Degenerative Phonology
Daniel Silverman

Part 1
Theory

Chapter One
Foundations

1. Many-to-many sound-meaning correspondence

At its most basic level of description, language structure involves a correspondence of sound (with a *form*) and meaning (with a *function*). Significantly though, no language possesses a one-to-one correspondence between sound and meaning—between form and function—such that phonetic elements are uniquely paired with semantic elements. Inevitably—and, as will be presently seen, for good reason—all languages possess (1) many-to-one correspondences between sound and meaning in the form of heterophone-maintaining alternations (a ubiquitous occurrence), and (2) one-to-many correspondences between sound and meaning in the form of homophone-inducing alternations (a rare occurrence). More specifically, this many-to-many relation between sound and meaning is *asymmetric*, in the sense that heterophonic alternations always far outnumber homophonic ones.

An asymmetric many-to-many sound-meaning correspondence being the *de facto* state of linguistic affairs, the tasks for the linguist include:

(1) Isolating the myriad pressures that interact on linguistic systems such that a one-to-one form-function correspondence is inevitably stymied,

(2) Motivating the fact that heterophone-maintaining alternations are rampant while homophone-inducing alternations are rare,

and, ancillarily,

(3) Determining if the sound components of sound-meaning correspondences are decomposed by language users into smaller elements that might combine and recombine with each other.

These research goals are pursued herein by investigating the manifold interactions between:

(1) The linguistic object, embodied as the product of conflicting pressures acting on (a) phonetics (form), and (b) semantics (function), and

(2) The linguistic subject, embodied as the product of conflicting pressures acting on (c) speakers, and (d) listeners (see Fig. 1.1).

