

# The medium-term dynamics of accents on reality television

Morgan Sonderegger<sup>\*1</sup>, Max Bane<sup>†2</sup>, and Peter Graff<sup>‡</sup>

<sup>1</sup>Department of Linguistics, McGill University

<sup>1</sup>Centre for Research on Brain, Language, and Music

Final preprint version of October 2016, to appear in *Language*

## Abstract

How flexible is an individual’s accent during adulthood, and how does this flexibility relate to longer-term change? Previous work has found that accents are remarkably flexible in conversational interaction, but predominantly stable over years, leading to very different views of the role of individuals in community-level sound change. This paper examines *medium-term* accent dynamics (days to months) by taking advantage of a ‘natural experiment’: a reality television show where contestants live in an isolated house for three months and are constantly recorded, forming a closed system where it is possible to both determine the dynamics of contestants’ speech from day to day and reason about the sources of any observed changes. We build statistical models to examine time dependence in five phonetic variables within individuals, in 14.5 hours of spontaneous speech from 12 English-speaking contestants. We find that time dependence in pronunciation is ubiquitous over the medium term: large daily fluctuations in pronunciation are the norm, while longer-term change over weeks to months occurs in a minority of cases. These patterns mirror the conflicting findings of previous work, and suggest a possible bridge between the two. We argue that time dependence in phonetic variables is influenced by contrast between sounds, as well as systematic differences between speakers in how malleable their accent is over time; however, we find only limited evidence for convergence in individuals’ accents. Our results have implications for theories of the role of individuals in sound change, and suggest that medium-term pronunciation dynamics are a fruitful direction for future work.\*

---

\*Corresponding author: [morgan.sonderegger@mcgill.ca](mailto:morgan.sonderegger@mcgill.ca)

†[max@thebanes.org](mailto:max@thebanes.org)

‡[graffmail@gmail.com](mailto:graffmail@gmail.com)

\*We thank Molly Babel, Lauren Hall-Lew, Tyler Kendall, Jane Stuart-Smith, Natasha Warner, and an anonymous reviewer for helpful comments on drafts of this paper, and the many other people who have provided useful feedback on this project, including Peter Auer, Rina Foygel Barber, Paul Brazeau, Alex Fine, James Kirby, Matt Goldrick, John Goldsmith, Dan Jurafsky, Robin Queen, Jason Riggle, Tyler Schnoebelen, Charlotte Vaughn, and Alan Yu; as well as audiences at CLS 2010, LSA 2011, CUNY 2011, NWAV 2012, LabPhon 2014, ICPHS 2015, and at USC, Indiana U., U. Michigan, McGill, U. Edinburgh, U. Ottawa, U. Minnesota, U. Toronto, Cornell, Northwestern U., and MIT. We thank N. Rothfels, M. Nelson, T. Knowles, M. Schwarz, M. Labelle, L. Bassford, R. Hwang, H. Hsieh, and M. Coumel for research assistance. We gratefully acknowledge permission from Channel

# 1 Introduction

Two of the most striking aspects of language are stability and change. Both aspects are especially clear for how sounds are pronounced, the focus of this paper. At the level of individuals, we intuitively know that a speaker’s accent can shift as a function of her interlocutor, and that some people’s accents seem to shift permanently (*accent change*), for example after moving to a new dialect region. However, we also intuitively know that many people’s accents are stable over time: this default assumption underlies our ability to identify where someone is from by their accent. At the level of speech communities, many aspects of the phonetic and phonological system are stable at a given point in time, yet cases of *sound change* are always occurring (e.g. Labov, 1994). This paper addresses three broad questions about accent dynamics and sound change: how plastic are the accents of individual speakers during adulthood, what are the sources of the dynamics of a speaker’s accent, and what is the relationship of accent plasticity in individuals to community-level sound change?

Previous work has addressed these questions on two timescales: *short-term* shifts in pronunciation during conversation and laboratory experiments (seconds–hours), and *long-term* accent change over the lifespan (years).

## 1.1 Short-term

One longstanding body of work emphasizes the flexibility of the sound systems of individuals, which are constantly updated over time as a result of interaction (Hockett, 1965; Paul, 1880; see Garrett & Johnson, 2013). For example, contemporary *emergentist* theories of phonetics and phonology (e.g. Johnson, 1997; Pierrehumbert, 2001), whose core tenets go back to the Neogrammarians, assume that a speaker’s cognitive representation of sounds includes all examples ever heard. This representation is updated with each new interaction, and used in speech production. Thus, emergentist models predict that pronunciation norms can and will shift over time, including during adulthood, as a result of interactions.

A primary source of evidence for accent plasticity in adults comes from short-term shifts in how one speaks (over seconds–hours), in conversations and in laboratory settings, under exposure to other people’s speech. These shifts (termed imitation, (phonetic) convergence, accommodation, etc.) have been documented using a variety of paradigms (Pardo, 2013). As assessed by perceptual measures, subjects shift their overall pronunciations towards a target voice in word shadowing tasks (Fowler et al., 2003; Goldinger, 1998) and towards their interlocutor in conversational interaction (e.g. Pardo, 2006). Specific acoustic variables typically also undergo short-term shifts in both interactive and non-interactive settings. For example, American English speakers increased the VOT of word-initial voiceless stops towards a model talker in an imitation task (Nielsen, 2011), and shifted their productions of five vowels (in F1/F2 space) towards a model talker in a shadowing task (Babel, 2011). Pardo (2009) also found American English speakers shifted vowel formants in conversational interaction, but in more complex ways. In these and other studies of shifts in phonetic variables, the existence of *some* shift, as a result of exposure to another person’s speech, is quite robust across speakers.

Short-term shifts in pronunciation are to some extent ‘automatic’ consequences of interaction (Delvaux & Soquet, 2007; Goldinger, 1998). They are also heavily modulated, by

---

4/Endemol to access footage from Big Brother 9 UK. Preliminary versions of the VOT results were given by Bane et al. (2010) and Sonderegger (2015), and by Sonderegger (2012) for all variables. This work was supported by grants from the Social Sciences and Research Council of Canada (#430-2014-00018) and the Fonds de Recherche du Québec Société et Culture (#183356).

three types of factors, which result in the significant variability between speakers in the size and directionality of shifts observed in most short-term studies:

1. *Social factors*, such as attitude towards the interlocutor and gender (Babel, 2010; Bourhis & Giles, 1977; Namy et al., 2002). Communication Accommodation Theory (Giles et al., 1991, *et seq*) proposes that these shifts result from individuals managing social distance using an *accommodation* strategy, such as *convergence* or *divergence* (e.g. to express solidarity or disapproval).
2. *Linguistic factors*. For example, Nielsen (2011) cites contrast maintenance to explain why English speakers will shift VOT of voiceless stops towards lengthened but not shortened VOT (c.f. Babel, 2011; Mitterer & Ernestus, 2008).
3. *Individual differences*, correlated with factors such as personality traits and gender (Yu et al., 2013; Namy et al., 2002).

In sum, plasticity in pronunciation, modulated by multiple factors, is the norm over seconds–hours for adults. This plasticity is in line with the ubiquity of *style shifting*: shifts in a speaker’s linguistic usage as a function of the addressee, topic, etc. (e.g. Bell, 1984; Eckert, 2000; Rickford & McNair-Knox, 1994) possibly many times over the course of a day (e.g. Coupland, 1980; Hindle, 1980).

**Change By Accommodation** The existence of substantial accent plasticity in interactions goes naturally with the view that change during adulthood is an important driver of sound change in communities. Versions of this view go back to the Neogrammarians (Paul, 1880), as summarized by Auer & Hinskens (2005), who call it the *change by accommodation model* (CBA). CBA explains the relationship between individual-level and community-level change as a sequence of three steps: (1) people accommodate to each other during interactions; (2) eventually, a given individual’s norms change as the result of accumulation of these interactions; (3) spread of the innovation in the wider community (as individuals undergo step 2). Different versions of CBA address dialect change/contact settings alone, or all sound change (e.g. Bloomfield, 1933; Garrett & Johnson, 2013; Labov, 1990; Trudgill, 1986, 2004), and differ in whether short-term shifts (and thus steps 2–3) are fundamentally automatic (every interaction makes the sound system of interlocutors more similar), or fundamentally social (as proposed by Communication Accommodation Theory). Regardless of their mechanism, short-term shifts are generally assumed to be “the driving force of language change” (Auer & Hinskens, 2005, p. 356). The fact that most short-term literature considers adults suggests that Steps 1–2 refer by default to adults. Thus, the link being made between accent plasticity in adulthood and the nature of sound change is: because adults’ pronunciations are so plastic in interactions (step 1), *long-term change* during adulthood (step 2) plays an important role in community-level *sound change* (step 3).

## 1.2 Long-term

On a timescale of years, the traditional view is that an individual’s linguistic system is largely fixed in adulthood (e.g. Chambers, 2003, p. 197). The assumption of accent stability over years underlies the widespread use of the *apparent-time construct* to study sound change in progress using a synchronic sample (Cukor-Avila & Bailey, 2013).

Previous work examines the extent of post-adolescent accent plasticity in two kinds of settings where accent change towards changing community norms seems intuitively likely:

individuals who remain in the same speech community, where some change is in progress (*panel studies*; Sankoff, 2005, 2013), and individuals who have moved between dialect regions (*dialect change* studies; Auer et al., 2005; Siegel, 2010).<sup>1</sup> For example, Sankoff & Blondeau (2007) examine change in rhotic realization ([r]→[ʀ]; a community-level sound change in progress) in Montreal French by a panel of 32 speakers recorded 11 years apart, and find relative stability of most individuals (72%), with the remainder increasing [ʀ] use. While most large-scale dialect change studies do not break down results by individual speakers, those that do suggest that significant accent change in adulthood is uncommon. For example, Foreman (2003) examines six American speakers who have settled in Australia for whom longitudinal data is available over 10–27 years; five of the six speakers show little or no significant change across six phonological variables. However, case studies have highlighted that dramatic accent change is possible. For example, Queen Elizabeth II’s realization of English vowels has changed significantly over 50 years of radio addresses, often paralleling community-level changes (Harrington et al., 2000, *et seq*); Yiddish folk singer Sarah Gorby shows a mixture of stability and change across phonological variables over 50 years, with change generally towards the standard (away from her native dialect) (Prince, 1987). More generally, long-term studies find that both variables and speakers (for a given variable) differ substantially in the degree of plasticity over years (Siegel, 2010, p. 51), for a variety of reasons. Though further work is needed given this heterogeneity, the picture which emerges from long-term studies to date across a range of languages and variables (e.g. Bowie, 2005; Brink & Lund, 1975; Nahkola & Saanilahti, 2004; Sancier & Fowler, 1997; Stanford, 2008) is that “the default for adults is apparently stability”, while a minority show significant pronunciation change, usually in the direction of community-level changes in progress or ambient norms. (Sankoff, 2013).

**Generational change** The finding that adults show limited and heterogeneous long-term plasticity in pronunciation goes naturally with the view that change during adulthood is not a primary driver of community-level sound change. An influential proponent of this view is Labov (1994, 2007), who argues that pronunciation changes internal to speech communities occur primarily via *generational change*: intergenerational *transmission* of norms from adults to children, and *incrementation* by children in the direction of changes-in-progress during childhood and adolescence. By contrast, the typical mechanism of changes rooted in dialect contact is *diffusion* in the course of interaction between adults who use different pronunciations. Diffusion is less important than transmission as a source of linguistic diversity *because* it is driven by adults, who show less and more sporadic “capacity to change their linguistic systems” compared to children and adolescents (Labov, 2007, 349). While Labov’s account has been subject to debate (e.g. Babel et al., 2013), it represents a common view: change in adulthood is not an important factor in the major source of sound change (internal change), while external change is a less important source precisely because change in adulthood is limited and sporadic.

### 1.3 Medium term

Empirical evidence from the short-term and long-term literatures broadly suggest that accents are both plastic during interactions and largely stable over the lifespan, and associated theoretical viewpoints reach conflicting conclusions about accent plasticity in individuals,

---

<sup>1</sup>We do not consider *age grading*, where a variable’s use changes during adulthood in a predictable way.

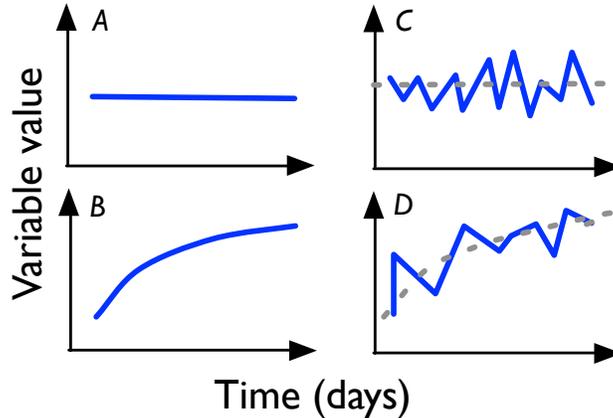


Figure 1: Schematic of possible types of time dependence in a phonetic variable within individual speakers: no change (A), time trend (B), by-day variability (C), by-day variability and time trend (D). Solid line indicates the variable’s value on each day; dotted lines indicate its mean value over several days.

and its relationship to community-level sound change. This conflict comes from extrapolating from short-term and long-term results to accent plasticity in individuals over the *medium term* (days–months), about which little is known.

Long-term studies suggest that accent stability is the norm in adulthood over years, but little is known about accent dynamics in adults on shorter timescales (c.f. Barden & Großkopf, 1998; Evans & Iverson, 2007; Pardo et al., 2012, where sampling points are separated by months). A reasonable null hypothesis for medium-term dynamics would be that an individual’s speech does not change over days to months; we call this the *stationarity hypothesis*. We can distinguish two versions of the stationarity hypothesis, thinking about the trajectory of an individual’s “baseline” use of a phonetic variable every day, over a period of months (schematized in Figure 1), after controlling for other factors (style, linguistic context, etc.). In the strong version, there could be no time dependence at all (Type A), either from day to day or over longer periods. A weaker version would be that individuals vary from day to day (*by-day variability*) but not over longer periods (Type C). There is almost no work examining whether an individual’s accent fluctuates from day to day (c.f. Heald, 2012; Pisoni, 1980, discussed below), though Nahkola & Saanilahti (2004) suggest this possibility.<sup>2</sup>

Short-term studies show that speakers regularly adjust their pronunciation due to the speech of others. The proposed link between short-term shifts in interaction and longer-term change is at the heart of CBA, and motivates many short-term studies (e.g. Babel, 2011; Delvaux & Soquet, 2007; Nielsen, 2011; Pardo, 2006). An important initial assumption being made in these literatures, which we call the *persistence hypothesis*, is that short-term shifts persist beyond individual interactions; they may then accumulate, and eventually lead to change in an individual’s pronunciation norms over the medium term. However, to our knowledge there is no work examining the persistence of short-term shifts in use of a phonetic or phonological variable for more than one hour.<sup>3</sup> The persistence hypothesis predicts that

<sup>2</sup>In contrast, Cukor-Avila (2015); Rickford & Price (2013) have examined daily fluctuations in individuals’ use of *morphosyntactic* variables, focusing on effects of style shifting and methodology.

<sup>3</sup>Goldinger (2000) shows persistence over one week of an overall percept, but acoustic variables are not examined.

if an individual engages in the same type of interactions sufficiently often, there should be steady change in her ‘baseline’ use of a phonetic variable over days to months—beyond any by-day variability—which we term a *time trend*. A phonetic variable could show a time trend without by-day variability (Figure 1: Type B), or both a time trend and by-day variability (Figure 1: Type D). By-day variability and time trends form two independent dimensions of medium-term time dependence: an individual’s pronunciation could show either, both, or neither.

## 1.4 Medium-term accent dynamics in a closed system

The disconnect between short-term and long-term plasticity and the paucity of empirical data on timescales in between motivates the current study, which investigates *medium-term* time dependence in speech production in individuals by taking advantage of a “natural experiment”: the reality television show *Big Brother UK*, whose structure is uniquely suited to investigating how and why an individual’s accent changes over the medium term. The show contains speech from the same individuals, recorded on a near-daily basis over three months, making it possible to examine whether aspects of contestants’ speech show by-day variability, time trends, or both. Contestants interact with each other constantly, without access to the outside world; the house is thus a linguistically “closed system” (Bane et al., 2010), a small community where persistence of short-term shifts in pronunciation might be expected to be especially likely, and where it is in principle possible to test whether change in an individual’s speech can be related to social interaction.

We first describe a corpus of spontaneous English speech from one season of this show, in which we examine five phonetic variables for 12 contestants over up to three months (Sec. 2–3). For each variable, we build statistical models to characterize its time dependence within each speaker, after controlling for other factors (Sec. 4). We then use the results of these models to address two research questions about accent dynamics over the medium term.

First, *what qualitative kinds of time dependence do phonetic variables show within individual speakers* over three months (Sec. 5)? We consider the types of time dependence shown in Figure 1, and find that medium-term time dependence of one kind or another is ubiquitous across speakers and variables, and that time dependence is primarily due to by-day variability and secondarily to time trends. Second, *to what extent can we account for the observed patterns of variability over time* (Sec. 6)? We consider the same types of potential sources which mediate short-term pronunciation shifts—social factors, linguistic factors, and individual differences—and find some evidence for a role of each one in shaping medium-term accent dynamics.

These findings allow us (Sec. 7) to evaluate the stationarity and persistence hypotheses, to situate medium-term dynamics between the disparate patterns individuals show on short-term and long-term timescales, and to evaluate the mechanism of accent dynamics in individuals and the possible relationship to community-level change.

## 2 The Big Brother Corpus

### 2.1 Show description

The Big Brother corpus consists of speech from the ninth season of the reality television show Big Brother UK (Channel 4/Endemol), which aired from June 5 to September 5, 2008

(93 days).<sup>4</sup> Twenty-one contestants (or *housemates*) took part in the show: 16 entered the Big Brother house on Day 1 and five more entered later in the season. Housemates were gradually evicted during the season by a combination of nomination by other housemates and voting by the viewing public. The last housemate remaining won a cash prize.

While on the show, housemates had essentially no linguistic input from the outside world: they did not leave the house, interact with people not involved with the show, or (with rare exceptions) have access to media. They spent most of their time interacting with each other in some form, e.g. in unstructured conversations or while participating in tasks set by Big Brother. The one important exception was the *diary room*, an isolated room where housemates could go to speak with “Big Brother”. Big Brother was in fact different people (male and female) at different times; he/she could see housemates, but could not be seen by them, and communicated only through audio. Housemates went to the diary room to talk about their feelings, events in the house, etc., and could either go to the diary room voluntarily or be called by Big Brother.

Housemates were recorded at almost all times, including by wearable microphones. During the season, the public could see video of housemates via a live feed or various produced shows, including daily ‘Highlights’ shows and weekly ‘Diary Room Uncut’ shows, which consisted of continuous segments from the house, presented without commentary.

Table 1 gives basic demographic information for the 21 housemates, who come from diverse dialect regions. Sixteen housemates are native speakers of British dialects. Five housemates reside in the UK, but are not native speakers of British dialects. Sara and Darnell are native English speakers from Australia and the USA, respectively; Sylvia and Mohamed were born in Sierra Leone and Somalia, respectively, and speak near-native English with light accents. Kathreya is a native Thai speaker whose English is heavily accented and frequently ungrammatical.

## 2.2 Corpus description

The corpus consists of all segments from the Highlight and Diary Room Uncut shows where a single housemate was in the diary room, which we call *diary room clips* (or *clips*). Audio from all clips was broadcast quality. The 749 clips contain roughly 14.5 hours of housemate speech. The corpus is unbalanced among housemates (Table 1), because housemates were on the show for different amounts of time: the less time a housemate spent in the house, the less frequently their speech was sampled (clips occur once per 0.8–5.4 days, for different speakers). Thus, to address our research questions about time dependence in individual housemates’ use of each phonetic variable—in particular, distinguishing between by-day variability and time trends—we can only consider housemates who spend a relatively long period in the house. Fifty days (of 93) was chosen as an arbitrary cutoff, leaving 12 housemates (of 21), listed in Table 1, who we refer to as the *core housemates*. The core housemates account for 85.6% of speech in the corpus, and are the focus of most analyses below.

The corpus is limited to diary room clips of single housemates in order to best address our research questions, given that it was not feasible to transcribe more than a fraction of the produced episodes (>100 hours). Limiting the corpus to one type of interaction in a relatively constant setting allows us to minimize differences in speaking style between different days

---

<sup>4</sup>Information about the show and housemates come from Wikipedia (2012a,b). This corpus is an expanded version (roughly twice as large) of that used by Sonderegger (2012).

Table 1: Demographic information, length of stay in the Big Brother house, and amount of data in the Big Brother corpus, for each of the 21 housemates from Big Brother 9 UK. The horizontal line separates the 12 ‘core housemates’ (top) and other housemates.

