The dynamics of near-merger in accommodation

Jennifer Nycz

This paper presents data on the near-merger of low back vowels in American English, and shows that neither classical OT nor the variable rule framework correctly predicts the phonetic properties of this near-merger. Both types of analyses fail for the same reason: each can only describe categorical facts of variation. A new account of the near-merger facts is presented within the framework of nonlinear dynamics. In this model, constraints are modeled as competing attractors in an attractor landscape; this formulation allows us to model the grammar's interaction with context and account for the gradient variable phenomenon of near-merger.

1. Introduction: ways of dealing with variation in phonology

There are two main treatments of variation in phonological theory. Within the rule-based model of SPE (Chomsky & Halle 1968), there is the variable rule, which allows both linguistic and extralinguistic factors to probabilistically affect rule application (Labov 1969, Cedergren & Sankoff 1974). Each relevant factor favors or disfavors application to some extent, and this is quantified in terms of application probabilities for each factor.

In Optimality Theory (Prince & Smolensky 1993), constraints may be variably ranked to yield different outputs from the same input[1]. There are several variants on this theme, for example the cophonologies of Anttila (1997), the stochastic OT of Boersma & Hayes (2001), and the floating constraints of Reynolds (1994). All of these, however, operate in essentially the same way. Given two competing constraints C1 and C2, either C1 outranks C2, yielding one output, or C2 outranks C1, yielding a different output. OT analyses have mostly been concerned with linguistic constraints on variation, but extralinguistic factors may also be incorporated in a limited way, for instance by globally shifting faithfulness constraints upward in the rankings for more formal speech styles (van Oostendorp 1997).

Both variable rules and OT can be, and have been, used to model variation between discrete choices. However, they cannot easily model phonetically gradient variation. This paper will describe one such type of variation – the variable near-merger of the low back vowels /□/ and /□/ in speakers of American English – and show why this data cannot easily be accounted for using either variable rules or variants of OT which operate over discrete outputs. It will be argued that this phenomenon is one facet of the more general problem of

[1] Coetzee (2004) presents a novel way of handling variation in OT: in his formulation, constraint ranking remains fixed, but surface variants of an underlying form may be chosen from among the "loser" candidates which are harmonically ordered by this ranking. Like the other OT systems discussed here, however, this mechanism is limited to the comparison of discrete outputs.

Proceedings of ConSOLE XIII, 2005, 273-285
http://www.sole.leidenuniv.nl

incomplete neutralization, which has been discussed at length in the phonetic and phonological literature. Ultimately, a model of this near-merger will be presented within the framework of nonlinear dynamics (as in Gafos 2003), and possibilities for future research will be discussed.

2. *The data: synchronic near-merger of the low back vowels*

The low back merger (i.e. merger of the vowels in the lexical sets exemplified by *cot* and *caught*) is a large-scale and rapidly spreading change in American English (Labov 1994). Speakers in large sections of the country do not distinguish between these vowels in production or perception, but the distinction remains robust in other areas, especially the Inland North and the Mid-Atlantic states. However, impressionistic data from the Mid-Atlantic region indicate that speakers in this area may be implementing a variable merger: speakers who natively possess the low back contrast sometimes produce their low back vowels distinctly, but sometimes seem to neutralize this distinction. This appears to be dependent at least in part on the merged or unmerged status of one's interlocutor. Someone from New York who ordinarily says [□□□] and [□□□] when speaking to other native New Yorkers may seem to say [□□□] and [□□□] when speaking to a transplant from California who does not share the contrast.

But are these vowels completely merging in the latter context? To examine the behavior of such speakers, a small lab study was carried out. Speakers from New York City who possess the low back vowel distinction were recorded in two contexts: first, conversing and completing tasks with an interlocutor who also has a two-phoneme distinction, and then later completing the same sort of activities with a different interlocutor who has a completely merged system. Acoustic analysis was completed using Praat (Boersma & Weenink 2004): F_1 and F_2 measurements were taken at the midpoint of all low back vowel tokens, and the results from one of the speakers, K, are presented in the table in (1) and the vowel plots in (2).

(1)

Context	Class		F_1 (Hz)	F_2 (Hz)
Distinct Speaker	□̣ (n=104)	Mean	766 (sd =65)	1378 (sd =77)
	□ (n=120)	Mean	708 (sd =66)	1247 (sd =93)
		$Mean_{/□/}$-$Mean_{/□/}$	58	131
Merged Speaker	□̣ (n=160)	Mean	697 (sd =59)	1312 (sd =66)
	□ (n=140)	Mean	669 (sd =53)	1217 (sd =79)
		$Mean_{/□/}$-$Mean_{/□/}$	28	95

Mean F_1 and F_2 values for the low back vowels produced by speaker K in two interlocutor contexts. The distance between the two vowels along each dimension is also calculated.

In fact, in both contexts, K's F_1 and F_2 means for the /□/ word class tokens versus the /□/ word class tokens were significantly different ($p<.05$, in independent samples t-tests), indicating that she continued to produce a contrast between the two vowels even while speaking with the merged speaker. However, K made *less* of a distinction in the merged speaker context: in this case, there is greater overlap between the realizations of these two categories, there is less distance between the mean F_1 and F_2 values for the two vowels, and the overall low back vowel space contracts.

