence of the tap in the canonical lenition context V_V.[5]

(8)   Constraints relevant to the realization of alveolar stops in medial syllables (American English)

    a.   FORTITION: Consonants are realized with increased closure duration at the onset of stressed syllables.

    b.   TAPPING: Alveolar stops are tapped in intervocalic contexts, where tap refers to: extra-short duration of closure, lack of a concomitant jaw raising gesture and lack of a glottal opening gesture.[6]

    c.   REDUCTION: A stressless vowel must be schwa.

Two of the constraints in (8) conflict: FORTITION and TAPPING cannot both be satisfied in the production of the same *VtV́* sequence. It is also clear that FORTITION outranks TAPPING since the *VtV́* sequences (including forms like *atómic, Satánic, Platónic*) are realized with unflapped, aspirated stops. Therefore FORTITION outranks TAPPING.

(9)   a.   FORTITION >> TAPPING

atómic	FORTITION	TAPPING
☞ i. [ə'tʰamɪk]		*
ii. [ə'ramɪk]	*!	

    b.   Effect of TAPPING when FORTITION is moot:

átom	TAPPING
☞ i. ['ærəm]	
ii. ['ætəm]	*!

Turning now to the Withgott effect, we note that the stop [t] in *militaristic* cannot be due to FORTITION since its syllable is stressless. Rather, it is due to

the constraint identified earlier, PU (STRESS: DURATION): the non-flapped [t] in *militaristic* is a partial preservation of the stress carried by its syllable in *military*. REDUCTION however (8c) conflicts with PU (STRESS: DURATION) and outranks it. In turn, PU (STRESS: DURATION) outranks TAPPING, as seen below:

(10) Constraint interaction producing stressless [tə] corresponding
to stressed [tæ] in base
Ranking: REDUCTION >> PU (STRESS: duration) >> TAPPING

	REDUCTION	PU (STRESS: DURATION)	TAPPING
☞ i. [ˌmɪlətəˈɹɪstɪk]		*	*
ii. [ˌmɪlərəˈɹɪstɪk]		**!	
iii. [ˌmɪlətæˈɹɪstɪk]	*!		*

The violation marks in the PU (STRESS: DURATION) column mark *degrees of durational equivalence* between the stressed syllables in the stem *militari-* of *militaristic*, as realized in individual candidates, and their correspondents in the isolation form *military*. We focus here on the syllable *-ta-*. The candidate that comes closest to achieving durational equivalence is (iii), the form in which the vowel is unreduced and the stop is untapped. This candidate, however, violates REDUCTION, the highest ranked constraint. The remaining two candidates are differentiated by tapping in the stressless *t[ə]*. The reduced but untapped [tə] is durationally closer to the stressed [tʰæ] of *military* than the [rə] of candidate (ii) and thus emerges as optimal. Although this analysis does not attempt to explain the relative ranking of TAPPING and REDUCTION, we note that the ranking may derive from conditions on the segmental properties signaling stress. REDUCTION is the correlate of the [±stress] distinction in English, therefore non-reduction (i.e. the ranking PU (STRESS) >> REDUCTION) will yield forms with the perceived stress contour *mìlitàrístic*.

We have seen that in deciding whether a PU constraint is satisfied by a given form, the grammar of English must enumerate the properties with respect to which the base stressed syllable differs from its correspondent in the derivative. At least one of these is universally non-contrastive: the durational difference between [t] and [r]. Judging from the only test of what can count as a phonological feature, this difference is not phonological. However, the presence of this property must be identified by a grammatical condition: and thus [t] and [r] cannot count as phonologically equivalent, despite the fact that the difference between them is classified as a matter of phonetic detail in a contrast-based feature theory.