Speaker	Gender	Age	Dialect region	Days on show	Number of clips	Housemate speech
Dale	M	21	N. England	65	43	0:43
Darnell	M	26	USA	93	73	1:24
Kathreya	F	31	L2 Thai	90	56	1:08
Lisa	F	40	N. England	86	53	1:02
Luke	M	21	N. England	58	49	1:09
Michael	M	34	S. Scotland	93	81	1:42
Mohamed	M	25	London	90	51	0:53
Rachel	F	24	S. Wales	93	69	1:03
Rebecca	F	21	W. Midlands	51	34	0:32
Rex	M	24	London	93	71	1:25
Sara	F	27	Australia	64	39	0:44
Stuart	M	25	N. England	57	32	0:38
Alexandra	F	23	London	14	17	0:25
Belinda	F	44	S. England	15	7	0:08
Dennis	M	24	S. Scotland	23	7	0:09
Jennifer	F	22	N. England	30	13	0:20
Mario	M	43	N. England	37	15	0:19
Maysoon	F	29	London	27	5	0:03
Nicole	F	19	S. England	22	19	0:24
Stephanie	F	19	N. England	9	9	0:08
Sylvia	F	22	London	23	6	0:09
Total					749	14:28

(i.e., clips), and thus better assess the stationarity hypothesis.<sup>5</sup> The only interaction in the clips is with Big Brother, whose role is usually limited to brief questions or answers. The register of the speech is generally casual and conversational, with characteristics of a sociolinguistic interview, but with the clear self-awareness and performativity expected given that the interaction may eventually be televised. The corpus allows us to examine how each housemate’s ‘baseline’ linguistic usage varies over time, abstracting away from short-term shifts which may occur during conversation with other housemates.

Each clip was segmented into speaker turns and orthographically transcribed by research assistants and the first author. The orthographic transcription and audio of each clip were then force-aligned using a version of the HTK-based aligner from FAVE (Rosenfelder et al., 2011; Young et al., 2006), customized for the Big Brother Corpus.

### 3 Data

We examine the dynamics of five aspects of pronunciation for speakers in the Big Brother house (*phonetic variables*): voice onset time (VOT), coronal stop deletion rate (CSD), and the quality of three vowels (GOOSE, STRUT, and TRAP; using the ‘lexical set’ notation of Wells, 1982 for a vowel’s realization in a given dialect).

These variables provide complementary evidence for our research questions about the existence and sources of medium-term time dependence in pronunciation. First, all variables differ greatly across varieties of English, meaning the possibility exists for contestants’ pronunciations to influence each other. VOT has been found to be very flexible over short-term timescales (e.g. Nielsen, 2011; Sancier & Fowler, 1997), making it a logical place to look for medium-term plasticity. Examining CSD lets us examine whether medium-term time change obtains for more categorical aspects of pronunciation, as well as continuous phonetic parameters. The three vowels examined differ along two dimensions which are possible sources of medium-term time dependence, and might affect the likelihood of convergence between housemates. Previous work has argued that the realization of STRUT and TRAP have high *social salience* across the UK, compared to GOOSE, which has been argued to constrain which variables can shift over the short and long term (e.g. Babel, 2010; Trudgill, 1986), and might affect a variable’s medium-term plasticity as well. GOOSE is undergoing *sound change* in communities across the UK (e.g. Docherty, 2010; Haddican et al., 2013; Hawkins & Midgley, 2005), while STRUT and TRAP are relatively stable. Speakers might show more flexibility in the realization of a vowel undergoing sound change, if their exposure to a greater variety of variants (outside the house) gives them a broader range of “self-exemplars” to draw on in converging towards other speakers.

We first introduce each phonetic variable, then describe annotation and data cleaning steps, and summarise the resulting dataset.

#### 3.1 Voice onset time

VOT, the time difference between the onset of a stop’s release burst and the onset of voicing in a following segment, is an important phonetic cue for the contrast between English ‘voiced’ and ‘voiceless’ stops (e.g. Docherty, 1992; Lisker & Abramson, 1967). VOT may

---

<sup>5</sup>Note that it was not possible to assess time of day (e.g. morning vs. evening) given information in the corpus; see Sec. 7.1 for discussion.

be positive (burst + aspiration duration) or negative (duration of voicing preceding burst). (Phonologically) voiceless stops are produced with variable degrees of positive VOT, while (phonologically) voiced stops are produced with either shorter positive VOT or negative VOT. Variability in VOT has been studied mostly in lab speech (see Auzou et al., 2000; Docherty, 1992), but some recent work examines spontaneous English speech (Chodroff et al., 2015; Yao, 2009), including Stuart-Smith et al. (2015)’s study of variability in positive VOT in spontaneous Glasgow vernacular—whose analytical choices we often follow, and who discuss the issues involved in measuring VOT in spontaneous speech. Like Stuart-Smith et al., we measured positive VOT (summed burst and aspiration duration) for every *word-initial* stop token produced with a burst or aspiration, without taking account of negative VOT or voicing during the closure. Thus, “voiced stops” throughout this paper always means *phonemically* voiced stops, as opposed to stops produced with negative VOT.

### 3.1.1 Annotation and dataset

We measured VOT semi-automatically for all word-initial stops in the corpus ( $n = 16784$  voiced, 13777 voiceless), using a procedure similar to Stuart-Smith et al. (2015).

First, an **automatic measurement** of VOT for each token was obtained by applying AutoVOT (Keshet et al., 2014; Sonderegger & Keshet, 2012). AutoVOT requires a classifier which has been trained on manually-labeled VOT data, as well as a window of time in the audio file for each token (specified in a Praat TextGrid; Boersma & Weenink, 2011) in which to search for the beginning of the VOT interval. Applying the classifier to each token in a file yields predicted VOT intervals, which are outputted on a new tier of the TextGrid. To predict VOT for voiceless and voiced stops, we used the classifiers for voiceless and voiced English stops distributed with AutoVOT. The window for each token was taken to begin and end 25 ms before and after the force-aligned segment boundaries. The parameters for minimum/maximum predicted VOT were set to 15/250 ms for the voiceless classifier and to 5/150 ms for the voiced classifier, and other algorithm parameters were kept at default values.

Second, in the **manual correction** phase, one of 3 phonetically-trained annotators (2/3 for voiced/voiceless stops), reviewed the predicted VOT intervals in Praat. For each interval, the annotator: (a) checked whether no burst was present (e.g. stop is realized as a fricative), in which case the token was marked for exclusion; (b) determined whether the VOT interval boundaries were where she would have placed them if annotating VOT manually; (c) manually corrected the boundaries, if this was not the case. Step (a) resulted in 3474 excluded tokens. Manual annotation (steps b–c) was performed in Praat. The left boundary was placed wherever the first “large” amplitude increase in high frequency frication occurred, established using the amplitude track and spectrogram; if there was a gradual rise in frication, the boundary was placed at the midpoint. When multiple clearly separate bursts occurred, the last one was used as the left boundary. The right boundary was determined primarily using the waveform: if periodicity was not present before the burst, at the zero crossing closest to the onset of periodicity; if periodicity was present throughout the burst, at the point where amplitude began to rise and general waveform shape changed abruptly.

A number of further exclusions were made. All 1511 tokens from Kathreya were excluded, since her data may show extensive transfer from Thai (which has a three-way VOT contrast for stops). 280 tokens with missing values for variables used in models below (listed in Table 2) were excluded. 80 tokens with VOT outside of 1–80 msec for voiced stops, or 8–175 for voiceless stops, were excluded as having extreme VOTs, with the cutoffs determined by

visual inspection of the distribution of VOT (separately for voiced and voiceless tokens). Because speech rate has a large effect on VOT relative to other variables, we excluded 62 tokens with speech rate greater than 10 syllables/second as having extreme speech rates (presumably due to forced-alignment errors), with the cutoff determined by visual inspection of the distribution of all tokens.

The final VOT dataset consists of 12908 voiced tokens (from 788 words, 678 clips) and 12246 voiceless tokens (from 964 words, 668 clips), across 20 speakers.

## 3.2 Coronal stop deletion

Coronal stop deletion (a.k.a. t/d deletion) is a variable process in English in which word-final coronal stops (/t/ and /d/) are sometimes deleted in word-final consonant clusters (e.g., *best* as [bɛs] vs. [bɛst]).<sup>6</sup> CSD has been examined in dozens of studies (usually) of spontaneous speech over the past 50 years across many varieties of English (reviewed by Hazen, 2011; Schreier, 2005; Tagliamonte & Temple, 2005). Fewer studies have examined CSD in British English varieties (e.g. Tagliamonte & Temple, 2005; Temple, 2009); particularly notable is Tagliamonte & Temple’s study of York English, whose analytical choices we often follow. Most previous work treats CSD as a binary variable, although “deletion” is clearly a categorical approximation for gradient realization of coronals (e.g. Temple, 2009). We use a categorical binary notion of “deletion” in order to make use of the results of previous work—including Tanner et al. (2016), where the CSD dataset described below is analyzed from a different perspective, unrelated to our research questions about time dependence.

### 3.2.1 Annotation and dataset

Annotation was carried out by 4 phonetically-trained annotators, for all speakers except Kathreya (20 speakers), whose data was excluded due to a near-categorical deletion rate (presumably because word-final clusters ending in t/d are phonotactically illegal in Thai). Each token ( $n = 14259$ ) for these 20 speakers whose underlying form ended in a t/d-final consonant cluster was manually annotated for the surface realization of the preceding and phonological context surrounding the final coronal, the length of any following pause (or other non-speech), the CELEX wordform ID (which was used to determine morphological class), as well as realization of the final coronal.<sup>7</sup>

Coronal stop realization was first annotated at a fine-grained level, using spectral and auditory criteria. Annotators chose from eight categories describing possible surface realizations of t/d (using the acoustic cues of burst, glottalization, sudden closure; realization as a glottal stop, an unreleased coronal, etc.). In cases where the underlying t/d was followed by a surface t/d (e.g. *want to*), the target t/d was only taken to be realized in cases where there was evidence for realization of a sequence of two distinct coronal stops. If the second t/d was clearly present following a closure (usually due to the presence of a burst), and no other evidence of a previous t/d realization was present (such as a glottal stop), the t/d was taken to be the realization of the following word, and the word-final t/d was annotated as unrealized.

---

<sup>6</sup>CSD is one type of final consonant cluster reduction (CCR) in English, and the most commonly studied. See Hazen (2011) on the relationship between CSD and CCR.

<sup>7</sup>We did not make exclusions such as frequent (e.g. *and*) or contracted (e.g. *don’t*) forms, as in much previous CSD work, instead accounting for lexical differences using by-word random effect terms.

The fine-grained annotation was then collapsed into a binary variable of present versus deleted, following previous work (e.g. Tagliamonte & Temple, 2005), where the presence of any unambiguous phonetic reflex of the underlying t/d indicates presence. We also followed Tagliamonte & Temple (2005) in discarding all tokens ending in /rt/ or /rd/ in rhotic varieties, given that many housemates are speakers of non-rhotic varieties.

We further excluded 141 tokens without a preceding consonant in their surface realization, 22 tokens where speech rate could not be reliably determined, 1 token with annotation errors, and 59 tokens at turn ends. This last step was necessary in order to include pause duration and following context as *independent* factors affecting CSD rate, following Tanner et al. (2016), where the motivation for this choice (rather than coding pause as an alternative following context) is discussed.

The final dataset consists of 12788 tokens (from 570 words, 667 clips), across 20 speakers.

### 3.3 Vowel formants

Accents of English differ primarily in their vowels (Wells, 1982). Within a given variety of English, vowel quality is determined primarily by the first two formants (F1, F2). We consider only variation in F1 and F2, for GOOSE, STRUT, and TRAP.

The realization of these vowels varies systematically across English dialect regions; this variation will be important for assessing whether factors such as social salience affect medium-term time dependence. We describe the expected pronunciations of each vowel in the English dialect regions for each of the 12 core speakers (Table 1)—Southern England (including London), Northern England, West Midlands, Southern Wales, Southern Scotland, General American, and General Australian (Cox & Palethorpe, 2007; Ferragne & Pellegrino, 2010; Kortmann & Schneider, 2004; Watson, 2007; Wells, 1982)<sup>8</sup>—as well as the range of realizations that might be expected for Kathreya, based on studies of Thai learner English, or transfer from her L1 (e.g. Tsukada, 2008). Figure 4 shows the approximate correspondence between IPA symbols used here and location in F1/F2 space.

GOOSE is subject to different degrees of fronting ([u]~[u:]~[y]) across varieties of English. In S. Scotland and General Australian, a fronted variant has long been the norm, while in Standard Southern British English (SSBE), Northern English, W. Midlands English, and many American varieties, fronting is a change in progress (e.g. Docherty, 2010; Haddican et al., 2013; Hawkins & Midgley, 2005). An exception is Welsh English, where a back [u] is the norm. We thus expect GOOSE to be fronted to different degrees for all native speakers, with the exception of Rachel (from Wales). We also expect Kathreya to show a less-fronted GOOSE than native speakers (close to [u]).

STRUT realization is one of the most characteristic features dividing varieties from the North and South of England. STRUT is merged with FOOT (characteristically [ʊ]) in Northern English, and is strongly socially marked (Wells, 1982); in accents of the West Midlands, the two are variably merged, with variable realizations ([ʊ]~[ɻ]~[ə]). They are not merged in other English varieties represented in our dataset, in which STRUT is realized as a lax mid-low vowel ([ʌ]~[ɐ]~[ə]). This is also the expectation for Kathreya.

TRAP is typically a front [æ] in General American and Australian. In the UK, accents other than SSBE considered here tend to use a more centralized pronunciation ([a]). Tra-

---

<sup>8</sup>“General American” is an approximation of the dialect of the single American speaker (Darnell), an African American who was born in the UK, grew up in St. Louis, Missouri, and moved back to the UK as an adult (Wikipedia, 2012a).

ditionally [æ] is used in SSBE, but younger speakers have shifted towards [a] (Hawkins & Midgley, 2005). Kathreya is expected to realize TRAP somewhere along the [æ]~[a] continuum.

In terms of social salience, previous work has argued that STRUT is highly salient (especially within England), while TRAP has medium social salience, and GOOSE has low social salience, in the sense of different degrees of fronting (e.g. Haddican et al., 2013; Wells, 1982).

### 3.3.1 Annotation and datasets

The vowel datasets consist only of data from the 12 core speakers. We considered all tokens produced by these speakers of the target vowels, defined as those whose reference pronunciation was [u], [ʌ], or [æ] in a word’s “primary pronunciation” in CELEX (Baayen et al., 1996). To avoid heavily-reduced vowels, we excluded tokens whose force-aligned durations were less than 30 msec, tokens from a list of highly-frequent words (“and”, “um”, “just”, “uh”, “to”), and tokens for vowels whose reference pronunciation is unstressed in the FAVE pronunciation dictionary. This procedure resulted in 4580 tokens for GOOSE, 6323 for STRUT, and 7903 for TRAP, for which we measured F1 and F2 semi-automatically.

In the **automatic measurement** step, we obtained F1 and F2 for each token, using the same customized version of the FAVE suite used for forced alignment (see Sec. 2.2). FAVE was set to use the `faav` method to determine measurement points,<sup>9</sup> and to perform ‘remeasurement’ within each file; other parameters were kept at default values. Following automatic measurement, F1 and F2 were Lobanov-normalized within each speaker (Lobanov, 1971), using the speaker’s entire automatically-measured vowel space, to control for physiological differences between speakers. We use normalized formants in all analyses below, usually referred to simply as F1 and F2.

**Manual correction** was carried out by two phonetically-trained annotators using FVR (Schwartz, 2015), a graphical program for filtering and remeasuring vowel formant data. The goals of manual correction were to exclude tokens realized using a “different” vowel from the target, including highly reduced tokens, and to verify and correct as necessary the formant measurements of the remaining tokens. All candidate tokens of a vowel from a given speaker were plotted in an F1/F2 display in which it was possible to hear each token’s corresponding audio. Each token was examined, perceptually judged, and either excluded or accepted. Tokens were excluded which were (perceptually) deleted or heavily reduced; produced during yelling, disfluencies, or with heavy glottalization; or were realized (perceptually) as anything besides a whitelist including all expected realizations in dialects represented in the house. All remaining tokens were accepted. If an accepted token’s formants were judged to be possibly mismeasured (judging by position in F1/F2 space, auditory judgement), they were checked in Praat, and manually remeasured if necessary. 1716 tokens were excluded for GOOSE, 1562 for STRUT, and 2482 for TRAP.

Finally, the first author examined F1/F2 plots of each speaker’s data for each vowel looking for any remaining extreme outliers in formant measurements, or words labeled with the wrong lexical class. This led to exclusion of 13 GOOSE, 9 STRUT, and 36 TRAP tokens.

The final vowel formant dataset consists of F1 and F2 measurements from 12 speakers for 2847 tokens for GOOSE (from 200 words, 588 clips), 4743 for STRUT (from 345 words, 591 clips), and 5365 for TRAP (from 434 words, 615 clips), across 12 speakers.

---

<sup>9</sup>Usually, this corresponds to one-third of the vowel’s duration—including for the three vowels considered here, except GOOSE after coronal consonants, for which the vowel beginning is used.

Table 2: Summary of static factors included in the static models of each phonetic variable.

<i>Variable</i>	<i>Static factor</i>	<i>Type</i>
VOT	Place of articulation	factor (bilabial, coronal, velar)
	Following segment type	factor (vowel, sonorant)
	Following vowel	factor (non-high, high)
	Syllable stress	factor (unstressed, stressed)
	Phrase position	factor (phrase-initial, phrase-medial)
	Speech rate deviation	continuous
	Speech rate mean	continuous
	Word frequency	continuous
	Gender	factor (male, female)
CSD rate	Following context	factor (vowel, other consonant, t/d)
	Preceding context	factor (other obstruents, sonorants, sibilants)
	Pause duration	continuous
	Morphological class	factor (reg. past, irreg. past, non-past)
	Word frequency	continuous
	Speech rate	continuous
	Gender	factor (male, female)
Vowel formants	Preceding segment	factor (oral apical, nasal apical, oral labial, nasal labial, liquid, obstruent+liquid cluster, palatal, velar, w/y, other)
	Following segment manner	factor (affricate, fricative, lateral, nasal, rhotic, stop, other)
	Following segment place	factor (alveolar, bilabial, interdental, labiodental, palatal, velar, other)
	Following segment voicing	factor (voiced, unvoiced, other)
	Vowel duration	continuous

### 3.4 Static factors

In building models of time dependence in each phonetic variable below, we control for a variety of factors besides time which affect phonetic realization (linguistic factors, social factors, properties of the utterance), which we call *static factors*, such as speech rate and the identity of surrounding segments. Table 2 summarizes the static factors for each phonetic variable. Because their effects are not directly related to our research questions about time dependence, previous work on static factors and how we control for them statistically are described in Appendix A.

## 4 Analysis

From an analysis perspective, it is easiest to divide the datasets of phonetic variables just described into nine *narrow variables*, each corresponding to a single phonetic parameter that may behave independently of the others: VOT for voiced and voiceless stops (2), CSD (1), and (normalized) F1 and F2 for each vowel (6). We now describe the analysis carried out on

the data for of narrow variable, culminating in one *dynamic model* per speaker per narrow variable, describing time dependence in the variable during the speaker’s time in the Big Brother house.

## 4.1 Method

The analysis of each narrow variable is conceptually similar. For each speaker, we fit four statistical models of the speaker’s realization of the variable, each assuming one of the types of time dependence shown in Figure 1, while controlling for static factors in the same way. These models are then compared to evaluate which qualitative type of time dependence best characterizes this variable, for this speaker. The best model is chosen, resulting in *one* model of time dependence per variable, per speaker. For example, we fit four models of VOT for voiceless stops for Rachel, controlling for the static factors listed in Table 2: one model assumes no systematic time dependence in VOT; one assumes some time trend (mean VOT changes over time) but no additional by-day variability; one assumes by-day variability in VOT around an unchanging mean value (no time trend); and one assumes both by-day variability and a time trend.

At a technical level, the analysis of each narrow variable uses both mixed-effects regression models (MEMs; e.g. Baayen et al., 2008; Gelman & Hill, 2007) and generalized additive mixed models (GAMMs; Hastie & Tibshirani, 1990; Wood, 2006), fitted using the `lme4` and `mgcv` packages in R (Bates et al., 2014; R Core Team, 2014; Wood, 2011). A MEM is first fitted for each variable for data from *all* speakers, as a function of static factors only. These *static models* allow us to determine, with maximum statistical power, which static factors should be controlled for in modeling time dependence within *individual* speakers. The results of the static models serve as input to GAMMs modeling the time dependence of each variable within individual speakers, which we refer to as *dynamic models*.

Before describing the modeling procedure for each variable, we discuss our choice of this particular analysis method, and briefly introduce GAMMs.

