Chapter 1

Introduction

Since its inception, the field of language variation and change has made great progress, moving from the question "can sound change be observed?" (Labov et al., 1972, pg. 6) to the question of what observing sound change in progress can add to our theoretical understanding of language and how theoretical linguistics can add to our understanding of variation and change. In this dissertation, I analyze a phonological change in progress, with the goal of using this change to illuminate aspects of phonology that are most visible in an analysis of phonology in flux. In so doing, I also highlight the usefulness of a structural analysis of language variation and change.

Because of the difficult nature of observing phonological change in progress, most hypotheses regarding phonological change are drawn from a post-hoc analysis, with evidence of the language's phonology preceding the change and following the change but sparse or no data from speakers during the change. This set of facts results in necessary speculation about what individual speakers must have produced in order to cause change in a language. This speculation is by no fault of phonologists or sociolinguists: phonological change is difficult to observe in real time, since it occurs relatively infrequently in comparison to phonetic change, and because large scale corpora of speech are, relatively speaking, new sources of data. Add to this the fact that the study of sound change in progress itself is a young field, it is unsurprising that studies of real-time phonological change within individual speakers are rare. This logistical problem of capturing phonological change in real time is eloquently articulated by Hockett's discussion of the phonemic

restructuring of /æ/ and /ɔ/ in early Middle English (Hockett, 1958, pg. 456–457, emphasis mine):

> Sound change itself is constant and slow. A phonemic restructuring, on the other hand, must in a sense be absolutely sudden. No matter how gradual was the approach of early ME [(Middle English)] /æ/ and /ɔ/ towards each other, we cannot imagine the actual coalescence of the two other than as a sudden event: on such-and-such a day, for such-and such- a speaker or tiny group of speakers, the two fell together as /a/ and the whole system of stressed nuclei, for the particular idiolect or idiolects, was restructured. *Yet there is no reason to believe that we would ever be able to detect this kind of sudden event by direct observation.*

Hockett points out that an abrupt change in phonological specification is an event so sudden and so difficult to observe that the chances of analyzing it are vanishingly small. In this dissertation, I attempt to do just that. Taking advantage of the large-scale and relatively new Philadelphia Neighborhood Corpus (PNC), I identify a phonological restructuring currently in progress in Philadelphia English /æ/. Using large-scale corpora as well as targeted interviews with the speakers most likely to be undergoing phonological change, combined with social evaluation experiments and a computational simulation of change, I attempt to provide a holistic sociophonological account of this allophonic restructuring.

I begin with a deceptively simple question: When phonological change occurs within a speech community, how do individual speakers contribute to that change? While different theories of phonology and phonological change make different specific predictions about the empirical outputs of individual speakers, it is only recently that our data sources have grown large enough to address this question for sound change; this dissertation represents one of the first large-scale investigations into phonological change in real time. The central drive of this project – determining how individual speakers drive community-wide phonological change – has in turn spawned its own related questions, which are the focus of Chapters 3, 5, and 6.

In §1.1, I outline the minimal theoretical assumptions necessary for my driving question. In §1.2 I describe the three primary theories of phonological change and the predicted outputs of

these mechanisms of change for individual speakers. In §1.5, I provide an outline of the chapters in this dissertation.

1.1 Phonological Change

The broad purpose of this dissertation is to investigate how individual speakers' productions drive community-level phonological change. This granularity of investigation represents somewhat of a break from tradition in quantitative sociolinguistics. The empirical study of language change originated as a concept that exists on the level of the community rather than on the level of the individual, as articulated in Labov et al. (1972):

> The general position that we have taken is that no useful distinction can be made between a change and its propagation (Weinriech et al., 1968) as long as we continue to consider language an instrument of communication. The language does not change if one man invents an odd form or develops an idiosyncrasy, even if people understand and evaluate his behavior; it does change when others adopt his idiosyncrasy and use it as a new social convention for communicating their intent.

Historically, the program of analyzing language change has taken as its primary focus the pattern of the speech community as a whole, as it is at this level that the language can be most clearly said to exist and change. Nevertheless, when a language or a dialect undergoes a change, it is through the individual speakers who produce language. Herein lies an apparent contradiction: while the sometimes idiosyncratic and non-prototypical language produced by an individual is not the same as language change, any change in the community is itself made of individuals producing a difference in their own language from that of the previous generation. In the decades since Labov et al. (1972) asked whether sound change can be observed, sociolinguists have documented many sound changes occurring in different speech communities in different languages in real time. As a field, we know quite a bit about how language change works on the level of the community, but not as much about how individual speakers drive that change along. Given the decades of work on how language change operates on the level of the community, we can now turn to the question

of how the production of individual speakers works in aggregate to produce the community-level change, which is the goal of the current dissertation.

