Hayes-Harb, R., Barrios, S., & Tripp, A. (2022, November 29). Whose Input Matters? The influence of socially-differentiated input sources in adult Lx phonetic learning.
<https://doi.org/10.31234/osf.io/t78cu>

Whose input matters? The influence of socially-differentiated input sources in adult Lx phonetic learning

Rachel Hayes-Harb and Shannon Barrios
University of Utah

Alayo Tripp
University of Minnesota

Abstract

Input is a necessary condition for language acquisition. In the language classroom, input may come from a variety of sources, including the teacher and student peers. Here we ask whether adult Lx learners are sensitive to the social roles of teachers and students such that they exhibit a preference for input from the teacher. We conducted an experiment wherein adult English speakers heard words in an artificial language. During an exposure phase, in one condition a “teacher” produced words with 25 ms of VOT on initial stop segments and a “student” produced the same words with 125 msec of VOT; in another condition the VOT durations were reversed. At test, participants judged productions by a different “student” and demonstrated a preference for the productions that matched the VOT durations of the teacher during exposure, providing evidence for an influence of social factors in differentiating input in Lx acquisition.

1 Introduction

The massive variability in language input presents a central obstacle to explaining how speech perception and learning proceed. To account for linguistic variation as a source of complexity in the language learning task, much attention has been paid to documenting differences between

languages, treating passive linguistic experience as the central driving force in determining language outcomes. To reconcile this account with what is known about the roles of motivation, social and cultural knowledge, and non-linguistic cognition, we must also consider the role of social milieu in shaping linguistic interactions and outcomes. Although listeners appear to perceive and acquire linguistic contrasts according to shared understandings regarding the distribution of social power, to our knowledge, no empirical data has previously been published demonstrating that adult language learners rely on knowledge of social roles to selectively attend to the fine grained phonetic variability of specific talkers.

In this study we ask whether extra-linguistic social cues (in this case, teacher and student roles conveyed by iconographic images) impact participants' memory of and preferences for speech input provided by different speakers, and consequently, their acquisition of the phonetics of a second or subsequent language (Lx; Dewaele, 2017).

In the next section we begin by discussing several kinds of language variation: that arising from phonological and phonetic differences between languages, or *language-specific* variation, and that arising from individual differences, or *talker-specific* variation. Given the literature suggesting that speech perception outcomes can also be associated with extra-linguistic, or non-linguistic social variation, we here pose the question of how labeling voices as teacher and student affects adult participants' judgment of their pronunciation. Given that language knowledge interacts with the ability to attend to differences between individual talkers (Goggin et al., 1991; Perrachione & Furbeck, 2019), and that this capacity may be adult-like around six years of age (Fecher & Johnson, 2018, 2021), it is not clear what role talker-specific perception plays in the learning of pronunciation. Knowledge of social roles may lead learners to selectively direct more attention to talker-specific linguistic variation when the talker can be

ascribed a powerful position. If so, social perception would have to be understood as partially gating the perceptual availability of language-specific information.

Given a preponderance of evidence associating deviation from “native” or “standard” speech behaviors with racial stigma and prejudice (Ramjattan, 2019c), distrust (Lev-Ari & Keysar, 2010), and impaired memory (Lev-Ari et al., 2018) on the part of listeners, coupled with the rhetorical importance historically placed on “native-speaker” status in language teaching (Llurda, 2005; Ramjattan, 2019a), it is important to investigate the mechanisms underlying how attention to detail in speech is influenced by social knowledge. We select voice onset time (VOT) as a readily manipulable acoustic feature which has been demonstrated to impact listeners’ representations of both talker-specific and language-specific speech variation, and can help us to learn how this variation is accounted for in the representation of Lx structure.

The results from an artificial language learning experiment show participants selectively generalizing their learning of pronunciation, indicating a greater preference for the pronunciations provided by talkers designated as teachers. Although it has been theorized that language learners necessarily rely on sociolinguistic distinctions when deciding who to learn from, most models of phonetic and phonological learning have not incorporated this type of social information (Tripp et al., 2021). The present data provide the first empirical evidence suggesting that in a classroom setting, perception of social authority can guide language learners’ attention to speech accent.

1.1 Kinds of linguistic input variability and their role in acquisition

Variability is a property of signed (Tyrone & Mauk, 2012) and spoken (Foulkes & Docherty, 2006) language production generally. Variability is readily accommodated by first language (L1) speakers in monolingual settings (Banks et al., 2015), and by multilingual speakers in

multilingual settings (Samuel & Larraza, 2015). Lx acquisition (Lx; Dewaele, 2017) can involve learning numerous novel language-specific contrasts. However, talker-specific contrasts present potential obstacles to perception in both the L1 and Lx, and the literature demonstrates that there is a tight link between the ability to recognize and interpret language-specific and talker-specific speech variation. This section briefly defines several types of linguistic variation, before moving to a review of the literature on their interaction in speech perception experiments.

Language-specific contrasts include phonological contrasts. For example, L1 English speakers, who have a voiceless glottal fricative phoneme /h/, must learn to contrast the glottal and pharyngeal voiceless fricatives /h/ and /ħ/ in Arabic in order to distinguish Arabic minimal pairs such as /s^hɑ:hir/ ‘watchful’ and /s^hɑ:ħir/ ‘charming’, whereas Spanish learners of Catalan encounter the novel Catalan /e/ and /ɛ/ contrast (e.g., /netə/ ‘granddaughter’ and /nɛtə/ ‘clean.fem’). However, even languages with the same (or similar) phonological contrasts can differ in how those contrasts are implemented phonetically, such that Lx acquisition also entails learning novel phonetics. For example, English learners of Spanish encounter voiced and voiceless coronal stops of Spanish (spelled <t> and <d> in both languages) produced with a dental place of articulation, in contrast to the English alveolar place of articulation. Moreover, the stop consonant voicing contrast present in both languages is also produced differently. In English, /d/ and /t/ are produced with short and long-lag voice onset time (VOT; the time measured from the release of the stop closure to the onset of voicing), respectively, whereas Spanish /d/ is prevoiced and /t/ is produced with a short-lag VOT (Lisker & Abramson, 1964).

Speech variability also includes *indexical* variation, or encoded social information about speakers (e.g., regarding speakers’ emotions, gender, social class, age, and dialect). Listeners are sensitive to the social information conveyed by speech variability, though their interpretation of

sociophonetic cues is not always accurate (see, e.g., Brown et al., 2021; Morgan & Ferguson, 2017; Moussu & Llurda, 2008). Importantly, listeners are also sensitive to extra-linguistic social cues (e.g., those conveyed through the association of dialect with ethnicity or national origin), which affect their processing of the speech signal (e.g., Drager, 2011; Kang & Rubin, 2014; McGowan, 2015; Szakay et al., 2016). There is also speech variation which is not language-specific, but which can be construed as more *circumstantial*. For example, the perception of a single talker's speech can be variable due to incidental factors such as whether the person's mouth is full (Samuel & Kraljic, 2009). Sources of variability, such as speech rate (Maslowski et al., 2019), phonological context (Baese-Berk, 2018), dialectal variation (Clopper & Bradlow, 2008), variation associated with multilingualism (Bradlow & Bent, 2008; Witteman et al., 2013; Xie et al., 2017), speech style (Smiljanic, 2021), and individual talker differences (Reinisch, 2016; Samuel & Kraljic, 2009), may complicate the mappings between the auditory input and underlying sound categories.