### 4.1.1 Choice of analysis method

Some discussion is warranted of our analysis method, as the steps of splitting up the data by speaker and dividing the analysis into two steps both carry risks. For our data, we would ideally build a *single* model for each phonetic variable, across all speakers, of both static factors and time. This was not possible because speakers turn out to show qualitatively different patterns of time dependence, making it necessary to consider four possible types of time dependence (Figure 1) for each speaker. Finding the best possible combination of types of time dependence for all 12 speakers within a *single* model would require comparing prohibitively many models ( $4^{12}$ ). Thus, we decided to compare possible models of time dependence within each speaker *separately*, for a given variable.

We considered two ways of doing this: (1) building four GAMM models accounting for both static factors and time dependence within an individual speaker, and (2) building four GAMM models of time dependence for an individual speaker using the residuals from a static model fit to all speakers. Method 1 has the disadvantage of less accurate and lower-power estimates of the effects of static factors, for which the data from *all* speakers is relevant; this could lead to controlling less accurately for the effects of static factors, and possibly underestimating time dependence within a speaker as a result. We use Method 1 for the CSD data, where Method 2 is conceptually impossible (as discussed below). Method 2 has

disadvantages associated with any case of “residual regression”, where the effect of one set of predictors  $X_1$  (here, static factors) on a dependent variable  $Y$  is “regressed out”, then the residuals are modeled as a function of another set of predictors  $X_2$  (here, time)—as opposed to building a single model,  $Y \sim X_1 + X_2$ . Residual regression can lead to bias in the estimated effect of  $X_2$ , which can either be conservative or anti-conservative depending on factors such as the correlation between  $X_1$  and  $X_2$  (Darlington & Smulders, 2001; Freckleton, 2002). Despite these issues, residual regression is widely used in language research, for example to control for the effect of word length on reading times in self-paced reading studies (e.g. Ferreira & Clifton, 1986; Fine et al., 2013), to remove by-speaker and by-item variation in acoustic parameters before examining the effects of predictors of interest on prosody (Breen et al., 2010), or in short-term and long-term studies analyzing change over time in VOT within speakers after controlling for other factors (Stuart-Smith et al., 2015; Yu et al., 2013). Here, we use Method 2 for the VOT and vowel formant data, in order to control for static factors with maximum statistical power.

#### 4.1.2 Generalized additive mixed models

While MEMs are now widely used in language research, GAMMs may be less familiar (for introductions see Winter & Wieling, 2016; Wood, 2006). GAMMs are an extension of MEMs: the dependent variable can still depend on factors and continuous variables (fixed effect terms), which may vary across groups (random effect terms), but it can now also depend on a *smooth* of one or more independent variables: a nonparametric function which locally fits the data, conceptually similar to a nonlinear smoother. GAMMs are closely related to Generalized Additive Models (GAMs), which similarly extend non-mixed regression models (e.g. linear and logistic regression) by allowing for smooth terms. Applications of these methods to linguistic data, to model variability across time and space, include event-related potentials (Kryuchkova et al., 2012), dialect variation (Wieling et al., 2014), and language evolution experiments (Winter & Wieling, 2016). The crucial aspect of GAMMs for our purposes is their ability to model largely arbitrary patterns of variability over time—here, in a phonetic parameter.

GAMMs allow for incorporation of two types of terms which conceptually correspond to the two types of time dependence which differentiate the four patterns in Figure 1: (1) a random intercept of `Clip`, which captures by-day variability (because different clips for a given speaker generally occur on different days); (2) a smooth of `Day`, which captures any time trend. Thus, we build four models per speaker/variable pair: one without (1) or (2), one with term (1) only, one with term (2) only, and one with both terms (1) and (2). These models are compared using the Akaike Information Criterion (AIC), a widely-used model selection criterion.<sup>10</sup> The model with the lowest AIC is chosen, resulting in a single model of time dependence for a given narrow variable per speaker, after controlling for static factors.<sup>11</sup>

---

<sup>10</sup>Best practices for model selection for GAM(M)s are not settled, especially with respect to selecting random-effect and smooth terms, as in the models compared here, for which various statistical issues arise (see Pinheiro & Bates, 2000; Wood, 2006). Comparison via AIC is one of several possible model selection procedures. A reviewer recommends model comparison via a  $\chi^2$  test on the difference in ML score, which can be more robust than AIC, but cannot be applied when comparing non-nested models (here, corresponding to types B and C of time dependence). We use AIC based on the discussion in the `gam.selection` page of Wood (2015), but note that methodology for model selection for longitudinal phonetic data is an important issue for future work.

<sup>11</sup>An important consideration in regression modeling of longitudinal data is the possibility of autocorrelated residuals due to non-independence of consecutive time points (see e.g. Pinheiro & Bates, 2000 Ch. 5); this often

## 4.2 Building dynamic models

We now summarize how the analysis was carried out for each variable, using the datasets described in Sec. 3.

**Voice onset time** The dependence of VOT on static factors was first determined by fitting two linear mixed-effects models of  $\log(\text{VOT})$  to the VOT data from the 20 speakers, for voiced and voiceless stops. These static models included (1) fixed-effect terms for all static factors (Table 2), including interactions based on exploratory data analysis, as well as annotator identity; (2) near-maximal by-speaker and by-word random effects (Barr et al., 2013). The model-fitting procedure and results are detailed in Appendix B.1.

We then built dynamic models of time dependence of VOT within the 11 core speakers, controlling for static factors, by modeling the *residuals* of the static models of VOT, offset for each speaker by the static model’s estimate of their “mean VOT” (the model intercept, plus the speaker’s fitted random intercept). These offset residuals, which we call *residualized VOT* have the interpretation for each token of “VOT, after controlling for static factors”.<sup>12</sup> For each speaker, for each type of stop (voiced, voiceless), we fit four dynamic models of residualized VOT, corresponding to the four possible types of time dependence. Each model is a GAMM with a gaussian link, and includes a by-word random intercept, as well as appropriate smooth terms (as described above). The model with the lowest AIC was selected from each set of four models. This procedure resulted in 22 models of time dependence (11 speakers  $\times$  {voiced, voiceless}) of VOT within individual speakers.

**Coronal stop deletion** Tanner et al. (2016) fit a mixed-effects logistic regression model of CSD realization for the dataset described above from 20 speakers. We use the results of this static model only to *motivate* the static factor terms included in the dynamic models of CSD (see below), rather than its residuals serving as input to the dynamic models (as for VOT); thus, only aspects of the model relevant for the dynamic models are summarized here. As fixed effects, the model included main-effect terms for all static factors (Table 2), as well as a term for annotator identity, and other main-effect and interaction terms motivated by the research questions of Tanner et al. (2016). Each variable was coded such that the model intercept has the interpretation of a grand mean (and this coding was retained for the dynamic models): factors were coded using Helmert contrasts, and continuous variables were standardized. Prior to standardizing, 0.01 seconds was added to pause duration and both pause duration and speech rate were log-transformed. Main-effect terms for six static factors—following and preceding context, word frequency, speech rate, pause duration, and annotator identity—reached significance, as well as several interaction terms, of which pause duration-by-following context had a much larger effect size than other interactions.

We built dynamic models of time dependence of CSD rate for the 11 speakers (12 core speakers, minus Kathreya), after controlling for these seven terms. Unlike for VOT, the residuals of the static model of CSD realization do not have the interpretation of “variable

---

occurs for linguistic data from laboratory experiments (van Rij et al., 2016). Examining residual autocorrelation is not straightforward for our data due to its complex temporal structure (irregularly spaced time points, many tokens per time point, i.e. per clip). Nonetheless, as a basic check, we examined autocorrelation plots of simplified versions of the predictions of each dynamic model, and found no evidence for significant autocorrelation of residuals, possibly because measurement points are spaced much further apart (days) than in experimental data.

<sup>12</sup>More precisely: after controlling for static factors, as well as variability among words, since the residuals factor out by-word random effects.

value, after controlling for static factors”, which cannot be calculated for individual tokens for a binary variable (see e.g. Agresti, 2002, Sec. 6.2). To model time dependence in CSD rate for a given speaker while controlling for static factors, we instead built models for that speaker’s data including terms for both time dependence and static factors. For each speaker, we fit four dynamic models of CSD realization, each one a GAMM with a logit link, including smooth terms capturing time dependence. Each model included the seven terms for static factors, and a by-word random intercept.<sup>13</sup> For each speaker, the model with the lowest AIC was selected, resulting in 11 models of time dependence (one per speaker).

**Vowel formants** Six linear mixed-effects models of normalized formants were first fitted to the data from the 12 speakers for which vowel formant data was obtained—two models (for normalized F1, F2) for each of GOOSE, STRUT, and TRAP. These static models included terms for all static factors (Table 2), parametrized in a different way than the VOT and CSD models, due to the relative sparsity of data (fewer tokens; see Appendix B.2). The models included (1) fixed-effect terms for the effects of following consonant voicing and vowel duration; (2) three *random intercept* terms for the effects of following consonant place, following consonant manner, and preceding segment class; (3) by-speaker and by-word random intercepts; (4) a by-speaker random slope for vowel duration. The model-fitting procedure and results are detailed in Appendix B.2.

We then built dynamic models of time dependence in F1 and F2 within the 12 core speakers, after accounting for static factors, in a manner similar to VOT, by modeling the residuals of the static models, offset for each speaker by the static model’s estimate of their mean formant value. These *residualized formants* have the interpretation within each speaker of “F1 for GOOSE, after controlling for static factors” (etc.). For each speaker, for each variable, the model with the lowest AIC was selected, resulting in 72 models of time dependence (12 speakers  $\times$  {F1, F2}  $\times$  {GOOSE, STRUT, TRAP}) of vowel formants within individual speakers.

### 4.3 Summary

The procedure just described resulted in 105 dynamic models (22/11/72 for VOT/CSD/vowels) describing the dynamics of each phonetic variable in each core speaker (with the exception of VOT and CSD for Kathreya, the non-native speaker). We present and discuss the results of these models in the next two sections, corresponding to our two research questions: what time dependence does pronunciation show within individual speakers over days-months (Sec. 5), and to what extent can we explain these patterns (Sec. 6)?

## 5 Time dependence

In this section, we describe each models’ predictions of the existence and nature of time dependence in each variable. For each variable, we plot the models’ predicted patterns of time dependence, in order to visualize the extent of time dependence beyond any effect of static

---

<sup>13</sup>Static factors were coded in the GAMMs as in the regression model of Tanner et al. (2016), with two exceptions. The the effect of speech rate was modeled using a smooth term (as for Day), based on exploratory plots suggesting that speakers showed qualitatively different speech rate effects. Pause duration was modeled as a log-transformed continuous variable—rather than discretized into four levels as in Tanner et al. (2016), which was motivated by the research questions of that study.

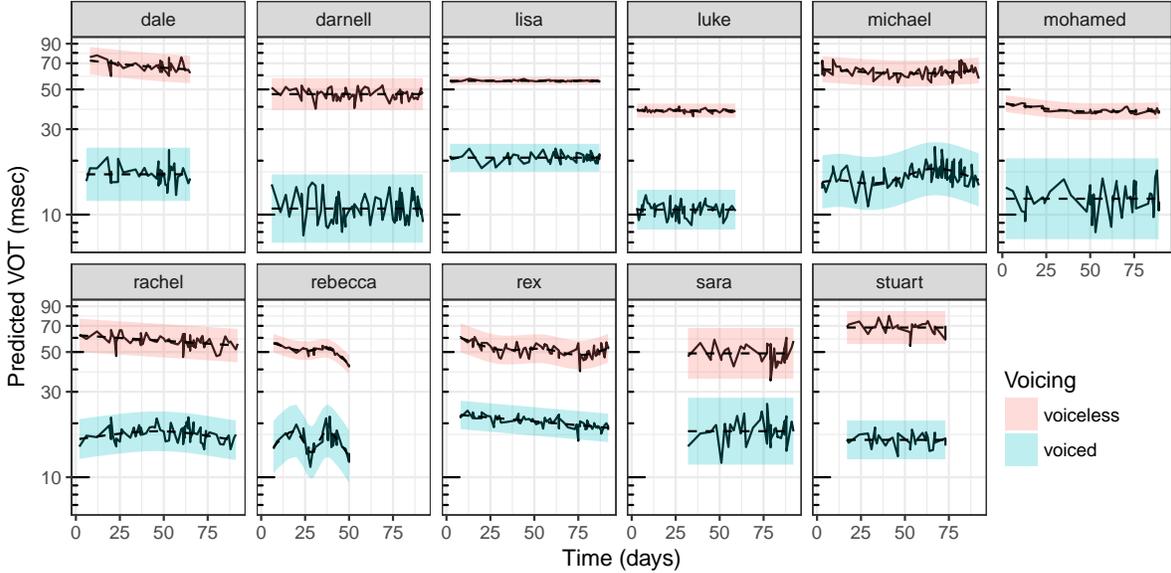


Figure 2: Predicted time dependence of VOT (in msec, on log scale) for each speaker, for voiced and voiceless stops. Solid lines indicate model predictions for each Day, including by-day variability. Dotted lines indicate predicted time trend of mean VOT (without any by-day variability). Shading indicates magnitude of by-day variability, between “low” and “high” days. (See text.)

factors or random noise.<sup>14</sup> For each variable, we also summarize time dependence in terms of the four-way categorization discussed above (Figure 1). A more detailed summarization of the predicted patterns of time dependence for each variable for each speaker, including quantitative measures, is given in *Supplemental Information B.2*, but is not assumed here.

## 5.1 Voice onset time

Figure 2 shows the model-predicted time dependence in residualized VOT for voiced and voiceless stops for each speaker, which we call *predicted VOT*. Because of how residualized VOT was computed, and the coding of static factors included in the static models (see Appendix B.1), predicted VOT can be interpreted as a speaker’s “baseline VOT” on a given day (averaging across words, spoken at the speaker’s average speech rate, etc.). The dotted line in each panel shows the predicted time trend: the mean predicted VOT on a given day, without taking by-day variability into account. A flat dotted line indicates no time trend. When a model predicts by-day variability, the estimated magnitude of these daily fluctuations is expressed by a fitted parameter in the model, denoted  $\sigma$ , such that roughly 95% of days (i.e., clips) are predicted to have  $\log(\text{VOT})$  within  $\pm 2\sigma$  of the mean.<sup>15</sup> The amount of by-day variability is shown in each panel by shading: the vertical range of shading shows the

<sup>14</sup>We plot model predictions rather than empirical means in all figures showing time dependence (Figures 2, 3, 5–8) in this paper because these predictions essentially *are* the results of the models which are relevant for our research questions—which the figures report in graphical form. (Unlike regression models such as MEMS, GAMMs do not include interpretable numerical coefficients for smooths or random effect terms.) Plots based on empirical means generally look very similar, but are not reported here due to space constraints.

<sup>15</sup>Formally,  $\sigma$  is the variance component for the by-clip random intercept.

Table 3: Summary of the qualitative type of time dependence each speaker shows for each variable, using notation from Figure 1: no change (A), time trend (B), by-day variability (C), time trend and by-day variability (D).

Speaker	VOT		CSD rate	GOOSE		STRUT		TRAP	
	Voiced	Voiceless		F1	F2	F1	F2	F1	F2
Dale	C	D	C	D	C	C	D	C	D
Darnell	C	C	C	C	D	D	C	C	D
Kathreya	–	–	–	C	C	D	D	D	D
Lisa	C	C	A	C	A	C	C	C	D
Luke	C	C	C	C	C	A	D	C	D
Michael	D	D	D	D	C	C	D	C	D
Mohamed	C	D	C	C	C	C	D	C	C
Rachel	D	D	C	C	D	C	C	D	C
Rebecca	D	D	B	C	D	C	D	D	C
Rex	D	D	C	C	C	C	C	C	D
Sara	C	C	B	D	C	C	C	D	C
Stuart	C	C	D	B	C	D	C	C	A

predicted range of VOT between a “high day” and a “low day” ( $\pm 2\sigma$  from mean, represented by the dotted line). Finally, it is also possible to extract from the model an estimate of the size of a fluctuation in VOT on a given day (the best linear unbiased predictor; Pinheiro & Bates, 2000, Sec. 2.2). Combining this estimate with the predicted mean gives the solid line, which shows the actual predicted VOT on a given day. The qualitative type of time dependence in VOT for each speaker is summarized in Table 3 (columns 2–3).

*Results:* All speakers show some time dependence in VOT, for both voiced and voiceless stops: by-day variability, a time trend, or both. In particular, every speaker shows by-day variability in VOT, for both voiced and voiceless stops. Of the 11 speakers, 6 (54.5%) show a time trend in predicted VOT. This corresponds to changes in both voiced and voiceless stops for 4 speakers, and in voiceless stops only for 2 speakers.

## 5.2 Coronal stop deletion

Figure 3 shows the model-predicted time dependence in CSD rate for each speaker (analogously to Figure 2), holding all static factors at their average values, and ignoring by-word random effects; the predictions can be interpreted as “average CSD rate” (for an average word, at a speaker’s mean rate, etc.). In each panel, the dotted line shows the predicted time trend in mean CSD rate (if any). The magnitude of fluctuations in CSD rate is parametrized by  $\sigma$  such that roughly 95% of days are predicted to have CSD rate (expressed in log-odds) within  $\pm 2\sigma$  of the mean. Shading shows the amount of by-day variability in CSD rate ( $\pm 2\sigma$  around the mean), and the solid line shows actual model-predicted CSD rate on each day. The qualitative type of time dependence in CSD rate for each speaker is summarized in Table 3 (column 4).

*Results:* Ten of the 11 speakers (91%) show some time dependence in CSD rate: by-day variability (6), a time trend (2), or both (2). 8 speakers (73%) show by-day variability. 4 of 11 speakers show time trends.

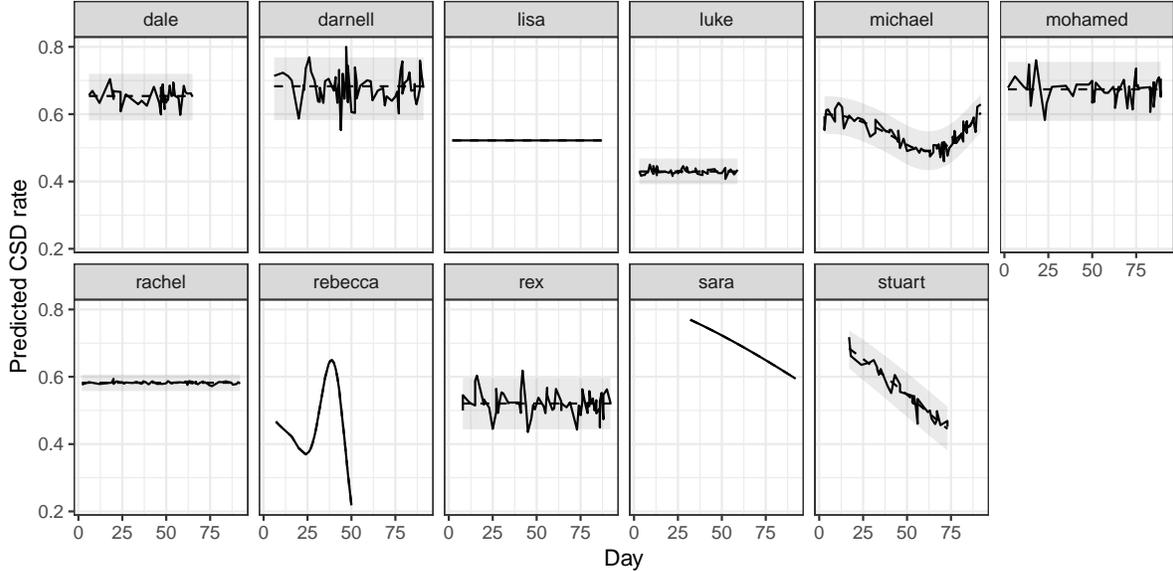


Figure 3: Predicted time dependence of CSD rate for each speaker. Solid/dashed lines and shading are defined as in Figure 2.

### 5.3 Vowel formants

Before turning to the results of the dynamic models for the vocalic variables, we examine the model-predicted mean locations of each vowel: how each speaker realizes GOOSE, STRUT, and TRAP, on average, after controlling for static factors (i.e., residualized F1 and F2). How speakers differ in realizing these vowels will be important in interpreting time dependence in F1 and F2 below. Figure 4 shows each speaker’s predicted mean F1 and F2 for each vowel, with superimposed IPA symbols showing the approximate phonetic realization corresponding to different parts of the F1/F2 space.

For GOOSE, the main division is between Kathreya, who shows a very backed [u], and native speakers (all others). Among native speakers, Rachel shows the most backed realization, as expected for a Welsh English speaker. All other speakers have markedly fronted GOOSE, to varying degrees. These realizations are all consistent with our expectations (Sec. 3.3).

