Chapter 7

Morphology As Lexical Organization

Joan L. Bybee

Most current conceptions of the apparatus behind linguistic behavior postulate separate components for rules and representations. Representations are static and fixed, the individual and idiosyncratic content of the morphology, while the rules are the "moving parts," the dynamic, the general statements that range over the representations. In the years of intensive research into the structure and domain of morphological and phonological rules since the mid-1960s, it has become apparent that many different types of rules exist and that they differ from one another in their freedom of application. Current work in modeling phonology and morphology attempts to adjust and divide the rule component in various ways in order to describe differences among rule types. Leaving aside autosegmental phenomena, representations and their properties have been the subject of interest only to the extent that they feed into the rules, which means that their form must be adjusted as rules and their mode of application are changed. This chapter takes a different perspective on representations, focusing on the lexicon directly and approaching rules as generalizations that arise from representations.

Among the descriptive frameworks for morphology utilized in the twentieth century, the one that has enjoyed the longest popularity and the widest application is the Item and Process (IP) model, currently best known as developed in generative phonology. The defining characteristic of an IP model is that it sets up one underlying form for alternating allomorphs and derives the surface forms by applying feature-changing rules to the underlying form (Hockett, 1954). Consider the stem vowel alteration in the present indicative of the following Spanish verbs, which are representative of a large lexical class of verbs:

(1)

	'to begin'				'to tell, to count'		
1sg	*empiézo*	1p	*empezámos*	1sg	*cuénto*	1pl	*contámos*
2sg	*empiézas*	2pl	*empezáis*	2sg	*cuentas*	2pl	*contáis*
3sg	*empiéza*	3pl	*empiézan*	3sg	*cuénta*	3pl	*cuéntan*

The stem has two forms in each case, *empiez-*, *empez-*, and *cuent-*, *cont-*; the choice is usually regarded as conditioned by the position of the stress. In an IP model, the stem is represented in the lexicon as a single form—either one of the surface alternants or a third form which does not occur on the surface but from which both alternants may be derived. Each verb has only one representation in the lexical component, and a series of rules in a separate component change the features in this underlying form to generate all the surface variants.

In contrast to the IP model, the Item and Arrangement (IA) model (which enjoyed some popularity during the 1940s and 1950s) would list all surface stem alternants, for example, *empiez-*, *empez-*, *cuent-*, *cont-*, in the lexicon and include with them a rule which states the distribution of each alternant (see Hockett, 1954). A third model, called the Word and Paradigm (WP) model, is a formalization of the intuition behind traditional grammar, which chooses the word (rather than the stem) as the unit of lexical representation and includes in the lexicon a full listing of paradigms (Matthews, 1972).

The main advantage of the IP model over the other two is that it allows the lexicon to contain a much smaller set of items, since the IA model must list all alternants and the WP model must list all words.[1] It is also sometimes argued (Kiparsky, 1968) that the IP model is superior because it allows alternations to be described in terms of very general rules that apply to many instances of the same alternation, while the IA model (in its early versions) requires a separate statement for each instance of an alternation.[2]

The IP model is the most used and the most extensively elaborated of these descriptive models, but it has also been subject to criticism. The major criticisms that have been leveled against it in its generative version are that it fails to distinguish between productive and nonproductive alternations, between morphologically and phonetically conditioned alternations, and that it allows underlying forms to differ too radically from the surface forms (Kiparsky, 1968; Vennemann, 1972, 1974; Hooper, 1976a; Hudson, 1980, etc.). Various refinements of the basic model have been proposed which introduce

[1]There seems to have been the feeling in the descriptive frameworks that lead to the IP model that a smaller number of lexical listings was preferable because it made for a simpler and more elegant description. This notion that the lexicon must be small and simple is disappearing as linguists realize the extent of lexical idiosyncracy, especially in morphological and syntactic properties. Moreover, considering the enormous storage capacity of the human brain, there is no reason to insist on a small lexicon.

[2]Of course, it is not impossible to formulate a version of the IA model that allows general statements regarding the distribution of allomorphs (Hudson, 1980).

these distinctions and constrain the relation between surface and underlying forms.

Despite the extensive investigations of the relation of the IP model to natural language data, and the many proposed constraints and refinements, there are still certain basic facts about the organization of morphological systems in the languages of the world that cannot even be described in an IP model of morphology, much less predicted or explained. These facts cannot be dismissed on the grounds that they are evidence from performance, for while some of the evidence does involve the results of psycholinguistic experimentation and child language acquisition, this evidence points in the same direction as the evidence from sources accepted even by those who wish to maintain a distinction between performance and competence, that is, evidence from synchronic distribution, historical change, and language universals.[3] In response to this accumulation of evidence against the IP model, I argue that the problem lies in IP's most fundamental tenet—that item and process are two distinct and discrete components of the description. I argue instead that the best exemplar of a rule and the best exemplar of a representation are two poles of a continuum, and that some rules have properties we associate with representations while some representations bear a resemblance to rules.[4] I propose some features of a model which can account for the rule-like nature of human language without forcing unnatural dichotomies.

1. Facts about Morphology That IP Models Neglect

It has often been pointed out that IP models have no way of representing paradigms, since each word of a paradigm is derived independently from an underlying form that may not correspond to any particular form of the paradigm. This would not be a serious problem were it not for the fact that some very strong cross-linguistic generalizations, as well as some significant tendencies in historical change and child language, can only be predicted in terms of paradigmatic relations. Here are some of the facts that an IP model misses because of its inherent structure:

First, the semantically least-marked forms of a paradigm are usually the morphologically least marked and usually have the highest text frequency. These unmarked forms are basic to the paradigm in the sense that children

[3]Theoreticians may limit the scope of their theories in whatever way they wish and thus ignore some sources of evidence which another theory may utilize. However, I see no reason to limit linguistic theories to theories of knowledge (or competence) when theories that integrate knowledge and use are within our grasp.

[4]In an IP model, a lexical representation may be thought of as a very specific rule, a rule that rewrites a semantic representation as a phonological one. This idea gets to the same basic point that I argue below, but I approach the question from the other direction.

learn them earliest and then use them to create the more marked forms. Moreover, the "analogical leveling" of diachronic change is the remaking of a more marked form using the less marked form as a base (Vennemann, 1972; Mańczak, 1980; Hooper, 1979; Bybee, 1985).