22.5 French C(ə)C

The second case we discuss involves the realization of French consonants in the alternating contexts CəC/CC. Morpheme-final schwa can delete in French. Rialland (1986) notes that when schwa is lost, the consonants left of its original syllable maintain allophonic qualities that would only be appropriate if schwa was still present. Rialland came to this conclusion by comparing pairs like *bas r'trouvé* [baʁtʁuve] 'stocking found again' (a realization of *bas retrouvé* [baʁətʁuve]) with *bar trouvé* [baʁtʁuve] 'bar found'. The postvocalic [ʁ] in *bas r'trouvé* is similar to an onset allophone ("fort et vibré") and has no lengthening effect on the preceding vowel. The post-vocalic [ʁ] in *bar trouvé*, an underlying coda, is lenited and doubles the duration of the preceding vowel (O'Shaughnessy 1981). Rialland formulates the comparison in syllabic terms stating that [ʁ] in *bas r'trouvé* is not resyllabified as a coda after the loss of schwa. If it had been, it would have been realized identically to the [ʁ] of *bar*. We will see however that the syllabic organization of the string resulting from schwa loss does not contribute to an explanation of the data. Rialland's observations are summarized in Table 22.1:

Table 22.1 *Rialland's observations:* bas retrouvé, bar trouvé, bas r'trouvé

	Syllable initial r	Coda r	r next to lost schwa
EXAMPLES	*bas retrouvé* [ba ʁətʁuve]	*bar trouvé* [baʁ tʁuve]	*bas r'trouvé* [ba ʁtʁuve]
PROPERTIES	Greater acoustic energy	Decreased energy	Greater acoustic energy
	Longer duration	Shorter duration	Longer duration
	No lengthening effect on preceding V	Preceding V lengthened by about 130%	No lengthening effect on the preceding V

The same study compared *pas d'rôle* 'no role' with *pas drôle* 'not funny': [ʁ] has a syllable-initial allophone in *rôle*, and a syllable-internal one in *drôle*. After schwa loss, [ʁ] maintains the syllable initial quality.

Rialland concluded that the string resulting from schwa loss is syllabified differently from strings lacking an underlying schwa. She suggests that syllables originally headed by schwa survive the loss of their nuclear vowel: thus *pas d'rôle* consists of the syllables [pa.dʁ.ʁol], with ambisyllabic and partly nuclear [ʁ]. This parse explains the observations made earlier: [ʁ] in *pas d'rôle* is different from [ʁ] in *pas drôle* because the former continues to stand in syllable-initial position in [ʁol], unlike the latter.

Table 22.2 *Rialland's observations:* pas de role, pas drole, pas d'role

	Syllable initial r	Onset r, syllable internal	r next to lost schwa
EXAMPLES	*pas de role* [padəʁol]	*pas drole* [padʁol]	*pas d'role* [padʁol]
PROPERTIES	Greater acoustic energy	Decreased energy in the higher formants	Greater acoustic energy
	Longer duration	Shorter duration	Longer duration

There is a basic obstacle to this interpretation: the syllable count depends in French on the presence of audible vowels, including [ə]. A verse such as *Il en est de pires il en est d'meilleurs* [il ɑ̃ e də piʁ, il ɑ̃ e d mɛjœʁ] ('There are worse and there are better ones')[7] scans correctly only when it contains exactly 10 syllables. If *d'meilleurs* 'of better ones' contains schwa, as [də mɛjœʁ], the verse is unacceptably long (cf. Dell 1989). Therefore we cannot adopt Rialland's suggestion that the invariant property in C(ə)C sequences is the number of syllables: when schwa deletes, the syllable count is correspondingly decreased. What will account then for her observations on allophone distribution?

Here too the answer invokes the effect of PU conditions on phonetic detail. Consider the phrase *de rôle* in *Il n'a pas d(e) rôle*. At least one component of this phrase—the noun *rôle*—can be produced in isolation. Its isolation form will necessarily contain a syllable initial [ʁ] allophone, characterized by longer duration and increased acoustic energy. It appears that such properties of the citation form are preserved in the reduced string *d'rôle*, after loss of schwa, whether or not the syllable structure is recomputed. Similarly, the citation forms of the components of *bas r(e)trouvé* are *bas* [ba] and *retrouvé* [ʁətʁuve], with schwa and syllable initial [ʁ]. It is the relative duration of *a* in [ba] and the syllable-initial properties of [ʁ] in [ʁətʁuve] that are preserved in the corresponding schwa-less phrase [baʁtʁuve]. The preservation of these properties of the citation or careful pronunciation can be attributed to constraints requiring the invariance of morpheme edges: all cases discussed by Rialland involve morpheme-initial or morpheme-final consonants whose quality remains relatively invariant in utterances with and without schwa. With this in mind, Cécile Fougeron and I have attempted to replicate Rialland's results through electropalatography (Fougeron & Steriade 1997). One group of utterances we studied is (11).