(2)

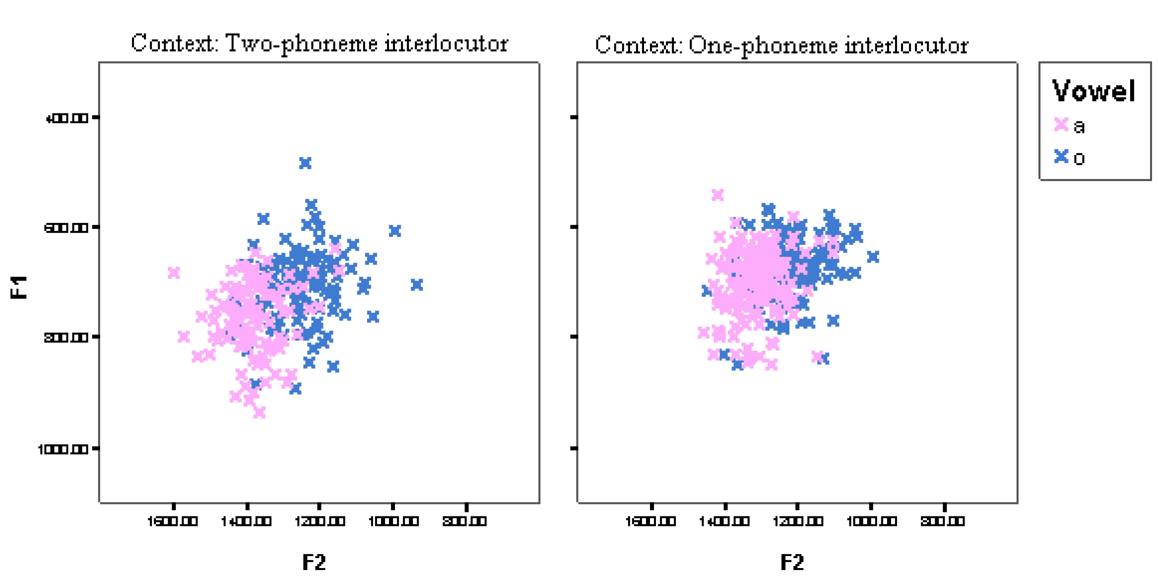

3. Why neutralization rules don't account for this data

One possibility is that this data can be described in terms of a neutralization rule; Herold (1990) alludes to such a rule in her account of the merger's spread among speakers in Pennsylvania. An important point to note, however, is that while K has adjusted her output while speaking to the merged speaker, she did not simply switch from a fully distinguishing grammar to a fully merged one. Instead, she seems to have moved to an intermediate point, where the essential contrast between categories is maintained, but the realizations become less distinct phonetically.

A neutralization rule cannot account for this behavior. Such a rule can only completely obliterate contrast: one vowel takes on the feature values of the other, essentially becoming that vowel for the purposes of phonetic implementation, and this should be reflected in an identical phonetic output.

A variable neutralization rule will also be inadequate, since it is limited to the categorical effects of application vs. non-application. This is illustrated in the diagram in (3). Some of the time, the rule does not apply, resulting in output like that in box (a). In such cases, the vowels in /□/ lexical set members such as *cot* are realized in the canonical '□ space', while /□/ set members such as *caught* have vowels which are realized in the '□ space.' The rest of the time, the neutralization rule applies, turning underlying /□/s into surface [□]s[2], which should then be realized in the □ space. This is shown in box (b). The overall output of the variable rule is simply a sum of the application and non-application cases. If speaker K is implementing a variable neutralization rule, surface realizations of underlying /□/s should occupy the same space as they would if no neutralization rule ever applied, while realizations of underlying /□/s should be spread over both the canonical □ space (in instances of non-

[2] Of course, it is also possible for /□/ to neutralize to [□], though this detail does not affect the argument given here.

application) and the □ space (in instances of application). Importantly, the overall low back space, and the two smaller spaces in which each vowel is realized, should remain stable. We do not expect to see the categories moving closer to one another in the vowel space, as we do in the data presented above.

(3)

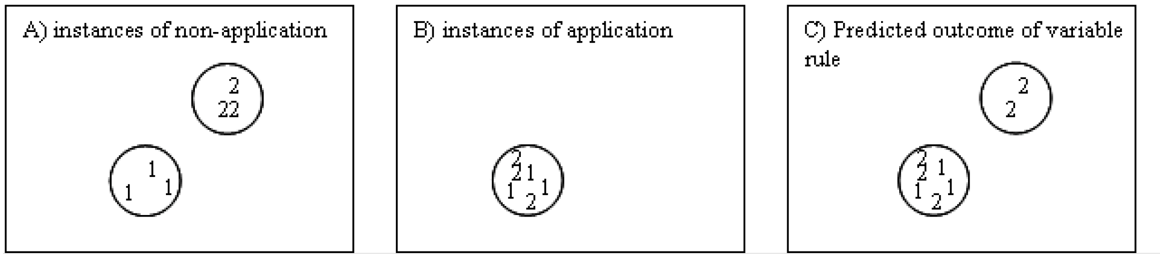

The expected phonetic outcomes of a) contrast maintenance, b) neutralization, and c) a variable neutralization rule. Numeral *1*s indicate lexical items which belong to the /□/ word class (e.g. *cot, top*), and numeral *2*s indicate items from the /□/ word class (e.g. *caught, talk*).