1.1.1 Modular, Feed-forward Separation of Phonology and Phonetics

Throughout this dissertation, I assume a modular, feed-forward architecture of phonological grammar, following terminology in Pierrehumbert (2002) (see also Bermúdez-Otero, 2007). The basic modular architecture is as shown in Figure 1.1: lexical representation is stored with underlying categorical phonological representation. For example, *mitten* is stored as a lexical entry with the underlying categorical phonemic representation of /ˈmɪtɛn/. This underlying representation then undergoes abstract phonological rules, which are also categorical in nature. Our example *mitten* would undergo /t/ allophony, producing a surface phonological form /ˈmɪʔn̩/ for many American English speakers. At this point, the lexical entry has two categorical phonological aspects: (1) the underlying representation and (2) the abstract rules that result in the surface representation. From this surface representation, forms then undergo gradient phonetic processes to finally produce a phonetic output. The modular aspect of this model separates each process into a distinct level, while the feed-forward aspect means that each stage can only "see" what was given to it by the previous stage; a phonetic process can only make reference to the surface phonological representation it has been fed – it cannot make reference to any underlying representations.

A number of variations on this main architecture have been proposed (see, e.g., Keating, 1985, 1990; Cohn, 1993). Here, I adopt a stratal version of this architecture (Bermúdez-Otero, 2007), as shown in 1.2, which breaks phonology into a stem-level, word-level, and phrase-level module. The underlying phonological representation then may undergo phonological processes at each of these levels, resulting in a phrase-level surface phonological representation that gets fed into the phonetics modules. Under this variation, there are four targets for phonological change: (1) the underlying phonemic representation, (2) abstract phonological rules which produce a stem level representation, (3) abstract phonological rules which produce a word level representation, or (4) abstract phonological rules which produce a phrase level representation. Notably, sociolinguists have often found that the abstract rules applying to each of these phonological levels are the same.

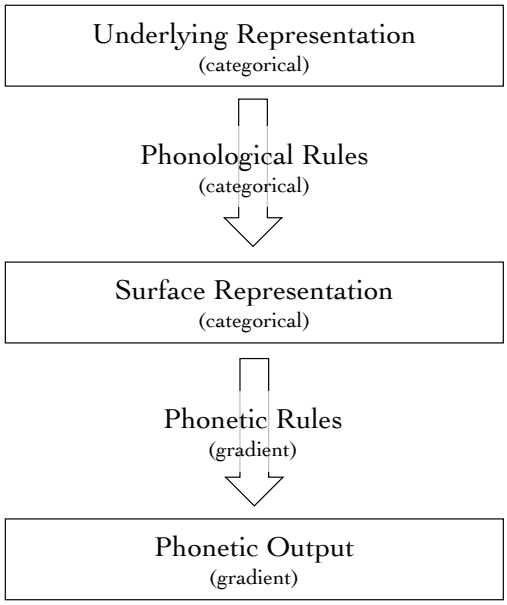

Figure 1.1: Modular, feed-forward phonology-phonetics interface.

In other words, a single rule may be repeated at each of these levels (see, e.g. Bailey, 2017, on /g/-retention in Mancunian English). However, because there are some processes which *only* apply at stem-level (e.g., the Scottish Vowel Length Rule Aitken, 1981) or at phrase-level (such as prosody), this must be representationally possible in the architecture.

The phonological modules in Figure 1.2 are boxed; any change occurring within one of the boxed modules constitutes phonological change. I note briefly that the phonetic components of the architecture are severely underdeveloped in the representation in Figure 1.2; this is fleshed out in several variations (Keating, 1990; Cohn, 1990), and often include distinct modules for language-specific phonetic processes and universal phonetic and articulatory processes. My exclusion of a more detailed phonetic framework in Figure 1.2 is not a theoretical stance, but rather intended to focus the dissertation on the levels of the architecture directly related to phonology. This architecture remains somewhat theory-neutral with regards to the formal specification of phonological processes. These specifications can be formalized under any theory of phonology compatible with

