

Chapter 1

Introduction

A thorough understanding of variation must by definition include an understanding of how it is heard and processed. Because language use fundamentally involves communication, every use implies both speaking and listening, encoding and decoding, expressing and interpreting. The study of variation has for the most part focused on the correlates of speakers' use of variation and comparatively little effort has been spent on listeners. Listener perceptions of variation are of inherent interest in that they are as fundamental an aspect of linguistic variation as correlates of speaker and situation. They are additionally valuable in offering answers to the many questions regarding the existence and nature of social meaning in relation to variation.

One of the central questions in the study of variation is how instances of variation connect to social structures. The concept of social meaning has helped us to understand the flexibility of socially linked variation. Variation is one of the many tools that people use to create meaning in their day-to-day lives. Seeing this has allowed us to explain many patterns of variable use that otherwise might remain a mystery (Eckert 2001b). Other questions are raised by this approach, however. Foremost among these is how to establish if variation actually carries meaning from speaker to listener. If so, how? Does all variation do so or only some? Are there limits on the meanings speakers can convey and where do they come from? When we speak of meaningful variation, what are we talking about in cognitive terms? Listener perceptions of variation offer insight into all of these issues and more.

This dissertation adapts existing methods to investigate social evaluations of a single variable, the English variable (ING). The work uses an adapted form of the Matched Guise Technique, using recordings digitally manipulated to create minimally paired stimuli, differing only in (ING). The stimuli were developed from recordings of spontaneous speech taken from sociolinguistic interviews with eight speakers, balanced by gender (four men and four women) and region (four from the South and four from the West Coast). Listener responses were collected in two phases, through group interviews and a survey. The results demonstrate the existence of social meaning, as indicated by observable changes in listener perceptions based on the manipulation of a single variable. They also demonstrate the crucial role of context in determining that meaning, the interconnected nature of social responses and the potential for variability among listeners as well as speakers.

This chapter provides the theoretical background on the conceptual issues of variation addressed by this project. First, I discuss the role of social meaning in understanding of sociolinguistic variation and introduce the important concepts in this tradition, including the co-constructed nature of meaning and the notion of indirect indexicality. I then discuss the significance of listener perceptions and briefly introduce work which has theoretically developed the role of listeners. Finally, I describe the structure of the rest of the dissertation in detail.

1.1 The problem of meaning

The earliest work on variation brought local meanings to the fore, examining the crucial role of local strife in the production of linguistic variants. Labov (1963) tied the use of locally salient vowel shifts on Martha's Vineyard to speakers' orientations towards the economic and social changes then underway on the island. Fischer (1958), which will be discussed in detail in Chapter 2, also addressed some level of social meaning, looking at differences in (ING) use and how they related to different identities or personalities. Fischer found that a "model boy" used more *-ing* while a "typical boy" favored *-in*. Much of the work that directly followed in the late 1960's and early 1970's stepped up to a larger scale, carrying out surveys and community studies

with large populations and random samples. These studies did not engage explicitly with meaning, but rather concentrated on issues of “social significance” (Labov 1966; Trudgill 1974) or how social position influenced linguistic performance. They did engage implicitly with social meaning, drawing on concepts such as stigma and prestige.

In recent years, researchers have been increasingly drawn to the issue of social meaning and in particular the local nature of such meaning. In trying to establish the theoretical groundwork for understanding the social meaning of variation, semantic notions of meaning are not very helpful. Silverstein (1976) discusses the wealth of ways in which language can carry meaning other than semantic, referential meaning.

To say of social behavior that it is meaningful implies necessarily that it is communicative, that is, that the behavior is a complex of signs (sign vehicles) that signal, or stand for, something in some respect. Such behavioral signs are significant to some persons, participants in the communicative event, and such behavior is purposive, that is, goal oriented in the sense of accomplishing (or in failing to accomplish) certain ends of communication, for example, indicating one’s social rank, reporting an occurrence, effecting a cure for a disease, and so forth. (Silverstein 1976:12)