Variation is handled differently in the OT framework. Here, variation between languages is the result of different constraint rankings; variation within languages and idiolects is accounted for in the same way. An OT analysis of the low back near-merger must involve the variable ranking of at least two relevant constraints: a markedness constraint favoring candidates which neutralize the contrast, and a faithfulness constraint favoring candidates which preserve the contrast. However, such a model cannot generate the type of gradient effects described above, *for the same reason* that the rule model cannot. In OT, two constraints C1 and C2 are always ranked with respect to each other such that C1 >> C2, yielding one discrete variant, or C2 >> C1, yielding another discrete variant. In the case of the low back vowels, either markedness will outrank faithfulness, resulting in identical outputs for the inputs /□□□/ and /□□□/, or faithfulness will outrank markedness, resulting in surface contrast. This is the OT counterpart of rule application vs. non-application: one may totally merge, or one may completely preserve contrast, but gradient approximation of two categories cannot be accounted for.

4. The general problem: incomplete neutralization

This phenomenon of near-merger is crucially not confined to the literature that self-identifies as sociophonetic. Many putative 'phonological neutralizations', on closer phonetic inspection, reveal themselves to be essentially cases of synchronic near-merger. One well-known example is syllable-final voicing neutralization, a salient phonological property of languages such as German and Dutch. Some representative data from German are shown in the table in (4). The words meaning 'association' and 'colorful' contrast underlyingly with respect to the voicing of their final obstruents. The traditional account of this pattern is that a neutralization rule applies to these obstruents in word-final contexts, causing the surface forms of both words to end in a voiceless stop. However, measurements of the phonetic properties of the resulting final voiceless obstruents (such as duration of the preceding vowel, duration of consonant closure, etc.) reveal that, in fact, the 'voiceless' surface obstruents which are underlyingly voiced are phonetically more voiced than the surface voiceless obstruents which have always been voiceless (e.g. Port, Mitleb & O'Dell 1981, O'Dell & Port 1983, Port &

Crawford 1989). Similar results have been found in studies of final-devoicing languages such as Polish (Giannini & Cinque 1978) and Catalan (Dinnsen & Charles-Luce 1984).

(4)

Underlying form	Surface form	
/□□□□/ ‘association’	[□□□□]	cf. [□□□□□□]
/□□□□/ ‘colorful(sing.)’	[□□□□]	cf. [□□□□□]

Neither rule- or OT-based mechanisms of neutralization as currently formulated are able to account for incomplete voicing neutralization, for the same reason that they cannot model the facts of low back near-merger: these mechanisms can only completely neutralize a contrast or completely maintain it. Moreover, as Dinnsen (1985) and others have pointed out, these facts are problematic for the traditional derivational view of phonology and phonetics. If a phonological neutralization rule yields two identical forms which are then passed on to phonetic implementation, phonetics is not supposed to ‘know’ to make differences in vowel length, consonant closure duration, etc. which correlate with the underlying voicedness of each obstruent. Port and Crawford ultimately conclude that, in German, ‘practical neutralization is a fact, but it is apparently not a rule’. Dinnsen and Charles-Luce, faced with similar data in Catalan, say that devoicing can be a rule, but this means that phonetic implementation rules that make reference to voicing must be ordered before the phonological neutralization rule. These pronouncements reflect the following problem: word-final devoicing is a qualitative fact about languages like German and Catalan which we would like to account for in the phonologies of these languages, but such placement is undermined by the quantitative facts, which partially flout complete neutralization in favor of preserving an underlying contrast.

5. The model: nonlinear dynamics

Gafos (2003) has described a model that is able to handle these types of facts. This model makes use of the mathematics of nonlinear dynamics, which is widely used throughout the natural sciences to model systems changing in real time (e.g., population growth and decay, economies, and aspects of motor coordination). This framework is now being applied in various ways to the study of cognitive systems such as memory, decision making, and language (see Port & van Gelder 1995 for a sampling of this research, as well as Benus 2005). Gafos uses this framework to model the dynamics of incomplete voicing neutralization described by Port & Crawford and others.

In this model, constraints on phoneme realization are modelled as *attractors* in a numerically-defined multidimensional state space. Each dimension reflects some continuous articulatory variable (such as tongue height or glottal aperture), and the entire state space encompasses all possible states of the system. Certain points in this space are attractors, reflecting preferred states of the system: that is, qualitative aspects of the grammar, such as the canonical realization of a given phoneme category. The figure in (6) illustrates this idea with a simple one-dimensional state space reflecting all the possible degrees of voicing[3]. The

[3] ‘Voicing’ is actually a multi-dimensional parameter of speech, incorporating several subparameters such as duration of closure, amount of glottal pulsing, etc. For simplicity’s sake, I illustrate the point with a one-dimensional space.

space contains two attractors, corresponding to the preferred values of voicing for voiced and voiceless obstruents.

