Chapter 6

The Inevitability of Phonological Change

Throughout this dissertation, my main objective has been to identify the mechanism of phonological change for the allophonic restructuring of /æ/ in Philadelphia. I've argued that this phonological change occurs via intraspeaker grammar competition between the abstract parameters of PHL and NAS, and furthermore that these abstract parameters are the subject of social evaluation. In this chapter, I turn to the question of how inevitable this change is. Many frameworks of phonology take articulatory and cognitive simplification to be a motivating factor for sound change. The allophonic restructuring from the phonologically complex PHL system to a simple surface-true NAS system appears on the surface to be a confirmation of the inevitability of phonological change via simplification. In this chapter, I conduct a computational simulation based on the Tolerance Principle to investigate whether this change was the result of an inevitable simplification. The work presented in this chapter is a slightly modified version of a collaboration with Josef Fruehwald and Charles Yang, which is currently under review.

6.1 The Role of Simplification in Sound Change

Simplification, whether cognitive, phonological or articulatory, is often appealed to as a major motivation for sound change. This notion can be found in a number of different theoretical frameworks, from European structuralism to generative phonology. While an appeal to simplicity is often considered intuitive, a definition of simplicity depends on the framework and what the target of simplicity is.

First, to phonological simplicity. The specifics of simplicity depend on the framework under consideration, but the primary cohesive factor is the idea that marked or dispreferred forms and systems are more "cognitively complex" (Givon, 1991) and therefore more susceptible to change. Cognitive complexity is, in itself, a somewhat slippery term to define. Writing in the functionalist tradition, Martinet (1952) appeals to the notion of structural harmony as a motivating factor in the history of a language. Here, structural harmony refers specifically to a linguistic, or phonemic, inventory that is maximally symmetrical and makes use of a limited number of active features, resulting in a cognitively efficient system. This idea is echoed in Feature Economy (Clements, 2003), in which the simpler systems are those that maximize the ratio of sounds to features. Under a featural phonological framework, a simpler system with simpler forms would be defined as a system needing fewer features to encode it than a complex system. The specifics of feature simplicity depend further on the framework involved, with Feature Geometry (Clements, 1985) and Contrastive Hierarchy (Dresher, 2011) providing a hierarchical account of active features, Government Phonology (Kaye et al., 1985) relying on nonlinear representations and classical Generative Phonology Chomsky and Halle (1968) calculating simplicity through binary feature bundles, just to name a few. Regardless of specific definition of feature simplicity, however, there is a shared notion across these frameworks that simplicity is a driving force in phonology.

If simplicity as measured by cognitive complexity is a main driving force in language change, we may be tempted to echo the question articulated by Martinet (1952): "How is it that after so many millennia of uninterrupted speech practice, patterns should still be in need of structural integration?" In other words, why, after so many thousands of years of speaking, have languages not settled on a maximally cognitively efficient system? Why do they continue to change?

One potential answer to this lies in the physical facts of using language. A cognitively perfect system must still pass through human articulators, whether oral or manual. This interface introduces a second type of simplicity which has been thought to have an effect on language change, namely, articulatory ease. The role of articulatory ease can be found hand-in-hand with cognitive complexity in nearly every framework: Martinet (1952) refers to the strain of physiology as a "germ of instability" within a linguistic system. While some markedness constraints in Optimality Theory refer to cognitive complexity, other markedness constraints refer to articulatory ease (see Haspelmath, 2005), whereby processes like coarticulation and consonant cluster reduction which may initially occur due to articulatory ease become phonologically encoded into the underlying system. An Exemplar Theoretic account simultaneously appeals to ease of articulation and ease of cognitive recall: developed from the observation that high frequency words exhibit reductive processes in production (Bybee, 1999) as well as faster recall (Segui et al., 1982; Grainger, 1990); many proponents of Exemplar Theory suggest that high frequency tokens will likewise exhibit distinct profiles of change (e.g. Hay et al., 2015). Blevins (2006) exemplifies this view of language, arguing that human perceptual and articulatory biases are the source of many of the phonological patterns found in languages today. The proliferation of framework-specific considerations outlined here highlight how the specific predictions of simplification-motivated sound change will depend on the framework used to define simplicity. Regardless of framework, however, the implicit notion is that simplicity in form and system will be preferable to speakers, and that given the choice between two plausible representations, speakers will select the simpler choice.

The change from PHL to NAS in Philadelphia English seems, on the surface, to be a case study in support of simplification as a driving factor in sound change. While the specific definition of complexity is framework-dependent, it is uncontroversial under any framework to state that PHL, with its disjoint set of phonological triggers and syllable structure references and lexical specificity, is simpler than NAS, a surface-true allophonic rule with little complexity. In this chapter, I delve into this question in detail, asking whether this change from PHL to NAS was the inevitable result of simplification.

