Chapter 11

Variation and Mental Representation

GREGORY R. GUY
New York University

KNOWING HOW TO SPEAK a language includes at a very basic level knowing the words and how to pronounce them—that is to say, knowing the lexicon and the phonology. Linguists have conventionally conceived of the lexicon as a set of mental representations of words that represent the phonological content of the word as a string of phonemes. Complex words that contain multiple morphemes can be treated as having lexical entries that reflect this internal structure, or as being the products of a derivational process that builds them up from the constituent morphemes.

Finally, actual utterances are generated by applying additional phonological and articulatory phonetic constraints or processes to strings of words to produce speech. Notably, such processes are frequently variable, so that a given word or phrase can have multiple surface realizations that do not change its lexical identity. In this model, lexical entries are static and invariant, while variability is a surface output of phonological processes. This model is illustrated in (1):

(1) Conventional model

	UR	Surface variation
and	/ænd/	[ænd ~ æn]
band	/bænd/	[bænd ~ bæn]
banned	/bæn#d/	[bænd ~ bæn]

Such an architecture accounts for many important facts about phonology and the lexicon. But complications arise in accounting for an elementary observation about the surface variability of lexical items: their distribution is significantly conditioned by properties of the word that are not incorporated into the simple string-of-phonemes representation. This fact has motivated alternative architectures, such as the usage-based phonology or exemplar theory articulated by Bybee (2001) and others. In this approach, the mental representation of words incorporates memory traces of previous productions of the word. This exemplar cloud will hence replicate the variability that the speaker has encountered, and the statistical distribution of the variants. New utterances of the word are generated by drawing stochastically from the exemplar cloud, rather than from the single underlying representation postulated in conventional models.

I will survey some of the evidence bearing on these issues that has been discovered in studies of phonological variation. Specifically, I will look at findings regarding lexical frequency, lexical exceptions, morphological constraints on variability, and priming. I will argue that the facts suggest a need for richer and fuzzier mental representations than the conventional model, but more constrained than the exemplar model.

Before turning to the data, however, I want to note that this problem has ancient roots in the history of linguistics. It is foreshadowed in the nineteenth century debate over the Neogrammarian hypothesis. The Neogrammarians argued for exceptionless sound change, operating on phonemic units; words were mere strings of phonemes, so they all necessarily underwent whatever change affected their constituent sounds. The contrary view was articulated in the slogan: "Each word has its own history." The conventional model of the lexicon that I just described is a direct descendant of Herman Paul's formulation of the Neogrammarian position, while exemplars invoke the idiosyncratic potential of each word.

Lexical Frequency

One kind of constraint that is a characteristic of the word as a whole, rather than as a string of phonemes, is its frequency of use. Some words are common, some are rare, and a variety of evidence shows that speakers treat rare words differently from high-frequency forms; for example, they are articulated more slowly, and are less likely to undergo lenition processes. With respect to linguistic variation, several studies have found significant frequency effects, such that higher-frequency forms behave differently from lower-frequency forms.

Two examples from my own work on English involve coronal stop deletion and the *-ing*/*-in* alternation. In Guy, Hay, and Walker 2011, looking at the New Zealand English Corpus at the University of Canterbury, we found a strong log linear frequency effect, such that more common words were deleted more often (see Figure 11.1). This is of course a lenition process, the kind of operation that occurs more often in faster speech.

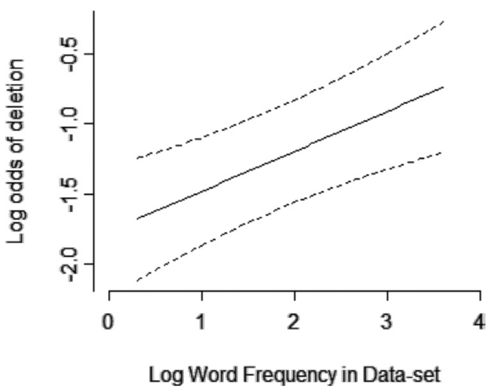

Figure 11.1. Coronal Stop Deletion Increases with Lexical Frequency. p=.0005.
Source: From Guy, Hay, and Walker 2008.