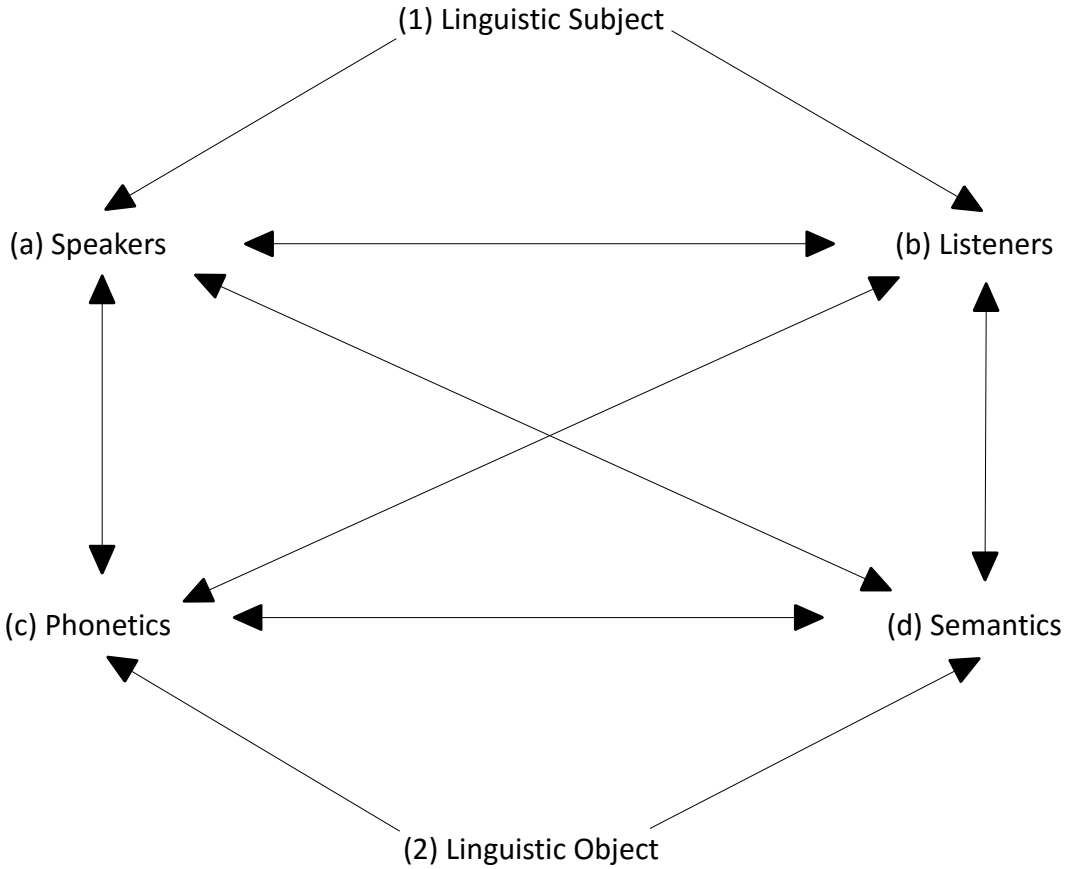

Figure. 1.1. Interaction between (1) the linguistic subject ((a) speakers and (b) listeners) and (2) the linguistic object ((c) phonetics and (d) semantics)

Briefly, linguistic structure in general, and *alternation* in particular, is herein argued to have its *indirect* origins in the interlocutionary act itself, in the successful conveyance—from speakers to listeners—of the elements of meaning (morphemes). The variation inherent to speech production, and the selectional pressures acting on this variation, culminates in a system that naturally and passively serves its communicative function. Morph selection—as embodied in alternation—is thus conditioned by the interaction of the pressures acting on the (1) linguistic subject (speakers and listeners), and (2) the linguistic object (phonetics and semantics).

Consider, for example, one common route to alternation (as we will see, there are others, too): under those particular circumstances in which meaning is successfully conveyed to listeners *despite minor articulatory simplifications* that are intermittently present upon morpheme concatenation—typically (though not exclusively), in the form of *assimilation* and/or *reduction*—then these simplified forms may be recycled by listeners as they themselves speak, eventually becoming conventionalized.

The end-product of these iterated scenarios may be both a simplification of the motor routines put in service to recurrent components of the speech code ("words"), their better phonetic separation, and concomitantly, their better separability (for listeners): frequently-required semantic content involves frequently-produced morpheme groupings (again, "words"), and thus involves frequently-produced phonetic content. Exactly due to their frequency and their consequent predictability, those phonetic

productions that are somewhat simplified (assimilated, reduced) in particular contexts may yet be successful in conveying the semantic content intended by speakers. The structural coherence of such frequently employed morpheme groupings may thus be cued in part by the very assimilatory patterns that are so prevalent within them: the less-common acoustic transitions that are characteristic of so-called word boundaries aid listeners as they parse the speech signal into its semantic components. Words and their phonotactic regularities, then, may passively emerge due exactly to recurrent strings of morphemes' repetitive use, and their context-dependent phonetic adjustments. The result may be a regularization of the motor routines put in service to encoding semantic content, and the concomitant emergence of alternations, as morphemes of different phonetic forms combine and re-combine with each other.

Thus, over the course of their early interlocutionary experience, listeners become better-practiced in deciphering a speech signal that is—and, as a consequence of its evolution, always has been—in a state that lends itself to just this decipherment. Practiced listeners may thus exploit as parsing aids the less-common acoustic patterns encountered at word boundaries, and the more-frequent acoustic patterns (and their accompanying limited inventory of motor routines) encountered word-internally.

Such patterns may prevail until listener confusion would set in: if word-internal motor routines would become too simplified, and hence, inevitably, too similar to each other, the requisite semantic clarity of the speech signal would become jeopardized, because semantically distinct primitives that are phonetically distinct in some contexts may become phonetically *non*-distinct in others.

This is a diachronic source of induced homophony, necessarily limited in prevalence due to the simple fact that an excess of such forms may interfere with the successful transmission of semantic content: only *successful* speech propagates. Unsuccessful speech is passively filtered out. The interlocutionary system thus has a built-in homophony-limiting mechanism.

The result is a system possessing both one-to-many and many-to-one correspondences between form and function of a specifically asymmetric character, since heterophonic alternations far exceed in number homophonic ones. Exactly because of its patterns of use and disuse, the system passively maintains a structure that facilitates both its function and its mastery.

There thus exist usage-based semantically-rooted diachronic pressures both promoting and, eventually, inhibiting an overall simplification of the phonetic content of the speech code. The product of these iterated interlocutionary tendencies is a linguistic system that naturally settles towards a semantically unambiguous state, a state whose semantic elements are combined and expressed with a limited inventory of motor routines, a state in which heterophony is rampant, and homophony is passively limited.