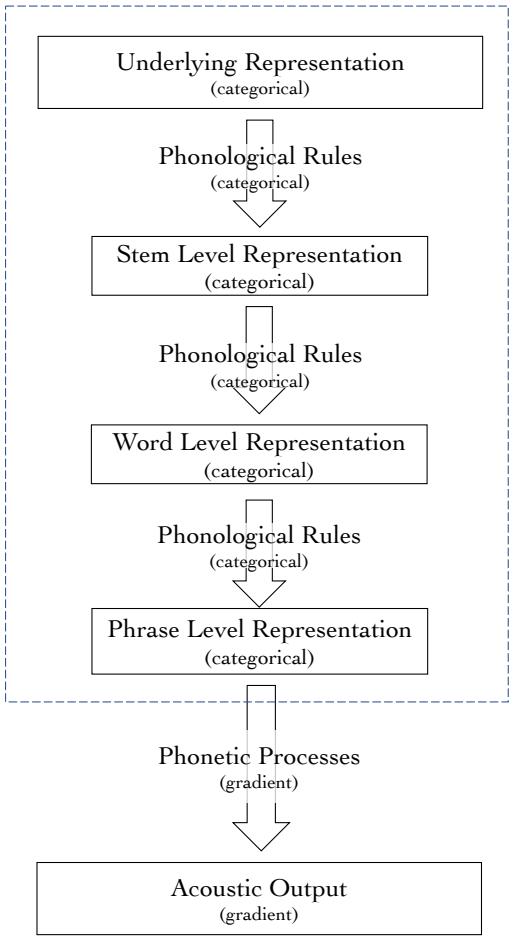

Figure 1.2: Modular, feed-forward phonology-phonetics interface with stratal phonology.

both underlying representations and categorical phonological processes.

Throughout the dissertation, I will refer to phonological processes as "rules," broadly adopting a broadly Generative Phonology framework. This terminology is not a theoretical stance; the phonological processes that take an underlying representation (like /t/) to the surface level phonological representation (like /ɾ/) can also be represented using most varieties of an Optimality Theoretic (Prince and Smolensky, 1993) as well as any version of an Exemplar Theory framework (Bybee, 2002; Pierrehumbert, 2001) that allows for categorical underlying specification which also undergo categorical processes (whether these features and processes are innate or emergent). Since I find rule-based notation easier to discuss, this is the terminology I adopt throughout the dissertation.

Defining Phonological Change

I consider phonological change to be any change to the phonological modules; this means either a change to (a) the underlying representation or a change to (b) any of the rules that produce a surface level representation.

1.2 Mechanism of Phonological Change

While the mechanism of phonological change is difficult to test in real time, there are three primary theories of how individual speakers contribute to community-level phonological change, which will be the focus of my dissertation. Here, I outline these three theories, the factors that govern them, and how they may be identified in the production of individual speakers.

1.2.1 Phonetic Incrementation

There is, to some degree, a level of conventional wisdom shared across a number of phonological frameworks which places the mechanism of phonological change on accruing errors in production or perception. This is the view espoused in Ohala (1981), which lays out a clear argument for the human body, rather than human cognition or abstract linguistic knowledge, as the locus of linguistic change. This is show in Figure 1.3, which provides a schematized illustration of a potential

phonological change from /ut/ to /yt/.

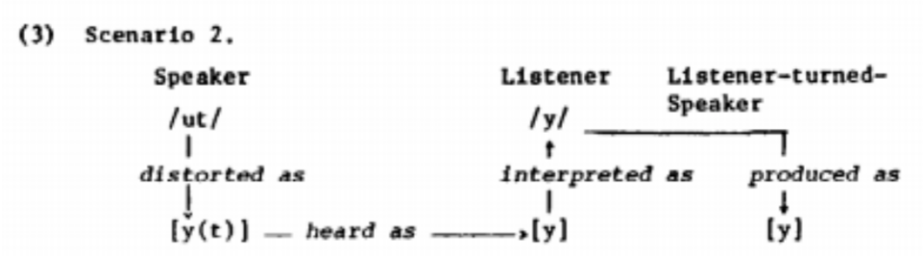

Figure 1.3: Accruing errors as the source of sound change. From Ohala (1981).

Ohala (1981) outlines a number of historical changes which can be accounted for by a perceptual bias of the surrounding phonetic environments that originally triggered such a change. This mechanism of sound change, however, still remains underspecified in terms of abstract linguistic properties. In Figure 1.3, the listener's failure to accurately account for the effects of coarticulation are shown in phonetic terms: the listener at this point has simply shifted their phonetic interpretation of the speaker's phonological content. At this point, phonological change as defined above cannot be said to have taken place. Furthermore, Ohala (1981) does not specify what the tipping point for phonetic incrementation turning into phonological change may be. Despite a lack of explicit specification of how or when this mechanism of sound change affects the abstract segments or rules, the mechanism of phonetic incrementation remains a possible driving force for phonological change; in the most general terms, this means that phonetic or perceptual processes drive sound change until it becomes phonologized either in the middle or at the end of the change (Kiparsky, 2015).