Variability in the speech signal associated with exposure to multiple talkers supports robust and generalizable acquisition of speech sound categories both in infants (Rost & McMurray, 2009, 2010) and in adults (Lively et al., 1993, 1994). Yet, Perrachione et al. (2011) found that the beneficial effect of multitalker input variability for adult Lx learners is moderated by their pre-existing perceptual sensitivity to the targeted sound contrast. Meunch and Creel (2013) further showed that exposure to two accents slowed adults' learning of new words relative to a single accent, demonstrating that multitalker exposure may not only facilitate, but may also interfere with speech processing. Relatedly, adult learners' exposure to so-called "mispronunciations" associated with L1-accented input from learner peers has been implicated in Lx learning difficulty. In such cases, variation from multitalker exposure could create

inconsistencies in the input that interfere with the establishment of lexical representations that effectively encode contrastive phonetic categories representative of the Lx (see, e.g., Llompart & Reinisch, 2021). One way that learners can navigate the complexity of linguistic variation is through tracking linguistic variation which is unique to particular talkers.

1.2 Variability in VOT

VOT cues word-initial stop voicing contrasts in many languages, and is thus relevant for defining phonemic categories and phonological contrasts. For example, English-learning infants and adult English speakers have been shown to leverage systematic variation in VOT in their input to learn relevant phonemic contrasts in the absence of other (explicit or implicit) information about the contrast (e.g., Hayes-Harb, 2007; Maye, 2000). This finding is suggestive of an exquisite ability of infant and adult listeners to attend to the statistical properties of their highly variable input and to make inferences about language-specific structure from those properties.

However, in addition to cueing language-specific contrasts, VOT also can be recruited to build talker-specific linguistic representations (Allen & Miller, 2004). Building on Allen, Miller, and DeSteno's (2003) finding that L1 English-speakers exhibit systematic *sub-phonemic* variability in VOT, Allen and Miller (2004) demonstrated that English-speaking listeners can track the characteristic VOTs of individual talkers. Following an exposure phase where two female talkers differed systematically in the VOT durations of voiceless stops, English-speaking adult listeners selected which of two VOTs is more characteristic of a given talker. Together, these findings suggest that not only do talkers differ systematically in their production of sub-phonemic variability, but that experienced listeners can learn sub-phonemic properties of the productions of particular talkers. It is unknown whether adult language learners similarly track

the phonetics of individual talkers in an unfamiliar language. One way Lx learners might acquire the phonetics and phonology of the target language is through selective attention to individual talkers whose input they expect to be representative of Lx users, and therefore maximally useful for learning language-specific representations.

Because listeners are sensitive to VOT as a cue to both language-specific and talker-specific variation, the case of subphonemic variation in VOT gives us an excellent place to begin probing the relationship between talker-specific knowledge, knowledge of social variation, and how learners create linguistic representations for a novel Lx.

1.3 Social influences on speech perception

Speech processing researchers often attempt to isolate phenomena of immediate interest from their social contexts by controlling for factors such as participant gender, age, social class, and language/dialect background. However, sociophonetic variation is an inherent property of spoken language (Foulkes & Docherty, 2006), and, as a result, speech input is never socially neutral. Indeed, the presentation of social information has been shown to significantly alter participant behavior in listening tasks. Niedzielski (1999) demonstrated that adults use knowledge of regional dialect in the interpretation of vowels, showing that the same talker elicited different responses when labeled as from Detroit or Canada. Similar contrasts in performance can be induced in measurements of talker intelligibility through the presentations of still faces (Babel & Russell, 2015; McGowan, 2015; Rubin, 1992). For a review of the literature on the attribution of social characteristics to talkers and its effects on speech perception see: (Babel & Munson, 2014; Tripp & Munson, 2022).

Effects on speech perception can be obtained through providing the participant with an enormous range of social and linguistic cues, including symbolic cues in the environment which

implicitly suggest salient social contrasts (Hay et al., 2006; Hay & Drager, 2010), social category labels explicitly describing the talker (D’Onofrio, 2018; Johnson et al., 1999), or through explicitly cuing the participant to listen for different named languages (Gonzales et al., 2019). Affiliative social information can also be gleaned from the acoustic signal alone: adults demonstrate the ability to use dialectal variation to broadly segregate groups of talkers using cues to regional dialect (Clopper & Pisoni, 2004). These effects may be associated with early and automatic neural responses, given evidence that listeners can pre-attentively and categorically discriminate between different dialects (Scharinger et al., 2011). Attention to these sociolinguistic cues can also have downstream effects on language processing, governing performance on tasks of semantic priming (Sumner et al., 2013), and predicting beliefs about non-linguistic characteristics and capacities of the talker (Kang & Rubin, 2009).

Importantly, listeners appear to recruit the same broad social categorization systems to aid their interpretation of speech in their various languages, using judgments of individual talkers to apply sociolinguistic expectations (Szakay et al., 2016). Speech processing outcomes are affected by both listener expectations regarding the identity of individual talkers and by the listener’s knowledge of stereotypes regarding the incongruency of pairings which combine specific linguistic and non-linguistic features, such as accents and racialized identities. For example, Rubin (1992) demonstrated that US college students exhibited more accurate comprehension of an audio lecture when it was accompanied by a photograph of a purported lecturer they identified as “Caucasian” than when the same audio lecture was accompanied by a person they identified as “Asian.” Accordingly, effects of manipulations to the social presentation of talkers is shown to have either facilitative or detrimental effects on learning (Babel et al., 2019).

In the present study we are interested in the effects of the social roles of teachers and students in the language classroom setting. Our research question is: do adult learners prefer the phonetics of a teacher's productions relative to those of a student peer in an artificial language classroom setting? We hypothesize that language learners expect teachers in the language classroom setting to be authoritative with respect to pronunciation of the target language, and thus demonstrate a preference for these talkers' phonetic input when constructing the categories required to judge the quality of speech in the unfamiliar language. While one may intuitively expect language learners to prefer teachers' pronunciations over those of fellow learners, to date, there is no empirical evidence demonstrating that presenting this social information is causally linked to learners' preferences for particular phonetic behavior.

To address our research question, we take advantage of the well-studied VOT dimension. We manipulated the VOT associated with initial voiceless stop consonants in an artificial target language such that the VOT associated with voiceless stops produced by individual talkers were either shorter or longer than typical VOTs reported for English speakers (Allen et al., 2003; Lisker & Abramson, 1964). By pairing talker's productions with visual information about the talker's role in the artificial classroom (using an iconographic image of a teacher and students; see Figures 1 and 2), we are able to examine the effect of talker role on participants' impressions of what the target language "should" sound like. If adult participants are able to track the VOTs of individual talkers (as previously reported by Allen & Miller, 2004), are sensitive to extra-linguistic social cues (Kang & Rubin, 2009; Rubin, 1992), and use these cues to guide the inferences they make about the phonetics of the target language, we expect to observe an effect of talkers' social roles on participants' preferences. Supposing teachers in the language classroom are preferred sources of input, we would expect that participants will value input from

the teacher over that of a student peer, showing a preference at test for long VOTs when the teacher produces long VOTs and a preference for short VOTs when the teacher produces short VOTs.

2 Methods

2.1 Researcher Positionality

The first and second authors are white L1 speakers of US English, and the third author is a racialized L1 speaker of US English; all three authors have studied additional languages (in our cases, elite multilingualism; Flores & Rosa, 2019), and the first two authors have experience as professional language teachers (of English and Spanish, respectively). Listeners' assumptions about the relative value of input from teachers and student peers is the phenomenon of interest in this research. As researchers, we are interested in these assumptions, and conduct this research for the purpose of investigating, not reinforcing or legitimizing, them.