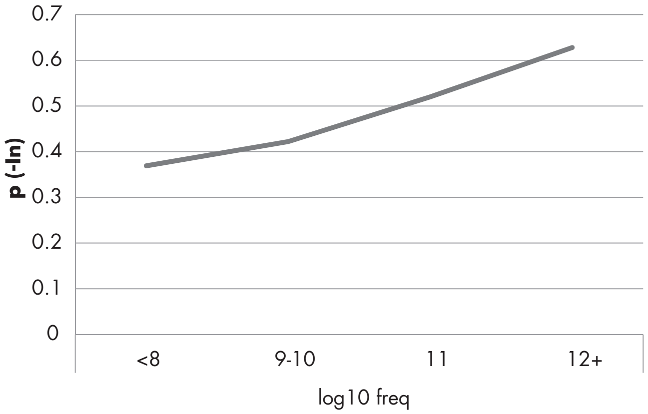

Figure 11.2. [-in] for *-ing* Increases with Lexical Frequency
Source: From Laturnus, de Vilchez, Chaves, and Guy (2016).

Similarly, the alveolar realization of the *-ing* suffix in English—*talkin'* instead of *talking*—is also favored in higher-frequency words. Figure 11.2 shows the frequency effect found in Laturnus et al. 2016. It is not clear that this is a lenition effect, although in a typological sense, alveolar nasals are certainly less marked than velars.

In these two cases, frequency effects were straightforward and significant. But in other studies, the results are different. Some studies of lexical frequency show no significant effects. The lenition of coda /s/ (sometimes characterized as "aspiration and deletion") is a common process in Caribbean Spanish. It involves shortening of duration and lowering the spectral makeup of the fricative noise. Erker's (2008) study of coda /s/ lenition in a corpus of Caribbean Spanish speech looked at both of these acoustic properties. However, neither is significantly associated with lexical frequency; Erker found a correlation coefficient of −0.02 (p = .74, n.s.) for center of gravity and 0.07 for duration (p = .136, n.s.).

What does this imply for the mental representation of words? Such results clearly indicate that quantitative information related to frequency of use is encoded in the lexical entry. Some phonological processes that are operative in the production of variable realizations must interact with or be sensitive to such information. I suggest that the most straightforward account of how the whole system works is effectively equivalent to a mixed effects probabilistic model, with fixed effects for the general constraints and random factors associated with individual words. Through experience, speakers learn that specific words are more or less likely to undergo a given phonological process and encode this in the lexical entry. Such individuation clearly requires an accumulation of evidence—hence it will be most evident in high-frequency words.

Bybee's usage-based phonology (2001, 2002) provides an alternative theory of frequency effects. Exemplar models preserve frequency information by storing in memory all exemplars of the word that the speaker has uttered or heard. But in such a model the interaction between frequency and a variable process requires an additional postulate. Bybee hypothesizes that lenition processes involve a kind of gradual erosion, advancing each time a word is uttered, so the cumulative effects are to make more frequent words appear more lenited:

> Given a tendency for reduction during production, the phonetic representation of a word will gradually accrue more exemplars that are reduced, and these exemplars will become more likely to be chosen for production, where they may undergo further reduction, gradually moving the words of the language in a consistent direction. The more frequent words will have more chances to undergo online reduction and thus will change more rapidly. (Bybee 2002, 271)

The problem with this model is that it overpredicts. Erker's study is a clear example of a gradual lenition process, and as we have seen, no frequency effect is evident. This suggests that "the reductive effect of articulatory automation" is not, in fact, automatic. Lexical frequency is a possible, but not automatic, constraint on phonological or phonetic processes.

Lexical Frequency and Morphology

Frequency also interacts with the second kind of constraint on variation that we will consider, namely morphology. It is well documented that many variable processes are sensitive to the morphological structure of forms. A well-known example is the *-in*/*-ing* alternation in English. Many studies (e.g., Houston 1985; Labov 1989; Tamminga 2014, 2016), have shown that the alveolar /n/ realization is much more common in verbal forms—*talkin*', *runnin*', and so forth—than in nominal forms (e.g., *ceiling*, *a building*). Laturnus et al. (2016) find that verbs are realized with alveolar *-in* in 55 percent of cases (N = 485, factor weight = 0.53), while nouns are realized with *-in* in only 38 percent of cases (N = 79, factor weight 0.32). The exceptional *something* and *nothing* exhibit alveolar *-in* 67 percent of the time (N = 81, factor weight = 0.51).

