THE ROUTLEDGE COMPANION TO SOCIOLINGUISTICS

Edited by
Carmen Llamas, Louise Mullany
and Peter Stockwell

2007

Routledge
Taylor & Francis Group
LONDON AND NEW YORK

3
TECHNIQUES OF ANALYSIS

I PHONOLOGICAL VARIATION

MATTHEW GORDON

Sociolinguists operate under the axiom that linguistic variation is not random but rather is shaped by social and linguistic factors. One of the goals of any sociolinguistic study is an account of the influence of these factors. Often this information is presented in the form of a statistical analysis that clearly defines the correlations among the linguistic forms and the various social and linguistic factors. When we read such an account, we rarely consider the process of analysis that went into creating it. Nevertheless, this process is in many cases the most time-consuming stage in the research.

It is easy to appreciate why this process takes so long when we remember that it typically begins with hours of recorded speech from several subjects. This is the pool of raw material in which the patterns of variation are to be found. To the casual observer, sociolinguistic variation can appear chaotic. It seems that some people use some forms more than other people and that some forms may be more common in certain words or contexts than in others, but firm generalizations are hard to deduce without a systematic analysis. That analysis is essentially a process of translating natural speech into data that allow comparison across speakers and linguistic contexts.

This chapter sketches the process for analysing phonological variables. We begin by considering issues related to how the variables to be analysed are defined. From there we turn to techniques for measuring the variation. The final section considers how variables are affected by their linguistic contexts.

DEFINING THE VARIABLES

A crucial early step in any variationist analysis is a definition of the **linguistic variables** to be examined (see Chapter 1). The research must specify the range of variation associated with each variable. This involves detailing the **phonetic variants** of the variable and their **phonological** distribution within the language or **dialect** studied.

It is common to think about phonological variation as coming in two basic flavours: *discrete* and *continuous*. Discrete variation involves phonetic variants that represent distinct alternatives. Some of the clearest examples of this type relate to the presence versus absence of a sound. Familiar examples from English include **H-dropping** in England (e.g. [hæt] ~ [æt] 'hat') and *r*-lessness in New York

City and elsewhere (e.g. [kɑr] ~ [kɑ] 'car'). Discrete variation often relies on a binary choice as in these two examples but may involve more variants. In a study of Newcastle speech, for example, Watt and Milroy (1999) identified four discrete variants of (o): [oʊ], [ɵː], [ʊə], and [oː].

In the case of continuous variation, there are no clear boundaries among the variants. Rather, the variable exhibits a range of realizations along a *phonetic continuum*. Many vocalic variables operate in this fashion. For example, in many varieties of American English, /æ/ is variably raised sometimes as high as [i]. Speakers do not simply alternate between [i] and [æ]. Instead they sometimes produce [æ], sometimes produce [i], and sometimes produce intermediate variants in the neighbourhood of [ɛ] and [e]; that is, they have available to them any phonetic value along the range from the low [æ] to the high [i].

In addition to defining the phonetic dimensions of the variation, the researcher must also specify the scope of the variable within the phonological system. If we think about variation as stemming from the application of rules, then the scope represents the contexts in which the rules may apply. Of course these rules may or may not apply in any given case and their likelihood of applying is influenced by social and linguistic factors (see below). Our concern here lies with defining where the variation might possibly operate. For phonological variables this is a relatively straightforward task as compared with the challenges presented by grammatical variables (see further Chapter 4).

In some instances the scope of a variable is simply that of the **phoneme** in which case the variation operates in all contexts containing that phoneme. Consider, for example, the raising of /æ/ noted above. In Chicago, Detroit, and elsewhere, this raising is part of a phenomenon known as the Northern Cities Shift (Labov 1994; Gordon 2001), a series of sound changes that lead to variable realizations of several phonemes. These changes apply across the board so that every word containing these phonemes may show the effects of the shift. Thus, the scope of the variable (æ) – parentheses are used to indicate **sociolinguistic variables** – is isomorphic with that of the phoneme /æ/.

Phonological variables often have a more restricted scope than the phonemes they involve. For instance, they are commonly subject to *phonological conditioning*. In the case of *r*-lessness, the /r/ is eligible for deletion only when it does not appear before a **vowel**. The raising of /æ/, while it applies across the board in the northern cities, operates within certain phonological restrictions in other locations, including Philadelphia and New York City (see Labov 1994). Thus in Philadelphia and New York the vowel may be raised when it appears before a **nasal** consonant (as in *man*) but not before a **voiceless stop** (as in *cat*).