Silverstein lays out a number of important concepts regarding indexicality and language. His use of the term *index* draws on the three-way divide in Peirce (1901) which divides **signs** into three categories: **icons**, in which the sign vehicle is associated with it signified by virtue of a similarity (e.g. a road sign with a schematic of an intersection); **indexes** in which the sign vehicle co-occurs with or predicts the signified in space or time (as smoke signals fire) and **symbols**, which constitute the rest of signs, in which the relationship between the signified and the signifier is arbitrary (as for most words). Although much of linguistics focuses on the semantic meaning carried by symbols, Silverstein points out that this is just one among many functions that language can take on. He formalizes the ways that linguistic indexes rely on aspects of the speech situation for their meaning. Further, he suggests a classification system based on two dimensions: the referential/semantic meaning carried by an indexical marker (e.g. tense is both referential and indexical while honorific markers

of speaker/hearer relationships are not) and the degree of performativity or creativity the markers allow for (e.g. locative deictics are fairly constrained while honorifics allow speakers to shift social situations within certain boundaries).

This work provides crucial theoretical structure by formalizing indexical meaning. It is extended to address indexing by linguistic variation by Ochs (1992), who looks specifically at linguistic indexing of gender. Her crucial insight is that “Few features of language directly and exclusively index gender” (Ochs 1992:340). Most linguistic behaviors associated with gender are used by both men and women, though perhaps to different amounts: “the relationship between language and gender is distributional and probabilistic” (Ochs 1992:340). Additionally, forms associated with gender tend to carry other social meanings and “the multiplicity of potential meanings allows speakers to exploit such inherent ambiguities for strategic ends” (Ochs 1992:340).

Ochs refers to these two observations as characterizing the “non-exclusive relation” between language and gender, the first of three characteristics she identifies. The third (temporal transcendence) does not concern us here. But the second characteristic is what Ochs calls the “constitutive relation” and is the basic idea of indirect indexicality, which will be important in the current discussion. Ochs proposes that linguistic resources may directly index gender or other social meanings such as speech activities (e.g. oratory), speech acts (e.g. ordering) and affective stances (e.g. coarseness). All of these meanings, including gender, work to constitute each other. Because the various social meanings all help to constitute each other, a given linguistic resource may be used in the production of gender or another meaning, even when it lacks a direct indexing relationship.

It is in this sense that the relation between language and gender is mediated and constituted through a web of socially organized pragmatic meanings. Knowledge of how language relates to gender is not a catalogue of correlations between particular linguistic forms and sex of speakers, referents, addressees and the like. Rather, such knowledge entails tacit understanding of (1) how particular linguistic forms can be used to perform particular pragmatic work (such as conveying stance and social action) and (2) norms, preferences and expectations regarding the distribution of

this work *vis-a-vis* particular social identities of speakers, referents, and addressees. (Ochs 1992:341-2)

The central idea is that a given resource may be used in the production of social meanings that it is not directly linked to. Instead, it may index that meaning indirectly, via an intervening meaning. Ochs' model introduces a sort of miniature semantic network, whereby a linguistic resource may be connected directly to a given social meaning which is itself connected to others. A correlational variationist approach might discover that women are more likely to use the word *please* than men, but Ochs's point is that associating the word with women directly would be a mistake, as is evident from the many examples where men may use the word without attempting to sound like women or being perceived as feminine. Instead, the word *please* may be directly linked to the notion of politeness, which is in turn connected to femininity or appropriate behavior for women.

This is the basic idea of social meaning for linguistic resources with which I am working. Social structures link together in a complex web, to which language connects at different points. The model Ochs presents is primarily aimed at situating gender and its relationship to language within the larger frame of mutually constitutive meanings. This theory applies more widely than gender, however.

Viewing linguistic variation as carrying social meaning rather than straightforwardly reflecting social address involves an important theoretical shift. If variation may be used strategically to convey particular kinds of meaning which impact a social situation, it means that rather than merely reflecting the social world, language is in fact central to constructing it. Identities like race or gender, as well as other aspects of the social world, do not simply exist as inherent qualities in people, places or situations, but are co-constructed as people interact with one another. The "co-" aspect of co-construction means that this process of construction is a joint project, one which all interactants contribute to. It does **not** mean that all contribute equally, that they engage with a spirit of cooperation or that a consensus is reached. We often also speak of social meanings being contested, to emphasize the competitive or even combative nature of some group processes of constructions (see (Rickford 1986) for a discussion of non-consensus based understandings of social class). This phrasing also

highlights the continual aspect of the process—meanings are continually contested, always under construction and never completed. The word **intersubjectivity** also removes the temptation to view meaning-making as necessarily cooperative, substituting the image of meanings occurring in the space between subjects (people) with no one able to claim final ownership.