Here, we again make use of the Tolerance Principle (Yang, 2016) as a method of diagnosing whether a proposed phonological rule would be plausible for a language learner. We apply the Tolerance Principle to the allophonic restructuring in /æ/, to investigate the likely route by which NAS is supplanting PHL in the community. In §6.2.1, we find that a child receiving entirely traditional input could not plausibly posit NAS as a productive rule. In §6.2.2, we further find that the change from PHL to NAS is unlikely to be the result of children positing incrementally simpler changes in PHL, removing a conditioning factor at a time until the speech community is left with NAS. Finally, in §6.3 we turn to the possibility that Philadelphian children have acquired NAS through receiving mixed input from both PHL and NAS speakers, in a situation of dialect contact.

6.2 Could Children have Endogenously Postulated NAS?

Recall that under a featural rule-based framework, PHL is described as in (25)[11]: tense before tautosyllabic anterior nasals and voiceless fricatives. In addition to a rule with relatively complex conditioning, PHL also requires speakers to memorize a list of lexical exceptions, as outlined in Chapter 3. In contrast to this, NAS (shown in 26) is a simple allophonic rule comprised of a single conditioning factor which typically requires no lexical exceptions.

(25) **PHL:** $\text{æ} \rightarrow \text{æh} / \_[\ \text{+anterior}\ ] \cap ([\ \text{+nasal}\ ] \cup \begin{bmatrix} \text{-voice} \\ \text{+fricative} \end{bmatrix})]\ \sigma$

(26) **NAS:** $\text{æ} \rightarrow \text{æh} / \_[\ \text{+nasal}\ ]$

Given that NAS is a surface-true generalized rule where PHL produces surface exceptions, one possibility to be addressed is whether Philadelphian children are spontaneously simplifying their traditional input into the new NAS system. In other words, a Philadelphia child, perhaps at some transient stage of language acquisition, may have postulated a NAS system despite receiving consistent PHL input: as we have seen, a statistical majority of the lexical items produced under the PHL system is in fact compatible with the NAS system, and children's tendency of regularizing inconsistent input to form a majority rule is well documented in naturalistic acquisition (Singh et al., 2004) and in artificial language learning experiments (Hudson Kam and Newport, 2009, 2005). The

[11] Because this chapter deals primarily with counting lexical exceptions under different versions of the regular rule, here I exclude a full list of lexical exceptions as part of PHL or NAS.

NAS system, once postulated, would of course encounter exceptions, but as outlined in Chapter 3, linguistic systems – including PHL – that have lexical specificity can still be stably acquired (Scobbie and Stuart-Smith, 2008; Payne, 1980; Roberts and Labov, 1995). Here, we are interested in whether the NAS system can become a viable endogenous response to the PHL system; if so, it would provide an example of simplification by children as a source of language change.

6.2.1 Can a NAS Postulation Tolerate PHL Input?

We begin first by asking the question "can a child who has posited NAS tolerate traditional PHL input?" To apply the Tolerance Principle to short-a systems in Philadelphia, assume a child is receiving input generated only by the traditional PHL system, with its disjunctive featural specification, syllabic sensitivity, and lexical exceptions. This learner could possibly hypothesize that their target grammar is simply 6, tense before nasals, producing tense æ in *ham, man*, etc. If they do so, they must somehow account for words they acquire that violates this generalization, such as lax æ in *bang*, or tense æ in *last.* If they maintain the generalization in 6, they must treat these and all other words that violate the "tense before nasals" generalization as stored lexical exceptions. If the number of such exceptions (e) is less than the tolerance threshold for that child's vocabulary size, then it is plausible that learners in Philadelhpia could endogenously hypothesize a NAS grammar given only PHL input. However, if the number of exceptions exceeds the tolerance threshold, then some other source of the NAS grammar in Philadelphia must be sought. As described in §3, N will be the entire set of æ words in a child's vocabulary, and e will be the list of words that violate R, where R = NAS.

We begin by using the CHILDES database (MacWhinney, 2000) to obtain a measure of the total N for a child's vocabulary. Each word type was coded for its realization under traditional Philadelphian input, under R = PHL, and under R = NAS. An example is shown in Table 6.1. Note that the mismatch between traditional input and PHL for *bad* reflects the fact that *bad* must be treated as a lexical exception, while PHL captures the regular phonological generalization.

This coding system allows us to measure the total number of exceptions produced by positing either PHL or NAS as a rule. Using Table 6.1 as a dummy lexicon with N = 5 words, we can see that a child positing R = PHL would have to list e = 1 exception to that rule, because the realization of *bad* under R = PHL does not match the child's input. Because 1 (fix), PHL emerges as a plausible rule for this dummy language. By contrast, a child positing R = NAS would have to list e = 4 exceptions, which does not pass the tolerance threshold of 3.11, rendering NAS an unproductive rule for the dummy language in Table 6.1.

Word	Traditional input	PHL	NAS
bad	Tense	Lax	Lax
hammer	Lax	Lax	Lax
cat	Lax	Lax	Lax
fast	Tense	Tense	Lax
bang	Lax	Lax	Tense

Table 6.1: Input realizations of /æ/ compared to expected /æ/ realizations for PHL and NAS. Mismatches between actual input and expected input (in gray) result in an exception.