The scope of some variables may be defined lexically, that is, in terms of particular words rather than phonological contexts. Often such variables involve alternations between phonemes (e.g. *either* and *neither* pronounced with /i/ or /aj/). In some cases the variation may apply only to a single word as in the American English examples of *aunt* which may appear with /æ/ or /ɑ/ or *ask* which may appear as [æsk] or [æks].

An exact definition of the linguistic variable is an essential prerequisite to any meaningful sociolinguistic analysis. The variation associated with a given variable is shaped by social and linguistic factors, but one cannot untangle the effects of those factors until the boundaries of the variation have been delineated. Failure to properly define the variable clouds the picture of the variation and may introduce serious bias to the results. Imagine a hypothetical study of *r*-lessness that failed to recognize the phonological conditioning of this variable and counted all instances of /r/ rather than just those potentially subject to deletion. If all of the examples from one group of speakers came from pre-vocalic /r/ (e.g. *ride, carry*), and all those from another group came from **postvocalic** contexts (e.g. *car, park*), the researcher might be led to the erroneous conclusion that the second group deleted /r/ much more frequently. Such an error is unlikely with a well studied variable like *r*-lessness, but the general caveat remains: clearly defining the variable helps ensure that one is comparing apples with apples.

MEASURING VARIATION

With the linguistic variable defined, the researcher can set about the task of measuring the variation associated with that variable. This task is essentially one of distilling the raw material of the recorded speech into usable data. These data will serve as the input for the later analysis in which the effects of social and linguistic factors are explored. Phonological variation is usually measured in one of two ways: (1) **auditorially** (by listening to the recordings), or (2) **instrumentally** (using spectrographic analysis of the acoustic signal).

Auditory coding

Over the last four decades of sociolinguistic research, the most common approach to measuring phonological variation has been to rely on the auditory judgements of the investigators who listen to the recorded speech to determine the variants used. With repeated listenings, researchers can train themselves to distinguish subtle phonetic variants reliably. A sample of speech from each subject is then reviewed and each instance of the variable under investigation is coded according to the variant produced. The goal of this coding is usually to produce some kind of statistical measure of each subject's usage.

In the investigation of discrete variables, this measurement is relatively straightforward. The researcher listens to and codes a certain number of instances or tokens of the variable, and then counts how frequently each variant appears. These raw numbers are easily translated into percentages by taking the number of tokens of each variant and dividing it by the total of all tokens. For statistical reliability it is recommended that at least thirty tokens of the variable be examined for each speaker (see further Guy 1993).

The coding of continuous variables is somewhat more complicated. The phonetic variation is too great for every variant to be transcribed. The researcher

must impose a classificatory system onto the variation to break the phonetic continuum up into steps. In effect this approach treats the continuous variable as if it were discrete. Returning to the example of /æ/ raising, the analysis might propose four degrees of raising represented by unraised [æ], the slightly raised [ɛ], the more raised [e], and the most raised [i]. The tokens would then be coded according to these steps in the raising process, keeping in mind that each step represents a piece out of the phonetic continuum. The phonetic codes can be converted to a mathematical index of each speaker's usage by assigning numerical values to each step. In this example, the conservative [æ] would be assigned a zero, [ɛ] would be one, [e] would be two, and [i] would be three. The index is calculated by taking the average of all the tokens coded. In this way a speaker with an index around zero shows very little raising while one approaching three shows consistent raising to the highest degree. Like the percentages calculated for discrete variables, these indexes allow for straightforward comparison across speakers.

Instrumental measurement

Phonological variation can also be examined using the instruments and techniques of **acoustic** phonetics. This approach was pioneered in sociolinguistics by Labov *et al.* (1972), and has become increasingly popular in recent years due in part to technological advances which make acoustic analysis possible on a personal computer. A thorough account of the theoretical underpinnings of these methods is beyond the scope of this chapter, but accessible introductions can be found in Kent and Read (1992) and Johnson (1997).