(5)

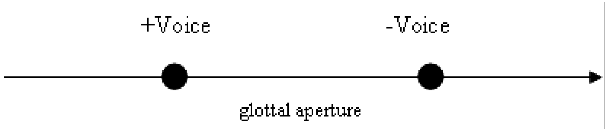

A more intuitive way to represent this is shown in the figure in (7)[4], which plots the potential function V(x) that describes this attractor landscape. The x-values of the minima of the function correspond to points of stability (in this case, the preferred states of voicing and voicelessness) which are relatively impervious to noise. One way to think of this is in terms of a ball moving in the potential function. If one puts the ball in the "voiced" well and shakes it a bit, the ball will ultimately land back in the voicing well. If the ball is placed on the summit at the origin and then shaken, it will fall into one well or the other. This potential function thus provides a good model of categorical perception: there will be a range of values for voicing which are perceived as qualitatively voiced (i.e. a range of points for which a ball placed at any of those points will land in the voicing well), but once a certain part of the voicing continuum is reached, judgments will suddenly become unstable (i.e. the ball is likely to fall into either well). Another important aspect of the model is that the shape of the potential function reflects the stability of each attractor. A well with steep walls is associated with a more stable attractor, insofar as realizations of this category will show little variation. A shallower well will be associated with a less stable attractor: realizations of this category will show more variation.

(6)

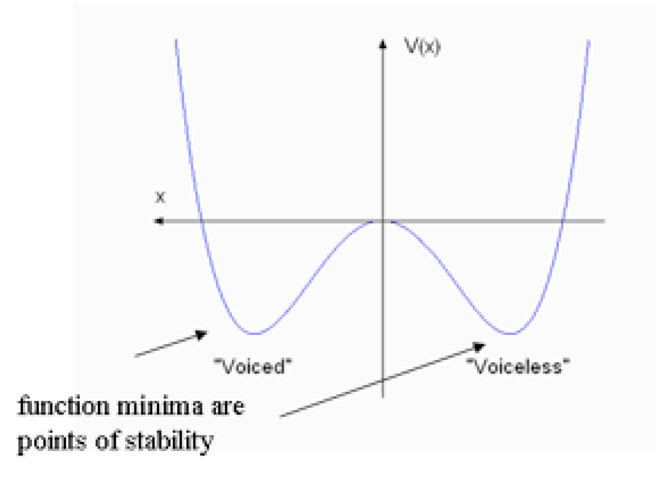

5.1. The model applied to incomplete voicing neutralization

In the case of German incomplete devoicing, there seem to be two different constraints at work, one similar to an OT markedness constraint ("Final obstruents are voiceless") and one

[4] Diagram in Section 5 reproduced from Gafos (2003)

corresponding to a faithfulness constraint ("Output obstruents voicing is faithful to the input voicing value"). As shown above, situating these OT-like constraints within a framework that evaluates discrete candidates will not yield the correct near merger output. Gafos, however, reformulates these constraints in terms of competing attractors in a continuous attractor landscape[5]. The lefthand figure in (8) shows the markedness constraint rendered in attractor form. Here we have a function with one stable point at the more voiceless end of the continuum, reflecting the coda devoicing rule in German. The faithfulness constraint will also be present as a potentially competing attractor in the state space. Given an underlying form /bund/, this attractor will be located at a more voiced value in the state space. Given /bunt/, the attractor will be at a value corresponding to a relatively voiceless value[6].

(7)

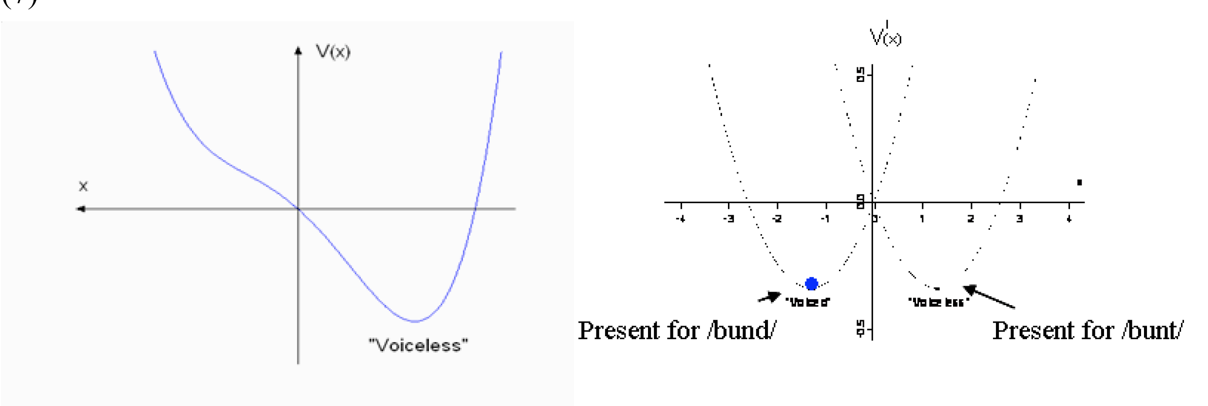

The easiest way to model the results of competition between the markedness and faithfulness attractors is by simply adding together the two potential functions. The result is a new function which has a stable point between that of the original two functions, as shown in the diagrams in 9. Here, the potential function describing the output has a minimum with an x-value which is slightly shifted towards more voiced values, in comparison to the potential corresponding to the markedness constraint which requires a certain level of devoicing.

[5] Gafos' terminology is different from mine. He refers to "grammar" and "intentions" rather than "markedness" and "faithfulness", respectively; I use the latter because these terms are likely to be familiar to the reader.

[6] While two faithfulness parabolas are plotted here on the same graph, only one will be present at a given evaluation, since only one underlying form (and one voicing value) is intended for a given utterance.