For STRUT, we can distinguish three groups of speakers. Dale, Lisa, Luke, and Stuart—the speakers from Northern England—have markedly higher and backer STRUT than other speakers, as expected given that they have merged STRUT and FOOT (usually realized [ʊ]). Rachel and Rebecca realize STRUT as more centralized on average than other speakers, while the last group—Darnell, Kathreya, Michael, Mohamed, Rex, Sara—realize STRUT as more low and back, something like [ʌ].

For TRAP, housemates vary roughly along a line from [æ] to [a]. We can nonetheless distinguish three groups, for purposes of exposition: Darnell and Kathreya, Sara/Rex/Mohamed, and all others (Dale, Lisa, Luke, Michael, Rachel, Rebecca, Stuart). The third group consists of speakers from regions where the norm is [a], all of whom indeed realize TRAP closer to [a] than [æ]. In the second group, Mohamed and Rex’s intermediate realization is expected given their dialect area (S. England), while Sara’s realization is unexpectedly far from [æ]. In the first group, Darnell’s TRAP is clearly [æ], as expected for an American speaker; Kathreya’s realization is similar, within the range of realizations expected for Thai-accented English.

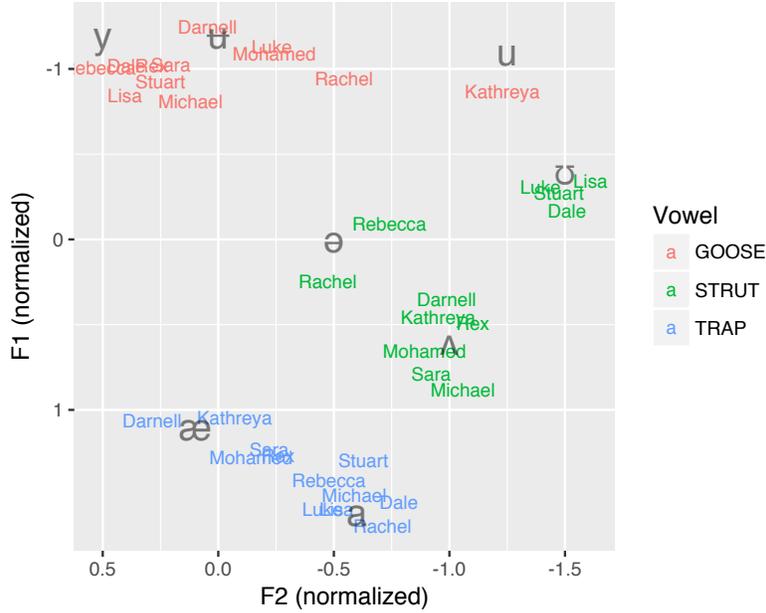


Figure 4: Predicted mean values of F1 and F2 (normalized) for GOOSE, STRUT, and TRAP, for each speaker (shown using names). Superimposed IPA symbols show approximate locations of main phonetic variants used for these vowels across dialects (see Sec. 3.3).

In sum, housemates’ mean realization of the three vowels fit our expectations given each speaker’s dialect background (Sec. 3.3).

### 5.3.1 Time dependence

We now turn to the dynamic models. Figure 5 shows the model-predicted time dependence in F1 and F2 for each speaker for each vowel, which we call *predicted F1/F2*. Analogously to the VOT and CSD cases, predicted F1 and F2 can be interpreted as a speaker’s “baseline F1 and F2” on a given day (at an average vowel duration, and averaging over possible types of preceding and following contexts). In each panel, the dotted line shows the predicted time trend in mean F1 and F2, vertical shading illustrates the magnitude of by-day variability between “high days” and “low days” (again characterized by the parameter  $\sigma$ , such that 95% of days have formant values within  $\pm 2\sigma$  of the mean), and solid lines show model-predicted formant values on each day. The qualitative type of time dependence in F1 and F2 for each vowel for each speaker is summarized in Table 3 (columns 5–10).

*Results:* All speakers show some time dependence in the realization of all three vowels (taking F1 and F2 together). In particular, all speakers show by-day variability in the realization of all three vowels: either F1, F2, or both are predicted to fluctuate from day to day around a mean value. For the vast majority of speaker/vowel pairs (32 of 36), both F1 and F2 show by-day variability. A subset of speakers show time trends in the realization of each vowel (in at least one of F1 and F2) during the season: 7 speakers for GOOSE (58%), 8 speakers for STRUT (67%), and 10 speakers for TRAP (83%).

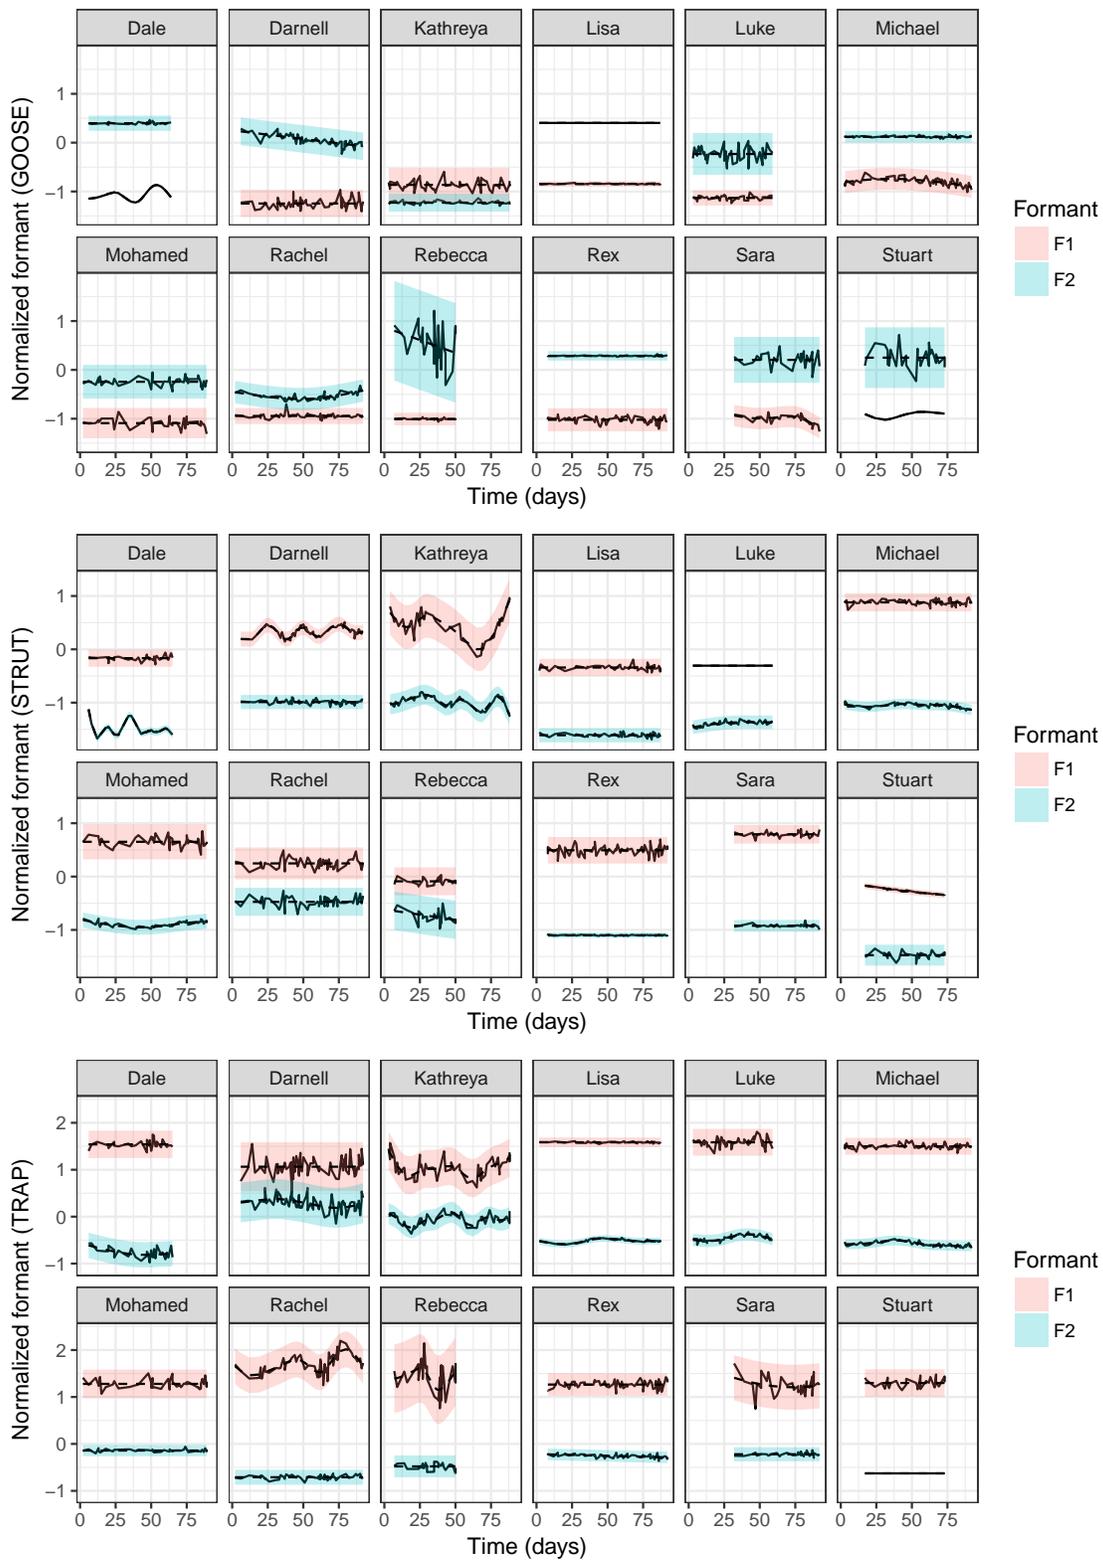


Figure 5: Predicted time dependence of (normalized) F1 and F2, for each speaker, for GOOSE, STRUT, and TRAP. Solid/dotted lines and shading are defined as in Figure 2.

Table 4: Number of speakers showing each type of time dependence (Figure 1), for each variable: none, time trend only, by-day variability only, or by-day variability and time trend.

Variable	Type of time dependence			
	None (A)	TT only (B)	BDV only (C)	BDV & TT (D)
VOT (voiced)	0	0	7	4
VOT (voiceless)	0	0	5	6
CSD	1	2	6	2
GOOSE	0	0	5	7
STRUT	0	0	4	8
TRAP	0	0	2	10

## 5.4 Summary

We can make several observations about medium-term time dependence in pronunciation based on Table 4, which summarizes this section’s results. First, across all variables, time dependence is ubiquitous: in all but one case (68 cases; 98.5%), a speaker’s pronunciation fluctuated daily, showed steady change in mean realization over time, or both. Medium-term time dependence takes on very heterogeneous forms across speakers and variables, but the *existence* of some time dependence is constant. Second, by-day variability in pronunciation is nearly ubiquitous (95.65% of cases). In particular, for every *gradient* variable (VOT, vowel formants), every speaker showed daily fluctuations in pronunciation; the two cases without by-day variability are for CSD. Third, time trends are often present, but are markedly less common than by-day variability (43.48% of cases), both within and across variables. For each variable, some speakers did not show steady change in mean pronunciation over time—at least not change that had a large enough effect size large enough for us to detect.

## 6 Sources of time dependence

So far, we have found that time dependence in phonetic variables over days to months is near-ubiquitous. The observed patterns of time dependence are also very heterogeneous (Figures 2, 3, 5). We now turn to our second question: can we account for any of the variability in patterns of medium-term time dependence? Short-term time dependence in phonetic variables is influenced by linguistic factors, social factors, and individual differences (Sec. 1.1), suggesting medium-term time dependence may be as well. We consider three possible sources of time dependence of these types: *contrast* between sounds corresponding to different variables; *convergence* in pronunciation of individual variables; and differences in *speaker plasticity*, across all variables.

### 6.1 Contrast

The contrastive role played by phonetic variables, including vowel formants (e.g. Babel, 2011), have been invoked to explain the results of short-term experiments where they are examined (Mitterer & Ernestus, 2008; Nielsen, 2011). Here, we consider one way in which variability over time in vowel formants might be constrained by their use as acoustic cues differentiating

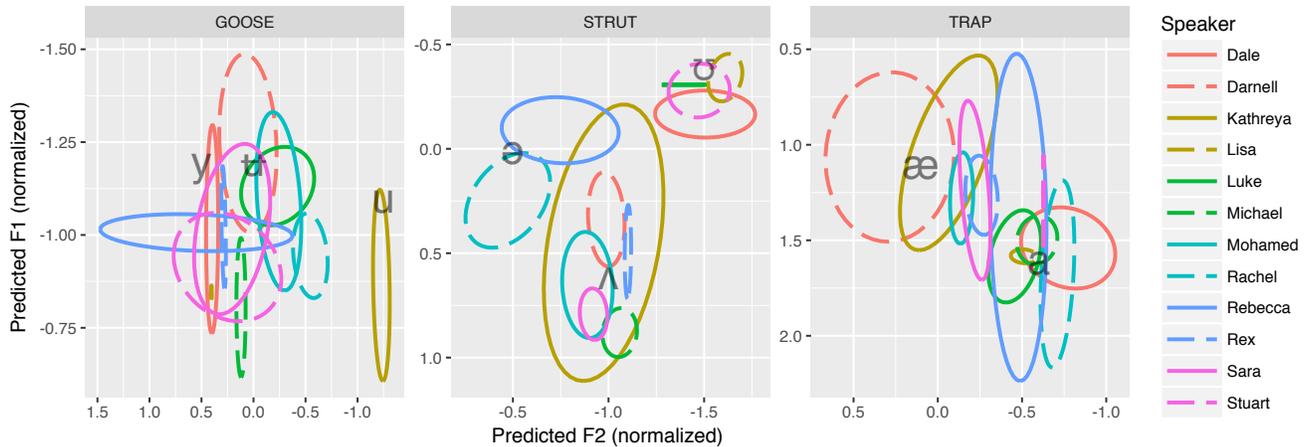


Figure 6: 95% confidence ellipses for predicted realizations of GOOSE, STRUT, and TRAP over time for each speaker. Superimposed IPA symbols show approximate locations of main phonetic variants used for these vowels across dialects.

speech sounds: variability in a speaker’s realization of a sound might be constrained by the variability in its realization across English dialects.

The realization of English vowels differs greatly among English varieties, including those represented in the house. If a speaker’s realization of a vowel changes enough over time, it could be realized as a variant not usually associated with the speaker’s variety. To examine whether this ever occurs, Figure 6 shows ellipses describing the range of each speaker’s predicted production during the season, for each vowel (ellipses contain 95% of predicted F1/F2 values, across clips), which we discuss with respect to the groups of speakers identified above for each vowel, according to their realization in the corpus and speakers’ dialect regions (Sec. 3.3, 5.3).

For GOOSE, the primary division was between Kathreya (clearly-backed [u]) and other speakers (different degrees of fronting). This division is maintained over time. The secondary division, among native speakers only (excluding Kathreya), was between Rachel (who showed the least fronting) and others. This division is largely maintained: although Luke’s GOOSE is also notably backed, on most days Rachel remains the native speaker with the most-backed variant.

For STRUT, the primary division was between speakers from Northern England and all others. This division is maintained over time: the ellipses for Northerners (Dale, Lisa, Luke, Stuart) and for other speakers do not overlap: the former always use something like [ʊ], while the latter never do. A second division, among non-Northerners, was between speakers who use variants closer to [ə] (Rachel, Rebecca) or closer to [ʌ]. Setting aside Kathreya, the [ə] and [ʌ] groups of native speakers retain distinct realizations of STRUT over time.

For TRAP, speakers showed a continuum of realizations from [a] to [æ], and we distinguished three groups for convenience. Speakers’ ranges of realizations of TRAP over time show more overlap, compared to GOOSE and STRUT. However, the three *groups* of speakers largely do not overlap: again setting aside Kathreya, there is little overlap between the ellipses for Darnell, Mohamed/Rex/Sara, and all other native speakers.

Thus, speakers do not generally use a markedly different phonetic variant for a given vowel over time, relative to the variants used in dialects represented in the house.

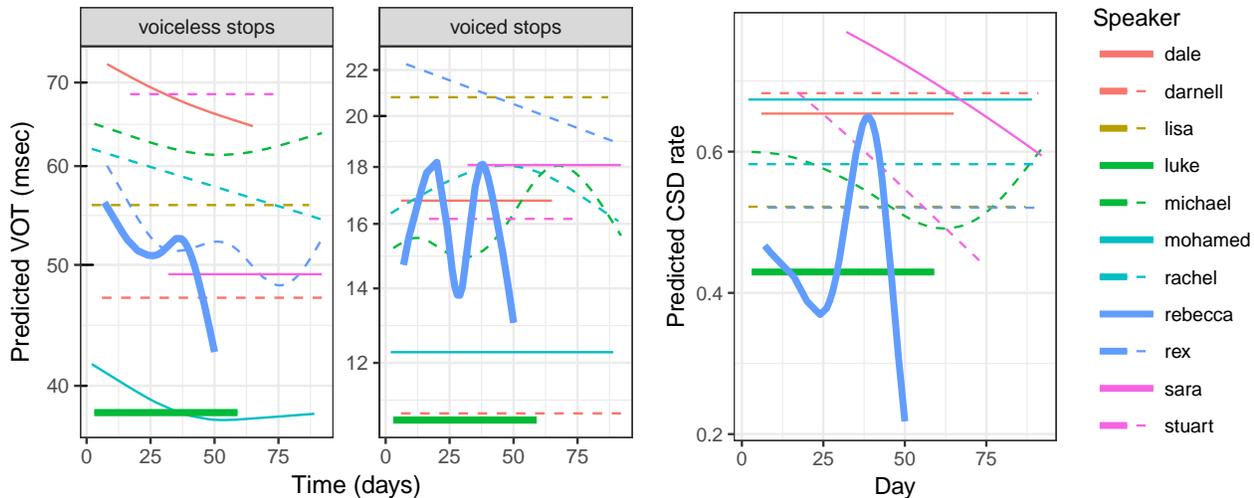


Figure 7: Trajectories of predicted mean VOT (left; msec, log scale) and mean coronal stop deletion rate (right), over time, for each speaker. Bolded lines correspond to Luke and Rebecca, two speakers who may show convergence.

## 6.2 Convergence

Convergence could take the form of converging time trends for a phonetic variable among all speakers (‘overall convergence’) or a subset of speakers (‘partial convergence’). We examine the evidence for each kind of convergence, using Figures 7–8, which show the predicted time trajectories of each variable, *without* taking by-day variability into account. These trajectories are the time trends for speakers whose dynamic models contain time trends, and otherwise a flat line at the speaker’s mean value.

**Overall convergence** The fact that only some speakers show time trends for each variable makes the strongest form of overall convergence impossible: speakers cannot all shift their mean pronunciation towards each other over weeks–months, since some of them do not shift at all over this timescale. However, it is possible that we have missed some time trends due to a lack of statistical power. We can still ask whether observed time trends are at least consistent with overall convergence, by asking for each variable whether time trends that *do* occur move towards each other.

Neither VOT (voiced or voiceless stops) nor CSD rate show clear evidence for overall convergence: in Figure 7, speakers with time trends do not systematically converge towards each other over the course of the season. The same is true in most cases for the vowels (Figure 8): F1 and F2 of STRUT and TRAP, and F1 of GOOSE. In the case of F2 for GOOSE, there is possible evidence for convergence: from about Day 50 on, the three speakers who do show time trends move towards a similar value. However, this pattern must be considered tentative, as it is based on only three speakers.