Using the full list of /æ/ word types in CHILDES, we calculated whether the number of exceptions a child would need to list under R =PHL and R =NAS would pass the tolerance threshold of e. We find that given the traditional Philadelphian input distribution, a child positing R =PHL would have to store e =39 lexical exceptions (mostly *mad, bad, glad*, strong verbs and function words), well under the tolerance threshold of 194.7. This, of course, is expected: children have been successfully learning PHL and its listed exceptions for well over 100 years (Labov et al., 2016, 2013). Turning to the question of whether NAS can be a productive rule given traditional input, we find that positing R =NAS requires listing a total of 324 exceptions (e.g. all tense /æ/ before anterior voiceless tautosyllabic fricatives), well over the tolerance threshold.

Thus, despite being a formally simpler rule, and in fact a featural subset of PHL, NAS is not a plausible innovation for Philadelphian children on the basis of only traditional Philadelphian /æ/ input. Positing NAS requires storing too many lexical exceptions for it to be productive.

6.2.2 Can NAS replace PHL incrementally over time?

It remains, however, that NAS is rapidly replacing PHL as the dominant allophonic rule for /æ/ in Philadelphia. Given the finding the r =NAS is not a plausible re-analysis of the traditional input, we can now turn to the question of incremental re-analysis. In other words, we ask whether it is possible that a child might posit an intermediate rule given traditional input, which might then be re-analyzed as a productive NAS rule by the subsequent generation of language learners. We take PHL, reproduced in 25, and break it down into its four constituent aspects. R =PHL can be spelled out as /æ/ is tensed when it precedes a (a) tautosyllabic (b) anterior (c) nasal or (d) voiceless fricative.

Using these four components of PHL, we construct six intermediate grammars between full PHL and NAS, beginning with excluding only one aspect of PHL at a time and ending with excluding two aspects of PHL. We do not analyze intermediate forms of PHL which consist of excluding the nasal trigger, since that would not produce an intermediate form between PHL and NAS; NAS being the result of excluding every component of PHL except the nasal constraint. In Table 6.2, these rules are described as PHL minus the components that have been excluded. We note that some rule exclusions result in the expansion of the set of triggering forms (as in PHL-ant). The set of triggering phonological contexts resulting from each intermediate rule is shown in the third column of Table 6.2. We note finally that NAS is the same as PHL minus the tautosyllabic, anterior, and voiceless fricative components.

Name	**Rule**	**Triggering Segments**
PHL-ant	$\text{æ} \rightarrow \text{æh} / \_[\ \text{+nasal}\ ] \cup \begin{bmatrix} \text{-voice} \\ \text{+fricative} \end{bmatrix}]\,\sigma$	m, n, ŋ, f, θ, s, $\int]_\sigma$
PHL-taut	$\text{æ} \rightarrow \text{æh} / \_[\ \text{+anterior}\ ] \cap ([\ \text{+nasal}\ ] \cup \begin{bmatrix} \text{-voice} \\ \text{+fricative} \end{bmatrix})]\,\sigma$	m, n, f, θ, s
PHL-fric	$\text{æ} \rightarrow \text{æh} / \_[\ \text{+anterior}\ ] \cap [\ \text{+nasal}\ ]]\,\sigma$	$\text{m, n}]_\sigma$
PHL-ant-taut	$\text{æ} \rightarrow \text{æh} / \_[\ \text{+nasal}\ ] \cup \begin{bmatrix} \text{-voice} \\ \text{+fricative} \end{bmatrix}$	m, n, ŋ, f, θ, s, ʃ
PHL-ant-fric	$\text{æ} \rightarrow \text{æh} / \_[\ \text{+nasal}\ ]]\,\sigma$	$\text{m, n, ŋ}]_\sigma$
PHL-taut-fric	$\text{æ} \rightarrow \text{æh} / \_[\ \text{+anterior}\ ] \cap [\ \text{+nasal}\ ]$	m, n

Table 6.2: Intermediate grammars between PHL and NAS.

In addition to testing the intermediate rules shown in Table 6.2, we also consider the effects of a smaller vocabulary. As mentioned in 3, smaller vocabularies are able to tolerate a higher proportion of exceptions. This is particularly relevant to the question at hand: perhaps children with smaller vocabularies would be able to plausibly posit NAS as a productive rule for their traditional input. To test this, we also test the plausibility of NAS and intermediate PHL forms on several subsets of the most frequent words in CHILDES, with at least 20, 50, and 100 mentions in the corpus, so as to provide a rough approximation of learners' vocabulary composition at progressive stages of language development. The results are shown in Table 6.3.

Rule	1 Mention N =1412 T =194.7	20 Mentions N =498 T =80.2	50 Mentions N =334 T =57.5	100 Mentions N=239 T =43.6
PHL	39	19	15	11
PHL-ant	244	60	42	31
PHL-taut	155	55	36	25
PHL-fric	155	64	48	38
PHL-ant-taut	273	94	63	45
PHL-ant-fric	237	93	67	51
PHL-taut-fric	240	92	65	50
NAS	324	121	84	63

Table 6.3: Exceptions required for each intermediate rule for vocabularies consisting of words with 1, 20, 50, and 100 mentions in CHILDES. Plausible grammars shaded.