Second, there are differing degrees of relatedness among the forms of a paradigm, depending on their semantic relations. For instance, the second and third singular forms of a present tense are more closely related to one another than each is to the same person form in a past tense. These differential relations are reflected in formal differentiation—the distribution of allomorphy as well as the propensity for one form to condition the restructuring of another in diachronic change (Hooper, 1979; Bolozky, 1980; Bybee and Pardo, 1981; Bybee, 1985).

A third set of questions for which IP morphology offers no account concerns the relative productivity of morphological and morphophonemic rules. Why are some rules more productive—better able to apply to new forms—than others? How does the extant set of forms a rule applies to determine its ability to apply to new forms? To what extent are nonce formations predictable? In an IP model these questions are not answerable, since a rule is structurally the same whether it applies to three forms, thirty, or three hundred.[5]

A fourth set of issues has to do with deviations from the one-to-one correspondence between meaning and form. An IP model takes the lack of allomorphy to be the simplest case: The combinatory rules can concatenate underlying forms and no processes need apply. The complications and the major work of the "processes" come about because of allomorphy. In such a model, unless allomorphic variation is phonologically motivated, it has no motivation whatever, and it is treated as though its occurrence were random. Such a model fails to describe a whole set of predictable phenomena: that stem allomorphy tends to follow the lines of major morphological categories and cut across minor categories (see Section 5); that irregularities tend to be greater in number in unmarked members of categories (in both stem and affix) (Greenberg, 1966); that allomorphy and suppletion are more likely in frequent lexical items than in infrequent ones. None of these generalizations can be stated in current versions of IP representation, nor can such a model explain why allomorphy persists over time rather than being eliminated rapidly.

The foregoing are a few reasons for entertaining thoughts about new and different models of representing morphology; these are important facts that

[5]Structural differences in rules, e.g., whether they contain a diacritic or a morphological feature, may correspond to some extent to productivity, since phonetically conditioned, transparent rules are usually productive, while morphologically conditioned rules are often not. But this correspondence is not direct enough to be predictive. For instance, the English vowel change rules for past tense would be structurally similar in an IP treatment, yet some of them (i.e., *string-strung*) are more productive than others (e.g., *bite-bit*).

can perhaps be represented more directly in a different sort of model. In addition I argue that there is motivation for rejecting the IP model specifically because of its major premise, which has gone completely unquestioned for decades: that rules and representations are discrete and distinct elements of the grammar. It is this point to which I now turn.

2. Rules and Representations

Morphological and morphophonemic rules cover the full range of types, from extremely productive and general (e.g., the suffixation of *-ing* to form participles and gerunds in English), to semiproductive (e.g., the "rule" which produces past forms such as *stung*, which sometimes applies to new forms), to minor rules (such as that governing the voicing alternation in *wives*, *leaves*, and so on), to "rules" dealing with admitted irregularities (such as *bring*, *brought*), and finally to suppletion (exemplified by *go*, *went*). The productive and general rules are the most independent of the representations to which they apply, but as we go down the scale, more and more information about particular representations has to be built into the rule. Thus a rule that applies to a particular lexical class (such as the class of nouns that has a final fricative that voices in the plural) must contain some particular signal—a diacritic feature or a phonological feature used as a diacritic—to match it with the particular representations to which it applies and prevent it from applying to forms which do not undergo the alternation (e.g., nouns such as *chief*). In other words, part of the representation has to be built into the rule. And of course, so-called rules governing suppletion are nothing more than representations.

A characteristic of alternations that are lexically and morphologically restricted is that they usually are not extendable to new lexical items. Consider the vowel alternations in Spanish verbs illustrated in (1). This diphthong/mid vowel alternation resulted from the diphthongization of lax mid vowels in stressed syllables, but this process is no longer phonetically conditioned in Spanish (indeed, the lax vowels that produced the diphthongs no longer exist). The alternation is usually approached as a problem in verb derivation. Although a large number of verbs undergo these alternations, there is also a large number of verbs that have nonalternating mid vowels, for example, *comer* 'to eat' with 1sg *cómo*, and *aprender* 'to learn' with 1sg *apréndo*. If in a generative treatment the mid vowel is taken as the underlying vowel, then it must be marked with a diacritic, since some mid vowels do not alternate (Harris, 1969). On the other hand, the diphthong could be taken as underlying, since there are very few verbs with a nonalternating diphthong, and these are all derived verbs, for example, *amueblar* 'to furnish' from *muebles* 'furniture' and *aviejar* 'to grow old' from *viejo* 'old.' This would suggest that a rule "a diphthong becomes a mid vowel in an unstressed syllable" could be formulated, and indeed an IP treatment along these lines could be made to work. However, additional facts indicate that these alternations are highly lexically restricted, that is, dependent upon the particular items to which they apply, and not extendable to new items.

Kernan and Blount (1966) used a nonce-probe task to test adult speakers of Mexican Spanish on the productivity of these alternations. They presented their subjects with a nonce form such as *suécha* in a context where it was clear that it was a 3sg of the present indicative, and then they asked the subjects to use it in a preterite context, where the stem is unstressed, so that the expected response is *sochó*, based on the rule stated above and the recurrent alternation pattern. However, their subjects uniformly answered *suechó*, producing a form which not only ignores the rule, but which also has an unstressed diphthong, which occurs only rarely in Spanish.

In Bybee and Pardo (1981) we tried even harder to get Spanish speakers to apply the rule. In our nonce-probe task, we presented nonce verbs in both diphthong and mid vowel alternates (e.g., *sochár* and *suécha*), but despite this very clear indication that the alternation exists for the nonce verb, in 25% of the responses the subjects produced a diphthong in an unstressed syllable. My interpretation of these results is that the rule governing these alternations is not independent of the existing lexical forms to which it applies. In some sense it is not a rule at all, but more a part of the representation of certain verbs.

Even rules that seem more independent, that is, that apply to new or nonce forms, have to emerge in acquisition from representations. In order to acquire rules, the child must extract them from the comparison of sets of related forms, which are entered in the mental lexicon. Studies by MacWhinney (1978) and Peters (1983) indicate that rule-like generalizations gradually emerge from stored rote forms, which are initially processed and stored as unanalyzed wholes.