In an acoustic analysis the measurements of the variation are taken instrumentally rather than by listening to the recordings. Today the process most often involves digitizing the recorded speech samples in order to enter them into a computer program to perform spectrographic analysis. This analysis produces a visual representation of the speech signal from which precise measurements can be taken. One of the most common representations used is the **spectrogram** in which shadings of light and dark are used to show degrees of acoustic energy at different frequencies across time. Research in phonetics has identified several measurable components of the acoustic signal that correlate with particular phonetic features. For example, the phonetic difference between **voiced** and **voiceless** sounds is indicated in a spectrogram by the presence or absence of a voicing bar, which appears as a dark band at low frequencies (see Figure 3.1).

Sociolinguists most commonly employ acoustic analysis in the study of vowels. This approach is especially useful in examining continuous vocalic variation such as in the case of /æ/ raising and other changes in the Northern Cities Shift. In a spectrogram, vowels appear as dark horizontal bands of energy known as *formants*. Formants are created by sound resonating in the mouth and pharynx. As the shape of the vocal tract changes by moving the tongue to produce different vowels, the sound resonates at different frequencies. For this reason, measuring the frequencies of the formants can provide indications of how the vocal tract is shaped;

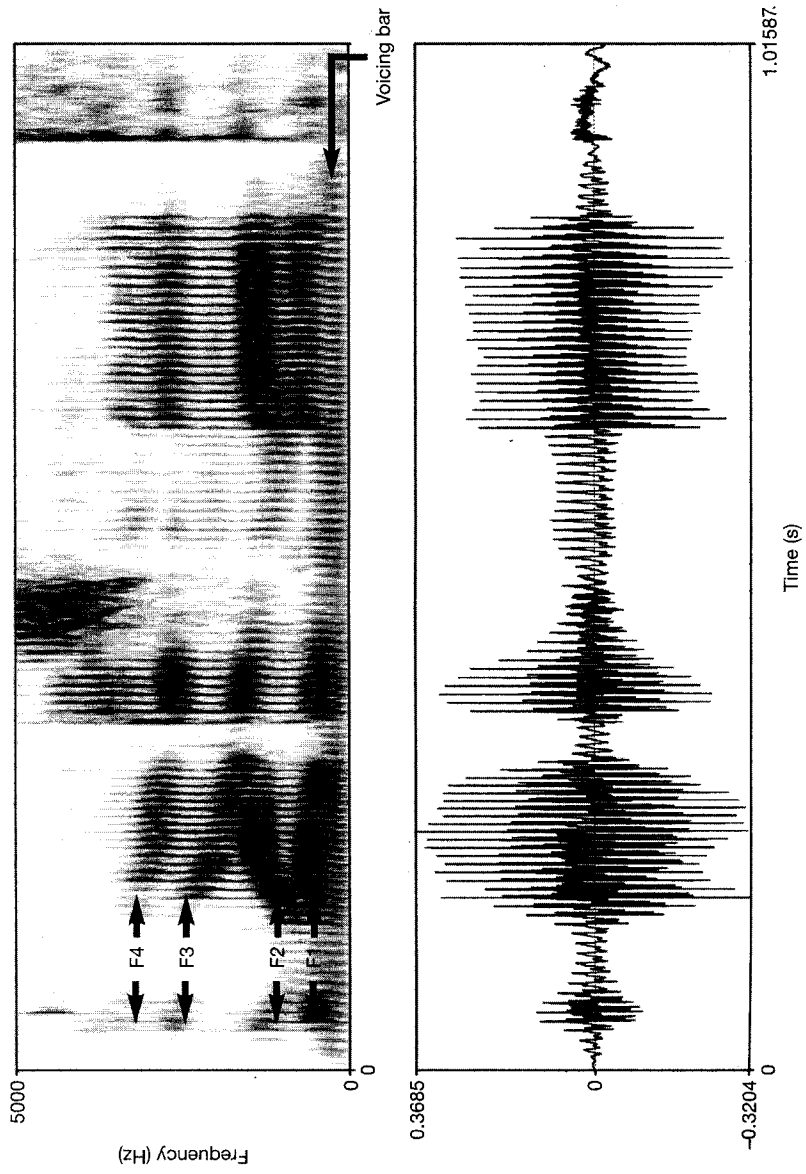

Figure 3.1 Spectrogram (top) and Waveform (bottom)