The process of co-construction depends crucially on the perception of meaningful social cues, including sociolinguistic variation. Research which adopts this theoretical framework draws more heavily on the processes involved in social perception. However, interpretation of socially meaningful variation or, indeed, variation in general, is studied significantly less often than its production. The research that has been done shows that social cues are bound up intimately with the interpretation of linguistic signals of all kinds. Although researchers traditionally place social judgments with semantic meaning at the “top” of the interpretation hierarchy, even the most basic processes of phonemic identification may be influenced by social factors. Phoneticians interested in the mechanisms of speech perception have discovered that listeners use a range of information regarding the speaker in order to normalize their perceptions, by creating a frame of reference to use during perception (Verbrugge *et al.* 1976; Assmann *et al.* 1982; Johnson 1990). Niedzielski (1996, 1998, 1999, 2001) established that in addition to these physical cues, purely social information also influences listeners’ phonetic perceptions at a very low level. Niedzielski provided listeners from Detroit with recordings of another Detroit speaker whose speech showed features of both the Northern Cities Chain Shift (NCCS) and Canadian Raising (CR). She had previously determined that while the NCCS did not figure largely in the metalinguistic awareness of Detroit residents, CR did, although they believed it was limited to residents of Canada. Listeners were told either that the speaker was Canadian or from Detroit and were asked to select synthesized tokens which most closely matched the vowels she produced. When hearing the vowels influenced by CR, which carried social awareness for these listeners, the speaker’s supposed nationality significantly influenced their selections. Those who believed they were listening to the Canadian speaker tended to select raised variants which were in fact closest to the actual token. Those who believed the speaker was from Detroit, however, more often selected

either standard or hyper standard variants. In the case of tokens influenced by the NCCS, the purported nationality of the speaker had no effect on listeners selections. Instead, listeners for the most part selected synthesized vowels that corresponded to the standard pronunciation of the segments, instead of the actual shifted variants they heard. Even when nationality is not implicated in the ideology, the social information of what token is appropriate or expected had a stronger effect than the sounds they actually heard.

What the current literature has neglected to date is the actual influence of specific variable on the interpretation of a linguistic performance. This question is crucial to treating linguistic variation as reflecting a system of social meaning. This becomes an issue because of the multiple associations any single variable may have. In real instances of use, a variable may respond to a range of situational and personal factors. Much early work on (ING) shows that it correlates with both socioeconomic status and the formality of the speech situation (Labov 1966; Trudgill 1974). Labov and Trudgill both suggest that variants associated with higher levels on the socio-economic hierarchy are more valued and thus more likely to be employed in situations where speakers are attending to the forms of their speech. A different, although related, explanation is that the meaning of (ING) is one which is of different use in different situations and also is valued or used differently by people of different socioeconomic classes. Eckert (2001a) suggests that (ING) primarily signals formality, with *-ing* being a more formal way of speaking and *-in* seen as more casual. Both the social and stylistic stratification might result from this kind of meaning, as different classes are likely to have different orientations towards formality and different situations involve different levels of formality as well.

The introduction of social meaning allows more powerful generalizations by sensibly linking divergent uses of the same variable. Even so, the range of uses found, particularly for stable and widespread variables such as (ING) makes it difficult to point to a single monolithic meaning. It is at this point that the concept of indirect indexicality becomes useful, as discussed previously. The possibility of indirect indexing allows us to explain a range of otherwise puzzling data. It also greatly expands the world of possibilities for linking language and social meaning. Because of this

expansion, we run the risk of throwing open our theoretical understanding to predict infinite flexibility. If any resource can mean anything, how do speakers choose what to say? How do listeners make sense of what they hear? What prevents anyone from combining any combination of variables to mean anything they wish? Since we know there are limits on the actual performances and interpretations we must insure our theory accounts for these.