Other evidence suggests that even though certain generalizations, especially over subregularities such as the English strong verbs, are recognized at some level by speakers, the generation of strong verb past tense forms is by lexical access rather than by feature-changing rule. We reported in Bybee and Slobin (1982) that experimentally induced errors involving vowel changes for past tense result in almost all cases not in the production of nonce forms such as the past tense of *heap* as **hept*, but rather in the replacement of one preexisting word for another, usually within the same semantic domain. Thus in 91% of the cases the wrong vowel-change response was a real English word; in 80% of the cases it was not only a real word, but also a verb; in 75% of the cases it was a past tense form of a semantically related verb. For instance, the past tense of *raise* was given as *rose*, the past of *seat* as *sat*, of *search* as *sought* (Bybee and Slobin, 1982).[6] These errors can only be explained by postulating that the production of strong verb past tense forms is a matter of accessing the lexicon rather than applying a rule to a base form to change the vowel.

These considerations lead to the conception of a model in which morphological rules and lexical representations are not separate from one another. Rather, morphological and morphophonemic rules are patterns that emerge from the intrinsic organization of the lexicon. Patterns that range over large numbers of lexical items are highly reinforced or strengthened and apply more readily to new items, while patterns that are found in a smaller number of items are correspondingly weaker and less apt to be productive. Thus the difference between major productive rules, minor rules, and suppletion is just a matter of degree, not a matter of qualitative difference.

The model I am proposing does not have a lexicon and a morphological component as separate compartments of the grammar. Rather the model has only a lexicon. The morphological facts of natural language are described in terms of independently necessary mechanisms of lexical storage: the ability to form networks among stored elements of knowledge and the ability to register the frequency of individual items and patterns. I discuss these mechanisms and their interactions under the headings of LEXICAL CONNECTION and LEXICAL STRENGTH.

3. Lexical Connections

It is uncontroversial that stored knowledge is organized, and that the lexicon is storage governed by multiple and diverse organizational patterns. Chief among these are semantic parameters by which morphemes are associated. Morphemes are connected via the semantic field they belong in (such as verbs for cooking, *boil*, *fry*, *roast*, *bake*), by the scripts they participate in (such as the restaurant script), by relations such as synonymy, antonymy, hyponymy (the relation between *rose* and *flower*), and many others.

Similarly, phonological connections exist among stored forms. The evidence for this is that we have a certain amount of access to the lexicon via the phonological shape of words: We can list words with particular initial segments; we can list words that rhyme; we can list words with a particular stress or tone pattern and a certain number of syllables. Some sort of phonological mapping function is required for speech perception, and speech

[6]Out of 46 errors, only three were nonwords: *hept* (as the past of *heap*), *snoze* (as the past of *snooze*), and *glew* (as the past of *glow*).

errors and punning behavior show that associations based on phonological similarity are accessible.

In this model, morphologically complex items are stored in the lexicon, and I refer to them as "words" although it is conceivable that some may be larger than traditional words, and some may be smaller, and there may even be typological differences among languages regarding the size of the lexical unit. Each lexical word is a pairing of a set of semantic features with a set of phonological features. Relations among words are set up according to shared features. For instance, the close relation between *dog* and *puppy*, shown in (2), can be derived from the fact that the semantic representations of these items have many of the same features. In fact, the semantic representation of the one can partially map onto the other. We may represent this as proximity in the lexicon, or, with the notation I have adopted here, as connecting nodes from one feature to the other. In the diagrams where phonological features are connected, for simplicity I have used solid lines to indicate that all features of two segments are identical, and broken lines to indicate that only some features are identical.

(2)

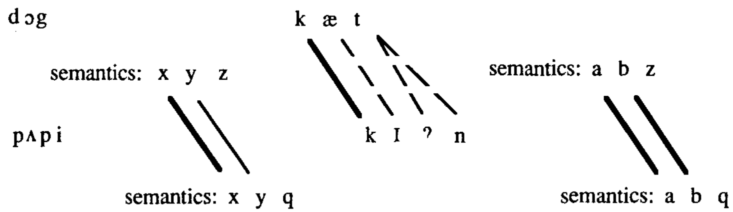

The pair *cat* and *kitten* have a similar semantic relation, since they share all features common to the species, but the latter term is restricted to the young of the species. A number of connections relate the semantic features of *cat* and *kitten*, and in addition, some connections exist between their phonological representations, since they share an initial consonant /k/, a front vowel, and an alveolar consonant. Now *kitten* and *puppy* are also semantically related, since they both designate the young of similar species, and they bear a phonological relation to one another also, since they both consist of a stressed obstruent-vowel-obstruent syllable, followed by an unstressed syllable. We could go on and show how the number of relations quickly multiplies and defies visual representation, except perhaps as something that resembles a bowl of spaghetti.

However, things are not entirely spaghetti-like. Sets of connections that parallel one another by running between the same two items may accumulate to form connections of varying strength, depending on the number of features connected, and in some cases on the content of the features.

Sets of connections such as those between *cat* and *kitten* are the basis of morphological analysis and morphological relations, for morphological relations are semantic and phonological connections that run in parallel. Consider the word *cats*. It forms both a semantic and a phonological connection with the singular form *cat* as shown in (3).

(3)

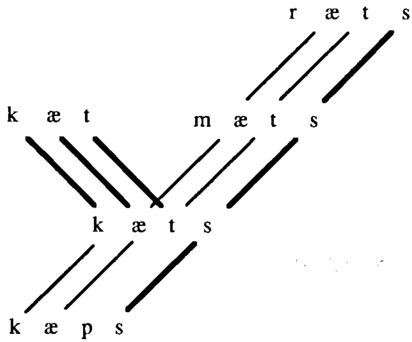

It also forms connections with other plurals, such as *mats*, *nats*, *laps*, *naps*, *tips*, *mits*, and so on, on the basis of the shared semantic feature plural and the identity of the final fricative. Thus where semantic and phonological connections coincide exactly, morphological identity can be established. This is drawn in (3), where parallel semantic and phonological connections are indicated by darkened lines, and connections that are only phonological are indicated with lighter lines.

In explicating this model, I first show how phenomena associated with the segmentation of words into morphemes are accounted for and argue that this model resolves a number of traditional problems in this area. I then introduce the notion of LEXICAL STRENGTH and the range of phenomena associated with this concept. Finally I return to the question of morphological rules and show how what appears to be rule-like behavior can be derived strictly from relations among representations.

4. Segmentation

Consider now the actual process by which we, both as speakers and linguists, discover that a word consists of more than one morpheme: We find a relation of phonological and semantic identity or similarity between some subpart of the word and a subpart of another word. This process is represented quite directly in this model: As shown in (3), *cats* forms connections with *cat*, as well as other plurals. The learning process involved here is very simple: When a new morphologically complex word is learned, it forms connections with existing lexical material on the basis of its meaning and phonological shape. The word is not physically dismembered, but its parts are nonetheless identified.