Item (a) involves an unambiguous prevocalic onset [d] followed in the next syllable by a syllable initial ʁ. Item (b) involves *d'* from [də] followed by [ʁ].

(11) a. *Il n'a pas de rôle en ce moment* [ilnapadəʁolɑ̃səmomɑ̃]
'He has no role right now.'

b. *Il n'a pas d'rôle en ce moment* [ilnapadʁolɑ̃səmomɑ̃]
'He has no role right now.'

c. *Il n'est pas drôle en ce moment* [ilnepadʁolɑ̃səmomɑ̃]
'He's not funny right now.'

d. *Il voit le jade rose en ce moment* [ilvwaləʒadʁozɑ̃səmomɑ̃]
'He sees the pink jade now.'

It is the properties of this [d] that we focused on. Item (c) involves an underlying onset [dʁ] sequence. Item (d) involves a coda [d] followed by an onset [ʁ]. Our conjecture was that the *a#d#ʁo* sequence (item b) will be systematically different from both *a#dʁo* (item (c)) and *ad#ʁo* (item (d)). Moreover, on the dimensions that distinguish (b) from (c) and (d), the (b) tokens will be closer to (a), the morphologically related form. We did not expect complete identity between the allophone of [d] in (a) *a#də#ʁo*, and that of [d] in (b) *a#d#ʁo*: one is followed by a consonant while the other is followed by a vowel. However, any similarity between (a) and (b) that is unexpected based on the surface composition of the string should be attributable to the effects of Paradigm Uniformity.

Two speakers fitted with Kay pseudopalates produced 20 repetitions of each of the sentences in (11). The results indicate that the amount of linguopalatal contact (measured at the point of maximum closure) is greatest for the prevocalic [d] of (11a) and not significantly different for the paradigmatically related preconsonantal [d] of (11b). The [d]'s of (11c, d), which are lexically unrelated and not underlyingly prevocalic, have significantly reduced contact in comparison to those in (11a, b). The duration of dental closure shows the same pattern as the amount of linguopalatal contact: the [d]'s of (11a) and (11b) are significantly longer than those of (11c) and (11d). Finally the incidence of [d]-lenition (tokens lacking a linguopalatal seal) displays a comparable pattern: under 10% of the [d] tokens in (11a, b) were lenited in contrast to 30% of the (c) tokens (onset [d] in *drôle*) and 40% of the (d) tokens (coda [d] in *jade*).

These results confirm Rialland's original observations. They also establish that the syllable-based explanation she offered does not account for the similarity between lexically-related strings like *de* and *d'*: according to Rialland's analysis, the *d'* of *d'rôle* occurs in the syllable [dʁ], whose nucleus is the syllabic first half of [ʁ]. Therefore this [d] occurs in a segmental context that differs from that of the original string [də]. Nothing in Rialland's analysis predicts that this [d] will be identical to the prevocalic [d] of [də] in duration, amount of contact and lenition possibilities. Thus, even if we maintain the hypothesis of

syllable invariance, we would still be missing an explanation for the similarities between *de* and *d'* .
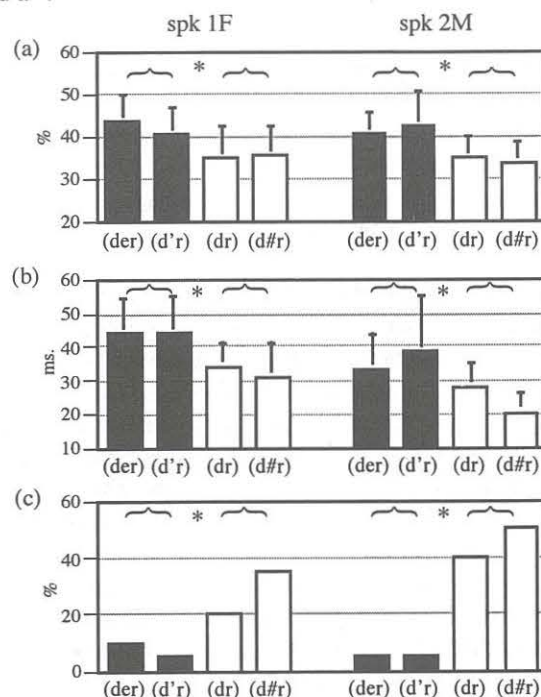