While the terrain laid out by indirect indexical meaning is flexible, it is not infinitely so. There is structure in the linguistic realm, as well as among the social meanings. This structure limits the choices speakers may make and shapes the processes of both speaking and listening. Meaningful linguistic variation is structured by markedness relations, by different levels of salience, by cognitive limits and by the habits of speaker/listeners form over time, to say nothing of internal linguistic constraints themselves. Social meanings bear structured relationships to each other and this structure is further shaped by ideologies and conscious beliefs. The more we understand about the structures of these two domains the better we will be able to understand the possibilities and the behaviors of the interface between them. I will briefly discuss each of these structural influences, beginning with the linguistic factors of markedness and salience.

Sociolinguistic variation is structured in such a way that certain variables and variants are more perceptible to speakers and listeners than others, leading to different patterns of use and different likelihoods of taking on social meaning. Linguistic variables may be more or less salient, for example the pilot work described in Chapter 2 revealed that listeners tend to be more conscious of (ING) than they are of the release or non-release of /t/. This may be due to the form the variable takes; a lexical item may be more quickly noticed from a single use than a phonological shift. Labov (2001) articulates the varying degrees of consciousness speakers have for different variables with a three-way divide: indicators, markers and stereotypes. **Indicators** are socially stratified linguistic behaviors which have no social awareness associated with them, are not topics for discussion and are difficult for listeners to detect. **Markers** are patterns with a degree more social awareness, a difference Labov hypothesizes results from them being further along in their process of change. In

this case speakers are more likely to display style shifting and to respond negatively to stigmatized forms. The highest degree of awareness is the **stereotype**, linguistic variation which is subject to overt commentary, potentially to the degree that speakers may discuss the form and its stigma without realizing they use it themselves (Labov 2001:196). The variable under study in this project, (ING), is clearly in the last category. Speakers discuss it easily, have conventionalized terms for referring to it (most commonly, “dropping one’s G’s”) and easily articulate ideologies concerning its use. The degree of consciousness that listeners have for particular variable is likely to affect the ways in which they respond to its presence in the speech of others.

While variables may be more or less salient, within a single variable, variants may be structured in terms of markedness. The concept of markedness refers to which member of a given paradigm is considered more natural, and less noticeable, either in general or with respect to contextual factors. In phonology, the unmarked member is the one which appears in absence of specific conditioning (Anderson 1985). In morphological paradigms, this default member is frequently the one with the least amount of explicit morphological marking. In a variable paradigm, the unmarked member might similarly be said to be the one which appears most often, which appears in absence of conditioning factors and which carries a lighter burden of social or situational meanings. In most settings, listeners seem to conceptualize *-ing* as the unmarked member of the paradigm of (ING), as I discuss in Chapter 7. Of course, a system of markedness being present in a variable does not mean that all speakers and listeners will agree on the distribution of markedness. Likewise, different variables will be more or less salient within different communities. One of the things which influences the distributions of markedness and salience is the habitual mode of speaking for a given speaker/listener, the patterns of language they are most accustomed to using and hearing. Although assessments of markedness do not depend solely on expectations derive from habit, it is likely that those variants which depart more strongly or unexpectedly from a listener’s customary experience are more apt to be noticed and assigned meaning than those which differ only slightly. As a result, patterns of markedness may be similar to social meanings themselves in being different for different people or groups and subject to contestation.

Just as linguistic variation is structured, so too is the social realm with which it is entangled. As part of constructing the social world, people form associations between certain concepts, considering them to be related, or likely to co-occur. These are built both from direct experience and from explicit beliefs about the structure of the world. Some of these patterns are stable and wide-ranging; researchers into language attitudes have found a consistent correlation between listeners' ratings of the speakers' intelligence, education and socioeconomic status across a variety of populations (Zahn and Hopper 1985). While this combination should surprise no one, it reflects a structuring of the social world, particularly when contrasted with the relative lack of correlation between these qualities and those typically associated with dynamism in this same literature, including qualities like forceful, persuasive, or self-assured. This structuring of listeners' perceptions influences the potential interface points for linguistic variation. Given the lack of correlation between status and dynamism and the association of regional dialect features with lack of education, this structuring makes it likely that dialect features will affect status-related qualities more than dynamism traits.