There are two major advantages to this method of representing segmentation. The first, which I have already mentioned, is that the mechanisms used, that is, the ability to form phonological and semantic associations among lexical elements, are necessary in any case, and probably differ little from the mechanisms used in storing nonlinguistic information. The second is that this model allows us to describe morphological relations in a gradient fashion, which accommodates the description of phenomena that must be ignored in an IP model.

The traditional definition of a morpheme as the smallest meaningful unit makes constancy of meaning the major criterion for identification of morphemes. In the ideal case, all the phonetic material in an utterance can be assigned to one meaningful unit or another. In such a case the semantic and phonological connections run parallel. But not all linguistic material is organized in this fashion. Deviations have been discussed for decades, for instance, the famous *cranberry* morph which results when *berry* is segmented from *cran*. Of course, this particular problem was solved by the makers of cranapple juice, but the general dilemma remains. Consider the days of the week: *Monday*, *Tuesday*, *Wednesday*, and so on all have the morpheme *day* in them, but what do *Mon*, *Tues*, *Wednes*, mean? With lexical connections we can associate the *day* sequence in these words with each other and the word *day* without requiring that the remainder of the word be meaningful. Rather the first syllable remains as a part of the whole word, but it has no connections to other items.[7]

Similarly, the formatives of generative grammar, such as the *-ceive* and *-cep-* of *receive*, *reception*, *deceive*, *deception* which Aronoff (1976, pp. 10–15) argues have formal behavior that identifies them as significant units, even if they are not meaningful, can be given some status in this model. *Deceive* is related to *deception* and *receive* to *reception* by both semantic and phonological connections. *Deceive* and *receive* are connected phonologically and by the fact that they are both verbs. Similarly, *deception* and *reception* are related phonologically and by the fact that they are both nominalizations. The semantic connections are weak between *deceive* and *receive* and *deception* and *reception*, but the phonological connections are strong.

Closely related to the *ceive* phenomena are phonaesthemes and other submorphemic units. In English there are sets of words with the same initial cluster that have a general semantic feature in common, such as having to do with the nose, in the case of *sneeze*, *snore*, *snort*, *sniff*, *sniffle*, *snivel*, *snoot*, *snout*, *snot*, and so on; or movement through air, as in the case of *flit*, *fly*, *flutter*, *fling*. These words may be connected phonologically in their initial clusters

[7]The sequence *Wednes* would seem to be perfect example of a uniquely occurring morph. However, I recently discovered that there is a town in England named *Wednesbury* and another called *Wednesfield*!

and semantically through one particular property. The relative weakness of such relations is due to the general nature of the semantic connection, as well as to the fact that the part of the word not involved in the connection is not meaningful, nor does it enter into any connections of its own. Phonaesthemes cannot be awarded any status in an IP model, nor can *cranberry* morphs. Formatives, such as *ceive*, can only be given status by relaxing the definition of the basic unit so that it need not be meaningful. The current model allows the description of the full range of phenomena and also allows for differential status according to whether both semantic and phonological connections are made, and how strong these connections are.

5. Degree of Relatedness

This model also allows a way of grading relations among words, so that we may say that certain pairs of words are more closely related than others. DEGREE OF RELATEDNESS is determined by the number of phonological and semantic connections. Why do we need a notion "degree of relatedness"? A series of experiments by Derwing and Baker (1977) shows that speakers can consistently rank pairs of words for semantic and phonological relatedness. In another experimental paradigm (Stanners *et al.*, 1979), the recognition of particular words can be speeded up more by priming with a closely related word than with a more distantly related one. For instance, *walked* increases the speed of recognition of *walk* more than *deception* increases the speed of recognition of *deceive.*

These experimental results point to the same conclusion as does the range of diachronic and cross-linguistic facts which I present in detail in Bybee (1985). There I argue further that it is not just the number of connections between forms that determines the degree of relatedness, but that the degree of relatedness can vary according to the semantic features involved.

If we consider first just words related through a shared stem, and in particular, inflectional paradigms, we can identify varying degrees of semantic relatedness depending on the meaning of the affix category. Some inflectional categories affect the meaning of the whole word more than others and thus produce forms that are less closely related semantically to the base form. As an example, consider the verbal category of aspect compared to person agreement with the subject. Aspect affects the meaning of the verb more, since it modifies the "internal temporal constituency" of the event or state described by the verb (Comrie, 1976). In fact, a change in aspect can produce quite a difference in the event described by the verb. Consider the example of the Spanish preterite/imperfect distinction, which is an aspectual distinction. The verb meaning 'sleep' in the preterite, *durmió* 's/he slept,' describes a completed event, which includes falling asleep and sleeping. The imperfect *dormía* translates approximately as 's/he was sleeping,' implying a state someone was in when something else occurred. For some verbs this aspectual distinction is large enough to produce distinctions that are expressed by separate verbs in another language. For instance, the preterite of the Spanish verb *saber* 'to know' translates into English as 'found out.'

Now compare aspect to person agreement. The function of agreement is to index the participants in the state or event described by the verb and has nothing to do with the inherent meaning of the verb. Two verb forms that differ only by person are much the same semantically if their tense, aspect, and mood are the same. The semantic differences produced by person markers are not the type that would be expressed by entirely different verbs. Thus such forms are more closely related than forms that differ in aspect.

Degree of relatedness is diagrammed by morphophonemic alternations: The more closely related two forms are semantically, the more likely they are to be similar morphophonemically. This means, for example, that stem changes in verbs are more likely to distinguish aspects or tenses than to distinguish person forms across aspects or tenses. Thus in Spanish there is a set of irregular verbs that have stem changes for the preterite aspect. The verb *saber* 'to know' has the 3sg preterite form *supo*, as shown in (4); similarly *tener* 'to have' has the 3sg preterite form *tuvo*; *querer* 'to want' has the 3sg preterite form *quiso*, and so on.

(4)

	Imperfect 'knew'				Preterite 'found out'		
1sg	*sabía*	1pl	*sabíamos*	1sg	*supe*	1pl	*supimos*
2sg	*sabías*	2pl	*sabíais*	2sg	*supiste*	2pl	*supisteis*
3sg	*sabía*	3pl	*sabían*	3sg	*supo*	3pl	*superion*

These irregular stems occur throughout the preterite person forms, and thus set off the preterite from the present and imperfect. On the other hand, there are no stem changes in Spanish that set off, for example, all first person forms in all aspects and tenses from all other person forms. Indeed, the hypothesis is that such a situation would be very rare. This hypothesis has been tested on a sample of 50 unrelated languages and was not disconfirmed. On the contrary, it was found that stem alternations of consonants or vowels are extremely frequent where aspectual distinctions are concerned and extremely rare where person distinctions are concerned (Hooper, 1979).