Phonological change via phonetic incrementation is also at the heart of many Exemplar Theoretic accounts of sound change (Bybee, 2002; Pierrehumbert, 2001; Hay et al., 2015). Here I set aside versions of Exemplar Theory that reject the notion of cohesive exemplar clouds altogether (e.g. Bybee and McClelland, 2005), and use the term Exemplar Theory to denote those frameworks that include some level of cohesive phonological identity, which in practice function as phonemes. Under this type of framework, the driving force of a sound change is also placed on listener misperception; here, the variation in the speech signal is caused both by physical reductive processes,

such as increased coarticulation and decreased duration, as well as more abstract cognitive reductive processes such as decreased lexical retrieval found in more frequent words (Grainger, 1990; Goldinger, 1998).

While the specific motivation for misperception varies by framework, the crucial driving force for phonological change in both cases is that some level of phonetic misperception accrues, which at some point results in a phonological change.

Phonetic Mitigation

It is worth briefly drawing attention to the difference between phonetic incrementation and what I term *phonetic mitigation*. Phonetic mitigation here refers to a process by which speakers change their phonetic production in response to social stigmatization. Speakers are often found to produce unsystematic phonetic mitigation of stigmatized forms, particularly in settings that are more formal or induce higher attention paid to speech (Labov, 1989, 2001). The crucial distinction between phonetic incrementation and phonetic mitigation for the purposes of this dissertation is in the community-based outcome of the acoustic output: while the acoustic production of phonetic mitigation may look very much like the production of a speaker during phonetic incrementation, the main distinction between the two is in whether or not that output *drives* sound change in the community. While phonetic incrementation drives phonological change in the community, phonetic mitigation is a response to change or evaluation from the community.

To determine whether a speaker's production is phonetic mitigation or phonetic incrementation, a speaker's social environment and peer sociophonological production must also be taken into account. If we find phonetically mitigated output in a number of speakers in a subset of a speech community where the cohort of speakers older than them produce unambiguously non-mitigated tokens and the younger cohort of speakers produce a phonological change, we can conclude that sound change via phonetic incrementation has taken place. If, on the other hand, we find phonetically mitigated output in a speaker whose subset of the speech community already produces the new phonology, we can conclude that the phonetic mitigation of the outlier speaker is not *driving* sound change but rather is the socially motivated response to a change that has already happened.

1.2.2 Spontaneous Phonologization

The second theory of phonological change provides a dramatic foil to change via phonetic incrementation. As argued by Janda and Joseph (2003), this "Big Bang" mechanism of phonological change places the phonologization at a very early stage of the change, wherein speakers innovate phonological and sociolinguistic conditions on a pre-existing (but brief in timespan) phonetic condition. This is taken up more strongly in Fruehwald (2013), who argues that phonologization may occur even *before* any perceptible phonetic conditioning has occurred. This spontaneous phonologization, if independently innovated by enough speakers in a speech community, would then be able to acquire phonetic correlates of the already existing phonological innovation and become a sound change on the level of the community (Ringe and Eska, 2013).

In considering the mechanism of community-wide change, it is important to differentiate between spontaneous phonologization as the solution to the Actuation Problem (reproduced in (3)) and spontaneous phonologization as the solution to the Transition Problem (reproduced in (4), both from Weinriech et al. 1968)

(3) **Actuation Problem**: What factors can account for the actuation of changes? Why do changes in a structural feature take place in a particular language at a given time, but not in other languages with the same feature, or in the same language at other times?

(4) **Transition Problem**: [...] the intervening stage which defines the path by which Structure *A* evolved into Structure *B*

These problems can be thought of as the split between an individual change and a change on the level of the community. The *actuation* of a change asks what causes a change to be innovated by individual speakers. The *transition* of a change asks by what path does structural change then become propagated throughout the community. As a solution to the Actuation Problem, the mechanism of spontaneous phonologization defines how speakers may come to posit idiosyncratic structural changes, and it is largely in this vein that Janda and Joseph (2003) and Fruehwald (2013) discuss spontaneous phonologization. This does not prohibit speakers from also spontaneously positing multiple structural analyses for their input data, which may in fact be a critical aspect of

the transition mechanism of competing grammars, which is discussed below. Here, I use the term *spontaneous phonologization* to describe the profile of Transition via spontaneous phonologization, and remain agnostic as to the Actuation of phonological change.

Under community level change via spontaneous phonologization, individual speakers posit a single phonological system and produce that system throughout their speech. In the beginning of the change, very few speakers in a given age cohort will have posited the change, but as time goes on, more speakers in each age cohort will produce the new system rather than the old system. As a mechanism of community-wide change, this predicts that what may look on a community scale to be intermediate productions between System A and System B is actually the result of some speakers producing A and some speakers producing B.