(8)

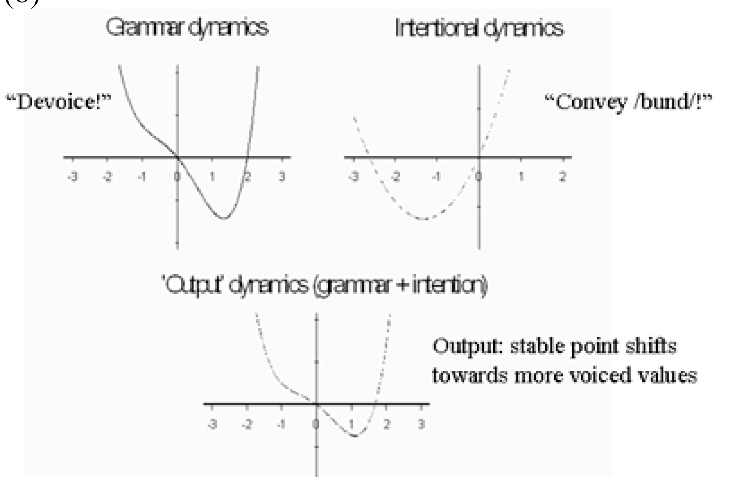

Adding the potentials that correspond to the markedness and faithfulness attractors results in an output potential bearing an intermediate voicing value: V(output) =V(mark) +V(Faith)

Gafos is also able to model the effect of pragmatic factors on the degree of incomplete neutralization. Port & Crawford (1989) describe a series of experiments in which minimal pairs such as German *bund* and *bunt* were produced by German speakers in several contexts. In one context, these forms were embedded in a running text such that speakers did not explicitly contrast the pairs. In another, the speakers read the pairs in contrastive sentences which semantically disambiguated the forms ("I said 'bunt' as in 'colorful', not 'bund' as in 'association'"). In the last context, speakers were instructed to read these forms in semantically-ambiguous contrastive sentences ("I said 'bunt', not 'bund'") to a German-speaking assistant, who took dictation. While Port & Crawford observed a significant difference between underlyingly contrastive stops in all of these contexts, the difference was most pronounced in the dictation case. There is an intuitive explanation of this behavior: when speakers had a greater intent to convey the underlying contrast (in this case, due to the presence of the assistant taking dictation), and could not semantically disambiguate the forms with additional words, their remaining option was to partially override the phonological constraint which obscures the distinction. Indeed, the degree of neutralization decreases as speakers' intent to convey the underlying form increases. Gafos incorporates these facts directly into the model by weighting each of the potentials in the output equation as shown in (9). If the speaker has a greater intent to maintain the contrast, the faithfulness potential will have more weight, with the result that coda obstruents will be even more voiced. Gradually increasing the weight associated with the faithfulness potential (Gafos' 'intentional dynamics') results in a gradual increase in the degree of voicing in the output.

(9)

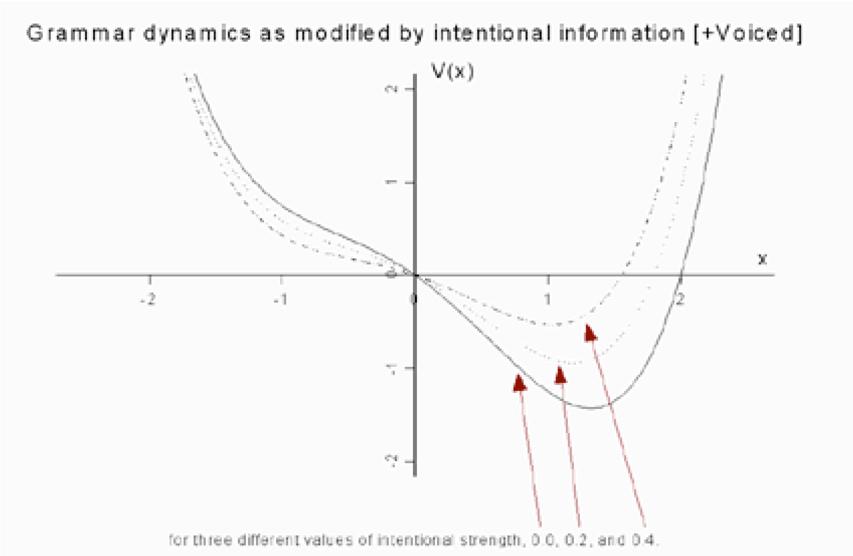

Changing the values of the F and B coefficients in the output equation V(output) = BV(mark) + FV(Faith) changes the relative strength of each attractor, and thus the relative effect that each attractor has on a given output.

6. Modelling accommodation

Gafos' model neatly accounts for the gradient effects of interaction between markedness and faithfulness constraints. As we know, however, the phonetic properties of linguistic output can also be affected by external input to the system: in the case of accommodation, this input consists of the forms produced by an interlocutor. In this section, I show how the attractor model can easily be extended to account for the effect of this input.

6.1 Identifying attractors

It is first necessary to identify the attractors of the system, corresponding to the observable qualitative states; I will illustrate this using the case of speaker K, presented in Section 2. K maintains two underlying vowel categories, /□/ and /□/. This means there will be an attractor corresponding to each of these vowels in the attractor landscape. When producing a particular lexical item, only one of these attractors will be present, reflecting the particular vowel present in that word. The figures in (10) show each of these attractors within a toy landscape; these are K's 'native' attractors, and the potential function used to describe each of them is a simple parabola[7]. These reflect how K will produce these vowels when under no particular influence from an interlocutor (in citation form, for instance).