As shown in Table 6.3, NAS does not emerge as a plausible analysis of traditional input, even with a limited vocabulary. However, we see that traditional input can be plausibly re-analyzed as any of the three intermediate rules that result from deleting one component of PHL. For example, a child could plausibly posit a phonological rule tensing /æ/ before all nasals and voiceless fricatives, including ŋ and ʃ (PHL-ant) without having to list more exceptions than the threshold. Given the plausibility of *at least some* children positing these intermediate grammars, we must now turn to the question of whether these intermediate children could plausibly contribute enough examples to the linguistic environment that in turn favors NAS, resulting in wholesale change for all subsequent language learners. To do so, we introduce the model of rule learning under heterogeneous input.

6.2.3 Rule learning under a mixture of PHL and Intermediate Grammar input

As stressed in 3, the Tolerance Principle applies on individual learner's lexicon composition; even if a representative sample of words (e.g., the 498 that appear at least 20 times per million) can be expected to support an intermediate grammar (e.g., PHL-fric, which is expected to have 64 exceptions an thereby falls below the tolerance threshold of 80), it remains possible that some learners may learn from a somewhat skewed sample, whose lexicon fails to support an intermediate grammar. Thus, if the endogenous emergence of NAS is achieved through successive generations of learners, we must consider the situation in which learners are exposed to a mixed input: some produced by speakers who happened to successfully acquire an intermediate grammar and some produced by speakers who have retained the traditional PHL grammar. The question of whether NAS is a plausible reanalysis must then be reframed as "what proportion of intermediate input does a child need in order to plausibly posit NAS?"

To answer this question, we simulate a child's acquisition given dialect contact between PHL and each intermediate rule, in the following way. First, we let m represent the proportion of input from the intermediate grammar that a child receives during acquisition, and 1-m the proportion of traditional input. We then construct a simulation of the plausibility of positing NAS, for values of m between 0 and 1 in steps of .01 for each of the three intermediate rules. We begin with the assumption that a child will store one form for each word type. For each run of the simulation, we generate a full mixed lexicon according to m. Each word is assigned lax or tense /æ/ on the basis of an intermediate rule or traditional input, according to m. For example, if m =24, each word in the lexicon will have a 24% chance of its /æ/ allophone being determined by an intermediate rule. This assumption is motivated by empirical studies of how children deal with mixed input where each lexical item is subject to probabilistic variation at the level of token frequency. In a series of studies (Hudson Kam and Newport, 2005, 2009), children were found to regularize mixed input to the statistically dominant variant. In the present case of mixed input with the level of m, we assume that each word type has an m probability of beign internalized in the child learner's vocabulary as an example of the intermediate grammar, and a $1-m$ probability of being internalized as an example of PHL. That is, the child regularizes a *probilistic* mixture of word tokens in the input as a *discrete* mixture of word types representing the two variant grammars: this is implemented by stochastically assigning each word type into one of two grammars with the associate probabilities. We then evaluate the viability of the two grammars on the basis of the resulting lexicon.

It is worth stressing several important features of the learning model. First, it is crucial to note that this is an acquisition model of how a single learner evaluates rules given variable input. This is clear from the description of the model, where the sample lexicon for the learner is stochastically drawn from the mixture distribution in the environment. By running the model many times, we can understand the outcome of learning for the speech community at large. Second, the model is agnostic as to the real-world source of the variable input. Both dialect contact scenarios and endogenous innovation scenarios are treated identically by the model. An individual learner evaluates rules on the basis of the lexicon they acquire from the mixed environment, and it is immaterial how such a mixture is introduced in the environment in the first place; see Yang (2000) for additional discussion and applications to syntactic change. Third, the model also does not imply any particular time course for change. For a given mixture of input data, it estimates the probability that PHL or NAS may be a plausible grammar for a speaker, but does not predict what the rate of use of either grammar would be for a speaker who has successfully acquired both systems. In other words, this model does not predict m for the next generation of learners. Fourth, we stress that this model does not address how a child may *generate* a possible rule, it is simply a model of how a child evaluates possible rules that have already been generated.

We calculate whether an input lexicon comprised of a mixture of PHL and intermediate grammars would allow NAS to be a productive rule for each trial. 1000 trials were run for each value of m between 0 and 1 in steps of .01, for each intermediate grammar.

Figure 6.1 presents the results of this simulation, with rates of m plotted along the x-axis and the proportion of trials that pass the tolerance threshold along the y-axis. It is important to note that the y-axis represents only the predicted proportion of children whose input would allow them to evaluate PHL (in stars) or NAS (in circles) as a plausible grammar for each value of m; it does not represent the predicted production of these children. Each intermediate rule was tested for whether NAS passed the tolerance threshold for each value of m (circles) and for whether PHL