1.2.3 Competing Grammars

The third mechanism of phonological change is an adaptation of syntactic grammar competition to phonology. Grammar competition accounts for the optionality that arises when mutually exclusive parameter settings coexist within the grammar of a single speaker, as in Kroch (1989). While competing grammars grew out of analysis of syntactic change, here I apply this concept to phonological change as well. Under a competing grammars framework, the structured optionality found within each speaker results straightforwardly from variation in a single abstract parameter, providing empirical support for a theory of generative syntax with abstract functional heads. Kroch (1989) demonstrates abstract competition between two variants of a parameter for a number of changes crosslinguistically, including the replacement of *have* by *have got* in British English, the rise of the definite article in Portuguese possessive noun phrases, the loss of verb-second word order in French, and the rise of English periphrastic *do*.

The rise of periphrastic *do* in English provides strong support for a theory of syntactic change through competing grammars, partially due to the large amount of data and partially because analyzing this change as competition in an abstract syntactic parameter provides an explanatory account for a number of distinct surface phenomena which can be best be explained as underlying variation between an abstract syntactic parameter (Kroch, 1989; Pintzuk, 1996).

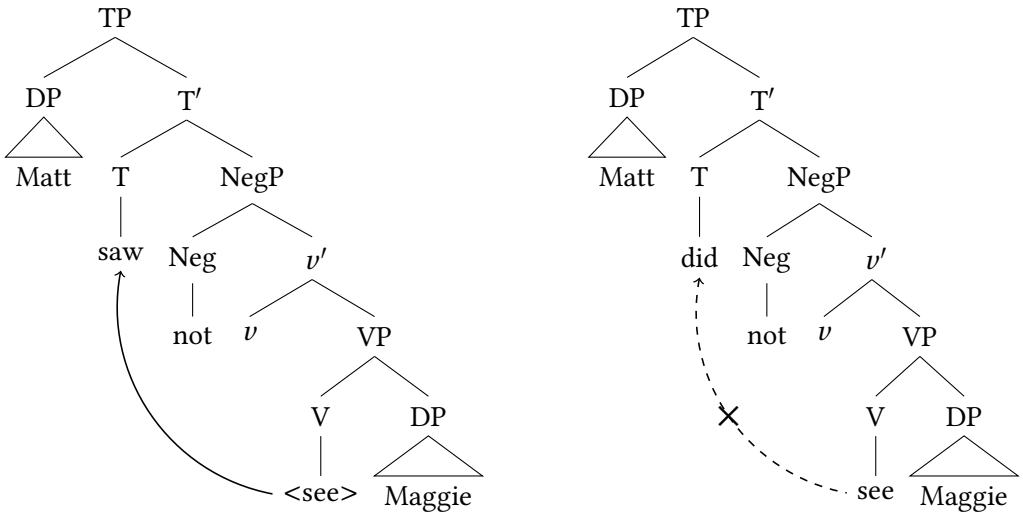

(a) Middle English V-to-T Raising.

(b) Early Modern English V-to-T blocked.

Figure 1.4: V-to-T parameter resulting in DO-support.

The structural analysis of periphrastic *do* in English is analyzed as a consequence of the loss of V-to-T raising (see Figure 1.4b) in English. This abstract structural parameter can be most clearly seen in contexts with an intervening element, such as negation or subject-auxiliary inversion. The evidence for *do*-support arising from the loss of the abstract verb raising parameter in English is also supported by what appears on the surface to be unrelated changes. If verb raising is lost in English, this makes specific predictions about the placement of adverbial forms like *never*. In Modern English, *never* precedes finite verbs (as in *I never found that article*); a pattern that falls out straightforwardly from the loss of V-to-T raising. In a diachronic analysis, Kroch (1989) finds all contexts of V-to-T raising exhibiting the same rate of change (referred to as the Constant Rate Hypothesis), which stands in contrast to the previously received conventional wisdom that syntactic change proceeds context by context.

The concept of competing grammars has, to some extent, been present in the study of phonological change from the beginning of modern sociolinguistics. Empirical Foundations for a Theory of Language Change (Weinriech et al., 1968, pg. 184), describes the transition problem as occurring through speakers with *heterogenous systems*:

Context	Old	New
Intervening Negation	John saw not the cat	John didn't see the cat
Subject-Auxiliary Inversion	Went he to the store?	Did he go to the store?
Subject-Auxiliary Inversion	Where went Matt?	Where did Matt go?
Intervening Adverb	He eats always broccoli	He always eats broccoli

Table 1.1: Some contexts exhibiting differences between V-to-T raising and the loss of V-to-T in English.