---

[7] The contrast between the low back vowels /□/ and /□/ is also in reality a multidimensional one, involving at least height and backness/rounding, as shown by the significant differences in F_1 and F_2 between these vowels. In this abstract model, only the relative locations of the attractors along the relevant dimension is important, though the equations for the arbitrarily-chosen parabola used are given in the diagrams.

(10)

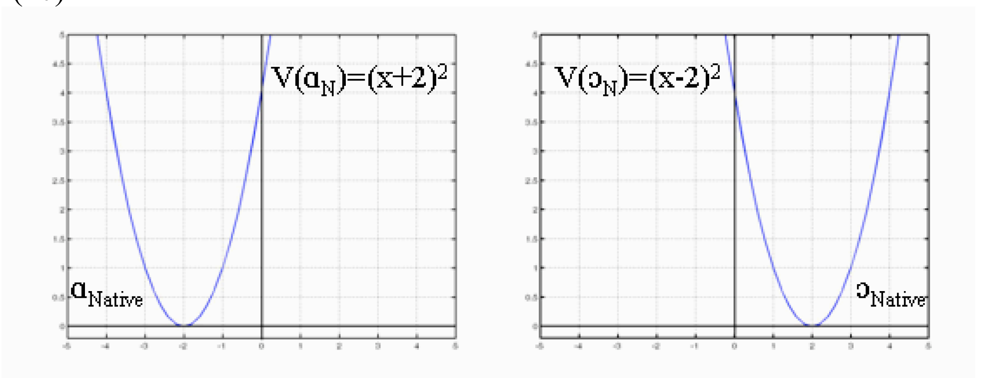

The data show that K's productions vary in a way that is dependent on the productions of her interlocutor. We can incorporate this effect by positing "heard" attractors corresponding to the productions of K's interlocutor: there will be a "heard" attractor for /□/, reflecting how the interlocutor produces tokens of this vowel (as perceived by K), and a similar attractor for /□/ based on the same heard evidence. In the case of a merged interlocutor, the posited attractors will be very close if not overlapping in the state space, since such a speaker produces no distinction between the relevant word classes (11).

(11)

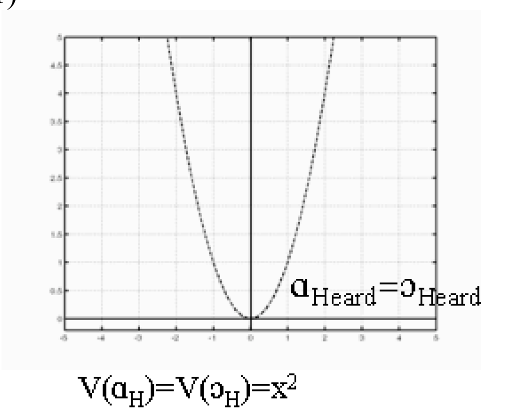

6.2 Competition

K's output for each vowel will be mediated by two constraints. One the one hand, there will be a pull to produce the relevant phoneme as given by the native system, and on the other hand, there will be a pull to approximate the productions of the interlocutor. The results of the interaction of these two constraints are described with the additive model of competition. For a production of *cot*, for instance, the function describing the native /□/ attractor and the function describing the heard /□/ attractor are added together, yielding an output potential which has an intermediate value as shown in the figures in (12). The main result of this competition will be that K's vowels, while they remain distinct in the context of the merged speaker, start to converge (13). This prediction is borne out by the data presented in Section 2.

(12)

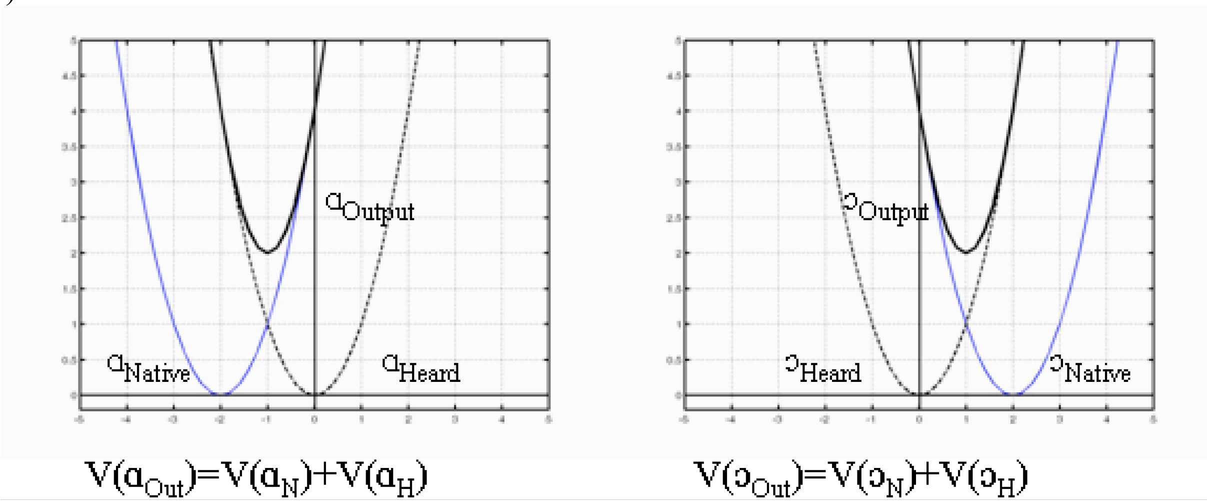

(13)