**Partial convergence** Partial convergence is more difficult to assess, because of the possibility of spurious results. For each variable, there are certainly pairs of housemates whose time trends seem to move towards each other over the course of the season, (e.g. Michael and Sara for CSD). However, such pairs would also exist if speakers showed random time trends

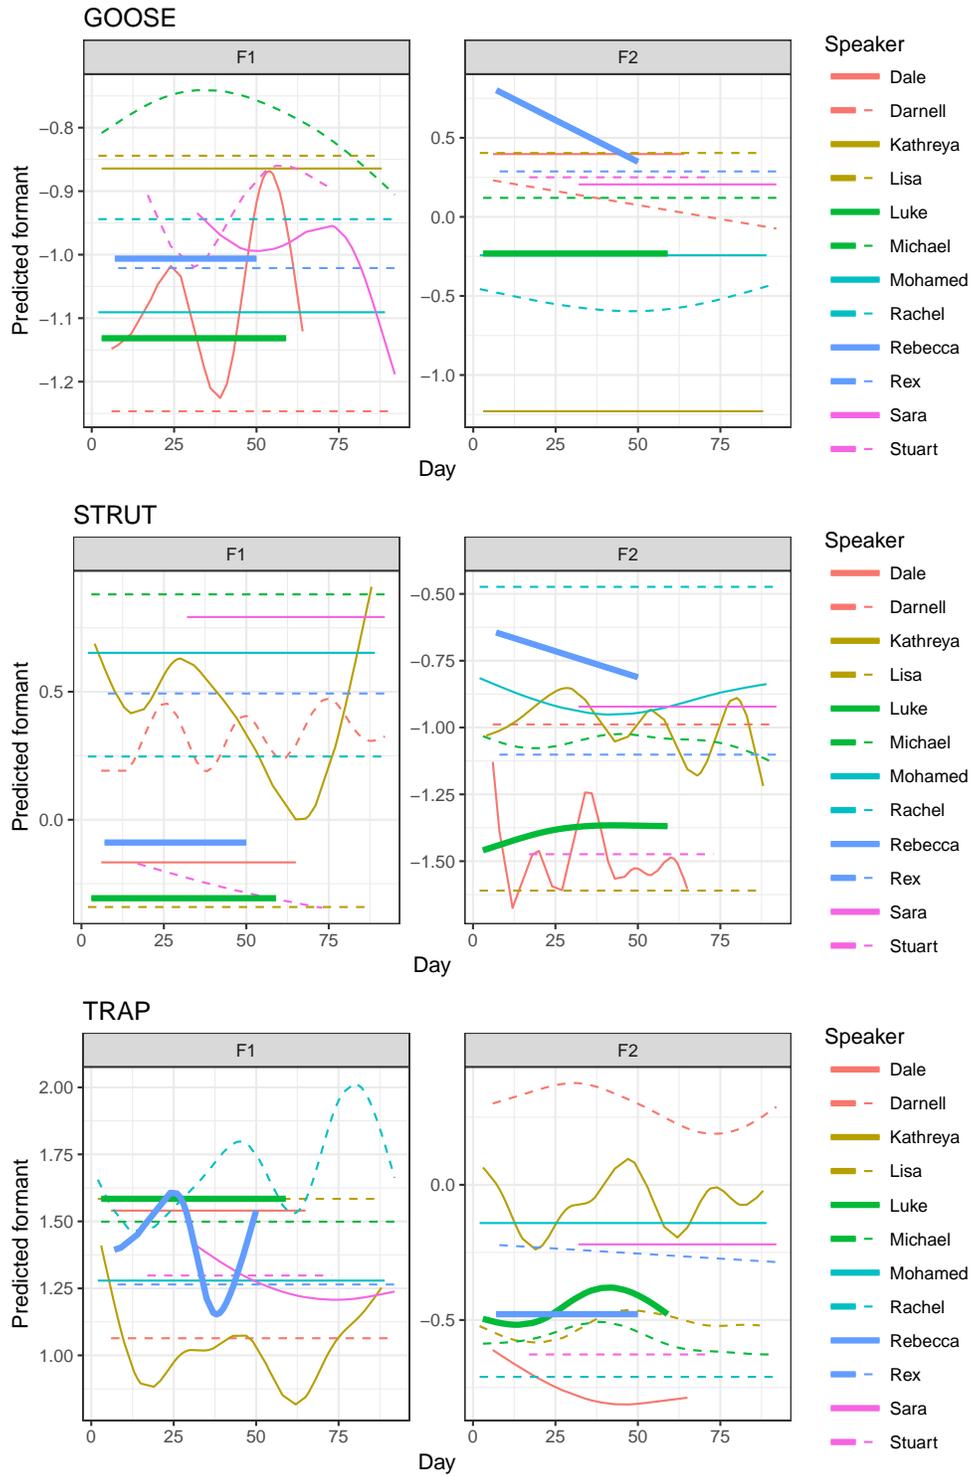


Figure 8: Trajectories of predicted mean (Lobanov-normalized) F1 and F2 over time, for each speaker, for GOOSE, STRUT, and TRAP. Bolded lines correspond to Luke and Rebecca, two speakers who may show convergence.

in phonetic parameters. The most solid evidence for partial convergence would be observing the same pattern of convergence across several phonetic variables, for the for the same subset of housemates, for whom there is a plausible reason for convergence to occur. We consider one such case.

Which pair of housemates would be most likely to converge? Following different theories of how short-term shifts may accumulate into accent change (Sec. 1.1), we might expect the most convergence between housemates who interact most often with each other, or who feel most positively towards each other. Both expectations arguably point towards the same housemates: Luke and Rebecca. These two housemates quickly become friends after entering the house, and interact progressively more with each other over the course of the season, until Rebecca is evicted, on Day 51. During this time, they develop romantic interest in each other, but at first only discuss this with other housemates. On Day 37, they admit their mutual interest to each other, and afterwards form an inseparable couple. Luke and Rebecca’s bond appears to have been genuine: after the show ended, they moved in together and remained in a relationship for over a year. No other pair of housemates on the show formed such a bond during the season, or interacted with each other as often, making Luke and Rebecca seem the most likely pair of housemates to converge.

To assess whether they do, the bolded lines in Figures 7–8 plot Luke and Rebecca’s predicted mean value of each variable over time. Rebecca’s GOOSE is progressively less fronted, moving steadily towards Luke’s norm. For STRUT, Luke’s [ʊ]-like production and Rebecca’s centralized production move slightly, but steadily, towards each other over the season. For the remaining variables, the pair’s pronunciations generally become more similar following around Day 37: Rebecca’s VOT shifts downwards towards Luke’s value, her pronunciation of TRAP shifts (back) towards his norm, and her CSD rate shifts radically downwards, reaching and then surpassing Luke’s low value. On the whole, the pair’s realizations of phonetic variables become markedly more similar over time—especially during their period of most intense interaction, after their relationship qualitatively shifts (Days 37–51)—suggesting convergence.

### 6.3 Speaker plasticity

Anecdotally, our results so far suggest another possible factor influencing patterns of time dependence: speakers may differ in how variable their accent is, across variables (*phonetic plasticity*)—whether this variability is due to a specific cause (e.g. convergence) or to random variation.

The clearest example is Lisa, who shows strikingly little variability over time compared to other speakers. For vowels, this is visually clear from Figure 6; she is also the only speaker to show no time dependence in CSD rate (Figure 3), and shows little time dependence in VOT (Figure 2). Luke and Rebecca provide another example: it is primarily time trends in Rebecca’s usage which are responsible for the pair’s convergence across variables, and the magnitude of variability over time generally seems greater for Rebecca than Luke, across variables (Figures 2, 3, 5). Both observations suggest Rebecca’s speech may be *systematically* more variable over time than Luke’s.

**Calculating plasticity over time** To test whether there are systematic individual differences in phonetic plasticity, we compute a single measure of plasticity per variable, per speaker, which takes both by-day variability and time trends into account, for each of

Table 5: Overall phonetic plasticity for each speaker. (See text.)

Lisa	Luke	Michael	Rex	Dale	Rachel	Mohamed	Stuart	Sara	Darnell	Kathreya	Rebecca
1.33	4.17	4.50	4.83	6.17	6.17	6.83	6.83	8.17	8.33	9.33	10.00

6 variables: VOT (voiced stops), VOT (voiceless stops), CSD rate, and the realization of GOOSE, STRUT and TRAP. We calculate these measures using the “residualized” predictions of the dynamic models (residualized VOT, etc.), which describe time dependence in each speaker’s use of a phonetic variable after controlling for static factors. We compute a measure of dispersion for each variable for each speaker, as follows:

- *VOT*: For each VOT subset (voiced, voiceless), for each speaker, we take the standard deviation of the predicted log(VOT) value for each of the speaker’s clips.
- *CSD rate*: For each speaker, we take the standard deviation of the predicted log-odds of deletion for each of the speaker’s clips.
- *Vowels*: For each vowel, for each speaker, we first extract the predicted F1 and F2 for each clip. We then compute the centroid of these values (the center of the ellipses shown in Figure 6) and the euclidean distance of each clip from the centroid (in F1/F2 space), and take the average of these distances.

This procedure results in six values per speaker (*Supplemental Information* Table SI4), describing the degree of variability in their realization of each variable over the season, which we call *plasticity* values.

**Testing for differences in plasticity** To test whether speakers differ significantly in plasticity, across all variables, we carry out a permutation test (see e.g. Good, 2013). The test statistic is the mean of rank correlations (Kendall’s  $\tau$ ) between pairs of plasticity values, which we denote  $\mu$ .<sup>16</sup> The observed value of  $\mu$ , denoted  $\mu_{obs}$ , is 0.265, suggesting a weak tendency for speakers to have similar plasticity across variables. The permutation test assesses how likely it is that we would observe a value of  $\mu$  this large, under the null hypothesis that each speaker’s degree of plasticity is randomly distributed across speakers, independently for each phonetic variable (i.e., speaker plasticities for different variables are unrelated).

We carry out this test by (a) randomly permuting plasticity values among speakers, independently for each phonetic variable; (b) re-calculating  $\mu$ ; (c) repeating steps (a) and (b)  $n = 100000$  times. The  $n$  values of  $\mu$  approximate its distribution under the null hypothesis. The  $p$ -value for this test is the probability (using this distribution) that  $\mu \geq \mu_{obs}$ , approximated as the proportion of the  $n$  runs for which  $\mu \geq \mu_{obs}$ .  $p = 2.8 \times 10^{-4}$  in our case, suggesting that the null hypothesis is very unlikely to be true, and thus that a speaker’s plasticity values for different variables are positively related. That is, there are “more plastic” and “less plastic” speakers.

We quantify each speaker’s degree of plasticity by their mean plasticity value (rank), across variables, (Table 5). As expected, Lisa is the least plastic speaker, while Luke is much less plastic than Rebecca, who is the most plastic speaker.

<sup>16</sup>These correlations are shown in *Supplemental Information* SI5.

## 6.4 Summary

Our findings in this section suggest that linguistic factors and individual differences account for some of the observed patterns of pronunciation change, and suggest a possible role for social factors. For vowels, housemates generally do not produce a markedly different phonetic variant for GOOSE, STRUT, and TRAP over time, with respect to the possibilities represented in the house. There are marked individual differences in phonetic plasticity, across variables. There was little evidence for convergence across all housemates in any phonetic variable, with one possible exception. However, we observed convergence across variables between between a pair of housemates with an unusually close bond. Social factors may also help explain patterns of variability in vocalic variables, as discussed below.

## 7 Discussion

Our results on medium-term dynamics in phonetic variables in the Big Brother house allow us to address the questions raised in the introduction: what trajectories do phonetic variables show within individual speakers over three months, and how do the results bear on the persistence and stationarity hypotheses? How do medium-term accent dynamics fit in between short-term dynamics and long-term dynamics, and how does the answer bear on the striking discrepancy between the plasticity of accents on these two timescales? What factors affect medium-term accent dynamics in individuals, and what is the relationship of these dynamics to the mechanism of community-level sound change?

### 7.1 Medium-term plasticity

Medium-term time dependence was ubiquitous: in 98.5% of cases, a speaker’s pronunciation of a variable varied systematically over time. Fluctuations between recording sessions on different days were the norm (by day-variability; 95.65% of cases); steady change over longer timescales was less common, but still occurred often (time trends; 43.48% of cases). Our results suggest that medium-term variability over time in pronunciation is primarily due to short-term fluctuations (days), and secondarily to longer-term change (weeks–months).

**Stationarity and persistence** Our results allow us to assess the stability of phonetic variables within individual speakers over days–months, about which little is known. Two reasonable null hypotheses, based on extending the stability many speakers show over years to shorter timescales, were that speakers do not systematically vary over time, or over a timescale longer than days (the “strong” and “weak” stationarity hypotheses).

Pronunciation is almost never static from day to day, suggesting that the strong stationarity hypothesis can be rejected. Previous sociolinguistic work (Nahkola & Saanilahti, 2004) has suspected the existence of such variability between recording sessions, but we are aware of only two laboratory studies examining variability within the same speaker’s speech over days (Heald, 2012; Pisoni, 1980); both find no evidence for by-day variability in productions of English vowels by a small number of speakers. Our contrasting result could be due to our larger sample size, or to greater variability in spontaneous versus laboratory speech.

An important caveat to our finding of ubiquitous variability among recording sessions is that some of this variation may not be a function of *time* per se (whether due to accumulated short-term shifts or to cognitive factors such as fatigue: Heald, 2012). We have treated a

housemate’s diary room clips (speech to Big Brother) as representing “baseline” use of each variable in the same social context, abstracting away from short-term shifts which may occur during interaction with other housemates. However, there could still be variability in speaking style between clips, e.g. due to style shifting or accommodation to different speakers acting as “Big Brother”. Such stylistic and interviewer/interlocutor effects are of concern in all linguistic studies using data from multiple time points (e.g. Gregersen et al., in press), and indeed are sometimes the focus (e.g. Rickford & Price, 2013; Gregersen et al., to appear). Like most panel and dialect change studies, we take the existence of a systematic pattern of time dependence—in our case, the near-ubiquitous *existence* of by-day variability—to reflect at least in part change over time, rather than factors not controlled for. More work is needed using speech from a controlled setting to establish how large daily fluctuations are for phonetic variables when factors such as style shifting are controlled for. Further work could also better establish the timescale of variability. While our data from diary room clips (which largely occur on different days, within a given speaker) lets us establish variability over time on two timescales (time trends, fluctuations between clips), it does not let us distinguish between by-day variability and shorter timescales. It is possible that what we are calling “by-day variability” actually reflects fluctuation on a scale of hours or minutes, instead of or in addition to daily variability. This possibility does not affect our main conclusions below, but accent dynamics on multiple timescales is an interesting direction for future work.

Pronunciation is also not in general static over weeks-months, suggesting the weak stationarity hypothesis can be rejected. The existence of time trends is notable, because it offers less ambiguous evidence that speakers’ “baseline” use of a variable changes over time, compared to by-day variability. Time trends are unlikely to be due to style shifting or accommodation to different Big Brothers, since a speaker would need to incrementally shift in the same direction over many clips. Indeed, it is hard to think of reasons why a speaker’s accent would systematically change over time—except for long-term accommodation to other housemates’ speech. Thus, the existence of time trends suggests that *medium-term persistence* of accommodation effects is possible—if very sporadic, across speakers and variables. More speculatively, it is also possible that some of the by-day variability in pronunciation is due to persistent accommodation effects.

## 7.2 The relationship of medium term to short and long term

We now turn to the the relationship between medium-term accent dynamics to shifts on shorter and longer timescales, to understand how individuals vary on different timescales.

In short-term studies, the existence of shifts in pronunciation on a timescale of seconds–hours are relatively robust across speakers and variables, and the magnitude and direction of these shifts are heavily modulated by social factors, linguistic factors, and individual differences. In our data, time dependence on a timescale of days is near-universal, across speakers and variables. We argue that medium-term time dependence in our data is primarily due to *by-day variability*, and can be partially accounted for by social factors, linguistic factors, and individual differences. In this sense, daily fluctuations in pronunciation (in our data) look qualitatively similar to shorter-term shifts (in previous work). By contrast, long-term studies suggest that accent stability is the norm in adulthood over a timescale of years, while a minority of adults show significant accent change. In our data, on a timescale of weeks–months, *time trends* occur more sporadically than by-day variability: across all speakers and variables, there is steady change in pronunciation in less than half of cases. In this sense, accent change over weeks-months (in our data) looks qualitatively similar

to accent change over years (in previous work). The similarity of different aspects of our results to the divergent results of the short-term and long-term literatures suggests that the disconnect between short-term and long-term accent plasticity is already present on a timescale of months: plasticity is widespread on timescales of up to days, but more sporadic over longer periods.

Why these similarities? We cannot say for sure, since we don't have data from the Big Brother house on short-term shifts (i.e., in interactions) or change over years. Nonetheless, we can sketch a possible account of the relationship between accent plasticity in individuals on different timescales, which addresses the mismatch between the ubiquity of short-term accent plasticity on the one hand, and the heterogeneity of long-term accent change on the other.

In the course of **interactions**, some short-term shift nearly always occurs. We assume that at least some of the ubiquitous by-day variability in our data can be ascribed to accumulated short-term shifts—that is, that speakers “bounce around” in their use of each phonetic variable as a result of short-term shifts in conversation, and thus have different “baseline” values when recorded in a constant environment (the diary room) on different days. From **day to day**, short-term shifts for a particular variable and speaker usually build up enough to “change” her pronunciation. The amount of build-up is highly variable, because of all the factors modulating short-term shifts. Next, note that in the vast majority of cases (for a given variable and speaker), the existence of a time trend implies the existence of by-day variability (Table 3). We hypothesize a relationship between these two types of time dependence: time trends are at least in part the result of accumulated by-day variability, but this build-up often does not occur (cases where there is no time trend), or is too small to be detected. The sporadicness of change over weeks–months might be because speakers differ in how phonetically “plastic” they are (see below); this could be capturing individual differences in size of short-term shifts, or how systematically they accumulate. Regardless of the source, the result is that over **weeks–months**, accumulation of day-to-day shifts into longer-term change in pronunciation norms occurs for only a minority of speakers and variables—although day-to-day shifts *could* accumulate into change over weeks–months in pronunciation norms, this only occurs in a minority of cases. Accent change over **years** is then also (correctly) predicted to be very heterogeneous.

### 7.2.1 Change by accommodation

Our medium-term results bear on longstanding Change By Accommodation theories, which propose a link between the short-term and long-term: short-term shifts in conversation (step 1) can accumulate into long-term accent change in individuals (step 2), and eventually community-level sound change (step 3). CBA theories have been difficult to test given the complexity of real-world speech communities and a lack of empirical data on medium-term dynamics (between steps 1 and 2); our medium-term results from a linguistically-closed system provide one of the first (indirect) tests.

The simplest version of CBA is that accent change and sound change are primarily driven by ‘automatic’ convergence during conversation: the more interactions take place in a group, the more similar group members’ speech should become. This idea, which has been repeatedly proposed (e.g. Bloomfield, 1933; Delvaux & Soquet, 2007; Paul, 1880; Trudgill, 2004), makes sense as a null hypothesis (Labov, 2000, p. 506). Our results suggest that this null hypothesis can be rejected: we see little evidence for overall convergence in how a set of housemates from diverse dialect regions speak over three months, despite constant interaction. This

(non-)finding bears most directly on steps 1–2 of CBA, which address accent dynamics in individuals. The combination of widespread medium-term time dependence in speakers’ accents with a lack of overall accent convergence is consistent with the heterogeneous results of longitudinal studies examining accent change over years, e.g. in college roommates or people who move between dialect regions (Barden & Großkopf, 1998; Evans & Iverson, 2007; Pardo et al., 2012; Wilson, 2010). Both our results and previous work suggest that the dynamics of phonetic variables over time within individual speakers, even in settings of intense social contact, are highly complex. More empirical data on longitudinal variation within individuals is needed before the intuitively-plausible link between social interaction and accent change during adulthood can be established, and to clarify the nature of this link.

More complex versions of CBA hold that accent change and sound change are mediated by a variety of factors, because short-term shifts (step 1) are also mediated by these factors (e.g. Auer & Hinskens, 2005; Garrett & Johnson, 2013). Our results, and the account sketched above of the link between dynamics on different timescales in individuals, are consistent with this position: the heterogeneous patterns of medium-term time dependence in phonetic variables we observe can be partially explained in terms of social factors, linguistic factors, and individual differences; the same types of factors which modulate the direction and extent of shifts observed in short-term studies.

### 7.3 Individual accent dynamics and community-level change

Medium-term time dependence in phonetic variables is both ubiquitous and heterogeneous, across speakers and variables. Our results allow us to assess several possible sources of accent dynamics in individuals—social factors, linguistic factors, and individual differences—and bear on the relationship between short-term shifts and longer-term change.

#### 7.3.1 Vowel dynamics: contrast and social salience

Our results for the vocalic variables suggest that time dependence in phonetic variables is constrained. For GOOSE, STRUT, and TRAP, despite extensive time dependence in their pronunciations, housemates generally do *not* produce markedly different phonetic variants with respect to the possibilities represented in the house, instead staying within the phonetic bounds expected given their dialect region. We suggest that the repertoire of possible realizations constrains variability in vowel realization: speakers generally do not produce variants that would “sound different”, for linguistic or social reasons, an interpretation similar to the role of *salience* invoked in both the short-term and long-term literatures (e.g. Auer et al., 1998; Kerswill & Williams, 2002; Trudgill, 1986).

Time dependence in vowel pronunciation may be constrained by contrast between speakers’ pre-existing phonetic categories. Findings from the short-term literature suggest that speakers’ productions in imitation tasks (e.g. Babel, 2011; Kim et al., 2011; MacLeod, 2012) or auditory feedback tasks (Katseff & Houde, 2008; Katseff et al., 2012) are strongly constrained by their existing “phonetic repertoire” (or “self-exemplars”; Garrett & Johnson, 2013). The observed pattern of a speaker’s vowel productions staying near a ‘characteristic pronunciation’ over days–months could simply result from using her phonetic repertoire over this timescale.