Another much-studied example is coronal stop deletion (Guy 1980, Guy and Boyd 1990). Guy (1991) finds that underived words systematically have the highest rates of deletion (38.1%, N = 658), and regular past tense forms have the lowest rate (16%, N = 181). For most English speakers, the irregular past tense forms like *kept*, *told*, *lost*, *have* intermediate rates of deletion (33.9%, N = 56).

The interaction between morphological structure and frequency emerges when we look at frequency effects on each morphological category separately. Table 11.1 shows the results of such an analysis of coronal stop deletion, from Myers and Guy 1997. When we look at just the monomorphemes, lexical frequency has a significant

Table 11.1. Frequency interacts with morphology: coronal stop deletion in English

	Monomorphemic Words		**Regular Past Tense Verbs**	
	N	% deleted	N	% deleted
Low frequency	151	18.5	96	7.3
High frequency	573	33.9	220	8.2
	p < .01		p > .70	

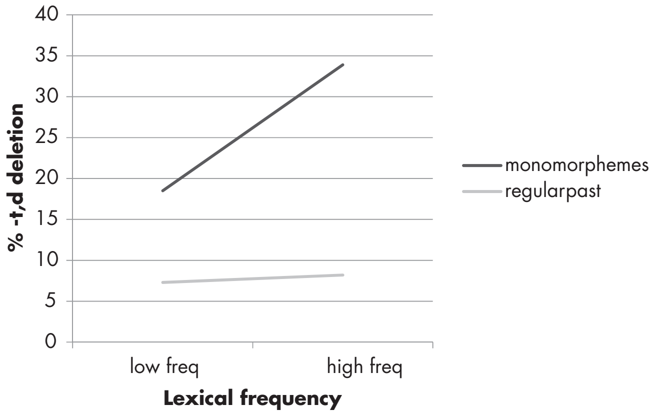

Figure 11.3. Morphology Interacts with Frequency: Coronal Stop Deletion in English

effect, as we saw earlier. But when we consider just the past tense verbs, there is no frequency effect.

Figure 11.3 shows a graphical display of these results. The line for the past tense verbs is basically flat with frequency, while the monomorphemes show the increase in deletion rate that was seen in Figure 11.1.

A continuous treatment of frequency yields a more refined picture of this effect. An unpublished study by Fruehwald, shown in Figure 11.4, examines coronal stop deletion rate in four lexical classes according to frequency treated as a continuous

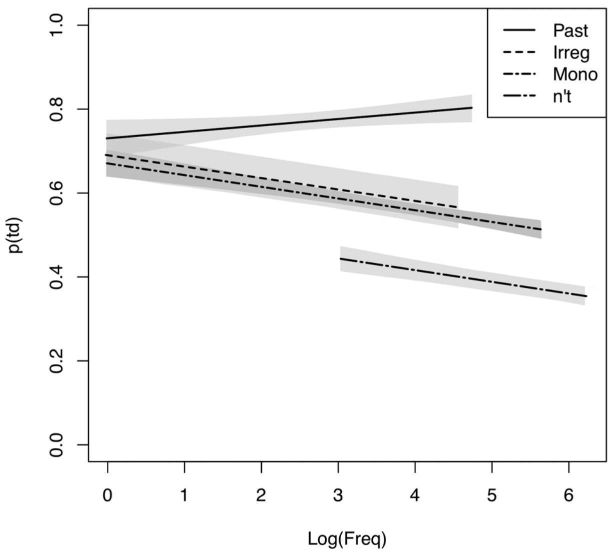

Figure 11.4. Frequency Interacts with Morphology

Source: From Fruehwald (n.d.).

variable. These data are presented in terms of stop retention instead of deletion, so the directionality of the graph is inverted: more deletion = lower down in the figure. The past tense verbs are the top line, showing modestly decreasing deletion rates as they get more frequent. But the monomorphemes, third from the top, show the same behavior as in the Myers and Guy study, with deletion increasing with higher frequency.