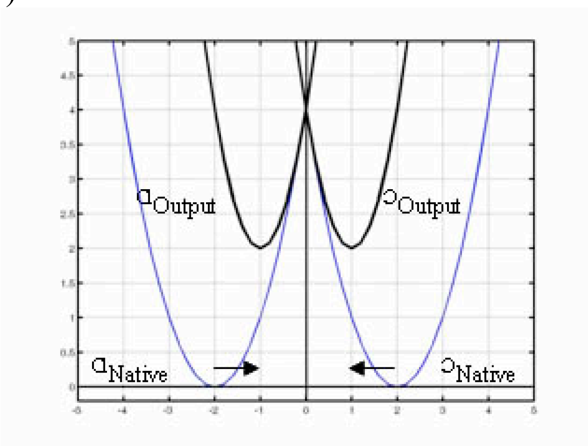

Of course, accommodation is not merely a matter of meeting an interlocutor halfway. Speakers accommodate to a greater or lesser extent, over long and short time scales. Over the course of the lifespan, developmental stages provide coarse ranges of weighting possibilities: while children are in one sense striving to be the ultimate accommodators during the language acquisition process, adults are less likely to pick up the speech patterns of those around them. Within a particular developmental stage, attitudinal factors will lead people to accommodate more or less between conversations, and even within conversations. The effects of developmental factors and attitude are united by and incorporated into the model via the type of weighting coefficients described here. As the figure in (14) shows, an increase in the weight of the heard attractor with respect to the native attractor (corresponding to some increase in the desire to move towards the interlocutor) will result in a realization that is even closer to this heard value.

(14)

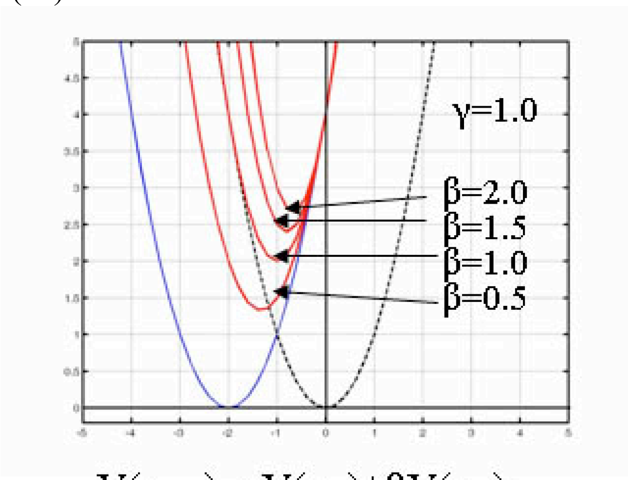

$$V(a_{Out})=\gamma V(a_N)+\beta V(a_H)$$

7. Conclusions and directions for future research

This paper presented data on synchronic near-merger which cannot be accounted for by traditional mechanisms of neutralization, either in rules-based or OT-based models. Gafos' dynamics approach to the similar problem of incomplete voicing neutralization was introduced, and it was shown how this model can be extended to incorporate the gradient effect of interlocutor input on a speaker's output.

In addition to accounting for the data presented here, the dynamics approach to accommodation opens up several possibilities for future research. Most significantly, this approach offers a way of smoothly integrating grammatical and sociolinguistic constraints on language usage within one model, and accounting for their interaction and gradient effects. In the toy model presented here, speakers posit attractors on the fly during the course of conversation with a particular interlocutor. However, the model is crucially not limited to responsive behavior. As a large body of work beginning with Bell (1984) has shown, there is a significant initiative dimension to style shift: though speakers from Social Group X may not be present, speakers can convey solidarity with X through the productive use of socially-indexed linguistic variables. For example, work by Eckert (2000), Labov (1972), and Schilling-Estes (1998) has shown that raised variants of the diphthong /□□/ are indexed with various salient local identities, depending on location - Detroit "burnout" teenagers, fishermen on Martha's Vineyard, and natives of Ocracoke Island, respectively. For speakers who deploy this feature, the choice between non-raised and raised /□□/ is not binary: there is a continuum of variation between these poles, and greater degrees of raising correlate with stronger intentions to convey the relevant social meaning. These facts are straightforwardly accounted for by giving a greater relative weight to the "identity" attractors in a given production. The dynamics model will be especially adept at handling the interaction between multiple, possibly conflicting identity attractors, as well as the competition between these attractors and linguistic constraints: social motivations may pull one towards raising /□□/, but one can only do so to a certain extent if the linguistic context disfavors raising.

This approach also gives us a way to model gradient change in speakers' systems over time. Phonological categories are stable, such that one conversation with a speaker of a different dialect or language will not significantly alter one's native categories in the long run, but they are also flexible, such that repeated exposure to different productions can cause them to change. This model thus accounts for the sociolinguistic concept of a 'norm enforcement mechanism': if a speaker is embedded in a community of people who all speak a certain way, then categories will tend to stabilize towards 'average' productions of this community. Isolated exposures to different idiolects will cancel each other out as noise, but sustained exposure to a particular accent that is different will result in change.