Social salience may also constrain variability over time, if speakers avoid productions that would “sound different” in the sense of having social salience. GOOSE shows extensive overlap in how fronted each native speaker’s realization is over time, while TRAP and STRUT show less

overlap. GOOSE is fronting over time in many British varieties, and previous work has argued that this change has low social salience. Native speakers will thus have exposure to very different degrees of fronting, and will not differentiate between degrees of fronting for social reasons; they are thus free to vary in GOOSE realization over time, as long as a fronted variant is used. This is exactly what is observed. STRUT realization is highly socially/regionally-marked (with [ʊ]-like realizations strongly associated with Northern England). Thus, native speakers cannot vary their production into a region of F1/F2 space associated with a different variant, without producing a socially-salient difference; this constrained variation is what is observed. Finally, TRAP realization has medium social salience; this could constrain speakers' pronunciations to not vary too greatly along the [a]~[æ] continuum, as is observed. In sum, it is possible that social salience constrains what medium-term change in pronunciation is possible in a variable, echoing short-term and long-term studies which invoke social salience to explain which variables shift more/less (e.g. Auer et al., 1998; Babel, 2010; Kerswill & Williams, 2002; Kim et al., 2011; MacLeod, 2012; Trudgill, 1986). This possibility could be more rigorously tested in future work by using a larger number of variables and a sample of speakers who are more likely to share evaluative norms, such as economic migrants or university students who move from one dialect region to another (Barden & Großkopf, 1998; Evans & Iverson, 2007; Wilson, 2010).

### 7.3.2 Convergence

In their simplest forms, both CBA and Communication Accommodation Theory predict that the more a set of speakers interacts, the more similar their pronunciation norms will become (Auer & Hinskens, 2005; Pardo et al., 2012). The Big Brother house seems like an ideal setting to observe accent convergence: constant interaction between speakers with very different accents, in a linguistically-closed system. However, we found strikingly little evidence for *overall convergence* among housemates for any of the five phonetic variables examined, in the sense of converging time trends. In only one case (F2 of GOOSE) was there a pattern which could be tentatively interpreted as convergence, with the caveat that 9/12 speakers showed *no* change in mean F2 over time. However, we did find clear *partial convergence* for one pair of housemates, across all five phonetic variables. These results are consistent with a recent study of five pairs of college roommates by Pardo et al. (2012)—similarly motivated by the idea that intense interaction should lead to phonetic convergence—who found that “overall levels of detected convergence... was modest, even after 3.5 months of relatively continuous cohabitation” (197), with a single pair showing much clearer convergence than others.

Why might we not have found overall convergence? The simplest possibility would be that convergence takes longer than three months (the duration of the show). Indeed, longitudinal studies of individuals who move between dialect regions have found that accent change is more widespread after 1–2 years than after three months (Barden & Großkopf, 1998; Evans & Iverson, 2007). We did find likely convergence for the pair of housemates who were arguably most likely to converge (Luke and Rebecca)—as measured either by frequency of interaction or social affinity—thus, convergence between housemates was *possible* over the timescale of the show. Why was it not more widespread? Luke and Rebecca may have simply interacted more frequently than other housemates, but this cannot be the whole story: they interacted frequently for 51 days, but it is only in the last couple of weeks when convergence is clear, across variables. We propose that convergence between them was also driven by their extremely strong social bond, cemented over these two weeks, when they

enter into a romantic relationship. Indeed, Communication Accommodation Theory (Giles et al., 1991, *et seq*) would predict Luke and Rebecca to express their increased emotional affinity towards each other in this period by converging during conversations, and that these shifts would accumulate over time. Other housemates, who do not have as close a bond, do not accommodate towards each other as consistently, hence the lack of overall convergence. In short, accent change over three months may require social motivation, rather than just frequent interactions.

Another reason why we might not observe overall convergence is the performative aspect of the show. Because housemates are under constant observation, the show entails an extreme version of the ‘observer paradox’ (Labov, 1972, p. 113). Housemates are in a permanent, performative interaction with *viewers*, in which context they may (un)consciously manipulate their accent. Between different clips, phonetic variables may undergo short-term shifts as part of a speaker projecting a certain image (Le Page & Tabouret-Keller, 1985), or depending on who the perceived audience is (Bell, 1984). These shifts may eventually change an individual’s pronunciation norms, as predicted by Auer & Hinskens’s (2005) “identity projection model”, a CBA-like theory with the crucial difference that short-term shifts in conversation result from the speaker expressing a particular persona by shifting his usage towards what he thinks that persona sounds like, not relative to the interlocutor’s speech, *per se*. This source of time dependence may be unusually strong in the context of a televised show, swamping ‘overall convergence’ that might otherwise occur. (For example, housemates may explicitly want to sound distinct from other housemates, as part of their performed persona.) If this is the case, overall convergence would be more likely observed in other close-knit temporary communities where speakers are not constantly recorded, such as a summer camp.

### 7.3.3 Individual differences

We found strong evidence for one additional source of accent dynamics: housemates differ in overall “phonetic plasticity”, showing a characteristic degree of variability over time across phonetic variables. To our knowledge, the quantitative demonstration that some speakers are more variable than others across different aspects of pronunciation is novel, in line with a long-observed fact about interspeaker variation by sociolinguists—that some speakers are ‘more variable’ than others, in the sense of having a greater range of styles and variable realization (e.g. Eckert, 2000; Rickford & Price, 2013). Housemates’ different degrees of plasticity could reflect individual differences in the cognitive mechanisms of perception and production, in the magnitude of style shifting, or something else. Establishing the extent and mechanism of individual differences in phonetic plasticity is an interesting direction for future work.

Regardless of their source(s), individual differences in phonetic plasticity may be relevant for community-level sound change. Any community-level change must be actuated by some speaker(s) deviating from the norm, then picked up by other speakers; the puzzle is why this happens at some times, and not others, given that most sound changes are rooted in presumably-universal phonetic bias factors (Garrett & Johnson, 2013). Much previous work suggests traits which make some people more likely to innovate and adopt new variants, who Milroy & Milroy (1985) call *innovators* and *early adopters*: social factors, including the range of variants individuals are exposed to due to life circumstances (Labov, 2000); and lower-level individual differences, such as in cognitive processing style (Yu, 2013) or the magnitude of coarticulation (Baker et al., 2011). Individual differences in phonetic plasticity may be another such factor. More plastic speakers are more likely to *randomly* generate new variants

over time, making them a more likely innovator (*ceterus paribus*). Thus, the more plastic an individual’s speech, the more likely she is to use a variant which could seed a sound change in an interaction with an “early adopter”. While speculative, this account adds to the list of systematic reasons why some individuals may be more likely sources of community-level sound changes.

## 8 Conclusion

We have used the “natural experiment” of a reality television show to conduct the first systematic study of medium-term accent dynamics, over days–months, filling in an empirical gap between studies of short-term shifts during conversation and long-term accent plasticity over years. We considered five phonetic variables in a corpus of 14.5 hours of spontaneous speech, and built statistical models of time dependence in these variables within each of 12 speakers over three months, after controlling for a range of confounds. Our main findings are that time dependence in an individual’s pronunciation over the medium term is near-ubiquitous and is primarily due to daily fluctuations, although systematic change over weeks–months occurs frequently as well. The combination of ubiquitous daily variability with more sporadic longer-term change mirrors the different extents of accent plasticity previously observed in conversation versus over the lifespan, suggesting that speakers show progressively less accent flexibility over progressively longer timescales. One question for future work is whether other aspects of an individual’s linguistic system (e.g. lexical choice, morphology, syntax) show similar medium-term dynamics, as might be expected given that pronunciation is thought to be among the *least* flexible aspects of an individual’s linguistic system during adulthood (Sankoff, 2005; Siegel, 2010).

We argued that the heterogeneous patterns of medium-term time dependence in our data can in part be accounted for, by the same kinds of factors which influence the degree and magnitude of short-term shifts in speech production: linguistic factors, social factors, and individual differences. Speakers systematically produce phonetic variants for vowels which remain within the boundaries expected given their dialect regions. We found evidence for accent convergence between two speakers with an exceptionally close social bond, but no strong evidence for accent convergence across all speakers. Finally, speakers show characteristic degrees of ‘plasticity’ over time across phonetic variables; more plastic speakers may be more likely to seed community-level change. However, to a large extent, the question of what factors determine the dynamics of an individual’s pronunciation over time remains open, and is the primary question raised for future work. Addressing this question through the study of medium-term pronunciation dynamics is key for understanding why accent change occurs (or does not) in individuals, the mechanism of sound change in communities, and the link between the two.

## References

- AGRESTI, ALAN. 2002. *Categorical data analysis*. 2nd edn. Hoboken, NJ: Wiley.
- ALLEN, J. SEAN; JOANNE L. MILLER; and DAVID DESTENO. 2003. Individual talker differences in voice onset time. *The Journal of the Acoustic Society of America* 113.544–552.
- AUER, PETER; BIRGIT BARDEN; and BEATE GROSSKOPF. 1998. Subjective and objective parameters determining ‘salience’ in long-term dialect accommodation. *Journal of Sociolinguistics* 2.163–187.

- AUER, PETER, and FRANS HINSKENS. 2005. The role of interpersonal accommodation in a theory of language change. In Auer et al., 335–357.
- AUER, PETER; FRANS HINSKENS; and PAUL KERSWILL (eds.) 2005. *Dialect change: convergence and divergence in European languages*. Cambridge: Cambridge University Press.
- AUZOU, PASCAL; CANAN OZSANCAK; RICHARD J MORRIS; MARY JAN; FRANCIS EUSTACHE; and DIDIER HANNEQUIN. 2000. Voice onset time in aphasia, apraxia of speech and dysarthria: a review. *Clinical Linguistics & Phonetics* 14.131–150.
- BAAYEN, R. HARALD; DOUGLAS J. DAVIDSON; and DOUGLAS M. BATES. 2008. Mixed-effects modeling with crossed random effects for subjects and items. *Journal of Memory and Language* 59.390–412.
- BAAYEN, R. HARALD; RICHARD PIEPENBROCK; and LEON GULIKERS. 1996. *CELEX2 (CD-ROM)*. Philadelphia: Linguistic Data Consortium.
- BABEL, MOLLY. 2010. Dialect divergence and convergence in New Zealand English. *Language in Society* 39.437–456.
- BABEL, MOLLY. 2011. Evidence for phonetic and social selectivity in spontaneous phonetic imitation. *Journal of Phonetics* 40.177–189.
- BABEL, MOLLY; ANDREW GARRETT; MICHAEL J HOUSER; and MAZIAR TOOSARVANDANI. 2013. Descent and diffusion in language diversification: A study of Western Numic dialectology. *International Journal of American Linguistics* 79.445–489.
- BAKER, ADAM; DIANA ARCHANGELI; and JEFF MIELKE. 2011. Variability in American English s-retraction suggests a solution to the actuation problem. *Language Variation and Change* 23.347.
- BANE, MAX; PETER GRAFF; and MORGAN SONDEREGGER. 2010. Longitudinal phonetic variation in a closed system. *Proceedings of the Annual Meeting of the Chicago Linguistics Society* 46.43–58.
- BARDEN, BIRGIT, and BEATE GROSSKOPF. 1998. *Sprachliche Akkommodation und soziale Integration [Speech accommodation and social integration]*. Tübingen: Niemeyer.
- BARR, DALE J; ROGER LEVY; CHRISTOPH SCHEEPERS; and HARRY J TILY. 2013. Random effects structure for confirmatory hypothesis testing: Keep it maximal. *Journal of Memory and Language* 68.255–278.
- BATES, DOUGLAS; MARTIN MAECHLER; BEN BOLKER; and STEVEN WALKER. 2014. *lme4: Linear mixed-effects models using Eigen and S4*. R package version 1.1-7.
- BEDDOR, PATRICE S. 1982. *Phonological and phonetic effects of nasalization on vowel height*. University of Minnesota dissertation.
- BELL, ALLAN. 1984. Language style as audience design. *Language in Society* 13.145–204.
- BLOOMFIELD, LEONARD. 1933. *Language*. New York: Henry Holt.
- BOERSMA, PAUL, and DAVID WEENINK. 2011. Praat: doing phonetics by computer (Version 5.2.17) [Computer program].
- BOURHIS, RICHARD Y, and HOWARD GILES. 1977. The language of intergroup distinctiveness. *Language, ethnicity, and intergroup relations*, ed. by Howard Giles, 119–135. London: Academic Press.
- BOWIE, DAVID. 2005. Language change over the lifespan: A test of the apparent time construct. *University of Pennsylvania Working Papers in Linguistics* 11.45–58.
- BREEN, MARA; EVELINA FEDORENKO; MICHAEL WAGNER; and EDWARD GIBSON. 2010. Acoustic correlates of information structure. *Language and Cognitive Processes* 25.1044–1098.
- BRINK, LARS, and JØRN LUND. 1975. *Dansk rigsmål: lydudviklingen siden 1840 med særligt henblik på sociolekterne i københavn*. København: Gyldendal.
- CHAMBERS, JACK K. 2003. *Sociolinguistic theory: Linguistic variation and its social significance*. Malden, MA: Wiley-Blackwell.
- CHO, TAEHONG. 2005. Prosodic strengthening and featural enhancement: Evidence from acoustic and articulatory realizations of /a, i/ in English. *The Journal of the Acoustical Society of America* 117.3867–3878.

- CHODROFF, ELEANOR; JOHN GODFREY; SANJEEV KHUDANPUR; and COLIN WILSON. 2015. Structured variability in acoustic realization: a corpus study of voice onset time in American English stops. *Proceedings of the 18th International Congress of Phonetic Sciences*, ed. by The Scottish Consortium for ICPHS. Glasgow: University of Glasgow. Paper number 632.1–5.
- COLE, JENNIFER; HEEJIN KIM; HANSOOK CHOI; and MARK HASEGAWA-JOHNSON. 2007. Prosodic effects on acoustic cues to stop voicing and place of articulation: Evidence from radio news speech. *Journal of Phonetics* 35.180–209.
- COUPLAND, NIKOLAS. 1980. Style-shifting in a Cardiff work-setting. *Language in Society* 9.1–12.
- COX, FELICITY, and SALLYANNE PALETHORPE. 2007. Australian English. *The Journal of the International Phonetic Association* 37.341–350.
- CUKOR-AVILA, PATRICIA. 2015. Growing INTO and IN African American Vernacular English. Paper presented at *UKLVC 2015*.
- CUKOR-AVILA, PATRICIA, and GUY BAILEY. 2013. Real time and apparent time. *The handbook of language variation and change*, ed. by Jack K. Chambers and Natalie Schilling, second edn., 237–262. Malden, MA: Wiley.
- DARLINGTON, RICHARD B, and TOM V SMULDERS. 2001. Problems with residual analysis. *Animal Behaviour* 62.599–602.
- DELVAUX, VÉRONIQUE, and ALAIN SOQUET. 2007. The influence of ambient speech on adult speech productions through unintentional imitation. *Phonetica* 64.145–173.
- DOCHERTY, GERARD. 1992. *The timing of voicing in British English obstruents*. Berlin: Foris.
- DOCHERTY, GERARD. 2010. Phonological innovation in contemporary spoken British English. *The Routledge handbook of world Englishes*, ed. by Andy Kirkpatrick, 59–75. London: Routledge.
- DOCHERTY, GERARD J; DOMINIC WATT; CARMEN LLAMAS; DAMIEN HALL; and JENNIFER NYCZ. 2011. Variation in voice onset time along the Scottish-English border. *Proceedings of the 17th international congress of phonetic sciences*, ed. by Wai-Sum Lee and Eric Zee, 591–594. Hong Kong: City University of Hong Kong.
- ECKERT, PENELOPE. 2000. *Linguistic variation as social practice*. Malden, MA: Blackwell.
- EVANS, B.G., and P. IVERSON. 2007. Plasticity in vowel perception and production: A study of accent change in young adults. *The Journal of the Acoustical Society of America* 121.3814–3826.
- FERRAGNE, EMMANUEL, and FRANÇOIS PELLEGRINO. 2010. Formant frequencies of vowels in 13 accents of the British Isles. *The Journal of the International Phonetic Association* 40.1–34.
- FERREIRA, FERNANDA, and CHARLES CLIFTON. 1986. The independence of syntactic processing. *Journal of Memory and Language* 25.348–368.
- FINE, ALEX B; T FLORIAN JAEGER; THOMAS A FARMER; and TING QIAN. 2013. Rapid expectation adaptation during syntactic comprehension. *PloS One* 8.e77661.
- FOREMAN, ANNIK. 2003. *Pretending to be someone you're not: A study of second dialect acquisition in Australia*. Melbourne, Australia: Monash University dissertation.
- FOURAKIS, MARIOS. 1991. Tempo, stress, and vowel reduction in American English. *The Journal of the Acoustical Society of America* 90.1816–1827.
- FOWLER, CAROL A; JULIE M BROWN; LAURA SABADINI; and JEFFREY WEIHING. 2003. Rapid access to speech gestures in perception: Evidence from choice and simple response time tasks. *Journal of Memory and Language* 49.396–413.
- FRECKLETON, ROBERT P. 2002. On the misuse of residuals in ecology: regression of residuals vs. multiple regression. *Journal of Animal Ecology* 71.542–545.
- GARRETT, ANDREW, and KEITH JOHNSON. 2013. Phonetic bias in sound change. *Origins of sound change: Approaches to phonologization*, ed. by Alan C.L. Yu, 51–97. Oxford: Oxford University Press.
- GELMAN, ANDREW, and JENNIFER HILL. 2007. *Data analysis using regression and multi-level/hierarchical models*. Cambridge: Cambridge University Press.

- GILES, HOWARD; JUSTINE COUPLAND; and NIKOLAS COUPLAND. 1991. Accommodation theory: Communication, context, and consequence. *Contexts of accommodation: Developments in applied sociolinguistics*, ed. by Howard Giles, Justine Coupland, and Nikolas Coupland, 1–68. Cambridge: Cambridge University Press.
- GOLDINGER, STEPHEN D. 1998. Echoes of echoes? an episodic theory of lexical access. *Psychological Review* 105.251–279.
- GOLDINGER, STEPHEN D. 2000. The role of perceptual episodes in lexical processing. *Proceedings of the workshop on spoken word access processes*, ed. by Anne Cutler, James M. McQueen, and Rian Zondervan, 155–158. Nijmegen: Max Planck Institute.
- GOOD, PHILLIP. 2013. *Permutation tests: a practical guide to resampling methods for testing hypotheses*. New York: Springer.
- GREGERSEN, FRANS; TORBEN JUEL JENSEN; and NICOLAI PHARAO. 2017. Comparing speech samples: on the challenge of comparability in panel studies of language change in real time. *Panel studies of language variation and change*, ed. by Suzanne E. Wagner and Isabelle Buchstaller. Routledge. To appear.
- GREGERSEN, FRANS J.; NORMANN JØRGENSEN; and JANUS SPINDLER MØLLER. 2016. Sideways: Five methodological studies of sociolinguistic interviews. Copenhagen: University of Copenhagen, MS.
- HADDICAN, BILL; PAUL FOULKES; VINCENT HUGHES; and HAZEL RICHARDS. 2013. Interaction of social and linguistic constraints on two vowel changes in northern England. *Language Variation and Change* 25.371–403.
- HARRELL, FRANK E, JR. 2014. *rms: Regression modeling strategies*. R package version 4.2-0. Online: <http://CRAN.R-project.org/package=rms>.
- HARRINGTON, JONATHAN; SALLYANNE PALETHORPE; and CATHERINE WATSON. 2000. Monophthongal vowel changes in Received Pronunciation: An acoustic analysis of the Queen’s Christmas broadcasts. *The Journal of the International Phonetic Association* 30.63–78.
- HASTIE, TREVOR J, and ROBERT J TIBSHIRANI. 1990. *Generalized additive models*. Boca Raton, FL: CRC Press.
- HAWKINS, SARAH, and JONATHAN MIDGLEY. 2005. Formant frequencies of RP monophthongs in four age groups of speakers. *The Journal of the International Phonetic Association* 35.183–199.
- HAZEN, KIRK. 2011. Flying high above the social radar: Coronal stop deletion in modern Appalachia. *Language Variation and Change* 23.105–137.
- HEALD, SHANNON L.M. 2012. *Tuning the targets of speech production: Effects of articulatory and perceptual experience on speech production*. Chicago: University of Chicago dissertation.
- HILLENBRAND, JAMES; LAURA A GETTY; MICHAEL J CLARK; and KIMBERLEE WHEELER. 1995. Acoustic characteristics of American English vowels. *The Journal of the Acoustical Society of America* 97.3099–3111.
- HILLENBRAND, JAMES M; MICHAEL J CLARK; and TERRANCE M NEAREY. 2001. Effects of consonant environment on vowel formant patterns. *The Journal of the Acoustical Society of America* 109.748–763.
- HINDLE, DONALD. 1980. *The social and structural conditioning of phonetic variation*. Philadelphia: University of Pennsylvania dissertation.
- HOCKETT, CHARLES F. 1965. Sound change. *Language* 41.185–204.
- JOHNSON, KEITH. 1997. Speech perception without speaker normalization: An exemplar model. *Talker variability in speech processing*, ed. by Keith Johnson and John W. Mullenix, 145–165. San Diego: Academic Press.
- KATSEFF, SHIRA, and JOHN HOUDE. 2008. Partial compensation in speech adaptation. *UC Berkeley Phonology Lab annual report*, 444–461.
- KATSEFF, SHIRA; JOHN HOUDE; and KEITH JOHNSON. 2012. Partial compensation for altered auditory feedback: a tradeoff with somatosensory feedback? *Language and Speech* 55.295–308.