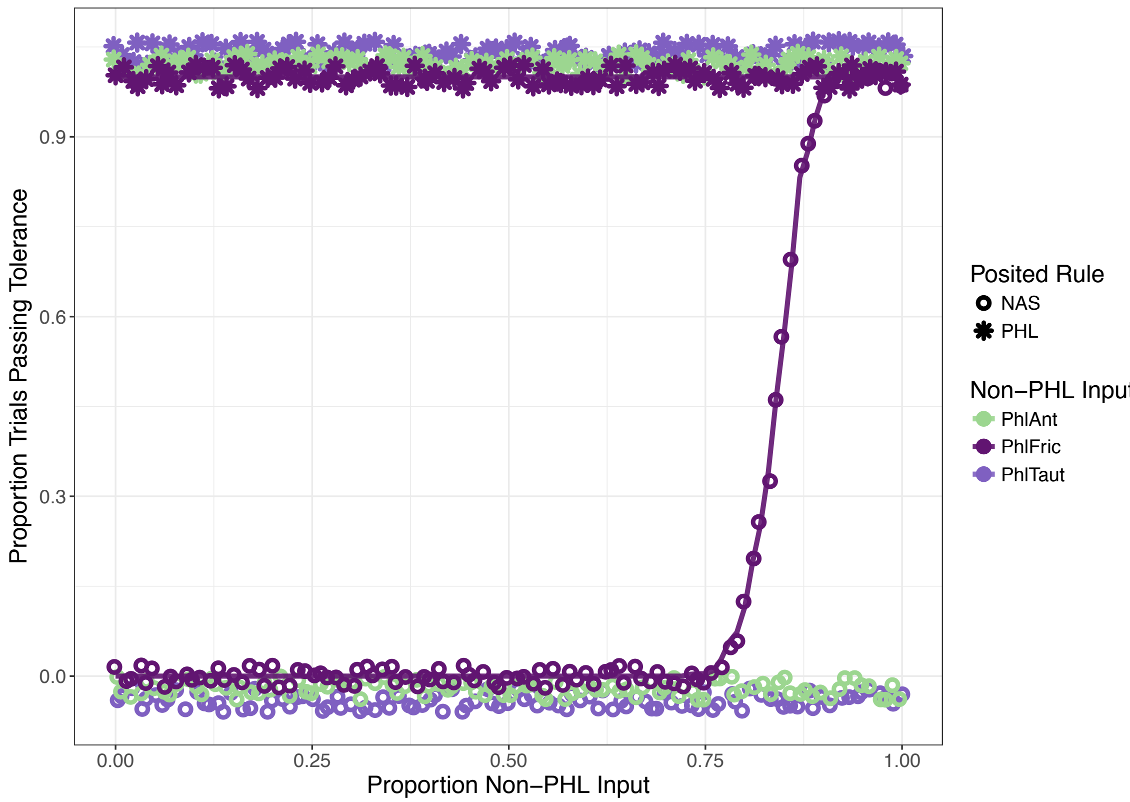

Figure 6.1: Proportion trials which pass the tolerance threshold for each proportion of intermediate rule input for positing NAS or PHL.

passed the tolerance threshold (stars).

We find two striking results. First, PHL is a plausible reanalysis of every intermediate rule, for all proportions of intermediate input, including 100% intermediate input. This speaks to the history of stability of PHL in Philadelphia; even if speakers have been spontaneously positing intermediate rules throughout the history of the /æ/ split in Philadelphia, these intermediate rules can still be reanalyzed as PHL by the next generation of speakers. Second, of the three intermediate rules that are plausible given traditional input, it is only the PHL-fricative rule that will allow NAS to be a plausible reanalysis of the intermediate rule. And this is only possible when children are receiving approximately 73% PHL-fricative input, which is the point at which the probability of accepting NAS becomes non-zero. That is, if at least 73% of Philadelphian children lost the voiceless fricative conditioning, then NAS can endogenously emerge as a consequence. We note that this possibility mirrors the argument in Ash (2002), who models the change from PHL to NAS in central New Jersey as occurring via an intermediate step of PHL-fricative.

However, we find this route of change to be highly implausible for Philadelphia, given the results of an empirical search for speakers exhibiting a PHL-fricative type grammar. Only 1 speaker out of 184 who had enough data to allow such an investigation was found[12]: Jake S, our outlier from Chapter 2. As I have argued in Chapter 2, Jake's social profile suggests he developed a PHL-fricative grammar as a *result* of NAS contact, rather than as an endogenous modification of the PHL system. Jake was born in 1992, and attended the elite Masterman middle and high school, then went on to graduate from the University of Pennsylvania. Most of Jake's peers – speakers born around 1992 who attended Masterman – produced NAS. Given the data, this suggests that language learners positing PHL-fricative was not the route by which NAS came into Philadelphia. In addition to a social profile that renders Jake's production of PHL-fricative an unlikely step in the change to NAS for his own subset of the speech community, it is also noteworthy that finding only one speaker out of 184 falls well short of the 73% PHL-fricative speakers necessary for NAS to be plausible for the following cohort of speakers.

[12]Using data from the PNC and the IHELP corpus , we analyzed every white speaker who produced at least 5 /æ/ tokens in both the fricative environment and lax nasal environment. The search was restricted to white speakers, as African American and Hispanic speakers in Philadelphia traditionally produce a neutral /æ/ system, produced as a raised lax form [ɛː] for all phonological categories (Fisher et al., 2015; Labov and Fisher, 2015).