There are two mechanisms behind this cross-linguistic pattern, one of which concerns the order of affixes. First, since aspect (and other categories that affect the meaning of the verb more) tend to occur closer to the verb stem than agreement affixes, they are more likely to produce phonological alternations in the stem. Second, when morphophonemic alternations are eliminated in analogical leveling, this takes place preferentially among more closely related forms, often leaving alternations intact in less closely related forms. Thus a leveling is more likely to occur among person-number forms of the same aspect or tense, than among forms in different aspects or tenses.

Further evidence for a hierarchical ranking among grammatical categories is the distribution of forms in suppletive or split paradigms. When inflectional paradigms split and realign, forming suppletive paradigms, the splits occur more often among forms that are less closely related semantically, than among forms that are more closely related. Rudes (1980) studied suppletive verbal paradigms in a large number of languages and found that splits occur in general along aspect or tense lines, as with *go* and *went*, and along person agreement lines only in the present tense, the most frequent tense. (The relation between frequency and suppletion is discussed in the following two sections.)

To summarize, the degree of relatedness among words is primarily determined by the number and type of semantic features shared. The degree of phonological similarity often parallels the degree of semantic relatedness.

6. Lexical Strength

The other theoretical construct that I propose for the lexicon that distinguishes this model from an IP lexicon, or even a WP lexicon, allows for the gradient representation of lexical strength. Previously the lexicon has been conceived of as the mental counterpart of a dictionary, a list of forms set down once and for all. I propose a more dynamic representation in which not all forms have the same status, but rather in which forms are affected by use or disuse. Frequently used forms gain in lexical strength and forms that are not used lose lexical strength. Lexical strength, then, is an index of word frequency, and the main reason for proposing the introduction of this notion into a model of the lexicon is to account for the psycholinguistic, historical, and cross-linguistic effect of frequency on morphology.[8]

One of the strongest and best-known effects on lexical access is the word-frequency effect. In various sorts of tasks, words that are more frequent are more quickly accessed or recognized. This alone is enough reason to build an index of frequency into a model of the lexicon. But there is also plenty of evidence outside of experimental contexts that frequency is an important dimension in the lexicon and in morphology. Most of this evidence is well known, but it has been ignored in models that emphasized rules and paid less attention to representations.

[8]For the moment, it is sufficient to define lexical strength as based solely on token frequency, but I would like to leave open the possibility that other factors may be involved. If lexical strength is based only on frequency, then it is not equivalent to the notion of "autonomy" as used in Zager (1982), Bybee and Brewer (1980), and Bybee (1985).

First, we need a notion of lexical strength to account for the maintenance of irregularity and suppletion in high-frequency forms. Conversely, the proposal that infrequently used forms fade accounts for the tendency to regularize infrequent irregular forms, for an irregular form that is not sufficiently reinforced will be replaced by a regular formation. The correlation of irregularity with high frequency can be documented in almost any language, but the historical mechanism behind the correlation is also easily demonstrated. For instance, as I pointed out in Hooper (1976b), the average frequency of a past tense Old English strong verb that has remained strong is more than 20 times greater than the average frequency of a strong verb that has regularized.

Second, lexical strength accounts for the tendency for lexical and inflectional splits to occur more often among high-frequency words. Lexical split describes the diachronic process by which previously related words lose their morphological relatedness as the originally derived word takes on a nonpredictable semantic representation. In a study of words with the prefix *pre-* in English, Pagliuca (1976) found a strong correlation between the frequency of the prefixed word and loss of a transparent semantic and phonological relation to its nonprefixed base. This cannot be accounted for in a model in which all words with the same affix are derived in exactly the same way. It requires a model in which a particular word, despite its morphological complexity, can be autonomous and develop semantic and phonological peculiarities. The same argument applies to inflectional splits, which, as is well known, occur only among the most frequent lexical items.

7. The Interaction of Lexical Strength and Lexical Connection

Lexical strength interacts with lexical connection in some very interesting ways. The first interaction we consider concerns what I have called elsewhere the BASIC/DERIVED RELATION. There is considerable historical evidence that speakers construct unidirectional relationships (or lexical connections) between morphologically related stem forms in such a way that the semantically unmarked or basic form is also morphophonologically basic (Watkins, 1962; Vennemann, 1972; Bybee and Brewer, 1980; Bybee, 1985). To give only the simplest sort of example, consider the potential regularization of a verb such as *creep* or *weep*. In these verbs, the vowels [iy] and [ɛ] alternate. In order for the verb to regularize, the alternation must be eliminated, that is, the same vowel must occur in base and past tense forms. As is well known, there is no question about which vowel will prevail: it is always the [iy] of the base form, not the [ɛ] of the past form. The reason for this is that the regularization takes place precisely because the irregular past form is either not represented in the lexicon or is extremely weak, while the base form is stronger.

An inability to access the irregular past tense form leads to the formation of a regular past form.

Cross-linguistic evidence for an asymmetrical relation among forms can be found in the prevalence of zero-marking in semantically basic forms of a paradigm, and non-zero marking in semantically complex forms. Such marking gives evidence that one form, the stronger form, is autonomous—not analyzable in terms of other forms—while the remaining forms of a paradigm may be stored and analyzed in relation to that basic form.

I describe this asymmetrical relation in the following way: Words are acquired and stored much as other types of knowledge, by integration with knowledge already stored. A word that is morphologically simple and highly frequent is likely to be acquired more or less independently of other words and may also be acquired earlier than other related words. The more complex related words of lower frequency are learned and stored in terms of the simpler, more basic words that are already present in the lexicon. In (5) strength is indicated by a boldface representation.

(5)

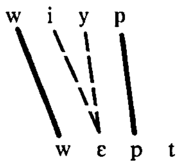

The continued frequency imbalance between the two forms will maintain the dependent relation of the more complex form on the simpler one.