These findings have further implications for mental representations. It is clear that variable processes do not see words as mere strings of phonemes. Rather, they are sensitive to aspects of the morphology. In previous work (Guy 1991), I have argued from the perspective of lexical phonology that variable phonological processes actually penetrate the morphological operations involved in derivation and word formation. But at the very least the mental representations of words that provide the input to production must distinguish between derived and underived forms. These facts about the interaction between morphology and frequency are particularly notable: they appear to be inconsistent with Bybee's usage-based phonology. Bybee's model of progressive lenition with more frequent repetition suggests that frequency effects should be orthogonal to and independent of morphology; every uttered word is stored as an exemplar and available as a model for new productions, so more frequently uttered words should presumably be candidates for greater lenition, regardless of their morphology. The data presented here are more consistent with the model advanced by Pinker and others, that derived forms are not stored in the mental lexicon, but rather generated on the fly. (We can allow that especially high-frequency derived forms do acquire lexical entries, as Bybee 1985 argues, by the same process that allows the acquisition of irregular derived forms, such as *leave-left*, *sell-sold*, *think-thought*. But these would count as lexical exceptions, as in the next section.)

Lexical Exceptions

Further evidence on the nature of mental representations can be obtained from the study of lexical exceptions. Various strands of research on how variation affects individual words have demonstrated the existence of clearly exceptional words, which undergo phonological processes in idiosyncratic ways. We will consider four such cases: the English conjunction *and*, the Portuguese first person plural verbal morpheme *-mos*, Salvadoran Spanish discourse markers, and the pronouns *I*, *my* in Southern American English.

The first case is the word *and* in English. Like all English words terminating in a consonant cluster with final -t or -d, *and* is susceptible to coronal stop deletion (CSD). But this word appears with the final /d/ absent at an extraordinarily high rate in all studies that have investigated it. Table 11.2 shows the figures from the Guy et al. New Zealand English study. As we saw in Figure 11.1, high-frequency words are

Table 11.2. Coronal stop deletion in Early New Zealand English: exceptional *and*

	N	% final stop deletion
and	597	80%
other words	3348	29%

Table 11.3. Coronal stop deletion in Early New Zealand English by following context (18 speakers from the ONZE corpus at University of Canterbury)

	Other words		and	
Following Context:	N	% del	N	% del
__C	1339	58.3	315	87.9
__V	1477	10.4	182	75.3
	Range:	47.9%	>	12.6%

expected to show high deletion rates, but *and* is exceptional even when we control for its high lexical frequency.

Notably, *and* is relatively insensitive to other constraints on the deletion process: the contextual constraints on CSD are significantly attenuated for this word. For example, the following segment in the speech stream is a powerful constraint on this process. Before a consonant, deletion rates are much higher than before a vowel: *west side* favors deletion, *west end* disfavors. But this constraint is much weaker for *and*: "*ham 'n' eggs*"—with deletion before a vowel—is almost as common as "*cheese 'n' crackers*," with deletion before a consonant. Table 11.3 provides the results of the New Zealand English study: preconsonantal tokens are almost six times more likely to be deleted in ordinary words, but only 15 percent more likely in *and*.

Similar exceptional lexical items can be found in many other variable processes, and they regularly show weaker effects of the phonological constraints on the process. Brazilian Portuguese and Caribbean Spanish both have processes of coda -s deletion, and both have prominent lexical exceptions. In Brazilian Portuguese it is the first plural verbal morpheme *-mos* that surfaces as *-mo* at an exceptionally high rate. In Table 11.4 we see the following context effects from a multivariate analysis

Table 11.4. Coda –s deletion in Brazilian Portuguese: following context constraint

Features of following C		Other words		1pl verb forms in ***-mos***
Voice/Manner:	sonorant	.69		.49
	voiced obstruent	.44		.58
	voiceless obstruent	.36		.44
	Range	.33	>	.14
Place:	labial	.32		.58
	coronal	.61		.53
	velar	.44		.39
	Range	.29	>	.19
N:		5880		1225
Goodness of fit (log likelihood)		–704.8		–791.5

for the *-mos* forms versus other words. In the *-mos* forms, the constraints are much weaker, with a smaller effect magnitude (indicated by the smaller range values), and a much poorer model fit.