To summarize the theoretical results so far, we have found that it is impossible for NAS to directly arise from a PHL system. It is conceivable that an intermediate grammar, specifically PHL-fricative, may eventually lead to NAS, but only if the vast majority of learners all converge onto that grammar under homogeneous PHL input. This, however, proves to be highly unlikely. Finally, although our simulation is intended to model the terminal state of language acquisition, it can also be used to understand the developmental time course of language acquisition in a single child/generation. It is clear that unless a child is nearly completely surrounded by PHL-fricative input (as indicated by the value m), it is virtually impossible for the grammar to survive until the stabilization of language acquisition (e.g., pre-teen years; Herold (1997); Johnson (2010); Johnson and Newport (1997)).

6.3 Acquiring NAS through dialect contact

Given the unlikelihood of and lack of empirical support for NAS emerging endogenously in Philadelphia, either through direct reanalysis of the original system or via a sequence of reanalyses, we now turn to the possibility of NAS emerging as the result of dialect contact between NAS and PHL.

6.3.1 Sociolinguistic background

The idea that Philadelphian children may be exposed to NAS speaking outsiders is not altogether unlikely. According to the Atlas of North American English (Labov et al., 2006), NAS has been found in the geographic area surrounding Philadelphia; it is not unlikely that some of these speakers may have access to and influence within Philadelphia. Ash (2002) also provides clear evidence of NAS gaining ground over both PHL and the New York split-/æ/ system in the Mid-Atlantic region in the region between Philadelphia and New York City. Furthermore, as outlined in Chapter 2, NAS is more likely to be found in graduates of elite non-Catholic high schools such as Masterman and Friends Central than in graduates of local diocesan schools. This pattern fits with an analysis of NAS as a change from above: the wealthier, more nationally-oriented schools adopt NAS early (perhaps via the influence of externally raised teachers), while the more locally-oriented neighborhood schools act as conservative forces holding on to PHL.

6.3.2 Theoretical analysis and predictions

Given that dialect contact with NAS speakers is a likely situation given the geographic and social patterns around Philadelphia, we now turn to the question of how much contact with NAS speakers is necessary for a Philadelphian child to accept NAS as a plausible system. Using the same simulation procedure described in §6.2.3, with NAS as the non-PHL input at the proportion of m, we tested which proportion of NAS input is necessary for a child to plausibly posit NAS. Figure 6.2 presents the results of this simulation, plotting the proportion of trials in which NAS emerged as a plausible rule (in circles) and PHL emerged as a plausible rule (in stars). Simulations were run for different sized lexicons, from words with one mention to words with 100 mentions in CHILDES, in order to capture the potential effect of differently sized lexicons. The full results are displayed in Table 6.4, which displays the proportion NAS input necessary for NAS and PHL to be viable at all as well as viable for 100% of trials.

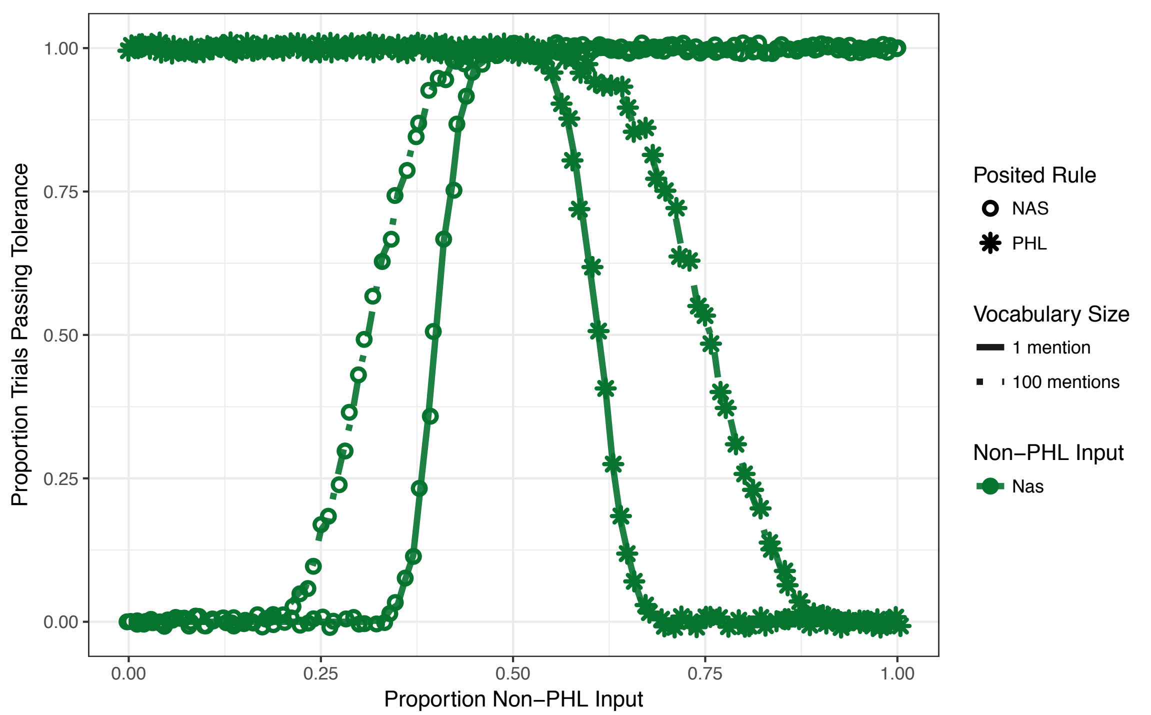

Figure 6.2: Proportion trials that pass the tolerance threshold for NAS (circles) and PHL (stars) for different proportions of NAS input.