It is not clear whether token frequency or morphological basicness is the main factor in establishing the directionality of such relationships. Evidence exists for both positions. In word-recognition experiments on Serbo-Croatian nouns, Lukatela *et al.* (1980) found that the least-marked form, the nominative singular, has the shortest response time even in paradigms where it does not have the highest token frequency. They argue for a satellite-entry model in which the basic form of a paradigm is the nucleus entry about which other forms cluster, which resembles in some respects the proposal I am making here.

On the other hand, Tiersma's (1982) paper on local markedness shows that frequency is an important factor in determining what is conceived of as the basic form. Tiersma discusses a number of interesting examples in which the semantically marked form serves as the basis of morphophonological regularization. For example, vowel alternations between singular/plural pairs in Frisian usually regularize with the use of the vowel of the singular for both forms, but a small set of nouns shows the opposite directionality—the vowel of the plural comes to be used in the singular. However, as Tiersma argues, these are all nouns in which the plural is more frequent than the singular, nouns that refer to objects that ordinarily appear in pairs or groups, for example, *arm*, *goose*, *horn* (of an animal), *stocking*, *tooth*, *splinter*, *thorn*, *tear.* For these nouns the plural is stronger than the singular. (See also Bybee, 1985, pp. 74–77.)

Lexical strength determines the directionality of morphological relations in the sense that the weaker words are learned and stored in terms of related stronger words. The more frequent words, even if they could be analyzed and stored in terms of other words, may be strong enough to be stored separately and may thus serve as the basis for innovations. A notion of lexical strength, then, is able to account for the basic/derived relation as well as local markedness. Moreover, lexical strength allows us to explain why irregularity and suppletion are characteristic of more frequent items. It allows us to explain how suppletion develops: if a form, even a morphologically complex one like *went* when it was the past tense of *wend*, can have its own representation, it can grow in frequency independently of *wend*, and gradually undergo a generalization of meaning that allies it semantically with *go.*

8. Affixes and Morphological "Rules"

In the diagrams presented so far, it can be seen that affixes are represented lexically, attached to their hosts, and that affixes form phonological and semantic connections with other instances of the same affix. What are represented as affixation rules in other models are in this model patterns of connections, such as that shown in (6) for the present participle suffix of English.

(6)

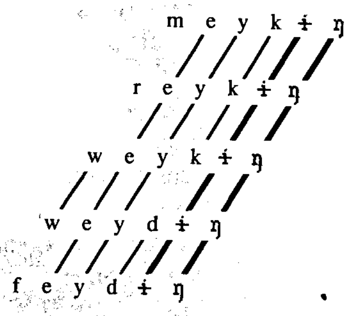

[. . . verb—iŋ] present participle

The same type of representation may be used for patterns that consist of internal changes in the stem rather than affixes. Diagram (7) shows a fragment of the pattern for the semiproductive strong verb class whose prototype is *strung.* In Bybee and Moder (1983) we showed that this class is fairly productive, given certain phonological features of the stem, including both the final and initial consonants.

(7)

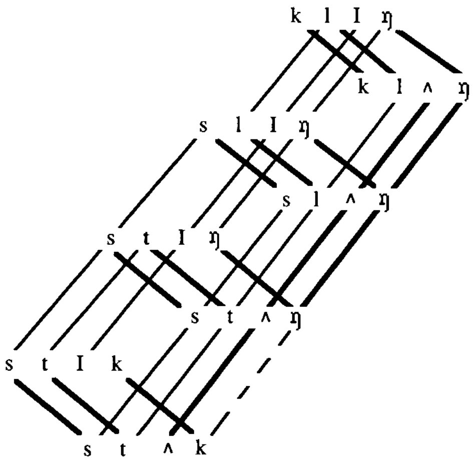

$$\left[\text{C (C) (C) ʌ} \begin{Bmatrix}\text{velar}\\ \text{nasal}\end{Bmatrix}\right]_{\text{past}}$$

I am suggesting, then, that what are usually thought of as morphological rules do not have a representation that is independent of the lexical items to which they are applicable. Rather, rules are highly reinforced representational patterns or schemas (Bybee and Slobin, 1982; Bybee and Moder, 1983). A schema may be thought of as an abstraction from existing lexical forms which share one or more semantic properties. The schema contains the features most strongly represented (i.e., represented most often) in existing forms in their positions of occurrence. Since the schema resembles a prototype, matching of form to schema works on a family resemblance basis rather than a categorical basis. Schemas may be complex, with some features being more strongly represented than others. Thus a schema may range over several allomorphs of the same morpheme, if the allomorphs have features in common. For example, the English plural suffix allomorphs may all be represented together, with the features that [s] and [z] share as the most strongly represented, the feature of voicing less strong and the features of the reduced vowel [ɨ] weaker still. While linguists today tend to think of schemas as being abstracted from lexical forms and stored in a separate component, I suggest that the evidence points toward schemas as tied ultimately to the forms from which they arise.

The tie to the lexicon is clearer for some schemas than for others. In particular, the schema that applies to fewer forms shows more evidence of being tied to the lexicon than the schema that applies to a large number of forms. The general, widely applicable morphological schemas (what are usually called "rules"), such as the regular suffixation for English past tense, appear to be free of the lexicon in the sense that they apply readily to new forms. In fact, such cases would appear to be qualitatively different from lexically restricted patterns because in acquisition they show real independence from particular lexical items. The usual sequence of acquisition for the English past tense is assumed to provide the strongest evidence that morphological rules are independent of the lexicon. This sequence of developments is as follows: In the first stage in which children produce past tense forms, they often produce the high-frequency, irregular forms, such as *went*, *came*, and *took*, correctly, as well as regulars, such as *looked*. In the second stage, however, they tend to regularize all past forms, producing *goed*, *comed*, and *taked*. This stage gives evidence for the formation of a rule, since forms are produced that are not present in the input. A third stage involves the correct relearning of the irregular form and the acquisition of the adult system.

Rumelhart and McClelland (1986) demonstrate that our assumption that this sequence of events can only be accounted for by positing the formation of explicit rules is erroneous. Rumelhart and McClelland used a parallel distributed processing model to simulate the acquisition of the English past tense. Their model and simulation did not involve the full problem of lexical storage, but only the problem with which we are concerned at the moment—the formation of rule-like patterns based on phonological features.

The main part of the model that interests us here is called the PATTERN ASSOCIATOR. The pattern associator contains two pools of units, which are phonological units of a type that I do not explain here. One pool of units is for the input and the other for the output. In this case, the input is the base form of a verb and the output is a past tense form. The pattern associator contains a modifiable connection linking each input unit to each output unit. Learning in this model involves modifying the weights of these connections. The model is given the base form of a verb and the target, which is the past form. From the base form, the model computes the past form it would generate, given the current connection strengths. It then compares its answer with the target and makes adjustments of the weights of the connections accordingly.