Similar findings emerge from Hoffman's (2004) study of Salvadoran Spanish, where the three common discourse markers ending in -s behave exceptionally: *entonces*, *digamos*, and *pues*. These are all realized without the final /s/ at exceptionally high rates, and they also show considerably weaker contextual constraints.

Another variable showing exceptional lexical items is the monophthongization of /ay/ in Southern American English. In this case the exceptional words are the pronouns *I* and *my*. The figures in Table 11.5, drawn from the work of Woods 2008, demonstrates that these two words are much more likely to undergo monophongization, and they are completely insensitive to following context (monophthongization is favored in the general vocabulary by a following voiced consonant), and they are also less sensitive to speech rate (typically more monophthongs are produced in faster speech with shorter segment durations). So these words stand out from the rest of the lexicon, in the same way as the previous three cases.

In all of these cases, drawn from three different languages, we find exceptional lexical items that differ systematically from the general vocabulary in two ways: they show exceptionally high rates of occurrence of the relevant variable process, and they appear relatively insensitive to contextual constraints on the process that are prominent for ordinary words. What does this suggest about mental representations?

As we argued previously, individual words, especially high-frequency words, may become associated in the mental representation with an idiosyncratic factor indicating their likelihood of undergoing some general phonetic or phonological operation. But this is not sufficient to capture the exceptional behavior of the cases we have just considered. A random factor in a mixed effects model will leave all the other

Table 11.5. Monophthongization of /ay/ in Southern American English: following context and speech rate constraints

	Other words		***I, my***
% monophthong	34%		53%
Following Context:			
__[+cons, +voice]	.76		(.51)
__[-cons]	.41		(.49)
__[+cons, -voice]	.17		(.48)
Range:	.59	>	.03 (n.s.)
Duration:			
shorter	.89		.68
longer	.49		.45
Range:	.40	>	.23

Source: From Woods (2008).

constraints on the process (the fixed effects) stable, and hence cannot capture the apparent weakening of constraints seen in the data. Rather, these data indicate that the underlying representation of an exceptional lexical item is not exclusively captured by the canonical full form of the word. Rather, each of these exceptional words must have an additional underlying representation that already incorporates the output of the variable process. That is to say, these words have underlying allomorphy. English *and* has an allomorph "*an*" or "*n*." When this allomorph is selected, the surface realization will lack a final /d/ regardless of contextual conditions.[1] The surface corpus includes a mixture of such unconditioned tokens with tokens of the underlying full forms to which deletion has applied, which has the mathematical consequence of appearing to attenuate the contextual constraints.

I suggest that this allomorph of *and* is what we represent orthographically in spellings such as "*rock n roll*." The same argument applies to the other cases we have looked at. The Portuguese first person plural verbal morpheme has an allomorph "*-mo*," lacking the final -s in underlying form. The Salvadoran Spanish discourse markers have underlying forms "*digamo*" and "*entonce*," also lacking final -s. And the words "*I*" and "*my*" in Southern American English have underlying representations with monophthongs.

Priming

The final issue I will consider is that of priming. This is the frequently observed phenomenon, sometimes known as persistence or perseveration, where specific variants tend to occur in clusters. In other words, when speakers make one selection from among the possible alternatives in a variable process, they tend to keep making the same selection, at least over a relatively short time span. This most likely reflects a cognitive property rather than a specifically phonological or linguistic one: the activation of a particular mental pathway or operation is heightened by an initial occurrence, and this heightened activation persists for a period of time, favoring its reuse on subsequent occasions. For present purposes, the utility of priming for diagnosing mental representations is that it involves like priming like. In other words, it provides a test of what mental representation is treated as being "like" another.

One exemplary variable for which priming has been well studied is the *-ing*/*-in* alternation in English. The results we will consider are drawn from the Laturnus et al. study cited earlier. In this study we found, like previous studies such as Houston 1985 and Tamminga 2016, that an occurrence of the velar form *-ing* favors another subsequent *-ing*, while the occurrence of an alveolar *-in* pronunciation favors more *-in* tokens in following words. Thus in Table 11.6 we see that after one token with an /n/, 79 percent of the time the next token also has an /n/, but when the priming word is realized with a velar /ŋ/, the /n/ realization occurs in only 21 percent of tokens. Notably, when there is no priming word—in other words, there hasn't been an *-ing* word in the previous few clauses—the data are split about 50–50 between *-in* and *-ing* realizations. This is an important point about the priming effect: when there is no prior priming event, then the current production gives the plain-vanilla output, the emergence of the unmarked.