Figure 22.1 (a) Amount of linguopalatal contact in [d]; (b) Duration of the lingual occlusion gesture of [d]; (c) Frequency of lenition of [d] in the 4 types of sequences: *de rôle* (der) ([e] = schwa), *d'rôle* (d'r), *drôle* (dr), and *jade rose* (d#r) (from Fougeron & Steriade 1997).

The analysis we propose for this data involves the interaction of constraints inducing durational reduction and eventually lenition with PU constraints. Our interpretation is that [d] is temporally reduced in preconsonantal position (or perhaps just before certain consonants such as [ʁ]). The diminution in extent of contact and [d]'s lenition are possible consequences of this durational reduction. We speculate that the PU constraint which ensures the close similarity between the articulation of [d] in *d'* and *de* refers only to the durational category of the consonant. Under this interpretation, a single correspondence constraint will be needed here, which imposes durational equivalence between the left edges of morphemes in careful and casual pronunciation.

(12) PU (LEFT: DURATION)
If two consonants, C and C', stand in correspondence and C is morpheme initial in the careful pronunciation of the relevant morpheme, C' is durationally equivalent to C.

This condition can shield the non-prevocalic [d] of *d'rôle* from lenition, if indeed lenition is a direct consequence of temporal reduction. If the temporal and spatial reductions observed are independent of each other, we will need to adopt a distinct PU condition—or an augmented (12)—that mentions equivalence for extent of contact. This analysis agrees with Rialland's on the point that some characteristic property of the citation or careful form is inherited by the phrase-medial schwa-less variant. Critically however, this property cannot be the number of syllables. We identify it as the durational category of the consonants and possibly their extent of contact: neither of these is an independently contrastive property and both display the sort of token-to-token variability that is said to characterize phonetic detail properties. Yet both of them are subject to paradigmatic leveling. This is then another phonetic analogy effect.

Other studies (Jun & Beckman 1993, and Manuel et al. 1995) observe effects comparable to those reported here: a process of gestural reduction or an increase in overlap between two gestures is inhibited to maintain lexical distinctions that are more prominently present in the careful pronunciation of the relevant forms. Thus the reduced pronunciation of *in the* may appear to sound just like *in a* but in fact maintains a systematically longer between-vowel interval and is identifiable as *in the* by listeners (Manuel et al. 1995). The loss of French schwa could be analyzed similarly, by claiming that a trace of the original schwa gesture is maintained, albeit reduced and completely masked by the neighboring consonants. However, what makes the case of French uniquely relevant to our discussion is the fact that schwa loss has phonological consequences: the syllable count depends on it. This is why all previous work on French schwa assumes that phonological rules and principles govern the occurrence of this vowel. Thus, while the processes discussed by Manuel et al. and Jun & Beckman may be interpreted as phonetic implementation rules, the case of French schwa indicates that phonetics and phonology are not easily separated.

22.6 Conclusions

The phenomena discussed here suggest that the realization of phonetic detail properties is governed by some of the same principles that must be invoked in studying phonological or potentially contrastive features. The family of principles discussed here involves paradigm uniformity. Earlier work (Docherty 1992, Keating 1984, Kingston & Diehl 1994, Pierrehumbert & Beckman 1987) has established that some aspects of phonetic implementation are speaker-controlled, rule-governed and possess language-specific characteristics, just like

the rest of grammar. The present study suggests a means of strengthening this result, by showing that some processes located in the 'phonetic implementation component' are qualitatively the same as the ones classified as phonological. Phonetic analogy is qualitatively the same process as cyclicity, the paradigmatic extension of contrastive properties. If this result can be maintained, then it should lead us to question the distinction between phonology and linguistic phonetics, i.e. speaker-controlled phonetic processes. The claim made here is that at least the *feature sets* of phonetics and phonology need not be distinct. A more conservative assessment of what has been shown is that if the phonological and phonetic feature sets are in fact distinct, they are not distinguished by potential contrastivity but in some other still unidentified way.