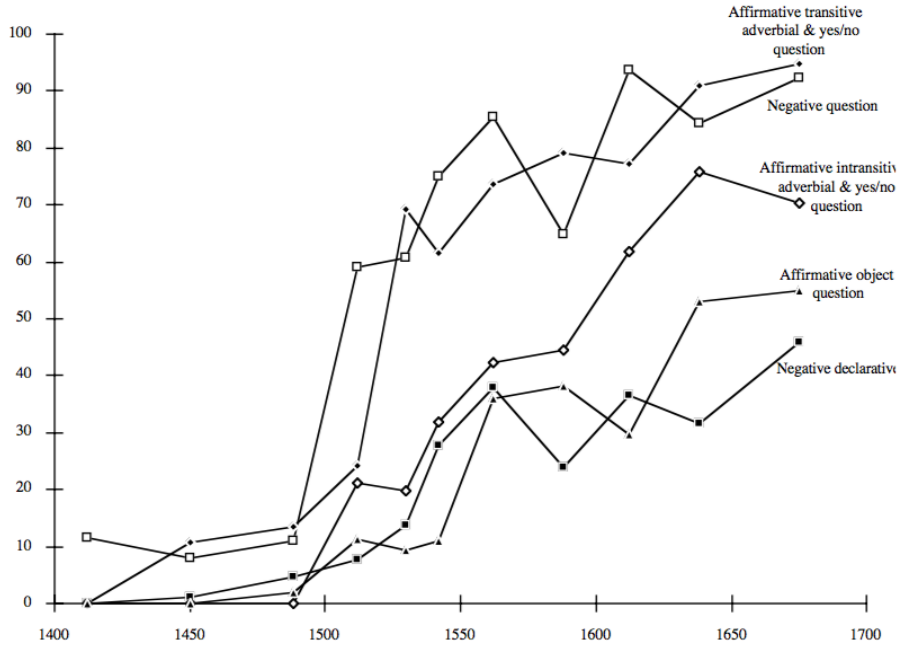

Figure 1.5: V-to-T loss increasing at the same rate across all syntactic contexts. From Kroch (1989).

> This transition or transfer of features from one speaker to another appears to take place through the medium of bidialectal speakers, or more generally, speakers with heterogenous systems characterized by orderly differentiation. Change takes place (1) as a speaker learns an alternate form, (2) during the time that the two forms exist in contact within his competence, and (3) when one of the forms becomes obsolete.

Using the framework of competing grammars more specifically, it becomes possible to make additional predictions about the time when two forms exist within a single speaker's competence. This has been done explicitly by Fruehwald et al. (2013), in an investigation of stop fortition in Middle High German. Using two corpora of written Early New High German, Fruehwald et al. (2013) find evidence for intraspeaker variation between a stop-fortition grammar and a non-stop-fortition grammar, which exhibits a Constant Rate Effect across all potential application contexts. In general terms, applying competing grammars to phonology as a mechanism of phonological change hypothesizes that variation on the level of the community may be the result of individual speakers exhibiting optionality between two options of a single abstract parameter.

Competing Grammars as a Single Parameter

In both syntactic change as well as phonological change, we conceive of the locus of variation being a single abstract parameter that governs surface-level output. Here, an example will be useful. Take, for example, the merger of the vowels in LOT and THOUGHT to LOT which is spreading geographically across the U.S. (Labov et al., 2006) as an example of phonological change to the underlying phonemic representation. A competing grammars mechanism of this change would consider there to be an abstract parameter governing the selection of LOT and THOUGHT for canonical THOUGHT words; within an individual speaker, this parameter would probabilistically select the LOT (merged) phoneme or the THOUGHT (unmerged) phoneme each time the speaker goes to produce a word. Different phonological contexts, such as following or preceding segment, are encapsulated under this single parameter. While the rate of usage across these contexts may differ, a competing grammars analysis requires that these contexts still exhibit the same rate of change, following the Constant Rate Hypothesis (see, e.g. Fruehwald, 2013, for an account of

phonological change analyzed using this Constant Rate Hypothesis).

That many contexts are classified under a single parameter which is realized in two competing ways is particularly important when the object of consideration encompasses multiple discrete contexts, as in the case of the phonological change I focus on in this dissertation. The use of *grammar* here in place of *parameter* has occasionally been the source of confusion for readers who are not operating under a Chomskyan theory of syntax, as it may be read to imply that the object under competition is a speaker's entire linguistic competency rather than a single parameter. I highlight here that the term *grammar* in the context of Kroch (1989) is drawn from a Principles and Parameters or Minimalist framework (Chomsky and Lasnik, 2008; Chomsky, 1995), in which syntactic items – both lexical and functional head – are selected by a merge function. In the case of syntactic change, *merge* has the option of selecting between two functional heads. Under this framework, the term *grammar* refers to the objects that are selectable by *merge* and not to a complete description of linguistic competency. In this dissertation, I use the terms *grammar* and *parameter* interchangeably.