In addition to implicit associations, listeners' images of the social world are structured by explicit beliefs and ideologies. Eagleton (1991:1) provides a list of potential definitions for the word ideology, including these two: "the medium in which conscious social actors make sense of the world" and "action-oriented sets of beliefs". The beliefs that people hold about the social world and its relation to language shape the choices that speakers make and the interpretations listeners construct. Niedzielski's work, discussed above, reflects not only the possibly implicit association of a specific linguistic pattern with Canadian speech, but also the ideological stance of Detroit residents that their speech is standard and therefore lacking this particular feature they perceive to be marked. This ideology is a weighty one because it does not limit itself to this single linguistic phenomenon, but reflects a widespread adherence to the notion of standard. It also requires a set of beliefs regarding what constitutes standard language and the belief that the speech in their region is perfectly reflective of that standard. Ideologies need not concern themselves directly with language in order to influence linguistic behaviors. For example, ideologies about education and

intelligence may lead to listeners' discounting the effect of language on a speaker's intelligence if the message content reflects a level of educational attainment.

At present, we know quite a bit about the structure of linguistic variation and our knowledge of the structure of the social meaning is growing rapidly. Research into the processes by which listeners interpret linguistic variation will help us learn more about the interface between these two. This, in turn, will help to eliminate the impact of this interface on the long term maintenance and change in patterns of linguistic variation. The research described in this dissertation will help us understand the ways in which listeners negotiate these linguistic, social and sociolinguistic structures to perform their share of the process of co-construction. Using listener perceptions to answer these questions is still somewhat unusual. In the next section, I will briefly discuss some of the many reasons that listener perceptions are useful objects of study.

1.2 Why do listener perceptions matter?

When contemplating the role of listener perceptions in the structuring of linguistic variation, it is apparent that speakers who differ in their linguistic preferences and social sense of meaning while speaking must also differ while listening. As a result, we would expect patterned variation in the responses that listeners give to particular variables, much as we witness patterned variation in the performances of speakers themselves. Indeed, not only are these likely to correspond, but the knowledge on which the performance is based must be to a large extent gained through observation of the performances of others. Further, to the extent that speakers are aware of this variation in their audiences or potential audiences, it will shape the performances they create.

Studies of listener perceptions have been carried out with a range of goals and a range of methodologies. Linguists may wish to investigate listener perceptions for the simple reason that listening is part of the overall process of language use and development and is thus as valuable a subject of inquiry as language production. Labov (1966), in his study of the English of New York City, speculated that the convergence of stylistic and socio-economic factors shown in his production study was

due to a common agreement as to the social status of the given forms. Labov designed the “subjective evaluation test” with which to verify this hypothesis. Rickford (1985) used a subjective reactions test to investigate the structure of the creole continuum in Guyana. Coming from a more technical perspective, Elman *et al.* (1977) demonstrated that the phonetics characteristics of segment perception are different for different languages, even within the same bilingual speaker.

In many cases, linguists may wish to study phenomena that are relatively rare in speech, making it more feasible to study perceptions. Although it is not commonly included in a discussion of listener perception work, the classic technique of eliciting grammaticality judgments in theoretical syntax is of this type, as are more recent approaches which take their cue from psychology (see, for example, Bard *et al.* (1996)). Both of these rely on the notion that certain phenomena may be too rare to investigate fully using production (Pullum 1990). A similar example is the elusive phenomenon of “Gay Speech”. Gay men, like straight men, present a range of personae and linguistic styles and thus there are no elements which tie the speech of gay men together and distinguish it from that of straight men (Kulick 2000; Podesva *et al.* 2001). Nonetheless, there is good evidence that there are some vocal cues which trigger a percept of “sounding gay” and that some of these cues are shared across some listeners (Gaudio 1994; Rogers and Smyth 2003; Levon 2005a). It is uncertain whether this style will turn out to consist of a particular set of variables, a set of practices in relation to the larger linguistic matrix or an ideologically defined concept which has no independent linguistic reality at all. Regardless, it is a pattern whose primary operational definition seems to be in terms of perception. I will discuss this phenomenon and its supporting literature in Chapter 5. Similar patterns may be seen in very early perceptual studies. Works such as Pear (1931) set out to find correlates between personality and speech. They succeeded in establishing that listeners are not very skilled at reading personality from speech, but that they tend to agree strongly with each other in their assessments. This suggests the possibility that the association of, for example, a deep voice with a commanding personality is one which exists mostly in perception rather than performance.