As expected, higher word frequency cutoffs produce shallower slopes; this is a reflection of the fact that these lexicons are smaller and therefore more proportionally tolerant of exceptions, resulting in a slightly higher proportion of trials that pass the tolerance threshold for each value of m. In contrast to the endogenously posited intermediate rules simulated in the previous section, we find that dialect contact between traditional input and NAS makes positing NAS a highly plausible solution for a child receiving both inputs. In other words, NAS becomes a plausible analysis of a child's input if that child is receiving at least 32% NAS input.

Vocabulary size	**NAS leaves 0% viable**	**NAS reaches 100% viable**	**PHL leaves 100% viable**	**PHL reaches 0% viable**
1 mention	.32	.48	.53	.7
20 mentions	.25	.46	.52	.82
50 mentions	.2	.47	.54	.86
100 mentions	.17	.48	.54	.9

Table 6.4: Proportion NAS input at which NAS and PHL become variable viable and categorically viable.

Again, although our model has been used to study contact-induced change, it is also applicable to children's developmental time course, and the sociolinguistic conditions of language acquisition. For example, as documented in detail (Johnson, 2010), young children may initially acquire the grammar of the parental input and then adopt a new grammar once immersed in their peer group under certain conditions. In the current study, the viability of PHL and NAS as a function of contact can be understood as follows. If there is a relatively weak presence of NAS in the environmet (e.g., m <.2), even if a child were to acquire NAS at home they will still end up adopting PHL. Likewise, if NAS is already quite dominant in a child's peer group (e.g., m >0.7), then the home PHL system will be abandoned in favor of NAS. In the region where m assumes an intermediate value, both systems are predicted to be viable. In other words, whichever system a child acquires at home, the linguistic environment of their peer group is sufficiently heterogeneous for these intermediate m values that either system will be sufficiently supported (i.e., neither will encounter an intolerable number of exceptions).

The above discussion is particularly applicable when the community network structure is taken into account. For instance, while m may be quite low over the entire speech community of Philadelphia, there may be local networks which may be geographically or socially defined, in which the concentration of NAS speakers is quite high, which may lead to the rise of NAS in specific groups before diffusing it to the wider dialect region. This is precisely the situation found in Labov et al. (2016) and reported in Chapter 2, which finds the highest concentration of NAS speakers amongst the graduates of elite public high schools, with other school networks lagging behind in the change to NAS.

6.4 Stability, Change, and Variation

So far, we have focused exclusively on what kind of input is necessary in what mixture for children to acquire a NAS grammar. However, the conclusion for the acquisition modelling is that across a broad range of mixtures, both PHL and NAS grammars are plausible. This raises two clear questions. First, is it possible that some learners acquire both PHL and NAS as a result of dialect contact? Second, once both grammars are in use within the speech community, is it inevitable that one should replace the other, as is being observed in Philadelphia?

Let's first consider the question of co-existing variation as the outcome of learning. There is considerable evidence that even for fully native bilingual speakers, one of the phonemic systems appears dominant (Cutler et al., 1989; Bosch and Sebastián-Gallés, 2003). The acquisition of the low-back merger system at the dialect boundary appears to be a case in point. At the beginning of this change, despite the presence of the merged system in their peer group, children retained the traditional distinct system. Once the merged system reached a certain level of prominence – above 23% – children acquired it en masse, resulting in the dramatic contrast in the vowel systems used by siblings separated by a few years as documented by Johnson (2010). However, the evidence provided in Chapter 4 suggests that for this allophonic restructuring in Philadelphia English, transitional cohort speakers do in fact acquire both PHL and NAS, and produce variation between the two systems as a whole. That Figure 6.2 finds both systems fully viable for such a wide overlap of NAS input (between 46% and 54% NAS input) provides a suggestion of the input data provided to the competing grammars speakers found in Chapter 4.

We now turn to the second question: if both NAS and PHL are viable, and speakers evidently acquire them, what is the long-term prognosis of this competition? Will one system necessarily replace the other, as it appears to be doing in Philadelphia? This question is quite different from the issues discussed so far in this chapter. We have mainly been concerned with the viability of a *single* system given a mixed environment. The Tolerance Principle based model has identified numerical conditions under which one grammar will replace the other as the terminal stage of language acquisition. It is a separate question entirely whether, having posited and acquired *two* competing systems, one will prevail. For the intermediate values of m, the learner can – and apparently does – acquire both systems, assigning a probabilistic distribution over them. Here, we have the more familiar sociolinguistic situation of variable rules, in which a speaker sometimes uses one variant of the allophonic system parameter and sometimes the other. The suitable mathematical model to study the dynamics of change is the variational model (Yang, 2000, 2002), where the terminal stage of language acquisition is a statistical distribution over two (or multiple) grammars. Language change is characterized by the change in this simulation over time, as governed by the differential utilities ("fitness") of the grammars in competition. Unlike the Tolerance Principle, which operates over *type* frequencies for rules and exceptions in the learner's lexicon and has a discrete outcome (whether a rule is tenable or not), the variational model presupposes the productivity of both rules and evaluates them on the basis of *token* frequencies.