2.2 Participants

Data from 264 participants (non-binary = 3, female = 154, male = 107, mean age = 29.32; range 18-62 years) are presented and analyzed here. All identified English as (one of) their “native” language(s), as well as a language that they spoke with parents and/or other childhood caregivers; 37 indicated one or more additional “native” or parental/caregiver languages (American Sign Language, Belizean Creole, Cantonese, Chinese, Dutch, Farsi, Finnish, Georgian, German, Greek, Hindi, Japanese, Korean, Lao, Mandarin, Polish, Portuguese, Punjabi, Russian, Samoan, Spanish, Tagalog, Tamil, Tongan, and Vietnamese), and 133 reported knowledge of one or more additional languages (American Sign Language, Catalan, Chinese,

Danish, Dutch, Farsi, French, German, Greek, Hindi, Italian, Japanese, Khmer, Korean, Latin, Lithuanian, Mandarin, Navajo, Old English, Russian, Serbo-Croatian, Spanish, Swedish, Tamil, and Vietnamese). Data collected from additional participants were excluded for the following reasons: participants who did not identify English as (one of) their “native” language(s), as well as a language that they spoke with parents and/or other childhood caregivers (16); those who were younger than 18 (1) or older than 65 years of age (5); those who self-reported speech, hearing, or neurological disorders (16) and those taking medications affecting motor skills (2); those who indicated familiarity with one or more of the voices in the study (the talkers were members of the department community; 9); and one participant whose response pattern indicated they were not attending to the task (1).

Participants were recruited via the [removed for review—university students enrolled in linguistics classes] participant pool (n=89) and Prolific (prolific.co; n=175), and received either course credit or financial compensation (at a rate of USD15/hour) for their participation. They were randomly assigned to one of eight conditions crossing Teacher-Talker (Talker 1, Talker 2), Teacher VOT (Short, Long), and Exposure Condition (Control, Experimental). Importantly, each of the two exposure talkers was associated with the teacher role in half of the conditions. In the Control conditions, the teacher and student talker both produced either long VOTs or short VOTs. As a result, the Control conditions allowed us to observe the effect of the VOT manipulation when there is no social group variability in VOT. In the Experimental conditions, the VOTs produced by the teacher and student during exposure differed in VOT duration, such that the teacher produced long VOTs and the student produced short VOTs or vice versa. As a result, the Experimental conditions allow us to observe the effect of the VOT manipulation in the presence of social group variability. Importantly, the talker who was assigned to the teacher role

was counterbalanced such that auditory input in the Experimental condition in which the Teacher-Talker is Talker 1 and produces long VOTs is identical to the condition in which the Teacher-Talker is Talker 2 and produces short VOTs, the same is true for the experimental conditions in which the Teacher-Talker is Talker 1 who produces short VOTs and the condition in which the Teacher-Talker is Talker 2 who produces long VOTs. The only difference between these experimental conditions is which talker is assigned the role of teacher.

There were a total of 33 participants in each of the eight conditions. The number of participants recruited via the participant pool and Prolific were roughly balanced for each of the eight conditions (median number of participants recruited via Prolific per Exposure Condition = 22; range = 18-25). Table 1 provides information about the participants in each exposure condition.

Teacher Talker	Teacher VOT	Student Talker	Student VOT	N	Mean age (range)	Gender identity: Non-binary/ Female/ Male
<u>Experimental Conditions</u>						
Talker 1	Long	Talker 2	Short	33	26.03 (18-60)	0/22/11
	Short		Long	33	27.88 (18-58)	1/14/18
Talker 2	Long	Talker 1	Short	33	29.09 (18-62)	2/23/8
	Short		Long	33	28.70 (18-50)	0/19/14
<u>Control Conditions</u>						

Talker 1	Long	Talker 2	Long	33	30.36 (18-55)	0/22/11
	Short		Short	33	28.82 (18-54)	0/15/18
Talker 2	Long	Talker 1	Long	33	29.91 (18-53)	0/20/13
	Short		Short	33	33.76 (18-59)	0/19/14
Total				264	29.32 (18-62)	3/154/107

TABLE 1. Participant characteristics by exposure condition.

2.3 Stimuli

Auditory and visual stimuli, along with experiment presentation files, are available at <https://osf.io/uw8y5/>. Auditory stimuli consisted of the 12 CVCV nonwords (4 /p/ initial, 4 /t/ initial, and 4 /k/ initial) shown in Table 2. Each word was produced several times by three women who were born in the US and self-identify as “native” speakers of American English in a quiet recording environment and recorded using Praat (Boersma & Weenink, 2020). Each word was produced with stress on the first syllable and an unreduced vowel in the second syllable. The experimenters manually annotated the long sound file to indicate the beginning and end of each nonword token, and a Praat script (Lennes, 2002) was used to extract each token from the long sound file and save it as a unique .wav file. Next, all tokens were scaled for intensity (root-mean-square amplitude) using a Praat script (Yoon, n.d.). We then selected four tokens of each nonword that were produced clearly and without noticeably lengthened vowels in the second syllables. Another Praat script (Crosswhite et al., n.d.) was then used to create textgrid files, and boundaries were manually inserted at the following points for each token: initial voiceless stop onset (burst), V1 onset, V1 offset, V2 onset, and V2 offset by consulting the waveform and the spectrogram. Vowel boundaries were determined by the intensity of F2. VOT was marked as the

period between the onset of the initial voiceless stop consonant and the onset of V1. A final Praat script (Schertz, 2015) was used to manipulate the VOT duration via Pitch-Synchronous Overlap-Add (PSOLA) for each auditory stimulus to create long and short tokens with 125ms VOT and 25 ms VOT, respectively. We did not otherwise control or manipulate the durations of the stimuli.

Table 2. The 12 nonwords and their assigned “meanings”

/pego/ - fork	/tefu/ - basket	/keto/ - helicopter
/pina/ - boot	/tiga/ - pan	/kinu/ - chicken
/pofi/ - corn	/tono/ - horse	/kole/ - rocking chair
/pule/ - flower	/tuli/ - piano	/kuba/ - spider

Visual stimuli were stylized pictures of a classroom setting consisting of a teacher, a blackboard, and two students. An image corresponding to each word’s “meaning” (line drawings from (Snodgrass & Vanderwart, 1980) appeared in the center of the blackboard each time an auditory word was played (e.g., /tefu/ was paired with a picture of a basket). Color was used to identify the talker during test and training. Exposure talkers (the teacher and one of the students) appeared in orange to indicate that the individual was “talking,” and the test talker (the student who was “talking” during the test) appeared in blue (See Figures 1 and 2 below).

2.4 Procedure

The study was completed entirely online. Participants accessed the study by a link provided on SONA (sona-systems.com) or Prolific (prolific.co). The link opened a Qualtrics (qualtrics.com) survey that was used to provide instructions for study participation and presented the consent document. The initial screen told participants that it was crucial that they complete the study on a

computer with Chrome or Firefox browser, and in a single attempt. They were additionally instructed to find a quiet environment free of distractions, to connect headphones, and to turn off any background music and close any browser tabs, programs, and apps which were non-essential to study participation. They clicked a button labeled “These conditions are met” when the study conditions were met. Next, participants read the consent document and provided informed consent by clicking a button labeled “Agree”.