- KENDALL, TYLER. 2013. *Speech Rate, Pause and Sociolinguistic Variation*. Basingstoke: Palgrave Macmillan.
- KERSWILL, PAUL, and ANN WILLIAMS. 2002. “Salience” as an explanatory factor in language change: evidence from dialect levelling in urban England. *Language change: The interplay of internal, external and extra-linguistic factors*, ed. by Mari C. Jones and Edith Esch, 81–110. Berlin: Mouton de Gruyter.
- KESHET, JOSEPH; MORGAN SONDEREGGER; and THEA KNOWLES. 2014. AutoVOT: A tool for automatic measurement of voice onset time using discriminative structured prediction [Computer program]. Version 0.91. Available at <https://github.com/mlml/autovot/>.
- KIM, MIDAM; WILLIAM S. HORTON; and ANN R. BRADLOW. 2011. Phonetic convergence in spontaneous conversations as a function of interlocutor language distance. *Laboratory Phonology*, 125–156.
- KLATT, DENNIS H. 1975. Voice onset time, frication, and aspiration in word-initial consonant clusters. *Journal of Speech, Language and Hearing Research* 18.686–706.
- KORTMANN, BERND, and EDGAR WERNER SCHNEIDER. 2004. *A handbook of varieties of English*, vol. 1: Phonology. New York: de Gruyter. Pages 47–67, 98–162.
- KRYUCHKOVA, TATIANA; BENJAMIN V TUCKER; LEE H WURM; and R HARALD BAAYEN. 2012. Danger and usefulness are detected early in auditory lexical processing: Evidence from electroencephalography. *Brain and Language* 122.81–91.
- KUZNETSOVA, ALEXANDRA; PER BRUUN BROCKHOFF; and RUNE HAUBO BOJESEN CHRISTENSEN. 2014. *lmerTest: Tests for random and fixed effects for linear mixed effect models (lmer objects of lme4 package)*. R package version 2.0-11. Online: <http://CRAN.R-project.org/package=lmerTest>.
- LABOV, WILLIAM. 1972. *Sociolinguistic patterns*. Philadelphia: University of Pennsylvania Press.
- LABOV, WILLIAM. 1990. The intersection of sex and social class in the course of linguistic change. *Language Variation and Change* 2.205–254.
- LABOV, WILLIAM. 1994. *Principles of linguistic change. Vol. 1: Internal factors*. Oxford: Blackwell.
- LABOV, WILLIAM. 2000. *Principles of linguistic change. Vol. 2: Social factors*. Oxford: Blackwell.
- LABOV, WILLIAM. 2007. Transmission and diffusion. *Language* 83.344–387.
- LE PAGE, ROBERT BROCK, and ANDRÉE TABOURET-KELLER. 1985. *Acts of identity: Creole-based approaches to language and ethnicity*. Cambridge: Cambridge University Press.
- LISKER, LEIGH, and ARTHUR S ABRAMSON. 1967. Some effects of context on voice onset time in English stops. *Language and Speech* 10.1–28.
- LOBANOV, BORIS M. 1971. Classification of Russian vowels spoken by different speakers. *The Journal of the Acoustical Society of America* 49.606–608.
- MACLEOD, BETHANY. 2012. *The effect of perceptual salience on cross-dialectal phonetic convergence in spanish*. Toronto: University of Toronto dissertation.
- MILLER, JOANNE L; KERRY P GREEN; and ADAM REEVES. 1986. Speaking rate and segments: A look at the relation between speech production and speech perception for the voicing contrast. *Phonetica* 43.106–115.
- MILROY, JAMES, and LESLEY MILROY. 1985. Linguistic change, social network and speaker innovation. *Journal of Linguistics* 21.339–384.
- MITTERER, HOLGER, and MIRJAM ERNESTUS. 2008. The link between speech perception and production is phonological and abstract: Evidence from the shadowing task. *Cognition* 109.168–173.
- MOON, SEUNG-JAE, and BJÖRN LINDBLOM. 1994. Interaction between duration, context, and speaking style in English stressed vowels. *The Journal of the Acoustical Society of America* 96.40–55.
- MORRIS, RICHARD J; CHRISTOPHER R MCCREA; and KAILEEN D HERRING. 2008. Voice onset time differences between adult males and females: Isolated syllables. *Journal of Phonetics* 36.308–317.
- MYERS, JAMES, and GREGORY GUY. 1997. Frequency effects in variable lexical phonology. *University of Pennsylvania Working Papers in Linguistics* 4.215–27.

- NAHKOLA, KARI, and MARJA SAANILAHTI. 2004. Mapping language changes in real time: A panel study on Finnish. *Language Variation and Change* 16.75–92.
- NAMY, LAURA L; LYNNE C NYGAARD; and DENISE SAUERTEIG. 2002. Gender differences in vocal accommodation: The role of perception. *Journal of Language and Social Psychology* 21.422–432.
- NEAREY, TERRANCE M, and BERNARD L ROCHET. 1994. Effects of place of articulation and vowel context on VOT production and perception for French and English stops. *The Journal of the International Phonetic Association* 24.1–19.
- NIELSEN, KUNIKO. 2011. Specificity and abstractness of VOT imitation. *Journal of Phonetics* 39.132–142.
- PARDO, JENNIFER S. 2006. On phonetic convergence during conversational interaction. *The Journal of the Acoustical Society of America* 119.2382–2393.
- PARDO, JENNIFER S. 2009. Expressing oneself in conversational interaction. *Expressing oneself/expressing one's self: Communication, cognition, language, and identity*, ed. by Ezequiel Morsella, 183–196. Mahwah, NJ: Lawrence Erlbaum.
- PARDO, JENNIFER S. 2013. Measuring phonetic convergence in speech production. *Frontiers in Psychology* 4.
- PARDO, JENNIFER S; RACHEL GIBBONS; ALEXANDRA SUPPES; and ROBERT M KRAUSS. 2012. Phonetic convergence in college roommates. *Journal of Phonetics* 40.190–197.
- PAUL, HERMANN. 1880. *Prinzipien der sprachgeschichte*. 1st edn. Halle: Max Niemeyer.
- PIERREHUMBERT, JANET B. 2001. Exemplar dynamics: Word frequency, lenition and contrast. *Frequency effects and the emergence of lexical structure*, ed. by Joan Bybee and Paul Hopper, 137–158. Amsterdam: Benjamins.
- PINHEIRO, JOSE, and DOUGLAS BATES. 2000. *Mixed-effects models in S and S-PLUS*. New York: Springer.
- PISONI, DAVID B. 1980. Variability of vowel formant frequencies and the quantal theory of speech: a first report. *Phonetica* 37.285–305.
- PRINCE, ELLEN F. 1987. Sarah Gorby, Yiddish folksinger: A case study of dialect shift. *International Journal of the Sociology of Language* 67.83–116.
- R CORE TEAM. 2014. *R: A language and environment for statistical computing*. R Foundation for Statistical Computing, Vienna, Austria. Online: <http://www.R-project.org/>.
- RICKFORD, JOHN R, and FAYE MCNAIR-KNOX. 1994. Addressee-and topic-influenced style shift: A quantitative sociolinguistic study. *Sociolinguistic perspectives on register*, ed. by Douglas Biber and Edward Finegan, 235–276. Oxford: Oxford University Press.
- RICKFORD, JOHN R, and MACKENZIE PRICE. 2013. Girlz II women: Age-grading, language change, and stylistic variation. *Journal of Sociolinguistics* 17.143–179.
- ROSENFELDER, INGRID; JOE FRUEHWALD; KEELAN EVANINI; and JIAHONG YUAN. 2011. FAVE (Forced Alignment and Vowel Extraction) Program Suite [Computer program]. Online: <http://fave.ling.upenn.edu>.
- SANCIER, MICHELE L, and CAROL A FOWLER. 1997. Gestural drift in a bilingual speaker of Brazilian Portuguese and English. *Journal of Phonetics* 25.421–436.
- SANKOFF, GILLIAN. 2005. Cross-sectional and longitudinal studies in sociolinguistics. *Sociolinguistics: An international handbook of the science of language and society*, ed. by U. Ammon, N. Dittmar, K.J. Mattheier, and P. Trudgill, vol. 2, 1003–13. Berlin: de Gruyter.
- SANKOFF, GILLIAN. 2013. Longitudinal studies. *The Oxford handbook of sociolinguistics*, ed. by Robert Bayley, Richard Cameron, and Ceil Lucas, 261–279. Oxford: Oxford University Press.
- SANKOFF, GILLIAN, and HÉLÈNE BLONDEAU. 2007. Language change across the lifespan: /r/ in Montreal French. *Language* 83.560–588.
- SCHREIER, DANIEL. 2005. *Consonant change in English worldwide*. New York: Palgrave MacMillan.
- SCHWARTZ, MISHA. 2015. Formant Visualization and Remeasurement (FVR) [Computer program]. Available at <https://github.com/mlml/fvr>.

- SCOBIE, JAMES M. 2006. Flexibility in the face of incompatible English VOT systems. *Laboratory phonology 8*, ed. by Louis M. Goldstein, Doug H. Whalen, and Catherine T. Best, 367–392. Berlin: Mouton de Gruyter.
- SIEGEL, JEFF. 2010. *Second dialect acquisition*. Cambridge: Cambridge University Press.
- SONDEREGGER, MORGAN. 2012. *Phonetic and phonological dynamics on reality television*. Chicago: University of Chicago dissertation.
- SONDEREGGER, MORGAN. 2015. Trajectories of voice onset time in spontaneous speech on reality TV. *Proceedings of the 18th International Congress of Phonetic Sciences*, ed. by The Scottish Consortium for ICPHS. Glasgow: University of Glasgow. Paper number 903.1–5.
- SONDEREGGER, MORGAN, and JOSEPH KESHET. 2012. Automatic measurement of voice onset time using discriminative structured prediction. *The Journal of the Acoustical Society of America* 132.3965–3979.
- STANFORD, JAMES N. 2008. A sociotonic analysis of Sui dialect contact. *Language Variation and Change* 20.409–450.
- STEVENS, KENNETH N., and ARTHUR S. HOUSE. 1963. Perturbation of vowel articulations by consonantal context: An acoustical study. *Journal of Speech and Hearing Research* 6.111.
- STUART-SMITH, JANE; MORGAN SONDEREGGER; RACHEL MACDONALD; and TAMARA RATHCKE. 2015. The private life of stops: VOT in a real-time corpus of spontaneous Glaswegian. *Laboratory Phonology* 6.505–549.
- TAGLIAMONTE, SALI, and ROSALIND TEMPLE. 2005. New perspectives on an ol' variable: (t,d) in British English. *Language Variation and Change* 17.281–302.
- TANNER, JAMES; MORGAN SONDEREGGER; and MICHAEL WAGNER. 2016. Production planning and coronal stop deletion in spontaneous speech. Montreal: McGill University, MS. Online: <http://people.linguistics.mcgill.ca/~morgan>.
- TEMPLE, ROSALIND A.M. 2009. (t,d): the variable status of a variable rule. *Oxford University Working Papers in Linguistics, Philology & Phonetics* 12.145–170.
- TRUDGILL, P. 1986. *Dialects in contact*. Oxford: Blackwell.
- TRUDGILL, P. 2004. *New-dialect formation: The inevitability of colonial englishes*. Oxford: Oxford University Press.
- TSUKADA, KIMIKO. 2008. An acoustic comparison of English monophthongs and diphthongs produced by Australian and Thai speakers. *English World-Wide* 29.194–211.
- VAN RIJ, JACOLIEN; MARTIJN WIELING; R. HARALD BAAYEN; and HEDDERIK VAN RIJN. 2016. *itsadug: Interpreting Time Series and Autocorrelated Data Using GAMMs*. R package version 2.2.
- WALKER, JAMES. 2012. Form, function, and frequency in phonological variation. *Language Variation and Change* 24.397–415.
- WATSON, KEVIN. 2007. Liverpool English. *The Journal of the International Phonetic Association* 37.351–360.
- WELLS, JOHN C. 1982. *Accents of English*. Cambridge: Cambridge University Press. 3 vols.
- WIELING, MARTIJN; SIMONETTA MONTEMAGNI; JOHN NERBONNE; and R HARALD BAAYEN. 2014. Lexical differences between Tuscan dialects and standard Italian: Accounting for geographic and sociodemographic variation using generalized additive mixed modeling. *Language* 90.669–692.
- WIKIPEDIA. 2012a. *Big Brother 9 (UK)* — Wikipedia, The Free Encyclopedia. [Online; accessed 21-April-2012]. Online: [http://en.wikipedia.org/wiki/Big\\_Brother\\_2008\\_\(UK\)](http://en.wikipedia.org/wiki/Big_Brother_2008_(UK)).
- WIKIPEDIA. 2012b. *Big Brother (UK)* — Wikipedia, The Free Encyclopedia. [Online; accessed 21-April-2012]. Online: [http://en.wikipedia.org/wiki/Big\\_Brother\\_\(UK\)](http://en.wikipedia.org/wiki/Big_Brother_(UK)).
- WILSON, JAMES. 2010. *Moravians in prague: a sociolinguistic study of dialect contact in the Czech Republic*. Frankfurt: Peter Lang.
- WINTER, BODO, and MARTIJN WIELING. 2016. How to analyze linguistic change using mixed models, growth curve analysis and generalized additive modeling. *Journal of Language Evolution* 1.7–18.

- WOOD, SIMON. 2015. *mgcv: Mixed GAM Computation Vehicle with GCV/AIC/REML smoothness estimation and GAMMs by REML/PQL*. R package version 1.8–7.
- WOOD, SIMON N. 2006. *Generalized additive models: An introduction with R*. Boca Raton, FL: Chapman and Hall/CRC.
- WOOD, SIMON N. 2011. Fast stable restricted maximum likelihood and marginal likelihood estimation of semiparametric generalized linear models. *Journal of the Royal Statistical Society (B)* 73.3–36.
- YAO, YAO. 2009. Understanding VOT variation in spontaneous speech. *Current numbers in unity and diversity of languages*, ed. by Manghyu Pak, 1122–1137. Seoul: Linguistic Society of Korea.
- YOUNG, STEVE; GUNNAR EVERMANN; MARK GALES; THOMAS HAIN; DAN KERSHAW; XUNYING LIU; GARETH MOORE; JULIAN ODELL; DAVE OLLASON; DAN POVEY; ET AL. 2006. *The HTK book (for HTK version 3.4)*. Cambridge University Engineering Department.
- YU, ALAN C.L. 2013. Individual differences in socio-cognitive processing and the actuation of sound change. *Origins of sound change: Approaches to phonologization*, ed. by Alan C.L. Yu, 201–227. Oxford: Oxford University Press.
- YU, ALAN C.L.; CARISSA ABREGO-COLLIER; and MORGAN SONDEREGGER. 2013. Phonetic imitation from an individual-difference perspective: Subjective attitude, personality and “autistic” traits. *PLOS ONE* 8.e74746.

# Appendices

## Appendix A Static factors

This section describes the major static factors (covariates unrelated to time) affecting each phonetic variable, which are incorporated into the statistical models described below (“static models”) of how each phonetic variable depends on static factors only. Table 2 summarizes the variables included in each static model, denoted in SMALL CAPS below.

We first define two static factors which are used as predictors for several phonetic variables: SPEECH RATE was calculated as the number of syllables per second (determined using the force-aligned transcriptions) within a phrase, defined as an interval of speech by a housemate bounded on each side by at least 60 msec of silence/non-speech. We follow Kendall (2013) in using a rate measure that excludes pauses and in the choice of 60 msec as a threshold. WORD FREQUENCY was taken to be the token count per million of the orthographic word in the corpus of transcriptions of housemate speech, log-transformed.

### A.1 Voice onset time

VOT is affected by a range of linguistic and social factors, reviewed by Auzou et al. (2000); Docherty (1992).

Whether the stop consonant is (phonologically) voiced or voiceless is the single largest factor affecting VOT (voiced < voiceless). Because VOT behaves qualitatively differently for voiced and voiceless stops, it makes sense to think of these as two different variables, and we build separate statistical models for VOT of voiced and voiceless stops. VOT is strongly affected by PLACE OF ARTICULATION, with VOT expected to be progressively larger for bilabial, alveolar, and velar stops (e.g. Docherty, 1992; Lisker & Abramson, 1967; Nearey & Rochet, 1994). VOT (in word-initial stops) is also influenced by the FOLLOWING SEGMENT

IDENTITY and FOLLOWING VOWEL HEIGHT in the host word: VOT is expected to be higher before consonants (in complex onsets) than before vowels (in CV syllables), and higher before high vowels than before low vowels (e.g. Docherty, 1992; Klatt, 1975; Nearey & Rochet, 1994). VOT is expected to be shorter for more frequent words (Stuart-Smith et al., 2015; Yao, 2009; Yu et al., 2013). VOT is expected to be greater in stressed or accented syllables, though this is more consistently found for voiceless stops than for phonologically voiced stops (Cole et al., 2007; Lisker & Abramson, 1967; Stuart-Smith et al., 2015). VOT strongly decreases with increased speech rate for voiceless stops, but for voiced stops, VOT shows no correlation with speech rate or only a weak trend (e.g. Miller et al., 1986; Stuart-Smith et al., 2015). In the static models we consider both a speaker’s mean speech rate (SPEECH RATE MEAN) and the deviation of speech rate from the speaker’s mean for each token (SPEECH RATE DEVIATION), following Stuart-Smith et al. (2015).

VOT also differs between speakers: varieties of English differ in the magnitude of positive VOT and the prevalence of prevoicing, including British varieties (e.g. Docherty, 1992; Docherty et al., 2011; Scobbie, 2006), and speakers of the same variety of English differ in their baseline VOT after controlling for differences in speech rate (Allen et al., 2003). SPEAKER GENDER may affect VOT (male < female; see Morris et al., 2008). By-speaker random effect terms in our models account for interspeaker differences in VOT beyond speaker gender.

## A.2 Coronal stop deletion

Deletion of word-final t/d in consonant clusters is affected by a variety of linguistic and social factors, reviewed by Hazen (2011); Schreier (2005); Tagliamonte & Temple (2005).

CSD rate is affected by the context the host word (ending in t/d) is produced in and properties of the host word. FOLLOWING CONTEXT strongly affects deletion rate, with deletion always found to be more frequent before consonants (C) than before vowels (V), and most frequent before another coronal stop (/t/ or /d/). Different studies find different effects of a following “pause”, which is usually treated as an alternative context (to C or V). Following Tanner et al. (2016), we instead consider PAUSE DURATION (following the t/d site), acting as a proxy for prosodic boundary strength, as a variable which can affect CSD rate independently of following context; deletion is expected to be less likely for longer pauses. (Pause duration was manually transcribed during the annotation process.) PRECEDING PHONOLOGICAL CONTEXT affects CSD, with deletion generally found to be more likely after sonorants than after obstruents, and most likely after sibilants (especially /s/). CSD, as a reductive process, is expected to be more likely for higher FREQUENCY words (e.g. Myers & Guy, 1997; Walker, 2012), at faster SPEECH RATES, and in more casual speech. MORPHOLOGICAL CLASS of the host word often affects CSD, with deletion progressively more likely in regular past-tense forms, irregular pasts, and non-pasts, though Tagliamonte & Temple (2005) notably failed to find any effect for York English.

While most social factors have weak effects on CSD rate, including speaker GENDER, different regional and ethnic varieties of English show very different deletion rates (Schreier, 2005). We thus expect the speakers in our dataset to differ greatly in overall CSD rate, due to their heterogeneous backgrounds, and account for speaker differences beyond gender in the static model of CSD rate using by-speaker random effect terms.