The adaption of the variational model to a case of allophonic restructuring is not entirely straightforward. By the traditional formulation, this model evaluates the proportion of input produced by the each grammar that can be parsed by the other. The inevitable "winner" will be that grammar that can parse more of the other grammar's production (i.e., receives a lower penalty probability). The idea here is that many utterances will be compatible with either underlying grammar that is in competition, but that the few utterances that are *not* compatible with one of the two possible grammars generates a penalty probability for that grammar. Whichever grammar receives the lowest penalty probability will eventually win. This is visually represented in Figure 6.3, which displays the overlapping production of two mutually incompatible grammars (G_1 and G_2). In this visualization, G_2 will eventually win out over G_1, because it can analyze a higher proportion of the other's output.

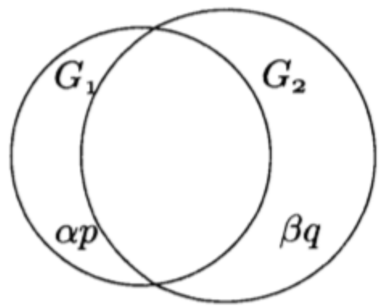

Figure 6.3: Two mutually incompatible grammars produce some proportion of ambiguous utterances. From Yang (2000).

This model has been successfully applied to syntactic parameters like the acquisition of a V2 grammar or pro-drop (Yang, 2000) as well as to phonological parameters like the LOT-THOUGHT merger (Yang, 2009), which produce assymetrical α and β values, resulting in an inevitable winning grammar. The problem of applying the variational model to the competition between PHL and NAS is that because any test token incompatible with PHL will be compatible with NAS and vice versa, here the penalty probabilities between the two grammars will be identical. While confusability can not be used here as a penalty probability, a potential direction for future research may lie in the social evaluation metrics reported in Chapter 5. PHL and NAS may be able to parse the exact same proportion of output, but they do not receive identical social evaluation scores. That structural sound change may be socially motivated has been a longstanding aspect of sociolinguistics (Labov, 1963); while future work may fruitfully apply the magnitude estimation scores of Chapter 5 to the variational model for the competition between PHL and NAS, this remains beyond the scope of the current dissertation.

6.5 Conclusion

The formulation of precise theoretical formulations such as the Tolerance Principle enables specific predictions, which hopefully lead in turn to theoretical advancement. In this chapter, we've demonstrated that applying the quantitative precision of the Tolerance Principle to the question of phonological change through language acquisition has allowed us to articulate a clearer model of the allophonic restructuring of /æ/ in Philadelphia in a way which would not be possible otherwise. Given a number of prima facie plausible hypotheses for the source of NAS innovation (grammar simplification, endogenous reanalysis, and dialect contact), we have been able to determine that only dialect contact is the likely source of this change.

We've investigated the possibility of NAS emerging in Philadelphia through regular transmission, finding that not only is NAS an implausible reanalysis of PHL input, but that it is also unlikely for NAS to have emerged through successive transmission simplifications of PHL. We've furthermore demonstrated that dialect contact is a far more likely source of NAS in Philadelphia, with the finding that NAS becomes a plausible analysis of mixed-environment input if that input is comprised of only 46% NAS. Importantly, it is not necessary for the *entire* speech community to be using NAS 46% of the time in order for NAS to make inroads into the speech community. Rather, it is only necessary for *some* learners to receive 46% NAS input.

This point bolsters the claim in Labov et al. (2016) that the shift from PHL to NAS is a change from above through dialect contact with NAS speakers who are unevenly distributed across social networks. Chapter 2 provides insight into the way educational systems in Philadelphia produce this uneven distribution, as well as the community level social characteristics that fit a classic change from above. Chapter 5 provides further evidence for this change as a change from above, with younger speakers in the Magnitude Estimation task rating all NAS-conforming tokens positively but rating tense PHL-conforming tokens negatively.

Finally, we've also found a relatively wide overlap in the tenability of PHL and NAS, with both systems completely viable when the input comprises between 46% and 54% NAS input. These findings are quickly turned into their own empirical predictions. We expect a child who is receiving less than 32% NAS input to posit PHL, and a child who is receiving more than 70% NAS input to posit NAS. A child receiving roughly 50% NAS input is expected to learn both systems and produce variation between the two. This predicts that a child with one NAS-speaking caregiver and one PHL-speaking caregiver who receives roughly equivalent input from both will emerge as a variable speaker, at least before they enter school and receive input from their peer group. We note that this prediction aligns neatly with the empirical results of Payne (1980), who found children with one PHL parent producing some /æ/ tokens that were inconsistent with PHL.