Rumelhart and McClelland simulated the learning of English past tense forms by supplying the model with verbs in a way that at least grossly duplicates the way verbs come into the child's system. For the learning trials they used a total of 420 English verbs, of which 20% were irregular. To begin with, they used the 10 most frequent verbs, presenting the model with these verbs for 10 trials each. This highest-frequency group had 8 irregulars and 2 regulars. These 10 learning trials were enough to produce good performance on these verbs. Then the remaining 410 verbs, mostly regular, were added to the first 10 and the trials continued. At this point, the performance on the original 8 irregular verbs reversed dramatically. For a short period, these verbs were regularized more often than produced correctly. Then gradually, correct responses for high frequency irregulars increased until they climbed over 90%. The simulation produced the same learning pattern that is observable in children, but without formulating an explicit rule. Rather the model adjusts the weights of connections, increasing or decreasing the probability of a certain output. Because the number of regular verbs so far outweighs the number of irregular verbs in English, the probability that the regular pattern will be followed is extremely high, but we need not assume that it is categorical.

Not only does the pattern associator learn the regular pattern and the irregular verbs it is given, it also learns the subregularities among the irregular verbs, so that when it is tested on verbs it has not encountered before, it behaves very much like an adult English speaker in a nonce-probe task. The majority of forms it produces have the regular suffix, including some that should not, such as *thrusted* and *sticked.* However, it also produces some correct irregulars, such as *bid* for the past tense of *bid*, *crept* for *creep*, and *clung* for *cling.* It also produces some incorrect, but highly probable vowel changes, such as the past tense of *slip* as *slept*, and some doubly marked pasts, such as *dripted* for *dripped.*

The connections made by the Rumelhart and McClelland pattern associator do not precisely parallel any connections in the model that I am proposing. The importance of their simulation is to show that, given fairly realistic input in terms of the frequency of regular versus irregular verbs, a model which only registers probabilities and never formulates an explicit rule produces behavior very similar to that of a human language learner. Let us see now briefly where these patterns show up in our model.

To begin with, recall that in our model connections are formed only between identical features, so that a pattern for the formation of the past *stung* will not be based on a connection between the [ɪ] of *sting* and the [ʌ] of *stung.* That this is correct can be seen from the fact that the members of the *stung* class of verbs do not all have the same vowel in the base. This is particularly true of the newer members, such as *strike*, *sneak*, and *drag.* Rather the basis for the formation of the pattern is an attempt to find a phonological schema with which to associate the semantics of past tense. Since there are many different shapes for past tense verbs in English, the pattern of connections is rather complex. It must be built up much as the Rumelhart and McClelland model builds up patterns: by registering probabilities for each feature paired with past tense according to how many times that feature appears in the representation of a past tense verb. As in the Rumelhart and McClelland model, the features that go into building up the pattern may be features of the stem as well as a suffix. Unlike that model, however, it is only the features of the past form that go into building the schema.

In Bybee and Slobin (1982) and Bybee and Moder (1983), we argued that the subregularities in English past forms could be accounted for by schemas, which specify the shape of a past verb form. The features of such a schema, we argued, are not categorical, but probablistic, like a Roschian natural category. The work by Rumelhart and McClelland shows that even the regular suffixation pattern should be treated in this way. The only difference between it and the subregularities is that the probabilities for suffixation are extremely high.

Thus the strongest features associated with past are alveolar and stop (in final position). The verbs that take the [ɨd] allomorph contribute the sequence "alveolar stop (the last consonant of the stem), reduced vowel, voiced alveolar stop" to the schema. Since the suffixation mode of forming the past tense involves so many verbs of such diverse phonological shapes, the rest of the schema is mostly unspecified. On the other hand, the schema for the *strung* class will have probability rankings for the initial consonants, the final consonant, and the vowel, since, as we showed in Bybee and Moder (1983), all of these features have at least some effect on the selection of this schema for new verbs.

Degree of productivity, then, is determined automatically in this system by several factors: First, how open the phonological definition of the schema is. The English past schema involving a final alveolar stop is associated with verbs of all phonological shapes, so it is open to extension to new verbs of any shape. A more strictly defined schema has correspondingly fewer chances to extend to new verbs. Second, the number of items participating in the schema. A larger number of distinct verbs participating in the same pattern will serve to strengthen it. Note that this is type frequency rather than token frequency. A verb of high token frequency will not serve to strengthen a schema; in fact, it appears that very high-frequency verbs have very little effect on productivity, since, as I mentioned in the preceding section, such forms seem to be processed without forming connections with other items. The third factor determining productivity may be termed CUE VALIDITY following Rosch. A more productive class has high cue validity, that is, most of the forms that fit the schema actually belong to the class. Thus the *strung* class has high cue validity since almost all of the verbs that could belong to the class do. On the other hand, the class that includes *tear, tore*; *bear, bore*; *wear, wore* has low cue validity, since so many regular verbs such as *snared, glared, aired, stared* exist.[9]

[9]Köpcke (1986), in a discussion of schemas for German plurals, suggests that the "salience" of a marker of a category may contribute to its productivity. SALIENCE refers to phonological size and shape: An affix consisting of two segments is more salient than a vowel change, such as umlaut, which affects only one or two features.

9. Conclusion

Of course, this chapter has not dealt with all the problems of lexical storage, of creating meaning-form correspondences, or of creating categories. I have proposed that lexical storage and organization involves the application of several principles, all of which appear to be well motivated and independently necessary: first, the ability to store strings of linguistic material in phonological and semantic representation; second, the ability to form, among stretches of this material, connections of a semantic and phonological nature; third, the accumulation of lexical strength due to token frequency; and fourth, the ability to organize sensory stimuli into categories. I have argued that these properties of lexical storage and organization allow us to account for a number of morphological phenomena that other models cannot account for.

First, this model allows us to conceptualize the internal structure of a word as a set of relations with other words, rather than as a string of discrete meaningful sequences, that is, morphemes. The problem of submorphemic units—*cranberry* morphs, formatives, and phonaesthemes—disappears, since the model allows the identification of a part of a word as a recurring meaningful unit without the necessity of assigning meaning to the remainder, and it allows the identification of phonological relations even in the absence of a clear semantic relation. Further, it postulates that morphological analysis proceeds directly from the discovery of relations among words, which is precisely how speakers and linguists accomplish morphological analysis.