Following consent, they were taken to the experiment controlled by PsychoPy (Peirce et al., 2019) and hosted on Pavlovia (pavlovia.org). The first screen provided an opportunity for participants to adjust their volume to a comfortable listening level as high-frequency English words (e.g., cow, horse, pig, etc.) produced by a male English speaker played. Next, participants read the following instructions: “You will observe a foreign language class where the students are learning new vocabulary. Sometimes the teacher will be saying the new words, and other times you will hear a student practicing the words. Your job is to learn as much as you can about the way this language sounds. You don’t necessarily need to memorize the words, but you should pay attention to the way they are pronounced.” The next screen read: “You will know the teacher is talking when you see the teacher highlighted in orange like this” while an image with the teacher highlighted in orange appeared on the screen. The same was done for the image indicating that a student talker is talking (See Figure 1).

FIGURE 1. Example visual stimuli for the words /pego/ ‘fork’ and /tiga/ ‘pan’. During the exposure phase, the purported “talker” is colored in orange (in the left panel, the teacher, and in the right panel, the student).

The procedure involved two exposure phases, each followed by a test phase. Each exposure phase consisted of the presentation of the twelve nonwords in a random order (1) produced by the teacher, (2) the student, (3) the teacher again, and (4) the student again, for a total of 48 exposure trials. Images appeared on the screen for 2 seconds and there was an inter-trial interval of 1 second before the presentation of the next item. Following each exposure phase, participants were told “Now you will take a test about the sounds of this language. You will hear a different student pronounce each word twice, and your job is to determine which of the two pronunciations is better. This student will be highlighted in blue” (See Figure 2).

FIGURE 2. Example visual stimulus for the word /tefu/ ‘basket’. During the test phase, the purported “talker” is colored in blue (here, the student on the left).

During the two test phases, participants saw the test student image with the student on the left colored in blue (See Figure 2) and heard two productions (one with short VOT and one with long VOT) produced by a third female talker, separated by an inter-stimulus interval of 1.5s. Participants indicated their response by pressing the “1” key if they preferred the first production, or the “2” key if they preferred the second production. Each test phase consisted of 12 trials, and the order of presentation of long and short VOTs for particular words was counterbalanced across Test 1 and Test 2 (e.g., participants heard [kole] with a long VOT followed by [kole] with a short VOT in Test 1 and [kole] with a short VOT followed by [kole] with a long VOT in Test 2). Following Test 2, participants completed a background questionnaire. Two test phases were administered to compare participant behavior before and after having their attention drawn explicitly to the discrimination task. The entire study took 15-20 minutes.

In the experimental conditions, participants heard the same distribution of phonetic variation. However, the talkers who produce the variation are associated with different roles. To

prevent a confound of social role and talker-specific voice information, the assignment of talkers 1 and 2 to teacher and student roles was counterbalanced across participants.

If we suppose that learners' social perception biases them to represent the behavior of the teacher as more authoritative or normative, then we should expect that utterances by the third talker are more often approved of when they phonetically pattern after the talker that was presented in the teacher role. If participants are not relying on social intuitions to guide their linguistic representations of the novel language, we should expect to see no preference for either production.

3 Results

Data and analysis code are available at <https://osf.io/uw8y5/>. Our data consisted of 24 observations (12 from Test1 and 12 from Test 2) from each of the 264 participants (6336 observations total). Prior to analysis, we first excluded any response that exceeded 4s (428 observations, 6.7% of the data). We then recoded "1" and "2" key responses as 'long' responses (1=long, 0=short) and computed the mean proportion long responses for each participant in each of the eight conditions at Test 1 and Test 2. The mean proportion long response corresponds to the proportion of trials in which a participant chose the auditory stimulus with a long VOT as the "better"-sounding pronunciation.

Figure 3 shows boxplots of proportion long responses by Exposure Condition (Control, Experimental), Test Number (Test 1, Test 2), Teacher VOT (Long, Short). Black dots indicate individual participant means, and the group means for each Exposure Condition are presented in superimposed text. The two left boxplots correspond to the two Control conditions in which the teacher and student both produced either long VOTs or short VOTs during the exposure phase. It was expected that participants in these conditions would favor the VOT that they heard from

both the teacher and the student during exposure, resulting in a difference in mean proportion long response at Test 1 and/or Test 2. As expected, the mean proportion long response for the participants in the Control_Long condition indeed exceeded that of the participants in the Control_Short condition at Test 1 (Control_Long=.66, Control_Short=.42) and Test 2 (Control_Long=.63, Control_Short=.42). This result demonstrates that the phonetic contrast was salient and reliably distinguishable by listeners at test when it was not associated with a contrast in social role of the talker.

These Control conditions can then fruitfully be compared to the Experimental conditions in which participants' auditory input included short and long VOTs in equal number, but were produced contrastively by the same talkers in different roles. If the teacher's productions have a greater influence on listeners' impression of "better" pronunciation, then we expect preference for the VOT that was associated with the teacher talker during exposure (preference for long VOTs when the teacher produced long VOTs, and vice versa) at Test 1 and/or Test 2. This is exactly what we observe. The mean proportion long response for the participants in the Experimental_Long condition indeed exceeded that of the participants in the Experimental_Short condition at Test 1 (Experimental_Long=.54, Experimental_Short=.50) and Test 2 (Experimental_Long=.59, Experimental_Short=.44).

If listeners do not attend to social roles in their phonetic learning, we should find no difference between the two Experimental Conditions, which each expose participants to an identical distribution of sounds. Instead, we find participants' prefer the productions which were more like those produced by the teacher, demonstrating that participants' responses were indeed affected by the social roles assigned to the talkers.

FIGURE 3. Boxplots of proportion ‘long’ responses by Exposure Condition (Control, Experimental), Test Number (Test 1, Test 2), Teacher VOT (Long, Short). Dots indicate individual participant means, and medians are presented as horizontal lines. Superimposed proportions show the group mean for the corresponding Exposure Condition.

We first conducted an omnibus analysis using generalized linear mixed-effects models using the lme4 package (version 1.1-28; Bates et al., 2015) of R (version 4.1.2). The model was fit with long responses as the dependent variable (coded as LongResponse = 1; ShortResponse = 0 for each observation) for which a logistic linking function was used. Following the procedures outlined in Linck and Cunnings (2015) for model construction and reporting, the analysis included contrast coded fixed effects for Teacher-Talker (Talker 1 = -.5, Talker 2 = .5), Exposure Condition (Control = -.5, Experimental = .5), Teacher VOT (Long = -.5, Short = .5) and Test

Number (Test 1 = -.5, Test 2 = .5) in a 2 x 2 x 2 x 2 factorial design and random by-participant and by-word intercepts and by-participant random slopes for Test Number. Models were fit using a maximum likelihood technique. Planned comparisons were conducted using simultaneous tests of general linear hypotheses with the multcomp package in R (version 1.4-18; Hothorn et al., 2008). P-values were adjusted using the single-step method and a family-wise error rate protection via Bonferroni correction. There was a significant main effect of Teacher VOT, and the interactions of Exposure Condition x Teacher VOT, Teacher VOT x Test Number, and Exposure Condition x Teacher VOT x Test Number were also significant. No other main effects or interactions were significant (see Table 3a). Importantly, the counterbalance Teacher-Talker variable exhibited no significant main effects or interactions, and was therefore removed from subsequent analyses.