Table 11.6. Priming effect for the *-in/-ing* alternation in English

Priming context	N	%[n]	weight
-n	129	79	.64
-ŋ	119	21	.27
null	401	54	.53

The data in Table 11.6 conflate all types of *-ing* words as priming contexts. But as we saw in Table 11.3, this variable is strongly differentiated by the part of speech of the *-ing* word: nouns have more velar realizations, verbs more alveolars. So, a more refined priming question is: do nouns and verbs prime each other, or does a verb have a priming effect only on another verb, a noun on a noun? Here the numbers give a nuanced answer. In Table 11.7, showing the factor weights from a multivariate analysis with Goldvarb (Sankoff et al. 2012), words of the same part of speech prime each other better than words of different parts of speech: /n/ realizations in the target word are strongly favored by an /n/ in the priming word (.70), and disfavored by an /ŋ/ in the priming word (.25). But, cross-priming is still strong and significant: target /n/ is about twice as likely after a prime with an /n/ than after a prime with an /ŋ/ (.52 and .25).

What does this tell us about mental representations? One possible analysis is that nominal and verbal *-ing* are actually two different suffixes: a verbal suffix that prefers the alveolar realization, and a nominal suffix that favors velar *-ing*. This would predict the different priming effects. This analysis is supported, as is well-known, by the historical sources of this suffix: Old English had verbal inflections containing /n/ (infinitival *-an*, participial *-ande/ende*) and a nominal suffix *-ung* or *-yng*. So in this interpretation, the differences between nominal and verbal *-ing* mean that we still have two discrete lexical entries, which are masked by the single standard spelling but revealed in the different priming effects and rates of occurrence of velar versus alveolar realizations.

The problem with this model is that it predicts more divergence between the forms than actually exists. Both nominal and verbal *-ing* words occur with both *-in* and *-ing* variants; the difference between them is probabilistic, not discrete. And their priming effect is also nondiscrete. It's not the case that they can't cross-prime, it's just

Table 11.7. English *-in/-ing*: within and cross-category priming effects

	Probability of /n/ realization in target		
Priming form	prime form /n/	prime form /ŋ/	none
same as target: V to V, N to N	.70	.25	
different from target: V to N, N to V	.52	.25	
no prime			.53

that cross-priming is weaker than same-to-same priming. So the quantitative facts suggest neither identical representations nor discretely different representations, but rather overlapping or partially similar representations. The lexical entry for *-ing* is blurry, not crisply focused on one unitary or two contrasting forms. One possible quantitative model of these data would treat the affix as a single entry with two allomorphs, each with a different probability of insertion into nominal and verbal morphological frames.

The Fuzzy Lexicon

Putting all this evidence together gives us a picture of mental representations that I will call the fuzzy lexicon. Words and morphemes are not stored in memory in crisply focused discrete forms, spelled out by contrastive phonemes, and devoid of any information about usage or past occurrences. Rather, lexical entries can take several forms, and can accrete various additional pieces of information. They may incorporate information derived from their history of usage—such as how frequently they have been used and to what surface realizations they have mapped. Some have allomorphs that are discretely different, like "*and ~ an*." They can overlap, be partially similar and partially distinct, or partially merged. Information about their morphological structure is sometimes stored in the lexical representation and sometimes not.

Collectively, these facts require a quantified element in mental representations. I think this is most straightforwardly captured by a probability function. Consider, for example, coronal stop deletion. Its effect is to produce surface forms with and without final stops, so the learner can postulate a phonological form for the word in which the final element is not as stable as the rest of the word; the stable parts of the word have a firm representation with probabilities approaching one, but the final segment is faded, with a probability of less than one. Lexical exceptions emerge when learners postulate that this segment may actually have a probability of zero under certain circumstances.