### A.3 Vowel formants

Variation in F1 and F2 is conditioned by many properties of context, the utterance, and the speaker. First, vowel formants are strongly affected by coarticulation. The strongest coarticulatory effects come from the consonantal context: the manner, place, and voicing of adjacent consonants affect vowel formants in complex ways which are well-documented for laboratory speech (e.g. Beddor, 1982; Hillenbrand et al., 2001; Stevens & House, 1963), as well as in sociolinguistic studies in (usually) spontaneous speech (Labov, 1994), making standard sociolinguistic practice appropriate for our own (spontaneous speech) dataset. Our models of F1 and F2 use a parametrization of consonantal context into four variables which is common in sociolinguistic studies: PRECEDING SEGMENT identity (10 levels), FOLLOWING SEGMENT MANNER of articulation (seven levels), FOLLOWING SEGMENT PLACE of articulation (seven levels), and FOLLOWING SEGMENT VOICING (three levels). We use the coding given by FAVE (Rosenfelder et al., 2011) for these variables, described in Table 2.

Prosody and speech style also strongly affect vowel formants, with vowels in more prominent positions or ‘clearer’ speech (indexed by slower speech rate, etc.) realized as less ‘reduced’ (e.g. more peripheral in formant space) relative to non-prominent positions or casual speech (e.g. Cho, 2005; Fourakis, 1991; Moon & Lindblom, 1994). We considered both speech rate and vowel duration as correlates of clear speech and prominence. Only log-transformed VOWEL DURATION was included in the final models of F1 and F2, as vowel duration and speech rate were highly correlated, and models including just vowel duration generally had lower AIC than models including both terms or just speech rate. We excluded tokens from unstressed syllables (Sec. 3.3.1).

Formant frequencies for a given vowel vary greatly by speaker: both across dialect regions (see Sec. 3.3), and among speakers within a dialect region, due to physiological differences and social factors (e.g. Hillenbrand et al., 1995; Labov, 2000) We do not explicitly model the effects of by-speaker factors on vowel formants, due to data sparsity, but account for speaker differences using by-speaker random effect terms.

## Appendix B Static models

This section describes the statistical models of how VOT and vowel formants depend on the static factors described above (Appendix A), as well as controlling for variability among speakers and words. Each static model is a mixed-effects regression model fitted using the `lme4` package in R (Bates et al., 2014). The static model for CSD comes from Tanner et al. (2016), and relevant aspects are summarized in Sec. 4.2.

### B.1 Voice onset time

Two linear mixed effects models were fit of  $\log(\text{VOT})$  for voiced and voiceless stops for the VOT dataset (Sec. 3.1.1). As fixed effects, each model contained main-effect terms for the nine static factors in Table 2 (top), as well as a term to account for annotator identity (ANNOTATOR), which was included as a factor with three levels for voiceless VOT and two levels for voiced VOT (the number of annotators for each subset) Each variable was coded such that the intercept of each model has the interpretation “VOT, when all variables are held at their average values”: factors are coded using Helmert contrasts, and continuous variables are *standardized* by centering and dividing by two standard deviations

(Gelman & Hill, 2007). `SPEECH RATE DEVIATION` was further coded as a nonlinear spline with two components for the voiceless stop data (using `rms` in the `rms` package; Harrell, 2014) and as a linear function for the voiced stop data, based on exploratory plots of the relationship between this variable and  $\log(\text{VOT})$  for each subset. The models contained interaction terms as fixed effects, determined by exploratory data analysis (exploratory plots, stepwise model selection). The voiceless VOT model contained interactions between `PLACE OF ARTICULATION` and each of `SPEECH RATE DEVIATION`, `STRESS`, and `FOLLOWING SEGMENT`; the voiced VOT model contained interactions between `PLACE OF ARTICULATION` and each of `SPEECH RATE DEVIATION`, `SPEECH RATE MEAN`, and `FOLLOWING SEGMENT`. Both models contained all possible by-speaker and by-word random effects (Barr et al., 2013), with the exception of terms for `ANNOTATOR`, but excluded correlations between random-effect terms, to aid convergence.

The results of the models are summarized in Table A1. Fixed-effect coefficients are shown with associated standard errors, degrees of freedom, test statistic ( $t$ ), and significances, calculated using the Satterthwaite approximation using `lmerTest` (Kuznetsova et al., 2014). Random-effect terms are not shown.

## B.2 Vowel formants

Six linear mixed-effects regression models of formant realization—F1 and F2 for `GOOSE`, `STRUT`, and `TRAP`—were fitted for the vocalic datasets (Sec. 3.3.1).

Each model contained terms for each static factor in Table 2 (bottom): fixed effect terms for `VOWEL DURATION` (log-transformed, then standardized) and `FOLLOWING SEGMENT VOICING` (coded using Helmert contrasts), and random intercept terms for `FOLLOWING SEGMENT MANNER`, `FOLLOWING SEGMENT PLACE`, `PRECEDING SEGMENT`. With this coding of static factor variables, the intercept of each model has the interpretation of a grand mean.<sup>17</sup> Each model also contained by-speaker and by-item random intercepts (where item = instance of a vowel in a particular word, e.g. first `TRAP` vowel in ‘grandad’), a by-speaker random slope for `VOWEL DURATION`, and a term for its correlation with the by-speaker random intercept.

This model structure differs from that used for the VOT and CSD models in two ways, which are motivated by data sparsity, while keeping in mind the goal of the static models: to account for the effects of static factors on vowel formants so that we can then examine whether there is time dependence in housemates’ pronunciations of vowels after accounting for these factors. First, some static factors are coded as *random effect* terms, rather than as fixed effects. We do this for the static factors which have more than two levels (`FOLLOWING SEGMENT MANNER`, `FOLLOWING SEGMENT PLACE`, and `PRECEDING SEGMENT`), because these factors are not independent and the data is highly unbalanced among different levels of these factors,<sup>18</sup> making stable estimates of fixed-effect coefficients for these factors impossible. Second, there are no random slope terms accounting for differences among speakers in the effects of static factors with the exception of `vowel duration`, because the data is simply too sparse to estimate differences between speakers in the effect of phonological context, given the large number of contexts. The models thus assume that the effects of context are the same for all speakers. This assumption should not affect the findings of the dynamic models with respect to our research questions.

---

<sup>17</sup>Note that random intercept terms index deviations from the grand mean.

<sup>18</sup>For example: there is no bilabial lateral in English, and among the 5365 `TRAP` tokens, only 4 are before a dental consonant.

Table A1: Summary of fixed-effect coefficients in the static models of log(VOT) for voiceless (top) and voiced (bottom) stops. Primes indicate different contrasts for factors with multiple levels, or spline components (for SPEECH RATE DEVIATION).

<b>Voiceless stops</b>					
Coefficient	Estimate	SE	df	<i>t</i>	<i>P</i> (> <i>t</i> )
Intercept	3.970	0.039	19.5	100.76	<0.001
Speech rate deviation	-0.214	0.020	50.7	-10.58	<0.001
Speech rate deviation'	0.034	0.009	33.1	3.91	<0.001
Phrase position	0.071	0.017	16.9	4.08	<0.001
Annotator	-0.014	0.006	12061.8	-2.35	0.019
Annotator'	0.024	0.003	12116.7	9.31	<0.001
Place of articulation (P vs T/K)	0.131	0.009	38.1	14.10	<0.001
Place of articulation' (T vs K)	-0.116	0.019	33.8	-6.17	<0.001
Stress	0.178	0.028	51.3	6.27	<0.001
Following segment type	0.168	0.024	30.8	6.92	<0.001
Vowel height	0.071	0.020	55.9	3.46	0.001
Frequency	-0.085	0.017	92.5	-5.10	<0.001
Speech rate mean	-0.041	0.081	18.1	-0.51	0.617
Gender	-0.094	0.078	17.7	-1.21	0.244
POA:Speech rate deviation	0.023	0.014	69.5	1.59	0.116
POA':Speech rate deviation	-0.027	0.017	575.0	-1.61	0.108
POA:Speech rate deviation'	-0.007	0.006	3209.1	-1.15	0.249
POA':Speech rate deviation'	0.007	0.009	160.6	0.82	0.411
POA:Stress	-0.057	0.020	34.4	-2.89	0.007
POA':Stress	0.103	0.033	102.1	3.17	0.002
POA:Following segment type	0.033	0.017	38.7	1.91	0.063
POA':Following segment type	-0.020	0.032	37.8	-0.62	0.54

<b>Voiced stops</b>					
	Estimate	SE	df	<i>t</i>	<i>P</i> (> <i>t</i> )
Intercept	2.766	0.050	18.4	55.62	<0.001
Speech rate deviation	-0.047	0.017	20.5	-2.82	0.01
Phrase position	0.096	0.033	30.1	2.94	0.006
Annotator	0.128	0.006	12777.4	22.10	<0.001
Place of articulation (B vs D/G)	0.225	0.011	26.1	20.07	<0.001
Place of articulation' (D vs G)	0.003	0.028	22.1	0.11	0.91
Stress	-0.035	0.039	23.6	-0.91	0.37
Following segment type	0.522	0.040	30.5	13.21	<0.001
Vowel height	0.080	0.026	37.1	3.15	0.003
Frequency	0.018	0.016	50.5	1.15	0.256
Speech rate mean	0.066	0.100	17.0	0.66	0.518
Gender	-0.150	0.072	22.0	-2.08	0.05
POA:Speech rate deviation	-0.015	0.009	11.6	-1.68	0.119
POA':Speech rate deviation	-0.029	0.022	20.9	-1.33	0.198
POA:Speech rate mean	-0.089	0.022	21.0	-4.03	<0.001
POA':Speech rate mean	0.039	0.056	19.4	0.70	0.494
POA:Following segment type	0.035	0.023	28.5	1.54	0.135
POA':Following segment type	-0.236	0.050	16.9	-4.75	<0.001

Table A2: Summary of static models of F1 and F2 for GOOSE, STRUT, and TRAP. *Top*: summary of fixed-effect coefficients. Primes indicate different contrasts for factors with multiple levels. *Bottom*: variance components (as standard deviations) for random effect terms corresponding to properties of adjacent consonants.

<b>Fixed effects</b>										
GOOSE	F1					F2				
	Estimate	SE	df	<i>t</i>	<i>P&gt;(&gt; t)</i>	Estimate	SE	df	<i>t</i>	<i>P&gt;(&gt; t)</i>
Intercept	-0.99	0.08	8.50	-13.21	0.00	-0.00	0.31	10.97	-0.02	0.99
Vowel duration	0.02	0.03	11.93	0.77	0.46	-0.07	0.05	12.25	-1.38	0.19
Following C voicing	-0.04	0.07	3.81	-0.55	0.61	0.08	0.29	4.41	0.28	0.79
Following C voicing'	-0.01	0.03	4.20	-0.56	0.60	-0.03	0.10	4.68	-0.27	0.80
<b>STRUT</b>										
Intercept	0.23	0.14	17.03	1.63	0.12	-1.09	0.14	23.01	-7.66	0.00
Vowel duration	0.35	0.05	11.60	7.38	0.00	-0.17	0.05	11.13	-3.15	0.01
Following C voicing	-0.12	0.09	8.45	-1.35	0.21	0.04	0.11	7.79	0.39	0.71
Following C voicing'	0.00	0.03	9.22	0.08	0.94	-0.01	0.04	8.27	-0.19	0.85
<b>TRAP</b>										
Intercept	1.37	0.20	238.54	6.85	0.00	-0.38	0.16	101.45	-2.42	0.02
Vowel duration	0.50	0.05	11.22	10.48	0.00	0.02	0.04	11.45	0.44	0.67
Following C voicing	-0.02	0.28	413.74	-0.06	0.96	-0.05	0.19	284.74	-0.25	0.80
Following C voicing'	-0.06	0.09	418.17	-0.65	0.52	0.03	0.06	290.06	0.49	0.63
<b>Random effects</b>										
	GOOSE		STRUT		TRAP					
	F1	F2	F1	F2	F1	F2				
Preceding segment	0.072	0.398	0.079	0.184	0.065	0.136				
Following segment manner	0.126	0.494	0.108	0.135	0.135	0.068				
Following segment place	0.003	0.159	0.074	0.106	0.055	0.103				

Table A2 summarizes the results of the models, with respect to the effects of static factors on F1 and F2 for the three vowels. Fixed-effect coefficients are shown with associated standard errors, degrees of freedom, test statistic ( $t$ ), and significances, calculated as for VOT using `lmerTest`. Also shown are the variance components for each random intercept corresponding to a static factor. Other random effect terms are not shown.

# Supplemental information

## 1. Dynamic model results

Tables SI1–SI3 supplement the results of the dynamic models reported in Sec. 5. Table SI1 summarizes predicted time dependence in VOT for voiced and voiceless stops for each speaker: predicted average VOT (a speaker’s “offset” in the static model), whether the model includes by-day variability and/or a time trend, and  $\sigma$ , which quantifies the degree of by-day variability. Table SI2 summarizes the predicted time dependence in CSD rate for each speaker (analogously to Table SI1): predicted average CSD rate (the intercept of a speaker’s dynamic model), whether there is by-day variability and/or a time trend, and  $\sigma$ , which quantifies the degree of by-day variability (when present). Table SI3 summarizes predicted time dependence in F1 and F2 for each vowel, for each speaker, similarly to VOT.

Table SI1: Summary of predicted time dependence in VOT for each speaker for voiced and voiceless stops. *Mean*: model-predicted average VOT. T/F values are boolean values describing which kind of time dependence a speaker shows (Figure 1): none, by-day variability (BDV), time trends (TT), or both.  $\sigma$  measures the size of by-day variability. (See text.)

Speaker	Voiced stops				Voiceless stops			
	Mean	BDV?	$\sigma$	TT?	Mean	BDV?	$\sigma$	TT?
Dale	16.79	T	0.17	F	66.96	T	0.09	T
Darnell	10.81	T	0.22	F	47.06	T	0.10	F
Lisa	20.79	T	0.09	F	55.84	T	0.03	F
Luke	10.66	T	0.13	F	38.07	T	0.04	F
Michael	16.50	T	0.17	T	62.58	T	0.08	T
Mohamed	12.27	T	0.26	F	38.47	T	0.05	T
Rachel	17.42	T	0.13	T	57.56	T	0.11	T
Rebecca	15.88	T	0.17	T	50.78	T	0.06	T
Rex	20.23	T	0.09	T	51.63	T	0.09	T
Sara	18.07	T	0.21	F	49.14	T	0.16	F
Stuart	16.17	T	0.12	F	68.52	T	0.11	F

## 2. Plasticity

Tables SI5 and SI4 give supplemental information for the phonetic plasticity analysis (Sec. B.2). Table SI4 shows the plasticity value for each variable/speaker pair. Table SI5 shows Kendall’s  $\tau$ , a measure of rank correlation, between the plasticities for each pair of variables (across speakers), with an associated p-value. (We use rank correlation because plasticity for different variables is on different scales.) With one exception (VOT for voiced and voiceless stops), all pairs of variables are positively associated ( $\tau \in [0.05, 0.67]$ ). However, these correlations are usually not significant ( $p > 0.05$ ; exceptions: STRUT/TRAP, GOOSE/CSD). Thus, while there is a consistent trend for a speaker who varies more in a given variable to vary more in another, the trend for individual pairs of variables is usually weak.

Table SI2: Summary of predicted time dependence in coronal stop deletion rate for each speaker. T/F columns are defined as in Table SI1.  $\sigma$  measures the size of by-day variability. (See text.)

Speaker	Mean CSD rate	BDV?	$\sigma$	TT?
Dale	0.62	T	0.15	F
Darnell	0.68	T	0.22	F
Lisa	0.52	F	–	F
Luke	0.43	T	0.08	F
Michael	0.51	T	0.12	T
Mohamed	0.67	T	0.20	F
Rachel	0.59	T	0.05	F
Rebecca	0.48	F	–	T
Rex	0.52	T	0.15	F
Sara	0.60	F	–	T
Stuart	0.61	T	0.13	T

Table SI3: Summary of predicted time dependence in (Lobanov-normalized) F1 and F2 for each speaker, for GOOSE, TRAP, and STRUT. T/F columns are defined as in Table SI1, with T in an ‘F1+F2’ column if either F1 or F2 is ‘T’ in the corresponding column.  $\sigma$  measures the size of by-day variability. (See text.)

Vowel	Speaker	F1+F2		F1			F2				
		BDV?	TT?	Mean	BDV	TT	Mean	BDV	TT		
				$\sigma$			$\sigma$				
GOOSE	Dale	T	T	-1.02	T	0.00	T	0.40	T	0.08	F
	Darnell	T	T	-1.25	T	0.14	F	0.05	T	0.14	T
	Kathreya	T	F	-0.86	T	0.17	F	-1.23	T	0.09	F
	Lisa	T	F	-0.84	T	0.03	F	0.40	F	–	F
	Luke	T	F	-1.13	T	0.08	F	-0.23	T	0.21	F
	Michael	T	T	-0.81	T	0.11	T	0.12	T	0.06	F
	Mohamed	T	F	-1.09	T	0.16	F	-0.24	T	0.17	F
	Rachel	T	T	-0.94	T	0.08	F	-0.54	T	0.11	T
	Rebecca	T	T	-1.01	T	0.06	F	0.52	T	0.51	T
	Rex	T	F	-1.02	T	0.12	F	0.29	T	0.05	F
	Sara	T	T	-1.03	T	0.11	T	0.21	T	0.23	F
	Stuart	T	T	-0.93	F	–	T	0.25	T	0.31	F
STRUT	Dale	T	T	-0.17	T	0.08	F	-1.51	T	0.03	T
	Darnell	T	T	0.35	T	0.07	T	-0.99	T	0.07	F
	Kathreya	T	T	0.45	T	0.21	T	-0.95	T	0.09	T
	Lisa	T	F	-0.34	T	0.08	F	-1.61	T	0.07	F
	Luke	T	T	-0.31	F	–	F	-1.39	T	0.06	T
	Michael	T	T	0.88	T	0.09	F	-1.06	T	0.05	T
	Mohamed	T	T	0.65	T	0.16	F	-0.89	T	0.07	T
	Rachel	T	F	0.25	T	0.15	F	-0.47	T	0.13	F
	Rebecca	T	T	-0.09	T	0.13	F	-0.74	T	0.18	T
	Rex	T	F	0.49	T	0.12	F	-1.10	T	0.02	F
	Sara	T	F	0.79	T	0.09	F	-0.92	T	0.06	F
	Stuart	T	T	-0.27	T	0.03	T	-1.47	T	0.10	F
TRAP	Dale	T	T	1.54	T	0.15	F	-0.78	T	0.14	T
	Darnell	T	T	1.06	T	0.26	F	0.29	T	0.21	T
	Kathreya	T	T	1.05	T	0.20	T	-0.07	T	0.11	T
	Lisa	T	T	1.58	T	0.05	F	-0.51	T	0.04	T
	Luke	T	T	1.58	T	0.14	F	-0.45	T	0.07	T
	Michael	T	T	1.50	T	0.09	F	-0.59	T	0.05	T
	Mohamed	T	F	1.28	T	0.15	F	-0.14	T	0.06	F
	Rachel	T	T	1.68	T	0.19	T	-0.71	T	0.08	F
	Rebecca	T	T	1.41	T	0.36	T	-0.48	T	0.12	F
	Rex	T	T	1.26	T	0.12	F	-0.26	T	0.06	T
	Sara	T	T	1.23	T	0.24	T	-0.22	T	0.08	F
	Stuart	T	F	1.30	T	0.15	F	-0.63	F	–	F

Table SI4: Plasticity value for each core housemate for each phonetic variable. (See text.)

Speaker	VOT: voiced (msec)	voiceless	CSD (% deleted)	GOOSE (normalized F1/F2)	STRUT	TRAP
Dale	0.118	0.072	0.111	0.094	0.082	0.123
Darnell	0.172	0.068	0.210	0.129	0.085	0.197
Kathreya	–	–	–	0.090	0.238	0.197
Lisa	0.060	0.008	0.000	0.006	0.049	0.028
Luke	0.097	0.022	0.035	0.119	0.035	0.098
Michael	0.139	0.065	0.173	0.062	0.051	0.063
Mohamed	0.179	0.038	0.164	0.108	0.100	0.080
Rachel	0.096	0.076	0.014	0.065	0.109	0.163
Rebecca	0.156	0.069	0.543	0.445	0.113	0.280
Rex	0.071	0.081	0.137	0.054	0.069	0.080
Sara	0.158	0.119	0.223	0.144	0.046	0.138
Stuart	0.088	0.077	0.297	0.171	0.068	0.075

Table SI5: Correlation across speakers in plasticity between each pair of phonetic variables. Below diagonal: rank correlation (Kendall’s  $\tau$ ). Above diagonal: associated  $p$ -values.

	VOT (vcd)	VOT (vl)	CSD	GOOSE	STRUT	TRAP
VOT (vcd)	–	0.648	0.121	0.165	0.359	0.121
VOT (vl)	-0.127	–	0.165	0.542	0.879	0.445
CSD	0.382	0.345	–	0.003	0.648	0.445
GOOSE	0.345	0.164	0.673	–	0.737	0.116
STRUT	0.236	0.055	0.127	0.091	–	0.014
TRAP	0.382	0.200	0.200	0.364	0.545	–