We next conducted a three-way analysis including contrast coded fixed effects for Exposure Condition, Teacher VOT and Test Number and random by-participant and by-word intercepts and by-participant random slopes for Test Number. The main effect of Teacher VOT, and the interactions of Exposure Condition x Teacher VOT, Teacher VOT x Test Number, and Exposure Condition x Teacher VOT x Test Number were significant; all other main effects and interactions were non-significant (see Table 3b).

Finally, we conducted pairwise comparisons of Teacher VOT for each Exposure Condition (Control, Experimental) at Test 1 and Test 2. The Test Number variable was an artifact of the study method and was not associated with the research questions; however, given the three-way interaction involving Test Number, Test Number is included in our analysis. At Test 1, while the effect of Teacher VOT was significant in the Control condition (estimate = 1.09, SE = .18, $z = 5.97$, $p < .001$) it was not in the Experimental condition (estimate = .12, SE =

.18, $z = .68$, $p = .98$). At Test 2, the effect of Teacher VOT in the Control (estimate = 1.06, $SE = .20$, $z = 5.29$, $p < .001$) and Experimental (estimate = .76, $SE = .20$, $z = 3.81$, $p < .001$) were both significant. We are limited in our ability to interpret the absence of significant differences in performance at Test 1, though we note that the trend is in the same direction as the significant differences found at Test 2. It may be that the additional variation between Teacher and Student created a more complex language learning task, and participants had not yet accumulated enough input by Test 1 to track differential VOT durations, or that the act of taking the first test cued them to the relevance of the VOT dimension, which they then tracked more carefully during the second exposure phase. Future research may help to clarify this issue. Nonetheless, the results demonstrate perception of social role and the overall diversity of the input both interact to determine how language-specific representations are constructed and used to complete a meta-linguistic task.

	<i>Fixed effects</i>				<i>Random effects</i>	
					by-participant	by-item
<i>Parameters</i>	Estimate	<i>SE</i>	<i>z</i>	<i>p</i>	<i>SD</i>	<i>SD</i>
(a) Omnibus analysis						
(Intercept)	0.11	0.09	1.32	0.19	0.78	0.22
Exposure Condition	-0.08	0.11	-0.68	0.5		
Teacher VOT	-0.76	0.11	-6.73	<.001		
Teacher-Talker	-0.09	0.11	-0.78	0.43		
Test Number	-0.03	0.07	-0.36	0.72	0.72	
Exposure Condition x Teacher VOT	0.63	0.23	2.81	<.01		
Exposure Condition x Teacher-Talker	-0.25	0.23	-1.11	0.27		
Teacher VOT x Teacher-Talker	-0.26	0.23	-1.15	0.25		
Exposure Condition x Test Number	0.05	0.15	0.36	0.72		

Teacher VOT x Test Number	-0.3	0.15	-2.07	0.04		
Teacher-Talker x Test Number	-0.14	0.15	-0.97	0.33		
Exposure Condition x Teacher VOT x Teacher-Talker	-0.09	0.45	-0.2	0.84		
Exposure Condition x Teacher VOT x Test Number	-0.68	0.29	-2.32	0.02		
Exposure Condition x Teacher-Talker x Test Number	0.32	0.29	1.09	0.28		
Teacher VOT x Teacher-Talker x Test Number	-0.05	0.29	-0.18	0.86		
Exposure Condition x Teacher VOT x Teacher-Talker x Test Number	-0.6	0.58	-1.02	0.31		
(b) Three-way analysis						
(Intercept)	0.11	0.09	1.32	0.19	0.79	0.22
Exposure Condition	-0.07	0.11	-0.66	0.51		
Teacher VOT	-0.76	0.11	-6.69	<.001		
Test Number	-0.02	0.07	-0.34	0.73	0.73	
Exposure Condition x Teacher VOT	0.64	0.23	2.81	<.01		
Exposure Condition x Test Number	0.05	0.15	0.35	0.73		
Teacher VOT x Test Number	-0.3	0.15	-2.05	0.04		
Exposure Condition x Teacher VOT x Test Number	-0.67	0.29	-2.29	0.02		
Note: All factors were coded using contrast coding, as follows: Exposure Condition (Control=-.5, Experimental=.5), Teacher VOT(Long = -.5, Short=.5), Teacher-Talker (Talker 1 = -.5, Talker 2 = .5), Test Number (Test 1 = -.5, Test 2 = .5)						
(a) Omnibus analysis model formula: LongResponse ~ Exposure Condition*Teacher VOT*Teacher-Talker*Test Number + (1 + Test Number Participant) + (1 Word), family=binomial, control = glmerControl(optimizer="bobyqa"))						
(b) Three-way analysis model formula: LongResponse ~ Exposure Condition*Teacher VOT*Test Number + (1 + Test Number Participant) + (1 Word), family=binomial, control = glmerControl(optimizer="bobyqa"))						

TABLE 3. (a) Mixed effects model omnibus analysis examining the effects of Exposure Condition, Teacher VOT, Teacher-Talker and Test Number and (b) mixed effects analysis examining the effects of Exposure Condition, Teacher VOT and Test Number

Recall that our research question is: do adult learners prefer the phonetics of a teacher's productions relative to those of a student peer in an artificial language classroom setting? Thus we are interested in what happens when participants are exposed to both short and long VOTs in equal numbers, but the short and long variants are produced by different talkers (as in the Experimental conditions). Importantly, by Test 2, participants who were exposed to a teacher who produced long VOTs and those who were exposed to a teacher who produced short VOTs differed significantly in their preference for long versus short VOTs at test, demonstrating that the social contrast communicated through the simple images and teacher/student labels was sufficient to influence pronunciation preferences. These data suggest that in order to guide their beliefs about pronunciation quality, learners not only recruit knowledge of individual talkers' phonetic detail, but also knowledge of the social roles of those talkers.

4 Discussion

In this study we asked whether adult learners prefer the phonetics of a teacher's productions relative to those of a student peer in an artificial language classroom setting. The results suggest that they do. Participants showed a preference for the VOT durations that were associated with the teacher during the exposure phase; that is, a preference for long VOTs when the teacher had produced long VOTs, and a preference for short VOTs when the teacher had produced short VOTs. This finding suggests that adult listeners are sensitive to (at least some types of) sub-phonemic variation, and that they can use extra-linguistic social cues to guide the inferences they make about the phonetics of a target language. This finding raises questions about how such socially-guided use of speech input might arise. In other words, how do language learners determine whose speech to attend to as informative about linguistic structure?

The literature on language acquisition provides some clues, suggesting that judgements of social belonging and linguistic representativity may be conflated in even very young first language learners. Although numerous experiments have shown infants can be induced to engage in phonetic learning from audiovisual stimuli, infants do more readily learn phonetic information in an unfamiliar language when interacting with live persons (Conboy et al., 2015; Kuhl, 2007). Infant phonetic learning has also been shown to be influenced by the perception of variation in physical features (Weatherhead & White, 2018), and may be guided by beliefs regarding the informant's knowledgeability and social group membership (Tripp et al., 2021). Toddlers similarly have been shown to use physical features of talkers to predict variation in pronunciation, but prefer to use affiliative behavioral cues when they are available (Weatherhead & White, 2021). Numerous experiments reveal a preference among young language learners for talkers with familiar accents (Kinzler et al., 2009, 2012; Liberman et al., 2017). These findings are all consistent with a theory of speech perception in which both linguistic and social representations are simultaneously encoded (Sumner et al., 2013).