Acknowledgements

Thanks to Diamandis Gafos, Lisa Davidson, Paul De Decker, Ken Lacy, Kevin Roon, John Singler, and members of the Ph-Lab at NYU for helpful comments and advice on this work.

Jennifer Nycz
New York University
jennifer.nycz@nyu.edu
http://home.nyu.edu/~jrn218

References

Anttila, A. (1997). Deriving variation from grammar. Hinskens, F., R. van Hout & L. Wetzels (eds.), *Variation, Change and Phonological Theory.* John Benjamins, Amsterdam/Philadelphia, pp. 35-68.

Bell, A. (1984). Language style as audience design. *Language in Society* 13:2, pp. 145-204.

Benus, S. (2005). *Dynamics and transparency in vowel harmony*. PhD. Dissertation, New York University.

Boersma, P. (1997). How we learn variation, optionality, and probability. Proceedings of the Institute of Phonetic Sciences of the University of Amsterdam 21, pp. 43–58.

Boersma, P. & B. Hayes. (2001). Empirical tests of the Gradual Learning Algorithm. *Linguistic Inquiry* 32, pp. 45-86.

Boersma, P. & Weenink, D. (2004). Praat. http://www.fon.hum.uva.nl/praat/.

Browman, C. & L. Goldstein. (1995). Dynamics and articulatory phonology. Port, R. & T. van Gelder (eds.), *Mind as Motion*. Cambridge, MA: MIT Press, , pp. 175-193.

Buder, E. (1991). A nonlinear dynamic model of social interaction. *Communication Research* 18:174-198.

Cedergren, H. & D. Sankoff. (1974). Variable rules: performance as a statistical reflection of competence. *Language* 50, pp. 233-55.

Chen, M. (1970). Vowel length variation as a function of the voicing of the consonant environment. *Phonetica* 22, pp. 129-159.

Chomsky, N, & M. Halle. (1968). *The sound pattern of English*. New York: Harper and Row.

Coetzee, A. (2004). *What it means to be a loser: non-optimal candidates in Optimality Theory*. Ph.D. Dissertation, University of Massachusetts, Amherst.

Dinnsen, D. (1985). A re-examination of phonological neutralization. *Journal of Linguistics* 21, pp. 265-279.

Dinnsen, D. & J. Charles-Luce. (1984). Phonological neutralization, phonetic implementation and individual differences. *Journal of Phonetics* 12, pp. 49-60.

Eckert, P. (2000). *Linguistic variation as social practice.* Malden, MA/Oxford, UK: Blackwell.

Gafos, A. (2003). Dynamics in grammar: comment on Ladd and Ernestus and Baayen. To appear in Goldstein, Whalen, Best, and Anderson (eds.), *Varieties of phonological competence,* Berlin/New York: Mouton de Gruyter.

Giannini, A. & U. Cinque. (1978). Phonetic status and phonemic function of the final devoiced stops in Polish. *Speech Laboratory Report.* Napoli: Istituto Universitario Orientale.

Guy, G. (2004). Language and time: Adaption, development, and diachrony. Paper given at *Sociolinguistic Perspectives on Age* workshop at New York University, November 12th, 2004.

Herold, R. (1990). *Mechanisms of merger: the implementation and distribution of the low back merger in eastern Pennsylvania.* Ph.D Dissertation, University of Pennsylvania.

Labov, W. (1969). Contraction, deletion and inherent variability of the English copula. *Language* 45, pp. 715-762.

Labov, W. (1972). *Sociolinguistic Patterns.* Philadelphia: University of Pennsylvania Press.

Labov, W. (1994). *Principles of linguistic change: internal Factors*. London: Blackwell.

Norton, A. (1995. Dynamics: an introduction. Port, R. & T. van Gelder (eds.) *Mind as motion*. Cambridge, MA: MIT Press, pp. 175-193.

Percival, I. & D. Richards. (1982). *Introduction to dynamics.* Cambridge University Press.

Port, R. (1996). The discreteness of phonetic elements and formal linguistics: response to A. Manaster-Ramer. *Journal of Phonetics* 24, pp 491-511.

Port, R. & P. Crawford. (1989). Incomplete neutralization and pragmatics in German. *Journal of Phonetics* 17, pp. 2257-282.

Port, R. & M. O'Dell. (1985). Neutralization of syllable-final voicing in German. *Journal of Phonetics* 13, pp. 455-471.

Port, R., F. Mitleb, & M. O'Dell. (1981). Neutralization of obstruents voicing in German is incomplete. *Journal of the Acoustical Society of America* 70, S10.

Port, R. & T. van Gelder (eds.). (1995). *Mind as motion: explorations in the dynamics of cognition*. Cambridge, MA: Bradford books, MIT Press.

Prince, A. & P. Smolensky. (1993). *Optimality Theory: constraint interaction in generative grammar,* unpublished manuscript, Rutgers University and University of Colorado, Boulder.

Reynolds, W. (1994). *Variation and phonological theory.* Ph.D dissertation, University of Pennsylvania.

Schilling-Estes, N. (1998). Investigating 'self-conscious' speech: the performance register in Ocracoke English. *Language in Society* 27, pp. 53-83.

van Oostendorp, M. (1997). Style registers in conflict resolution. Hinskens, F., R. van Hout & L. Wetzels (eds.), *Variation, change and phonological theory.* John Benjamins, Amsterdam/Philadelphia, pp. 207-229.