Similarly, I refer at times to the allophony of /æ/ under investigation here as an *allophonic system* as well as an *allophonic rule*. As I argue in Chapter 3, any allophonic rule also includes any lexical exceptions to that rule, meaning that *system* and *rule* are synonymous, both referring to one of the two parameters in competition.

1.3 Transition Cohort Speakers

Finally, here I briefly define the target research population of this dissertation, which is the *Transition Cohort Speakers*. I've defined phonological change as a difference in phonology between older speakers and younger speakers within a given speech community. In the time period before any change, every speaker in the community produces the old phonology; after the change is completed, every speaker produces the new phonology. It is the speakers acquiring language in between these two time periods who are of the most interest to the mechanism of phonological change. The phonetic outputs of these transitional cohort speakers are what, in the aggregate, produce the overall community shift. The primary question in this dissertation is whether the

transition cohort speakers produce a community-level phonological change via phonetic incrementation, spontaneous phonologization, or competing grammars.

1.4 Disambiguating Evidence

While the three proposed mechanisms of phonological change result in clearly distinct trajectories of a change, it is not necessarily straightforward to disambiguate between the three mechanisms by the production of a single speaker. In this section, I discuss some of the evidence that must be drawn on in order to disambiguate potentially ambiguous data.

1.4.1 Phonetic Evidence for Competing Phonological Parameters

It is occasionally assumed that phonological competing parameters will manifest in a phonetically obvious manner (see, e.g. Dinkin and Dodsworth, 2017). Unfortunately, this is not the case. Phonological change occurring via a mechanism of competing grammars refers *only* to variation in the abstract linguistic parameters. Assuming a modular and feed-forward model of phonology, as I do here (and in fact, as do Dinkin and Dodsworth 2017) means that the phonetic manifestation of the phonological input is not within the domain of phonology. It is possible, in other words, for competing parameters to be active in a speaker's cognitive representation of the language without that competition resulting in an easily measurable output. It could even be active without any difference at all in output, in a situation where the phonetics interpret two distinct surface representation as having the same phonetic output. Setting aside this case, which results in a theoretical distinction without an empirical difference and is therefore a moot point, it remains that phonological competing grammars may not be easily discrete. This is particularly true for phonological mergers and splits, which although produce structurally radical differences, may not be easily identifiable in the phonetic implementation of those abstract differences.

The main point here is that a theory of competing grammars makes no assumptions about the phonetic output of those competing grammars. Of course, a grammar competition that is completely imperceptible to other speakers will not last beyond the speaker(s) who innovated that change. Phonetically distinct but similar outputs, on the other hand, may require an extremely

large data set to analyze the underlying mechanism of change. One potential method of distinguishing between competing grammars and phonetic incrementation in a case where the phonetic targets of the two parameters are similar is in the expected standard deviations for conditioning factors under each theory. In general, we would expect change via competing grammars to exhibit higher standard deviations for each conditioning factor (because speakers are actually producing two targets) than we would expect for change via phonetic incrementation (where speakers produce only one target per conditioning factor). Unfortunately, the amount of data required to make a strong distributional case for competing grammars is out of reach for most phonological variables in current sociolinguistic corpora. While advances in recording technology are making it easier to obtain relatively large-scale data sets from speakers, the sheer volume of data needed to distinguish the significance of standard deviations of phonetically similar outputs is, at this point, prohibitive.

Fortunately, phonological mergers or splits are not the only type of phonological change that can be investigated. In this dissertation, I analyze the mechanism of phonological change for an allophonic restructuring in Philadelphia English. The nature of this restructuring means that both the old system (which I call PHL) and the new system (which I call NAS) produce outputs that are phonetically distinct. This means that the amount of data required to identify a competing grammars speaker is relatively small, compared to a merger or a split, making it opportune for investigating the mechanism of phonological change.

1.4.2 Social Evidence for Spontaneous Phonologization

In first-wave sociolinguistics, a speech community is generally thought of as a relatively monolithic entity exhibiting an "enigma of uniformity" (Labov, 2009). And in fact, generally speaking, the level of uniformity in both production of and evaluation of features found across millions of speakers in a single speech community is difficult to explain given speakers' lack of contact with the entirety of their speech community. Layered above this backbone of general uniformity, however, smaller communities of practice (Wenger, 1998; Eckert and McConnell-Ginet, 1992) introduce local socially defined loci of linguistic variation that is often itself socially meaningful in nature.