Frequency information in this model is incorporated by updating the probability function based on experience. I think this is cognitively more plausible than the exemplar model: rather than storing each token one hears, we store only the probability function. Frequency effects operate like an "elsewhere condition," in that high frequency favors lexically specific, marked, and exceptional outcomes over the unmarked general outcomes. In other words, items start with a neutral unmarked probability, and deviate from this only as required by repeated experience. Practice doesn't exactly make them perfect, but it does allow them to become more nuanced.

Morphological information is also at least partially quantified. The resistance of morphological markers to deletion can be restated as a high probability of realization. This was stated in my derivational model of -t,d deletion as a consequence of the stratal structure of the lexicon, but a probabilistic representation of morphemes and morphological boundaries can achieve the same ends.

Priming, as I said, is likely a general cognitive rather than specifically linguistic phenomenon, but it shows us that lexical representations are not necessarily discretely differentiated. The nominal and verbal *-ing* morphemes are partially differentiated, or from a historical perspective, partially merged.

In short, the mental representation of words and morphemes is fuzzy. Items do not have firmly delimited unique identities, but probability functions, with some elements in boldface, and others faded.

Note

1 If the deletion process applies to the full *and* allomorph at the same rate as other words, the probability of selecting the *an'* allomorph can be calculated at about .67.

References

Bybee, Joan. 1985. *Morphology: A Study of the Relationship between Meaning and Form.* Amsterdam: Benjamins.
Bybee, Joan. 2001. *Phonology and Language Use*. Cambridge: Cambridge University Press.
Bybee, Joan. 2002. "Word Frequency and Context of Use in the Lexical Diffusion of Phonetically Conditioned Sound Change." *Language Variation and Change* 14:261–90.
Erker, Daniel. 2008. "Rethinking Coda /s/ Lenition in Spanish: Continuous Descriptions and the Explanatory Value of Lexical Statistics." PhD Qualifying Paper, New York University.
Fruehwald, Josef. 2009. "Summer R Study Group Report: What I Couldn't Have Done Otherwise." Report presented at Spunch. https://jofrhwld.github.io/papers/rstudyhandout.pdf.
Guy, Gregory R. 1980. "Variation in the Group and the Individual: The Case of Final Stop Deletion." In *Locating Language in Time and Space*, edited by William Labov, 1–36. New York: Academic Press.
Guy, Gregory R. 1991. "Explanation in Variable Phonology: An Exponential Model of Morphological Constraints." *Language Variation and Change* 3:1–22.
Guy, Gregory R., and Sally Boyd. 1990. "The Development of a Morphological Class." *Language Variation and Change* 2:1–18.
Guy, Gregory R., Jennifer Hay, and Abby Walker. 2011. "Phonological, Lexical, and Frequency Factors in Coronal Stop Deletion in Early New Zealand English." Paper presented at Laboratory Phonology 11, Wellington, NZ.
Hoffman, Michol. 2004. "Sounding Salvadorean: Phonological Variables in the Spanish of Salvadorean Youth in Toronto." PhD diss., University of Toronto.
Houston, Ann. 1985. "Continuity and Change in English Morphology: The Variable (ING)." PhD diss., University of Pennsylvania.
Labov, William. 1989. "The Child as Linguistic Historian." *Language Variation and Change* 1:85–97.
Laturnus, Rebecca, Natalie de Vilchez, Raquel Chaves, and Gregory R. Guy. 2016. "Dialect, Priming, and Frequency Effects on (-ING) Variation in English." Paper presented at the 45th Annual Conference on New Ways of Analyzing Variation (NWAV-45). Simon Fraser University, Vancouver.
Myers, James, and Gregory R. Guy. 1997. "Frequency Effects in Variable Lexical Phonology." *University of Pennsylvania Working Papers in Linguistics* 4 (1).
Sankoff, David, Sali A. Tagliamonte, and Eric Smith (2012). Goldvarb Lion: A multivariate analysis application for Macintosh. http://individual.utoronto.ca/tagliamonte/goldvarb.html
Tamminga, Meredith. 2014. "Persistence in the Production of Linguistic Variation." PhD diss., University of Pennsylvania.
Tamminga, Meredith. 2016. "Persistence in Phonological and Morphological Variation." *Language Variation and Change* 28:335–56.
Woods, Laurie. 2008. "Whither the Southern Belle?: Language and Identity in the American South." PhD diss., New York University.