In adult Lx learning, beliefs about social belonging and knowledgeability may similarly influence beliefs about linguistic authority, leading learners to assumptions about whose talker-specific speech is most representative of the target language. Given that listeners demonstrated a preference for the VOT that had been associated with the teacher role during exposure, and that the teacher role was cued generically via iconographic image, the question is why the teacher's productions were preferred in this way. One possibility is that learners ascribe social power and domain expertise to the teacher role. They may additionally assume native-speakerness or native-likeness on the part of the typical language teacher.

A dominant ideology of preference for “native” or “native-like” input has led researchers to focus on acquisition of “native-like” perception and production as the target of Lx teaching. Defining successful Lx acquisition this way justifies characterizations of “non-native-accented” input as having a “negative effect” (Young-Scholten, 1995), as well as characterizations of variation in multilinguals’ productions as “mispronunciations” (e.g., Bosch & Ramon-Casas, 2011; Llompарт & Reinisch, 2021), as “less accurate” than native speakers’ productions (Hayes-Harb & Masuda, 2008) or “inauthentic” (Flege & Liu, 2001). This ideology has, of course, been challenged (e.g., by Falkert, 2016; Golombek & Jordan, 2005), often in the context of promoting intelligibility (as opposed to native-likeness) as a goal for Lx learners (e.g., Levis, 2005, 2020), though it is important to note that the intelligibility construct is also not devoid of social influence and bias (see, e.g., Babel & Russell, 2015; Ingvalson et al., 2017). Many scholars have further challenged the underlying native/non-native speaker dichotomy altogether (e.g., Cheng et al., 2021; Kabel, 2009; Moussu & Llurda, 2008), highlighting the ways in which racism and the racialization of “non-native” teachers of English underlie ideologies of preference for native-speaker language teachers (Ramjattan, 2019a, 2019b).

Evidence in fact demonstrates that language teachers for whom the target language is not their L1 do not differ from teachers who are L1 speakers of the target language in pronunciation teaching “effectiveness” (Levis et al., 2016), and indeed that some language learners express a preference for language teachers who share their L1 (Ma 2012). Nonetheless, language learners bring their own biases and expectations regarding pronunciation, which may include assumptions about the value of input from individuals with various social identities and roles such as talker gender, native or non-native speaker status (Ballard & Winke, 2017). To the extent that participants in the present study ascribe pronunciation authority to the teacher role, this may

explain their preference for the VOT durations produced by the teacher over those of a student peer.

The interpretation of this study's results is tempered by a number of methodological limitations, in particular having to do with ecological validity. First, the artificial classroom setting and the teacher and student roles were cued by means of iconographic images, limiting the generalizability of the present findings to real-world language learning settings. A further limitation is that the stimuli used here exhibited none of the within-talker variation in VOT that is typical of natural speech (here, VOT duration was fixed at either 25 ms or 125 ms). Indeed, some evidence from the perceptual learning literature suggests that listeners require and use within-talker variation to accommodate accented speech (Sumner, 2011). Sumner (2011) reported a shift in category boundary only for participants exposed to variable but not invariant VOTs during exposure. Perhaps most importantly, while we demonstrated an impact of the talkers' social roles on preference for specific talkers' phonetics, we have not directly addressed the question of whether such preferences actually guide Lx learning by adults through more ecologically valid communicative tasks with actual language learners. An additional concern here is that the observed effect sizes for the crucial comparisons were quite small (e.g., $\eta_p^2 = .092$ for test groups at Test 2), which may have resulted in part by difficulty perceiving the subtle VOT contrasts, potentially exacerbated by the online implementation of the study in various locations and varying equipment. Escudero and Kalashnikova (2020) found that infants are more attentive to phonetic detail when perceiving speech and learning words when the detail is relatively easy to perceive; it is possible that adults will show a similar effect with phonetic contrasts that are easier to perceive than the VOT contrasts used here.

5 Conclusion

We have presented a potentially promising method for examining the differential learning of speech input from speakers with different social roles. Building on this method, researchers could examine the effects of other, and more complex, cues to social position. For example, we provided participants with no (intentional) cues to the native or non-native speaker status of the “teacher”, though such status is likely to impact learners’ judgments of the value of the teacher’s input (see, e.g., Ballard & Winke, 2017; Moussu & Llorca, 2008). Relatedly, in displaying iconographic images instead of photographs of actual people, we attempted to avoid providing participants with cues to the students’ or teachers’ racialized identities, which may affect the perceived value of their speech input (Von Esch et al., 2020). Further, manipulating the types of interactions might impact learners’ assumptions of the value of different speakers’ input (e.g., corrective feedback from teacher to student may serve to reinforce the authority associated with the teacher role), or different communicative settings altogether (e.g., conversations among peers, between parents and children, etc.) are likely to affect learners’ perceived value of various input sources. These and other future manipulations will help Lx speech researchers and language teachers understand the ways in which language learners modulate their speech input in Lx learning.

References

- Allen, J. S., & Miller, J. L. (2004). Listener sensitivity to individual talker differences in voice-onset-time. *The Journal of the Acoustical Society of America*, 115(6), 3171–3183.
- Allen, J. S., Miller, J. L., & DeSteno, D. (2003). Individual talker differences in voice-onset-time. *The Journal of the Acoustical Society of America*, 113(1), 544–552.
- Babel, M., & Munson, B. (2014). Producing socially meaningful linguistic variation. *The Oxford*

- Handbook of Language Production*, 492, 308–325.
- Babel, M., & Russell, J. (2015). Expectations and speech intelligibility. *The Journal of the Acoustical Society of America*, 137(5), 2823–2833.
- Babel, M., Senior, B., & Bishop, S. (2019). Do social preferences matter in lexical retuning? *Laboratory Phonology*, 10(1). <https://doi.org/10.5334/labphon.133>
- Baese-Berk, M. (2018). Perceptual learning for native and non-native speech. *Psychology of Learning and Motivation*, 68, 1–29.
- Ballard, L., & Winke, P. (2017). The interplay of accent familiarity, comprehensibility, intelligibility, perceived native speaker status, and acceptability as a teacher. In T. Isaacs & P. Trofimovich (Eds.), *Second Language Pronunciation Assessment* (Vol. 107, pp. 121–140). Multilingual Matters / Channel View Publications.
- Banks, B., Gowen, E., Munro, K. J., & Adank, P. (2015). Cognitive predictors of perceptual adaptation to accented speech. *The Journal of the Acoustical Society of America*, 137(4), 2015–2024.
- Bates, D., Mächler, M., Bolker, B., & Walker, S. (2015). Fitting Linear Mixed-Effects Models Using lme4. *Journal of Statistical Software*, 67, 1–48.
- Boersma, P., & Weenink, D. (2020). *Praat: Doing phonetics by computer* (Version 6.1.37) [Computer software].
- Bosch, L., & Ramon-Casas, M. (2011). Variability in vowel production by bilingual speakers: Can input properties hinder the early stabilization of contrastive categories? *Journal of Phonetics*, 39(4), 514–526.
- Bradlow, A. R., & Bent, T. (2008). Perceptual adaptation to non-native speech. *Cognition*, 106(2), 707–729.
- Brown, K. M., Dahl, K. L., Cler, G. J., & Stepp, C. E. (2021). Listener age and gender diversity: Effects on voice-based perception of gender. *Journal of Voice: Official Journal of the Voice Foundation*, 35(5), 739–745.
- Cheng, L. S. P., Burgess, D., Vernooij, N., Solís-Barroso, C., McDermott, A., & Nambodiripad, S.