Second, this model provides for varying degrees of relatedness among words as the consequence of both phonological and semantic features. I have proposed that some semantic connections are stronger than others because of the nature and number of shared semantic features. This proposal in turn predicts that analogical reformation is more likely among certain pairs of related forms than among others, depending upon their degree of semantic relatedness. As a result, leaving aside very high frequency forms, the degree of phonological connection reflects the degree of semantic connection.

Third, the notion of lexical strength, an index of word frequency, has been built into the model to account for the effect of word frequency on lexical access, and additionally to account for two major facts about morphological systems in the languages of the world: (a) that irregularity and suppletion are more common among high-frequency words and paradigms (because high-frequency items are less prone to analogical reformation); (b) that the more-frequent form of a pair of closely related forms is the one that serves as the basis for the analogical reformation of the other.

Finally, recurring morphological patterns emerge as accumulations of similar or identical sets of connections and are described as SCHEMAS. The notion of schema accommodates a range of pattern types from the most lexically specific and idiosyncratic to the most general and productive. The shape of a schema and the likehood that it will influence the formation of new words is directly determined by the number and types of items over which it ranges and does not have to be indicated separately.[10]

The most important difference between this theory and previous theories is that in this theory the generalizations that in other theories are called "rules" are here part of the representations. They arise out of the organization of phonetic and semantic substance of the language, and they have no existence independent of the substance that brings them to life.

Acknowledgments

I acknowledge fruitful discussions on the model I present here with Uli Frauenfelder, Klaus-Michael Köpcke, Carol Lynn Moder, William Pagliuca, Revere Perkins, Sandra Thompson, and Lolly Tyler.

References

Aronoff, M. (1976). "Word Formation in Generative Grammar." MIT Press, Cambridge, MA.

Bolozky, S. (1980). Paradigm coherence: Evidence from Modern Hebrew. *Afroasiatic Linguistics* **7**, 2–24.

Bybee, J.L. (1985). "Morphology: A Study of the Relation between Meaning and Form" Benjamins, Amsterdam.

Bybee, J.L., and Brewer, M.A. (1980). Explanation in morphophonemics: Changes in Provençal and Spanish preterite forms. *Lingua* **52**, 271–312.

Bybee, J.L., and Moder, C.L. (1983). Morphological classes as natural categories. *Language* **59**, 251–270.

Bybee, J.L., and Pardo, E. (1981). Morphological and lexical conditioning of rules: Experimental evidence from Spanish. *Linguistics* **19**, 937–968.

Bybee, J.L., and Slobin, D.I. (1982). Rules and schemas in the development and use of the English past tense. *Language* **58**, 265–289.

Comrie, B. (1976). "Aspect." Cambridge Univ. Press, Cambridge.

Derwing, B., and Baker, W.J. (1977). The psychological basis for morphological rules. *In* "Language Learning and Thought." (J. Macnamara, ed.). Academic Press, New York.

Greenberg, J. (1966). "Language Universals." Mouton, The Hague.

Harris, J. (1969). "Spanish Phonology." MIT Press, Cambridge, MA.

Hockett, C. (1954). Two models of grammatical description. *Word* **10**, 210–231.

Hooper, J.B. (1976a). "An Introduction to Natural Generative Phonology." Academic Press, New York.

Hooper, J.B. (1976b). Word frequency in lexical diffusion and the source of morphophonological

[10]In models that have separate rules components, productive and unproductive rules can be assigned to different components. Such assignment, however, is completely ad hoc and allows no way of associating the relevant variable—how many words a pattern occurs in—with the degree of productivity.

change. *In* "Current Progress in Historical Linguistics" (W. Christie, ed.). North Holland, Amsterdam.

Hooper, J.B. (1979). Child morphology and morphophonemic change. *Linguistics* **17**, 21–50; also in "Historical Morphology" (J. Fisiak, ed.). Mouton, The Hague.

Hudson, G. (1980). Automatic alternations in non-transformational phonology. *Language* **56**, 94–125.

Kernan, K.T., and Blount, B.G. (1966). The acquisition of Spanish grammar by Mexican children. *Anthropological linguistics* **8**, 1-14.

Kiparsky, P. (1968, 1973). How abstract is phonology? *In* "Three Dimensions of Linguistic Theory" (O. Fujimura, ed.), pp. 5–56. TEC, Tokyo.

Köpcke, K. (1988). Schemas in German plural formation. *Lingua* **68**, in press.

Lukatela, G. *et al.* (1980). Representation of inflected nouns in the internal lexicon. *Memory and Cognition* **8**, 415–423.

MacWhinney, B. (1978). "The Acquisition of Morphophonology." Child Development Publication, Univ. of Chicago Press, Chicago.

Mańczak, W. (1980). Laws of analogy. *In* "Historical Morphology" (J. Fisiak, ed.). Mouton, The Hague.

Matthews, P. (1972). "Inflectional Morphology: A Theoretical Study Based on Aspects of Latin Verb Conjugation." Cambridge Univ. Press, Cambridge.

Pagliuca, W. (1976). *PRE-fixing*. SUNY at Buffalo Master's Project.

Peters, A.M. (1983). "The Units of Language Acquisition." Cambridge Univ. Press, Cambridge.

Rudes, B.A. (1980). On the nature of verbal suppletion. *Linguistics* **18**, 655–676.

Rumelhart, D.E., and McClelland, J.L. (1986). Learning the past tenses of English verbs: Implicit rules or parallel distributed processing? *In* "Mechanisms of Language Acquisition" (B. MacWhinney, ed.). Erlbaum, Hillsdale, NJ.

Stanners, R., Neiser, J., Hernon, W., and Hall, R. (1979). Memory representation for morphologically related words. *Journal of Verbal Learning and Verbal Behavior* **18**, 399–412.

Tiersma, P.M. (1982). Local and general markedness. *Language* **58**, 832–849.

Vennemann, T. (1972). Rule inversion. *Lingua* **29**, 209–242.

Vennemann, T. (1974). Words and syllables in natural generative grammar. "Natural Phonology Parasession," pp. 346–374. Chicago Linguistic Society, Chicago.

Watkins, C. (1962). "Indo-European Origins of the Celtic Verb." Dublin Institute for Advanced Studies, Dublin.

Zager, D. (1982). *A real time process model of morphological change*. SUNY at Buffalo dissertation.