Let me conclude by addressing two questions raised by the possible unification of phonological and phonetic features. The first involves the functioning of paradigmatic uniformity. Phonetic analogy—the paradigmatic extension of non-contrastive properties—is far less categorical than instances of paradigm leveling affecting the global distribution of contrasts. Thus, a given token of French *d'rôle* can be produced with a *d* that is identical in duration and extent of contact to the *d* of *drôle*. The PU effect in the French case accounts for the trend rather than for the quality of individual tokens. Similar variability is not reported in the study of 'phonological' cyclic effects. This may be due to the fact that phonological studies rely on introspective reports but a difference is likely to emerge even if the investigative methods are held constant. We must explain this difference in the categorical status of paradigmatic extensions. I would speculate here that any sound property or any cluster of properties may give rise to paradigmatic leveling but that the categorical or variable nature of the effect will depend on the perceptibility of the property being generalized through leveling. The less perceptible the contrast generated in this way, the harder it is to detect and enforce uniformity in each and every relevant token.

The second question takes us back to the beginning of the paper. The idea that some phonetic categories lack phonological relevance was inspired by the observation that lexical contrasts are limited in number, in any given language. How should this observation be handled if we fail to distinguish phonetic from the much smaller set of phonological categories? An answer to this second question must proceed from the distinction between *feature* and *contrast*. The grammatical object that phonologists refer to by the term 'the feature [±voice]' is a contrast, not a feature (cf. Keating 1984). It is a contrast implemented through a large number of features: closure duration, prevoicing, VOT, pitch etc. The contrast is robust across contexts and speech circumstances only when many features are jointly employed to distinguish its terms. This means that in order to have some optimally differentiated contrasts, a language must drastically limit their numbers, so as to minimize the featural overlap between contrastive categories (Lindblom 1990b, Flemming 1995). A primitive example of

the role played by this distinction between features and contrasts is the statement in (13), which requires every stop to possess a certain number of attributes that jointly identify it as belonging to either the 'voiced' category (a) or the 'voiceless' category (b).

(13) A stop must have either one of the following sets of properties:

    a. {short closure duration, voicing during closure, VOT value $< x$ ms, long V_1}

    b. {long closure duration, no voicing during closure, VOT value $> y$ ms, short V_1}

The condition in (13) requires voice-differentiated stops to differ with respect to four distinct features. In principle, this condition prohibits the contrastive use of any of the properties listed in (13) in stops: for instance, any stop with a long closure will necessarily possess all the other attributes of voiceless stops, the absence of closure voicing, the longer VOT, the shorter V_1. To do otherwise would violate condition (13). In practice, the surface effects of (13) will be a function of the interactions between it and competing constraints of the language. The statement in (13) thus amounts to a violable constraint on the expression of the voicing contrast. Compare now a contrast based on all the features in (13) with one based exclusively on closure duration, i.e. the smallish durational difference separating English *t* from *d*. The contrast implemented through only one feature is obviously worse than the one based on four: it is worse not because it uses the *wrong* feature but because it does not use enough features to differentiate its terms.

The suggestion then is that by exploiting constraints on contrast like (13) we eliminate a major reason to impose limitations on what should count as a phonological feature. The feature set need not be restricted in order to distinguish good contrasts from bad ones: a theory of contrast-goodness and specific constraints on contrast are sufficient for this task (Flemming 1995). We are therefore free to assume, if necessary, that all properties listed in (13) are grammatically relevant despite the fact that none of them is independently contrastive. This paper has shown that this is indeed necessary: non-contrastive features such as [extra-short closure] *are* grammatically relevant, in the sense that the evaluation of paradigmatic uniformity constraints requires the grammar to note distinctions based on these attributes.