- (2021). The problematic concept of native speaker in psycholinguistics: Replacing vague and harmful terminology with inclusive and accurate measures. *Frontiers in Psychology*, 12(1), 715843.
- Clopper, C. G., & Bradlow, A. R. (2008). Perception of dialect variation in noise: intelligibility and classification. *Language and Speech*, 51(Pt 3), 175–198.
- Clopper, C. G., & Pisoni, D. B. (2004). Some acoustic cues for the perceptual categorization of American English regional dialects. *Journal of Phonetics*, 32(1), 111–140.
- Conboy, B. T., Brooks, R., Meltzoff, A. N., & Kuhl, P. K. (2015). Social interaction in infants' learning of second-language phonetics: An exploration of brain-behavior relations. *Developmental Neuropsychology*, 40(4), 216–229.
- Crosswhite, K., Antoniou, M., & Qi, Z. (n.d.). *Text grid maker [Praat script]*.
http://web.mit.edu/zqi/www/uploads/1/4/8/9/14891652/text_grid_maker.praat
- Dewaele, J.-M. (2017). Why the dichotomy “L1 versus LX user” is better than “native versus non-native speaker.” *Applied Linguistics*, 39(2), 236–240.
- D’Onofrio, A. (2018). Personae and phonetic detail in sociolinguistic signs. *Language in Society*, 47, 513–539.
- Drager, K. (2011). Speaker age and vowel perception. *Language and Speech*, 54(Pt 1), 99–121.
- Escudero, P., & Kalashnikova, M. (2020). Infants use phonetic detail in speech perception and word learning when detail is easy to perceive. *Journal of Experimental Child Psychology*, 190, 104714.
- Falkert, A. (2016). The relevance of accent in L2 pronunciation instruction: A matter of teaching cultures or language ideologies? *International Journal of Pedagogies and Learning*, 11(3), 259–270.
- Fecher, N., & Johnson, E. K. (2018). Effects of language experience and task demands on talker recognition by children and adults. *The Journal of the Acoustical Society of America*, 143(4), 2409.
- Fecher, N., & Johnson, E. K. (2021). Developmental improvements in talker recognition are specific to the native language. *Journal of Experimental Child Psychology*, 202, 104991.
- Flege, J. E., & Liu, S. (2001). The effect of experience on adults' acquisition of a second language. *Studies in Second Language Acquisition*, 23(4), 527–552.

- Flores, N., & Rosa, J. (2019). Bringing race into second language acquisition. *Modern Language Journal*, 103, 145–151.
- Foulkes, P., & Docherty, G. (2006). The social life of phonetics and phonology. *Journal of Phonetics*, 34(4), 409–438.
- Goggin, J. P., Thompson, C. P., Strube, G., & Simental, L. R. (1991). The role of language familiarity in voice identification. *Memory & Cognition*, 19(5), 448–458.
- Golombek, P., & Jordan, S. R. (2005). Becoming “black lambs” not “parrots”: A poststructuralist orientation to intelligibility and identity. *TESOL Quarterly*, 39(3), 513–533.
- Gonzales, K., Byers-Heinlein, K., & Lotto, A. J. (2019). How bilinguals perceive speech depends on which language they think they’re hearing. *Cognition*, 182, 318–330.
- Hayes-Harb, R. (2007). Lexical and statistical evidence in the acquisition of second language phonemes. *Second Language Research*, 23(1), 65–94.
- Hayes-Harb, R., & Masuda, K. (2008). Development of the ability to lexically encode novel second language phonemic contrasts. *Second Language Research*, 24(1), 5–33.
- Hay, J., & Drager, K. (2010). Stuffed toys and speech perception. *Linguistics and Philosophy*, 48(4). <https://doi.org/10.1515/ling.2010.027>
- Hay, J., Warren, P., & Drager, K. (2006). Factors influencing speech perception in the context of a merger-in-progress. *Journal of Phonetics*, 34(4), 458–484.
- Hothorn, T., Bretz, F., & Westfall, P. (2008). Simultaneous inference in general parametric models. *Biometrical Journal. Biometrische Zeitschrift*, 50(3), 346–363.
- Ingvalson, E. M., Lansford, K. L., Federova, V., & Fernandez, G. (2017). Listeners’ attitudes toward accented talkers uniquely predicts accented speech perception. *The Journal of the Acoustical Society of America*, 141(3), EL234.
- Johnson, K., Strand, E. A., & D’Imperio, M. (1999). Auditory–visual integration of talker gender in vowel perception. *Journal of Phonetics*, 27(4), 359–384.
- Kabel, A. (2009). Native-speakerism, stereotyping and the collusion of applied linguistics. *System*, 37(1),

12–22.

- Kang, O., & Rubin, D. (2014). Listener expectations, reverse linguistic stereotyping, and individual background factors in social judgments and oral performance assessment. In J. Levis Alene (Ed.), *Social Dynamics in Second Language Accent* (Vol. 10, p. 239). Walter de Gruyter.
- Kang, O., & Rubin, D. L. (2009). Reverse linguistic stereotyping: Measuring the effect of listener expectations on speech evaluation. *Journal of Language and Social Psychology*, 28(4), 441–456.
- Kinzler, K. D., Dupoux, E., & Spelke, E. S. (2012). “Native” objects and collaborators: Infants’ object choices and acts of giving reflect favor for native over foreign speakers. *Journal of Cognition and Development: Official Journal of the Cognitive Development Society*, 13(1), 67–81.
- Kinzler, K. D., Shutts, K., Dejesus, J., & Spelke, E. S. (2009). Accent trumps race in guiding children’s social preferences. *Social Cognition*, 27(4), 623–634.
- Kuhl, P. K. (2007). Is speech learning “gated” by the social brain? *Developmental Science*, 10(1), 110–120.
- Lennes, M. (2002). *Save labeled intervals to wav sound files [Praat Script]*.
https://github.com/FieldDB/Praat-Scripts/blob/a2444ca3dd96223d0677f2d6db49828f7739d24a/save_labeled_intervals_to_wav_sound_files.praat
- Lev-Ari, S., Ho, E., & Keysar, B. (2018). The unforeseen consequences of interacting with non-native speakers. *Topics in Cognitive Science*, 10(4), 835–849.
- Lev-Ari, S., & Keysar, B. (2010). Why don’t we believe non-native speakers? The influence of accent on credibility. *Journal of Experimental Social Psychology*, 46(6), 1093–1096.
- Levis, J. M. (2005). Changing contexts and shifting paradigms in pronunciation teaching. *TESOL Quarterly*, 39(3), 369–377.
- Levis, J. M. (2020). Revisiting the intelligibility and nativeness principles. *Journal of Second Language Pronunciation*, 6(3), 310–328.
- Levis, J. M., Sonsaat, S., Link, S., & Barriuso, T. A. (2016). Native and nonnative teachers of L2