The study of listeners may be able to answer important cognitive questions about

processing of language. Niedzielski's work, discussed previously, demonstrated that social categorization of the speaker, in this case geographic region, operates at a very low level to literally affect the perception of the sounds of language. Strand (1999) and Johnson *et al.* (1999) show a similar phenomenon with respect to gender, drawing on the integration of linguistic and visual information first established by McGurk and Macdonald (1976). Plichta and Preston (2005) shows that listeners can process variables along a continuous basis and are capable of aligning that continuum with the social (or geographic) one. All of these studies answer questions about the cognitive abilities and/or habits of speaker/listeners and do so but a study of listening. In the case of vowel splits or mergers, perceptual results can shed light on whether listeners are maintaining distinctions that may be difficult to observe through traditional measurement methods (Labov *et al.* 1972; Di Paolo and Faber 1990).

Listener perceptions are perhaps most thoroughly studied within the covert study of language attitudes, using the Matched Guise Technique developed by Lambert and his colleagues (Lambert *et al.* 1965; Lambert 1967). This approach uses the study of perceptions of linguistic traits as a foil for investigating covert attitudes regarding groups of people. Because linguistic prejudice is frequently more socially sanctioned than other forms, Lambert and his colleagues hypothesized that respondents would express their opinions more openly if the responses were prompted by linguistic performances. Researchers use recordings created by the same person speaking in different manners to investigate listeners' beliefs and attitudes regarding the categories of people commonly believed to speak in those ways. By using the same speaker and varying specific traits, the aim is to eliminate irrelevant variation and guarantee responses which reflect only the differences under study. Since this is the technique used in this study, it will be discussed in detail in Chapter 3, along with the description of the methods used in the current study.

Lastly, there's a great deal of evidence that speakers take listeners and their processes into account when constructing their performances. Communicative Accommodation Theory takes the interaction of linguistic behaviors between participants as crucial to the overall construction of variation (Giles *et al.* 1973; Giles 1973). Accommodation Theory began in the early 1970s as a way of incorporating an understanding

of style, style shifting and code choice into sociolinguistic understanding. It centers around the ways in which speakers orient towards the linguistic practices, real or imagined, of their interlocutors, converging or diverging strategically (Giles and Powesland 1975; Thakerar *et al.* 1982). Accomodation, or convergence, is the common behavior that “at least one member of an interactive diad tends to adopt the speech patterns of the person to whom he (sic) is talking” (Giles and Powesland 1975:156). This may occur for a range of reasons, usually relating to a desire to increase one’s similarity to one’s interlocutor and therefore one’s attractiveness. Early experimental work established that increasing communicative similarity between speaker and listener also increased the listener’s evaluation of the speaker’s attractiveness, as well as their ability to predict and understand what the speaker is saying (Bishop 1979; Berger 1979; Triandis 1960). Convergence of this type relates to the speaker’s desire for social approval and so the less powerful participant in interaction is more likely to converge to their interlocutor than the more powerful, since they are more likely to need approval to accomplish their social goals (Thakerar *et al.* 1982). In other situations, participants may diverge in order to express anger or distance.

Accomodation Theory is notable for the degree to which it focuses on the actual sociopsychological processes involved in “recipency to talk”, breaking it down explicitly into perceptual, labeling, attributional and evaluative dimensions (Coupland and Giles 1988:178). Coupland and Giles also emphasize the importance of combining detailed linguistic analysis with a good understanding of the relevant ideological factors and of combining observations of natural data and various kinds with sophisticated experimental work. Much of the early work in the framework focuses on relatively straightforward concepts of convergence and divergence towards the speech style of another, but over time more complex approaches have been developed. Atkinson and Coupland (1988) discuss incorporating ideologies into Accomodation Theory, noting that “situated talk is both conditioned by, and in some sense constitutive of, some higher-order, structured value-system” (p. 321).