Notes

1 See also Burzio 1994, 1996, Flemming 1995, Benua 1998 and Kager 1998 for further discussion of these issues.

2 These points emerge also from a study by Cutler (1979) on subject preferences for nonce words using stress-neutral suffixes (*-able, -ment, -ness, -ish*) as against nonce forms built with stress-modifying suffixes:(*-al (N), -ial (Adj.), -ity, -ous*). Cutler's subjects did not evince any global preference for productive as against unproductive suffixes: but their comments did display awareness of the conflict between the considerations of base invariance and metrical well-formedness mentioned in the text. Thus several mentioned that "although *villagérial* might [...] be preferred to *víllagerish* as an English word, the latter would be more likely to get the message across [...]" (p.84). We speculate that those who found *villagérial* to be "aesthetically more pleasing" did so on the grounds of LAPSE avoidance (for which see below).

3 The reader will note that a form like *dìsciplináble* (or *dísciplinàble*))—with some level of stress on the suffix—satisfies both PU(STRESS) and *LAPSE. This fact reflects further conditions on affix invariance, dicussed by Burzio (1994).

4 Withgott rejects a cyclic analysis of the *capitalistic–militaristic* contrast on the grounds that an unexpected stop is also encountered in the word *Mèditerránean*, where no cyclic principle can predict it. I attribute the unflapped [t] in this unique form to the orthographic geminate 'rr', which is interpreted by speakers as an indication of secondary stress on the preceding vowel (Nessly 1977). In forms lacking a geminate, e.g. *mèri[ɾ]ocrátic*, the *t* is regularly flapped.

    The effect of tap suppression observed here does not obtain in syllables that directly follow the tonic: *statístic–stà[ɾ]istícian*. The interpretation of this data is not entirely clear but what seems certain is that very few instances of non-tapped t's in the *v́_V* context have been encountered so far. This may be due to the fact that constraints that induce tapping are more stringent (i.e. more highly ranked) in the immediate post-stress position than elsewhere. PU effects surface only when the tapping constraint is weaker.

5 See Kirchner 1998 on the reasons why this particular context induces closure shortening and lenition.

6 The articulatory properties of taps are reviewed by Banner-Inouye (1995). The absence of a jaw raising gesture in flaps is noted by Fujimura (1986). Voiceless taps do occur which indicates that the effect of shortening on glottal opening may have to be controlled by a distinct condition. The characteristic absence of jaw raising and glottal opening movements may be seen as consequences of the extra-short duration allotted to closure.

7 From a poem by Georges Brassens.

23

Commentary: Counting, connectionism, and lexical representation

GARY S. DELL

A colleague of mine once told me that grammars don't count. She meant that numbers—reals, integers, or anything that identifies quantity—are not a part of the formal system that characterizes linguistic knowledge. There are no rules that allow, for example, something to happen three, but not four, times. I later learned that this is not entirely true, at least for phonology, where a rule may permit two but not three syllables per foot, or where the number of morae is an important property of a syllable. Nonetheless, the idea that linguistic knowledge is stated solely in terms of nominal, rather than numeric, categories has congenially coexisted with my belief that linguistic *data*—whether or not a given form exists in the language—is also nonnumeric. We were taught during the classical period of the 1960s and 1970s that linguists build theories of competence (or generative grammars) and these theories accounted for linguistic judgments. Neither the theories nor the data involved counting. In contrast, those of us who studied language performance, such as psychologists and phoneticians, were allowed to hypothesize quantities and use numerical data. However, we were also told that the relation between our data and theories and competence theories was quite abstract. Consequently, we had little expectation that competence theories would stand or fall on our numbers.

    The papers that form the basis of this commentary, and more generally, those of the several Laboratory Phonology conferences illustrate the fact that the classical period is long over. I have never before seen so much counting at a linguistics conference! In the lexical representation session, what are being counted are lexical items. And the resulting quantities are being used to develop radically new approaches to phonological knowledge and the relation between knowledge and performance.