- pronunciation: Effects on learner performance. *TESOL Quarterly*, 50(4), 894–931.
- Lieberman, Z., Woodward, A. L., & Kinzler, K. D. (2017). Preverbal infants infer third-party social relationships based on language. *Cognitive Science*, 41 Suppl 3, 622–634.
- Linck, J. A., & Cunnings, I. (2015). The utility and application of mixed-effects models in second language research. *Language Learning*, 65(S1), 185–207.
- Lisker, L., & Abramson, A. S. (1964). A cross-language study of voicing in initial stops: Acoustical measurements. *Word & World*, 20(3), 384–422.
- Lively, S. E., Logan, J. S., & Pisoni, D. B. (1993). Training Japanese listeners to identify English /r/ and /l/. II: The role of phonetic environment and talker variability in learning new perceptual categories. *The Journal of the Acoustical Society of America*, 94(3 Pt 1), 1242–1255.
- Lively, S. E., Pisoni, D. B., Yamada, R. A., Tohkura, Y., & Yamada, T. (1994). Training Japanese listeners to identify English /r/ and /l/. III. Long-term retention of new phonetic categories. *The Journal of the Acoustical Society of America*, 96(4), 2076–2087.
- Llompart, M., & Reinisch, E. (2021). Lexical representations can rapidly be updated in the early stages of second-language word learning. *Journal of Phonetics*, 88, 101080.
- Llurda, E. (2005). *Non-native language teachers: Perceptions, challenges and contributions to the profession*. Springer Science & Business Media.
- Maslowski, M., Meyer, A. S., & Bosker, H. R. (2019). How the tracking of habitual rate influences speech perception. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, 45(1), 128–138.
- Maye, J. (2000). *The acquisition of speech sound categories on the basis of distributional information* [PhD]. University of Arizona.
- McGowan, K. B. (2015). Social expectation improves speech perception in noise. *Language and Speech*, 58(Pt 4), 502–521.
- Morgan, S. D., & Ferguson, S. H. (2017). Judgments of emotion in clear and conversational speech by young adults with normal hearing and older adults with hearing impairment. *Journal of Speech*,

- Language, and Hearing Research: JSLHR*, 60(8), 2271–2280.
- Moussu, L., & Llurda, E. (2008). Non-native English-speaking English language teachers: History and research. *Language Teaching*, 41(3), 315–348.
- Muench, K. L., & Creel, S. C. (2013). Gradient phonological inconsistency affects vocabulary learning. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, 39(5), 1585–1600.
- Niedzielski, N. (1999). The effect of social information on the perception of sociolinguistic variables. *Journal of Language and Social Psychology*, 18(1), 62–85.
- Peirce, J., Gray, J. R., Simpson, S., MacAskill, M., Höchenberger, R., Sogo, H., Kastman, E., & Lindeløv, J. K. (2019). PsychoPy2: Experiments in behavior made easy. *Behavior Research Methods*, 51(1), 195–203.
- Perrachione, T. K., & Furbeck, K. T. (2019). Acoustic and linguistic factors affecting perceptual dissimilarity judgments of voices. *The Journal of the*.
<https://asa.scitation.org/doi/abs/10.1121/1.5126697>
- Perrachione, T. K., Lee, J., Ha, L. Y. Y., & Wong, P. C. M. (2011). Learning a novel phonological contrast depends on interactions between individual differences and training paradigm design. *The Journal of the Acoustical Society of America*, 130(1), 461–472.
- Ramjattan, V. A. (2019a). The white native speaker and inequality regimes in the private English language school. *Intercultural Education*, 30(2), 126–140.
- Ramjattan, V. A. (2019b). Racist nativist microaggressions and the professional resistance of racialized English language teachers in Toronto. *Race Ethnicity and Education*, 22(3), 374–390.
- Ramjattan, V. A. (2019c). Raciolinguistics and the aesthetic labourer. *Indian Journal of Industrial Relations*, 61(5), 726–738.
- Reinisch, E. (2016). Speaker-specific processing and local context information: The case of speaking rate. *Applied Psycholinguistics*, 37(6), 1397–1415.
- Rost, G. C., & McMurray, B. (2009). Speaker variability augments phonological processing in early word learning. *Developmental Science*, 12(2), 339–349.

- Rost, G. C., & McMurray, B. (2010). Finding the signal by adding noise: The role of noncontrastive phonetic variability in early word learning. *Infancy: The Official Journal of the International Society on Infant Studies*, 15(6). <https://doi.org/10.1111/j.1532-7078.2010.00033.x>
- Rubin, D. L. (1992). Nonlanguage factors affecting undergraduates' judgments of nonnative English-speaking teaching assistants. *Research in Higher Education*, 33(4), 511–531.
- Samuel, A. G., & Kraljic, T. (2009). Perceptual learning for speech. *Attention, Perception & Psychophysics*, 71(6), 1207–1218.
- Samuel, A. G., & Larraza, S. (2015). Does listening to non-native speech impair speech perception? *Journal of Memory and Language*, 81, 51–71.
- Scharinger, M., Monahan, P. J., & Idsardi, W. J. (2011). You had me at “Hello”: Rapid extraction of dialect information from spoken words. *NeuroImage*, 56(4), 2329–2338.
- Schertz, J. (2015). *Duration manipulation via PSOLA [Praat Script]*.
<http://individual.utoronto.ca/jschertz/scripts.shtml>
- Smiljanic, R. (2021). Clear speech perception. In *The Handbook of Speech Perception* (pp. 177–205). Wiley.
- Snodgrass, J. G., & Vanderwart, M. (1980). A standardized set of 260 pictures: Norms for name agreement, image agreement, familiarity, and visual complexity. *Journal of Experimental Psychology. Human Learning and Memory*, 6(2), 174–215.
- Sumner, M. (2011). The role of variation in the perception of accented speech. *Cognition*, 119(1), 131–136.
- Sumner, M., Kim, S. K., King, E., & McGowan, K. B. (2013). The socially weighted encoding of spoken words: a dual-route approach to speech perception. *Frontiers in Psychology*, 4, 1015.
- Szakay, A., Babel, M., & King, J. (2016). Social categories are shared across bilinguals' lexicons. *Journal of Phonetics*, 59, 92–109.
- Tripp, A., Feldman, N. H., & Idsardi, W. J. (2021). Social inference may guide early lexical learning. *Frontiers in Psychology*, 12, 1516.

- Tripp, A., & Munson, B. (2022). Perceiving gender while perceiving language: Integrating psycholinguistics and gender theory. *Wiley Interdisciplinary Reviews. Cognitive Science*, 13(2), e1583.
- Tyrone, M. E., & Mauk, C. E. (2012). Phonetic reduction and variation in American Sign Language: A quantitative study of sign lowering. *Laboratory Phonology*, 3(2), 425–453.
- Von Esch, K. S., Motha, S., & Kubota, R. (2020). Race and language teaching. *Language Teaching*, 53(4), 391–421.
- Weatherhead, D., & White, K. S. (2018). And then I saw her race: Race-based expectations affect infants' word processing. *Cognition*, 177, 87–97.
- Weatherhead, D., & White, K. S. (2021). Toddlers link social and speech variation during word learning. *Developmental Psychology*, 57(8), 1195–1209.
- Witteman, M. J., Weber, A., & McQueen, J. M. (2013). Foreign accent strength and listener familiarity with an accent codetermine speed of perceptual adaptation. *Attention, Perception & Psychophysics*, 75(3), 537–556.
- Xie, X., Theodore, R. M., & Myers, E. B. (2017). More than a boundary shift: Perceptual adaptation to foreign-accented speech reshapes the internal structure of phonetic categories. *Journal of Experimental Psychology. Human Perception and Performance*, 43(1), 206–217.
- Yoon, K. (n.d.). *Normalize intensity-dB [Praat script]*.
- Young-Scholten, M. (1995). The negative effects of “positive” evidence on L2 phonology. In Eubank, L., Selinker, L., & Sharwood, S. M. (Ed.), *Current State of Interlanguage: Studies in Honor of William E. Rutherford* (pp. 107–121). John Benjamins Publishing.
- Linck, J. A., & Cunnings, I. (2015). The utility and application of mixed-effects models in second language research. *Language Learning*, 65(S1), 185–207.