that is, of how the tongue is positioned in making the vowel sound. In this way, formant frequency measurements serve as a corollary of the position of vowels in the mouth. Of particular interest are the first and second formants, labelled F1 and F2 here, which are usually taken as corollaries of vowel height and frontness respectively. Low vowels are characterized by a high F1 frequency and high vowels by a low F1 frequency as measured in Hertz (Hz). A typical male speaker might have a high [i] with an F1 frequency of about 300 Hz and a low [æ] with an F1 around 700 Hz. For the front–back dimension, high F2 frequencies characterize front vowels and lower F2 values characterize back vowels. Thus, an [i] produced by a male speaker might have an F2 of 2,000 Hz while his [u] might have an F2 of 1,000 Hz.

Researchers can use formant frequency measurements to create a picture of a subject's vowel space by plotting the data on a graph. Individual tokens of the variable can be plotted, but for a less cluttered picture the researcher may choose to plot mean values for F1 and F2 which have been calculated on the basis of several tokens of the variable. These vowel plots can be oriented in keeping with the traditional representations of vowel articulations (i.e. the *vowel quadrangle* with [i] in the upper left corner and [ɑ] in the lower right – see Figure 3.2).

Comparing the positions of the vowels as measured acoustically with their expected positions demonstrates the progress of vowel shifts. For example, a speaker who is advanced in the Northern Cities Shift might have a vowel plot with /æ/ in the high front position (indicating a low F1 and a high F2) very near his or

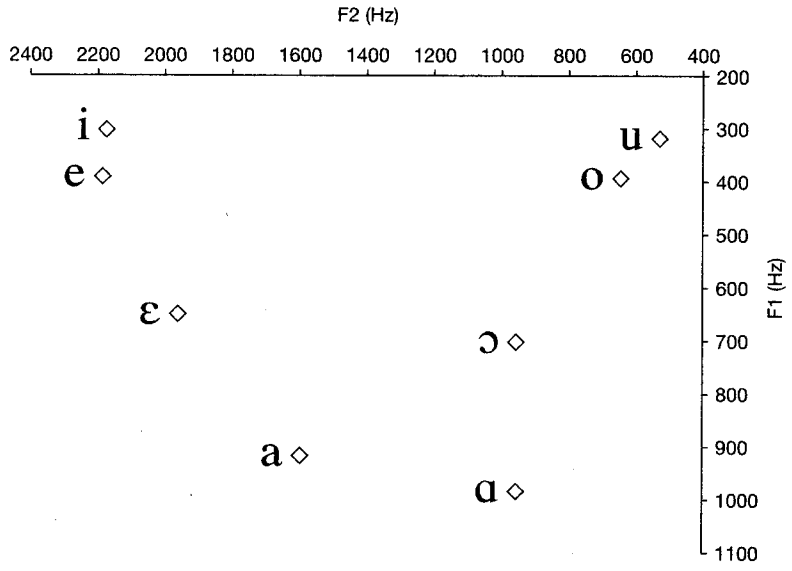

Figure 3.2 F1 and F2 plot of a male speaker's vowel space

her /i/, which is not affected by the shift (see Thomas 2001 for a collection of plots representing various American dialects).

One of the advantages of acoustic analysis is that it allows the researcher to measure variations that are too subtle to be reliably detected auditorially. Also, because the measurements are taken mechanically and therefore rely less on the researcher's judgement, they are felt to be less subjective. Still, acoustic analysis raises questions of its own (see further Milroy and Gordon 2003). One of the most important challenges posed by the use of acoustic data is cross-speaker comparison. Formant frequencies and other acoustic measures are affected by physiological differences in vocal tracts. Since no two subjects have identical vocal tracts, their formant measurements cannot be directly compared. This fact leaves the researcher with two choices: (1) compare speakers in terms of the relative positions of the vowels rather than the absolute formant frequencies, or (2) normalize the frequency measures by applying a mathematical formula.

The former approach can be useful in studying vowel shifts such as the Northern Cities Shift where progress can be measured, for example, by the position of /æ/ relative to /i/. To compare a large sample of speakers, the researcher might develop a coding system to mark degrees of shifting in much the same way as is done in an auditory analysis. Fridland (1999) presents this type of analysis in her study of another set of changes known as the Southern Shift.