This idea has been theoretically developed by others, as well. Drawing on work in Accomodation Theory, Bell (1984); Bell (2001) developed a model of audience design. Bell proposed that a comprehensive model of stylistic variation may be based

on individual speakers designing their performances for their audience, either real or imagined. The core of Bell's theory is that speakers form themselves into social groups which establish recognizable identities, familiar to themselves and others. As part of forming this identity, the social group differentiates its language from that of others through the exploitation of variation. As the larger community recognizes which language behaviors are associated with which groups, listeners may evaluate language used by others and assign social meaning to it. Because speakers in turn are aware of this, they "design their style primarily for and in response to their audience" (Bell 2001:143).

Bell's analysis crucially depends on the alignment between intraspeaker and interspeaker variation. In his assessment, intraspeaker style shifting derives directly from manipulations by the speaker made in response to potential audience evaluation, structured by interspeaker variation. Bell acknowledges the constitutive nature of certain kinds of style shifting, where speakers change their linguistic performance in order to change the situation, rather than responding to an external change. He labels this form of initiative **referee design** in which a social group which is not present in the interaction is invoked using linguistic means thus influencing the development of the interaction by being referenced within it. In this way he proposes a full account of stylistic variation based on the connection of linguistic resources to particular groups, while emphasizing the complexity of individual interactions.

Both of these approaches focus on the listener as an individual, influencing specific linguistic acts by being present in an interaction or referenced by a speaker. Another way to think about how speakers consider their audience is at a more general level. While speakers undoubtedly have models of actual interlocutors, they may also make linguistic choices on the basis of a broader ideological sense of how they are perceived. This is captured in the concept of a linguistic marketplace, put forward by Bourdieu (1982). This approach considers a setting for linguistic exchanges as a market, like a stock market or a vegetable market. Within this market, participants produce their linguistic "offerings" whose value is determined by a number of factors relative to qualities of the linguistic performances themselves as well as the vagaries of that particular setting. Different settings may assign value differently, for example

interactions within a university setting are likely to place greater positive value on the use of the standard language within a performance than might a factory, or a country music concert. As in other economies, the value of one's product determines and is determined by the amount one can get for it. For example, greater value might lead to a higher salary, greater social acceptance, more sales, or greater romantic success, to name a few. Because speakers will be, to varying degrees, aware of the market within which they are operating, they may be expected to craft their performances for maximum value, to the extent that they are capable of doing so. The concept of the marketplace allows us to understand one aspect of how large-scale patterns visible in interpretation feed back into those visible in production. Speakers' awareness of the expectations and beliefs of their listeners will lead them to construct their speech in such a way as to maximize their ability to achieve their particular goals in a given setting.

All three of these approaches speak to the crucial role of listeners in shaping the choices made by speakers. They, with the earlier points regarding the social and cognitive processes of listening itself, provide good reasons why the study of listener perceptions of socially linked variation is a useful endeavor. This provides some of the theoretical background and explanation for the purpose of this research. The following section describes the structure of the dissertation.

1.3 Organization of the chapters

This dissertation is intended to expand our knowledge of how listeners interpret linguistic variation and participate in the ongoing construction of social meaning, a process which shapes the face of linguistic variation. By understanding the factors which influence the interpretations that listeners construct, we can better understand the systems which connect small, idiosyncratic and context dependent decisions to large scale structures with broadly visible patterns. I explore the multiple and complex social meanings relating to articulateness, education, masculinity and regional difference which are all tied to (ING).