2. Degeneracy

Adapting terminology employed to characterize biological and other complex adaptive systems, the sort of system just outlined evinces both *degeneracy* and *pluripotentiality.* Degeneracy is present in a system when *single functions are subserved by multiple forms.* Pluripotentiality is present in a system when *single forms are responsible for multiple functions*. Herein, we conflate these two properties, subsuming the latter into the former, exploring in preliminary detail the proposed degenerate character of phonological systems. *Degenerative phonology* is thus both the subject (the theory) and the object (the data) of our investigation.

As with all systems that are degenerate in character, a degenerative phonology possesses elements that are at once (1) sufficiently impervious to insult such that they remain vital to the proper functioning of the system as a whole (culminating in the system's *robustness*), (2) sufficiently variable such that they might adapt to new conditions coming to act on their form (culminating in the system's *evolvability*), and (3) sufficiently interactive such that they enter into a hierarchical organization (culminating in the system's *complexity*). Indeed, robustness, evolvability, and complexity, are inherent properties, hence hallmarks, of any degenerate system (Whitacre 2010).

The linguistic system in general, and the morpho-phonological system in particular, is subject to myriad pressures—some in a state of antagonism, others in harmony—such that a one-to-one relationship between form and function is inevitably stymied, but stymied *not* as an incidental artifact of wholly independent pressures on the evolution of the system, but rather, stymied because *degeneracy is inherent and crucial to the system's functional efficacy*: *any* complex system that is subject to evolutionary pressures on its forms and its functions is likely degenerate by its very nature.

Figure 1.2 presents a fairly standard (and, as will be immediately argued, a somewhat incomplete) way of schematically portraying the many-to-many nature of degenerate systems. For present purposes, again, *function* refers to elements of meaning (morphemes), and *form* pertains to these elements' phonetic expression (morphs).

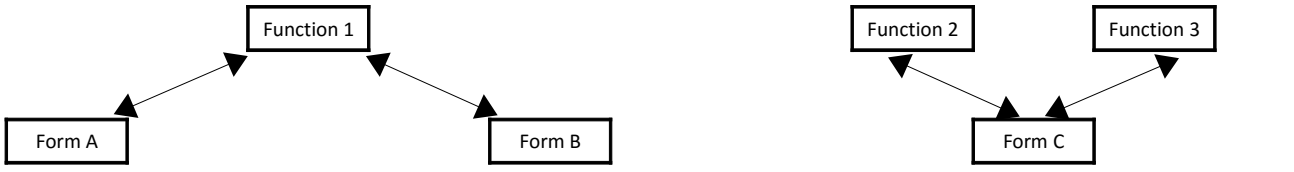

(a) Maintained heterophony (ubiquitous) (b) Induced homophony (rare)

Figure 1.2. Many-to-many form-function relationships.

In Figure 1.2a, a single morpheme (Function 1) is associated with two morphs (Forms A and B). Forms A and B are thus heterophonic alternants. In 1.2b, a single morph (Form C, one among more than one alternant) is associated with multiple morphemes (Functions 2 and 3). Form C is thus a homophone.

But despite its apparent straightforwardness, Figure 1.2 does not compellingly convey the degenerate nature of the system, primarily because the components that are being paired here—Function 1 with Forms A and B; Functions 2 and 3 with Form B—are considered in the absence of the contexts that induce the specific characters of their respective form-function relationships. Instead, in order to understand morph selection, it is vital to consider a larger domain, one that includes relevant morpheme-external content.

So consider the situation portrayed in Figure 1.3. In 1.3a, a morpheme has two alternants, again, Forms A and B, that acquire their partially distinct phonetic characters as a consequence of the phonetic properties of the morphemes that follow (here, Forms D and E): the end-spans of A and B are affected by the beginning-spans of D and E (and the beginning-spans of D and E are affected by the end-spans of A and B). That is, the non-final spans of the first forms may be determined in whole by morpheme-internal content, while their final spans are determined in part by (or *bonded* with) morpheme-external content. This establishes distinct *temporal spans of bonded material* between the

potentially stable flanking spans of A and D (Bond 1), and also between the potentially stable flanking spans of B and E (Bond 2); bonds are shaded. These phonetically distinct bonds thus embody A-B alternation. The same holds for Form C in 1.3b: as a consequence of Bond 3 (itself a consequence of Form F's placement), Form C is phonetically non-distinct from a form of some other morpheme.