With change rather than variation in mind, these local socially meaningful community of practice divisions in a broader speech community raise the possibility that phonological change may be introduced or innovated differently across communities of practice within the broader community. This is particularly of importance in trying to determine whether phonological change has occurred via spontaneous phonologization or not. In spontaneous phonologization, speakers produce either the old phonology or the new phonology. However, this output is also consistent with the beginning and end stages of phonetic incrementation and competing grammars. If change has occurred via phonetic incrementation or competing grammars but is affecting different subsets – or communities of practice – within the larger speech community at different times, the output of these speakers as a whole will show some speakers with the old system and some speakers with the new system. Taking only speakers' phonetic outputs into account will not allow us to disambiguate between different mechanisms of phonological change. Instead, the social divisions within a larger speech community must also be taken into account; if all speakers within a subset of the community produce only one system, this suggests that community of practice is not in flux and has either not undergone the change or has already completed the change. If, on the other hand, some speakers within a single community of practice produce the old system and some produce the new system, this suggests change via spontaneous phonologization.

It will therefore be necessary to obtain information on the relevant social divisions within a broader speech community in order to disambiguate whether phonological change has occurred via spontaneous phonologization or another mechanism of change.

1.5 Roadmap

In this chapter, my goal has been to outline the motivating theoretical question of this dissertation and the minimal theoretical assumptions I make. As highlighted in §1.4 above, a full investigation of the mechanism of phonological change must bring social, phonological, and phonetic evidence to bear, which is what I aim to do in this dissertation. The dissertation is organized as follows.

In Chapter 2, I describe the phonological change that serves as the case study in this dissertation, which is the allophonic restructuring of /æ/ in Philadelphia English. I outline the community-

level pattern of this change, specifically highlighting the meaningful social divisions produced by the educational system in Philadelphia. I argue that the social divisions produced by the educational system results in communities of practice that either promote the change (in the case of Special Admissions non-Catholic schools) or inhibit the change (in the case of Open Admissions Catholic schools). A bipartite network diagram visualizes the distinct fragmentation in Philadelphia's school system and the subsequent linguistic consequences. Chapter 2 also presents an analysis of the intergenerational pattern of change, finding that the allophonic restructuring of /æ/ occurs in three stages.

Because the allophonic status of /æ/ in traditional Philadelphia English has often been the topic of phonological debate (e.g., Ferguson, 1972; Labov, 1989; Kiparsky, 1995; Dinkin, 2013; Labov et al., 2016), I devote some considerable space in Chapter 3 to a theoretical account of traditional Philadelphia /æ/ as a productive allophonic rule with limited lexical specificity. I propose more generally in Chapter 3 that productive phonological rules, much like productive morphological rules, can tolerate a limited number of lexical exceptions. I specifically appeal to the Tolerance Principle formula from Yang (2016) as a way to define the upper limit of lexical exceptions that a productive process may tolerate. This solution provides a resolution for a number of phonological relationships that have been set aside as troubling or puzzled over as intermediate between phonemic and allophonic under the classic definitions of contrastiveness, without needing to add any additional categories such as *quasi-phonemes* or *fuzzy contrasts* to the phonological architecture.

In Chapter 4, which provides the main evidence for the mechanism of phonological change for the allophonic restructuring of /æ/ in Philadelphia, I take a close look at the speech of transitional cohort speakers to determine which mechanisms of change are at play. I find evidence that the change in /æ/ is occurring via competing grammars in Philadelphia, suggesting that phonological change and syntactic change proceed in the same manner. In this chapter, I also present evidence that the lexical exceptions discussed in Chapter 3 participate in the intraspeaker variation, supporting the claim in Chapter 3 that lexical exceptions are in fact stored as part of the productive phonological rule.

The findings in Chapter 4 suggest the existence of a single parameter governing the choice of

allophonic system: in Chapter 5, I investigate whether this abstract parameter may be the target of social evaluation. Using a Matched Guise task, I find Philadelphian participants rating a speaker with the old /æ/ system as more *accented* than a speaker with the new /æ/ system. I follow this with a modified Magnitude Estimation task, which finds Philadelphians evaluating the pronunciation of /æ/ under different conditioning factors in a surprisingly systematic (rather than phonetic) way. My results suggest that not only are speakers able to socially evaluate phonological structure, but that an investigation of evaluation during a period of phonological change may reveal an underlying abstract reason for apparent surface-level results.

Finally, in Chapter 6, I tackle the question of the inevitability of this change, asking whether the complex traditional /æ/ system was destined to be replaced by the simpler, surface-true nasal /æ/ system. Using a computational simulation of acquisition given mixed input, I find that Philadelphian children could not plausibly produce this change through a reanalysis of their input and that instead it is most likely through dialect contact with outside speakers that the nasal /æ/ system entered the Philadelphian speech community.

In Chapter 7, I provide some concluding remarks and directions for future research.