The research described here consists of an adapted version of the Matched Guise

Technique, using recordings of spontaneous speech digitally manipulated to create identical recordings differing only in (ING). Using these recordings, I first collected open ended qualitative data regarding listener perceptions of the speakers used in the recordings and of the influence of (ING) in particular. Using this data, I constructed the materials for a large-scale survey, which I carried out over the web, using students from two similarly prestigious but geographically distant schools. The quantitative and qualitative data show the paradox of broad scale relatively consistent and simple beliefs and responses regarding (ING), combined with intricate and curious individual differences based on the particular idiosyncrasies of the speech of the given speaker or the content of their utterance. I show that despite this seeming paradox, these idiosyncrasies are not exceptions to a larger norm, but rather the very behaviors which, in aggregate, make up these large-scale norms.

In Chapter 2, I survey the literature on (ING), highlighting the existing treatments of social meaning and social significance to find possible interpretations of the variable. The production literature has found a range of correlates for (ING), including socioeconomic status, situational formality, race, gender and educational background. Researchers have also found evidence linking (ING) to regional variation and specific performances of masculinity. I discuss all of these correlates and the suggestions they hold regarding the nature of (ING) and the possible meanings listeners might attach to it. After reviewing the literature, I describe a pilot study which investigated the effect of (ING) on listener ratings of speakers along the traditional axes of the matched guise formulations. This experiment demonstrated that it is possible to capture responses to linguistic variation using the matched guise technique. It also showed that that this approach, using the traditional metrics, does not provide sufficient complexity to truly understand listener behavior. I describe the methods used in the pilot and discuss the results. I then point out some of the shortcomings of this pilot study, setting the stage for the main study.

Chapter 3 lays out the methodology for my study. I review the history of the Matched Guise Technique and other work using listener responses to investigate linguistic variation. I also present the crucial methodological points of the current study,

drawing on this literature to support these choices. I then describe the actual procedures used in my study. The original recordings were created using sociolinguistic interviews. I used the software package Praat to manipulate tokens of (ING), creating minimal pairs of recordings. Both speakers and listeners were university students in the South and on the West Coast. I first conducted individual and group interviews eliciting general responses to the recordings, as well as specific reactions to the function of (ING). I then used the qualitative data to construct a survey instrument with ratings and adjective checklists and administered it over the web. The last section of this chapter describes the statistical procedures used in the analysis of the survey data.

Chapters 4, 5 and 6 discuss the results of the study. Chapter 4 takes the list of social correlates from the literature on (ING) which are discussed in Chapter 2 and investigates their relationship to (ING) in the results of the study. It shows that while (ING) impacts our range of social perceptions, it tends not to do so in a straightforward way. Only two responses were influenced by (ING) across all of the data: speakers were rated as more *educated*¹ when they used *-ing* and were also more likely to be described as *articulate* in that guise. Other results were more complex. This chapter steps through each correlate in turn and discusses the relevant results.

One fundamental influence on the role of (ING) was the regional background of the speaker. Chapter 5 explores the relationship of (ING) to region and to the concept of accent. It shows the striking importance of language ideologies in how listeners understand the role of (ING) in a specific situation. It also shows the flexibility of the concept of accent and how (ING) influences the perceived performance of different accents. Specifically, *-in* increases the perceived ratings of *accented* speech on the part of the Southern speakers, while decreasing how *accented* a different speaker sounded, one whose accent was associated with urban centers and alternative sexuality.

Chapter 6 moves away from speaker-related influences to the role of the listener in creating meaning. This chapter discusses different listeners' reactions to individual speakers and to the social constraints involved in sharing the explicit evaluations of

¹Descriptions in italics refer to responses on the matched guise survey, either checkbox or ratings or to exact quotes from respondents.

others. It shows that listeners with more or less positive reactions to a speaker may interpret a “social move” in very different ways. This chapter demonstrates that the limits on speaker agency are clearly provided by the agency of listeners who may assign a variety of meanings to a given performance, regardless of the original intent.

After going over the specific results of the study, I turn in the final chapter to the theoretical implications. Chapter 7 discusses the answers this research provides to the questions discussed in this chapter. It focuses particularly on the role of style as a cognitive tool, one which allows listeners to contextualize their linguistic and social expectations of a given speaker. After introducing the notion of style, I discuss the insight into its character given by the research presented here and describes some of the the many open questions regarding style and its role in the development of variation. Finally, I describe more generally the open questions are arising from this work.