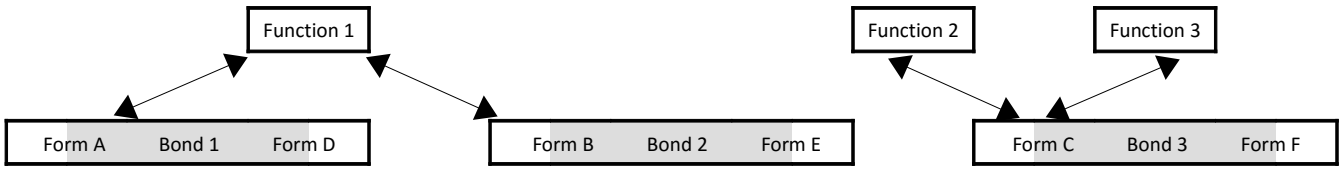

(a) Maintained heterophony (ubiquitous) (b) Induced homophony (rare)

Figure 1.3. Degeneracy in form-function relationships

As will be argued, Figure 1.3 highlights the proposal that bonding among morphemes is a crucial characteristic of the degenerative phonological system, with, as will be seen, major consequences for both the linguistic object (phonetics and semantics), and the linguistic subject (speakers and listeners).

For speakers, the bond embodies a simplification of the motor routines put in service to frequently used morpheme groupings (words), and further, it increases the speed of information encoding, since it efficiently organizes the semantic content of distinct morphemes by means of simultaneous phonetic cueing.

For listeners, the bond thus provides information about *both* morphemes: repeated encounters with Bonds 1 and 2 quickly come to unambiguously signal Forms A and B's identical semantic content, while also providing some "look-ahead" information about both the phonetic content of following Forms D and E, and (especially with high-frequency morpheme groupings) their distinct semantic content as well. As will be seen, bonding serves these functions (although to a far lesser extent) even across word boundaries.

In the vast majority of instances then, bonding assists in the establishment of *paradigmatic* relationships among forms (heterophonic alternants' semantic non-distinctness), and in the establishment of *syntagmatic* relationships among forms (parsing). Far from being a drag on efficiency, bonding plays a crucial role in the evolution of the interlocutionary system: information flow between speaker and listener is sped, enhanced, and clarified.

But now consider Figure 1.3b. Here, Bond 3 participates in the phonetic character of Form C, thus indeed providing some look-ahead information about both F's phonetic (and often semantic) content, and yet, the resulting structure actually subserves two functions, as it is homophonous with some other morpheme or morpheme alternant. Still, genuine ambiguity and hence listener confusion is unlikely to arise; recall, the system has a built-in mechanism inhibiting the pervasion of semantically ambiguous linguistic structures.

As will be discussed, understanding the crucial role of bonding in a degenerative phonology also enhances our understanding of stem-modifying, fusional, vowel-harmonic, and other sorts of non-concatenative morphological patterning. It will further be suggested that it is the evolved recyclability

of bonds and their attendant motor regularities that is source of *phonological productivity*.

To sum up, the frequent juxtaposition of particular morphemes manifests an symmetric many-to-many relationship between form and function; a degenerative phonology. Information-rich bonding content increases the efficiency of semantic *en*coding for speakers, and semantic *de*coding for listeners. Degeneracy serves to provide cues to both the paradigmatic and the syntagmatic relationships among linguistic elements, hence enhancing the system's structural and functional *robustness*, *evolvability*, and *complexity*; the hallmarks of any degenerate system.