The raw frequency data can be made comparable by applying a formula that normalizes the values. Several such formulas have been proposed, and their merits are the subject of ongoing debate in the field (see, for example, Adank *et al.* 1999). All of these normalization routines are intended to factor out the effects of physiological differences across vocal tracts. In this way, the normalized frequency values can be used to compare speakers through statistical analysis (see, for example, Labov 2001).

Analysing linguistic factors

When we think of important sociolinguistic studies, we tend to concentrate on findings related to the social significance of linguistic variables rather than on those related to phonological conditioning. From Labov's study of New York department stores (1966), for example, we might remember that **rhotic** pronunciations were more common among the employees of the high-end store Saks though it is harder to recall that rhoticity was also more common at the end of a word than before another consonant. This is understandable since it is the focus on the social functioning of language that distinguishes this work from other areas of linguistics. We should keep in mind, however, that sociolinguistic analysis also involves a thorough exploration of linguistic factors shaping usage. Indeed, generalizations about the social distribution of a speech form cannot be made without an understanding of its internal linguistic patterning.

Phonological variables are often influenced by elements of their linguistic context. Labov's finding that /r/ is more often deleted in pre-consonantal contexts

(e.g. *fourth*) than in word-final position (e.g. *floor*) illustrates the phonological conditioning that commonly shapes usage of a variable. The importance of phonological conditioning was noted above as a key element in defining a variable. The kind of conditioning at issue here differs in that it does not apply categorically but rather it contributes to tendencies. Thus, in this section we are interested in the kinds of phonological factors that make a given variant more or less likely to occur. With *r*-lessness, /r/ is more likely to be deleted pre-consonantally, but it does not have to be, and conversely it is less likely to be deleted word-finally, but it certainly can be.

The elements of the phonological context that serve as conditioning factors will vary according to the type of variable. Similar to *r*-lessness in New York, **glottalization** of /t/ in many British varieties is influenced by the position of the consonant in the word as well as by whether it is followed by a vowel or consonant (see, for example, Docherty and Foulkes 1999). In the case of vowels, adjacent consonants often play a role. The raising of /æ/, for example, has been found to be promoted by the appearance of a following nasal consonant (e.g. *ham*, *hand*) and retarded by a preceding **liquid** (e.g. *laugh*, *rat*) (Labov 1994). Such findings about phonological conditioning are arrived at by simply comparing usage across contexts as measured either auditorially or instrumentally. In the same way that one might calculate an index for an individual speaker's usage of (æ), one can break the data up by phonological context and calculate an index for all the tokens involving a following nasal, a preceding liquid, and so on.

In addition to phonological context, a variable might be influenced by individual lexical items; that is, use of some variant might be more or less common in a particular word. Ash (1997) reports a case of lexical conditioning for the process of /l/ **vocalization**, in which /l/ is pronounced as a vowel such as [o] or a **glide** such as [w] (as in [fɪo] ~ [fɪw] for 'fill'). Among the Philadelphians Ash studied, she found vocalization was not common when /l/ appeared between vowels except in the word *Philadelphia* itself.

Recognizing the influence of linguistic context, researchers often take steps to ensure they sample a range of contexts for each subject. Thus, they might limit the number of tokens of any given word to be coded. This helps to reduce the potential for skewed results from a particular lexical item or phonological context. Such measures are necessary in order to conduct a reliable comparison across speakers and thus to examine the influence of social factors.

This chapter has sketched out some of the major components of a sociolinguistic analysis of phonological variation. Readers seeking a fuller treatment of these issues may wish to consult Hudson (1996) or Milroy and Gordon (2003). My goal here has been to give a sense of the analytical process that lies behind the results we encounter when reading a sociolinguistic study. Understanding the methodological choices that the researcher has made is essential to evaluating the validity of the study's conclusions.

FURTHER READING

Chambers, J.K. (2003) *Sociolinguistic Theory* (2nd edition), Oxford: Blackwell.

Hudson, R.A. (1996) *Sociolinguistics* (2nd edition), Cambridge: Cambridge University Press.

Johnson, K. (1997) *Acoustic and Auditory Phonetics*, Oxford: Blackwell.

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

Milroy, L. and Gordon, M. (2003) *Sociolinguistics: Method and Interpretation*, Oxford: Blackwell.