3. Non-Compositionality

The recurrent motor routines characteristic of speech (whatever their form may turn out to be) are not, in and of themselves, linguistic primitives. This is because, quite simply, any particular motor routine that might be both isolable and recyclable does not typically pair with any particular semantic function. Thus, although they obviously constrain speech patterns in particular ways in particular languages, still, the absence of any regularity in form-function pairing between these recurrent motor phenomena and recurrent semantic phenomena precludes their candidacy as genuine linguistic primitives. And after all, the functional relevance of any phonetic component of the linguistic system is established exclusively by its role in maintaining distinctions in *meaning*, not by maintaining distinctions in sound itself.

Instead, it will be argued herein that morphemes themselves (or, rather, their respective inventories of alternants) are the genuine elements of phonological structure, since it is the morph-morpheme correspondence that embodies the elemental pairing of form and function, of sound and meaning. Consequently, just as there is no compelling evidence to support the proposal that recurrent and recyclable motor routines are componential linguistic primitives, neither is their compelling evidence that sub-parts of these routines are componential (again, however recurrent and recyclable they are purported to be), precisely because such proposed structures do not directly participate as independent players in the degenerate system of form-function relations.

4. Preliminary exemplification

Consider the examples of Spanish nasal assimilation in Table 1.1 (adapted from Nathan 2008).

form:	function:
ũn-õmbre	"a man"
ũm-beso	"a kiss"
ũŋ-faktor	"a factor"
ũŋ-gato	"a cat"

Table 1.1. Spanish nasal assimilation exemplified.

Based on the preliminary discussion that precedes, it should be clear that the symbol-by-symbol rendering of the sound-and-meaning pattern in Table 1.1 conveys neither the true nature of the form-

function relations here, nor the degenerate character of the system in which these patterns are embedded. The transcriptional changes to the indefinite article suggest that when one morpheme comes to abut (though, counter-factually, not bond with) another, there is a phonetic switch-out of sub-morphemic content (here, the nasal), but the sequenced morphs themselves remain otherwise unchanged, readily distinct from one another, and readily separable from one another. The presence of the so-called "morpheme boundary" symbol ("-") hammers home this flawed characterization, thus nailing shut the possibility of conveying the degenerate alternative.

Herein then, the International Phonetic Alphabet is augmented by a simple system of underscoring and overscoring, in a preliminary attempt to graphically suggest the truer-to-nature form of morpho-phonological bonding: The first typographically-sequenced morph is underscored, the second overscored. Underscoring and overscoring highlight (rather roughly) the distinct morphs as their phonetic content is distributed in the speech stream. Components of the structure that embody the bond thus possess both underscoring *and* overscoring, as in table 1.2. (When morphs appear embedded in context, the en dash is *not* intended to represent a so-called "morpheme boundary". Rather, it is intended as a *variable*, suggesting that there is additional phonetic material beyond the typographic frontier that, although it varies as a consequence of context, is nonetheless crucial to the phonetic character of the morph(s) with which it is affiliated.

form:	function:
ũn-õmbre	"a man"
ũm-beso	"a kiss"
ũŋ-faktor	"a factor"
ũŋ-gato	"a cat"

Table 1.2. Spanish nasal assimilation exemplified with morpheme affiliation indicating by under- and overscoring.

The overscoring in Table 1.3 properly indicates that the virtual entirety of the morpho-phonological complex (the word) contributes phonetic cues to the root. Consider ũm-beso. Clearly, the assimilated nasal conveys information about the following functional element (the root)—*it is part of the root, as much as it is a part of the affix*—informing listeners that this root is labial-initial, and thus serving to narrow listeners' lexical search. Nasal assimilation is tolerated here because homophony is rarely if ever induced, and thus even these articulatorily simplified variants may yet be successful in conveying listeners' intended meaning.

But also, vowels that flank consonantal spans are mutually influential, and, necessarily, also affect (and are affected by) the spectral properties of the intervening consonantism itself (Öhman 1966). The so-transcribed ũ thus actually bears the phonetic mark of the nasal-stop-vowel span, since the onset transitions as the lips close for m-b are influenced by the offset transitions as the lips open again into e. Again, the virtual entirety of the affix is bonded with the root, thus possessing some phonetic "look-ahead" information about the semantic element that is to follow.

So, when bonding with another root, as in, say, ũŋ-gato, this same affix contains modified phonetic properties as influenced by the different phonetic shape of the root. Thus, here too, due to bonding, the affix conveys "look-ahead" information about the form (and oftentimes the function) of the root itself. In short—and typographically misleadingly—the ũ of ũm-beso is both phonetically distinct from, and may serve to convey partially-distinct semantic content of, the ũ of ũŋ-gato.

To be clear, whenever the indefinite article appears with roots of different shapes, it inevitably engages in a phonetically and semantically informative alternation that encompasses a significant majority of its temporal span: the bond here encompasses almost the entirety of the article, but also encompasses a non-trivial temporal span of a following root. The root, meanwhile is affected well into its initial vowel, thus reinforcing the phonetic properties of the article itself. Recurrent experience with ũm, ũɱ, ũn, and ũŋ quickly inform learners that their phonetic differences are semantically inert with respect to the article, but are semantically active with respect to the following morpheme, thus providing information about both the paradigmatic and the syntagmatic properties of the speech stream. It is a consequence of their *frequency of use* that bonds emerge, and thus those less-frequently employed phonetic routines at so-called word boundaries evolve weaker bonds, and so, by dint of their rarity, their more perspicuous phonetic discontinuities come to serve as parsing aids, cueing to listeners that a new morphological complex (a new word) has begun.

Referring to the Spanish morphological system as *concatenative* in nature—one that is qualitatively different from so-called non-concatenative systems of various sorts—may thus be revealed to exemplify a specious distinction. Rather, the different word-formation systems found in the world's languages are perhaps better seen as residing somewhere on a sliding scale, with different languages evolving towards different degrees of bonding, but not differing from one another in a genuinely qualitative way.

Summary

It may be a jarring realization, that morphs are not self-contained phonetic entities, but rather, in an organic-like way, they spread, they merge, they interact, and they overlap with the phonetic content of other morphs, and thus the speech stream simultaneously encodes information about multiple semantic entities. But to the extent that the linguistic system is degenerate in character, this organicity should not be surprising at all. Indeed, the morpho-phonological system may bear a remarkable likeness to organic systems of growth and development not by coincidence, but because the deep pressures and principles that affect its structure are qualitatively non-distinct from those that affect any and all complex adaptive systems.

Still, there is nothing particularly novel about this approach to morpho-phonological structure:

(1) It acknowledges that phonological systems are put in service to encoding (and, largely, ensuring) distinctions in meaning, just as any and all phonemic/segmental approaches do.

(2) It acknowledges that alternations have consequences not only for phonological structure, but for the encoding of meaning as well, just as any phonological theory that draws a distinction between so-called allophonic (necessarily heterophone-maintaining), and neutralizing (potentially homophone-inducing) alternations.

(3) It acknowledges that juncture phenomena are not merely worthy of observation, but are linguistically relevant, just as virtually all theories have always done.

Where degenerative phonology departs from other approaches is in its proposed *locus of explanation* for the phonological patterning that all scholars of linguistic sound structure investigate. As will be argued, a degenerative phonological analysis operates under the assumption that the linguistic system is qualitatively non-distinct from other complex adaptive systems in terms of its organizing principles, and thus requires no special, domain-specific theoretic machinery for its operation.

Regarding the remainder of Part One ("Theory"), Chapter Two ("Background") briefly discusses the history of the concept of degeneracy. Chapter Three ("Origins") considers the possible pre-linguistic origins of the degenerate system, proposing that the pressures and principles responsible for its phylogenetic emergence are also responsible for its maintenance . Chapter Four ("Bonding"), considers the sources and functions of assimilation, lenition, and other sorts of boundary-blurring, bond-creating phenomena, and briefly considers their relevance to parsing. Finally, Chapter Five ("Morphology"), explores how degeneracy may manifest itself in a variety of ways, as concatenative, partially-concatenative, and non-concatenative morphological systems.

Regarding Part Two ("Data"), in Chapters Six ("Heterophony"), a number of case studies are considered that apply the conclusions of Part One. Chapter Seven ("Homophony") considers linguistic data supporting the claim that induced homophony is inevitably limited in its prevalence. Finally, in Chapter Eight ("Contrast"), a case will be made for the non-compositionality of morphemes, that is, for the phonological unanalyzability of morpheme-internal content.