Journal of Experimental Psychology: Learning, Memory, and Cognition

Gradient Phonological Inconsistency Affects Vocabulary Learning

Kristin L. Muench and Sarah C. Creel

Gradient Phonological Inconsistency Affects Vocabulary Learning

Kristin L. Muench and Sarah C. Creel
University of California, San Diego

Learners frequently experience phonologically inconsistent input, such as exposure to multiple accents. Yet, little is known about the consequences of phonological inconsistency for language learning. The current study examines vocabulary acquisition with different degrees of phonological inconsistency, ranging from no inconsistency (e.g., both talkers call a picture /vig/) to mild but detectable inconsistency (e.g., one talker calls a picture a /vig/, and the other calls it a /vIg/), up to extreme inconsistency (e.g., the same picture is both a /vig/ and a /dIdʒ/). Previous studies suggest that learners readily extract consistent phonological patterns, given variable input. However, in Experiment 1, adults acquired phonologically inconsistent vocabularies more slowly than phonologically consistent ones. Experiment 2 examined whether word-form inconsistency alone, without phonological competition, was a source of learning difficulty. Even without phonological competition, listeners learned faster in 1 accent than in 2 accents, but they also learned faster in 2 accents (/vig/ = /vIg/) than with completely different labels (/vig/ = /dIdʒ/). Overall, results suggest that learners exposed to multiple accents may experience difficulty learning when 2 forms mismatch by more than 1 phonological feature, plus increased phonological competition due to a greater number of word forms. Implications for learning from variable input are discussed.

Keywords: word learning, accent processing, accented speech, learning under variability

A major challenge in language acquisition is identifying the relevant sources of variability in a complex signal. A listener hearing a novel language must rapidly determine which of numerous differences from sentence to sentence, or from word to word, indicate a change in meaning and which changes are irrelevant to meaning. Further, recent research indicates that high levels of variability can lead to language change (e.g., Hudson Kam & Newport, 2005, 2009; Smith & Wonnacott, 2010): Learners tend to regularize their input, making morphosyntactic forms more consistent than their input would dictate (see also Kirby, Cornish, & Smith, 2008).

However, most research on the role of variability in language acquisition has focused on morphosyntax and semantics, leaving open what happens in situations of phonological input variability. This is an important consideration, because most listeners experience some degree of phonological inconsistency within single word forms due to *accent variability*: Children in South Carolina

must learn that their parents mentioning a pen (/pIn/) and a newscaster mentioning a pen (/pɛn/) refer to the same object. Children whose parents are native Spanish speakers of English may hear their parents call a pen something closer to "pain" /peIn/ due to the different distributions of vowel sounds in English and Spanish (e.g., Fox, Flege, & Munro, 1995). Later in life, adult learners of a second language may encounter this mapping problem as they are forced to reconcile phonetically different words referring to the same object in the new language. These apparent sound changes may seem to learners to violate the basic principle that a particular word will always contain a single sequence of speech sounds. How, then, does phonological inconsistency affect vocabulary acquisition?

Our goal in the present study was to explore the effects of phonological inconsistency during vocabulary learning. We consider two possibilities. First, listeners may encode a single representation, discarding accent variability as irrelevant. Second, listeners may encode multiple forms of a word, essentially storing different lexical representations for different pronunciations of the same word (see Goldinger, 1996; Pierrehumbert, 2002, for related accounts). These representations might be stored simply as different words, or they might be encoded as belonging to a particular context rather than varying freely. In the next sections, we review evidence on processing of linguistic variability generally and phonological variability specifically. We then present two experiments in which listeners were asked to learn phonologically variable words.

Tuning Out Variability in Language

Researchers have explored the effects of variability as a force in language change (e.g., Hudson Kam & Newport, 2005, 2009; Smith & Wonnacott, 2010; Wonnacott, 2011; Wonnacott, New-

Kristin L. Muench and Sarah C. Creel, Department of Cognitive Science, University of California, San Diego.

Kristin L. Muench is now at the Department of Pediatrics, Emory University School of Medicine, and the Marcus Autism Center, Children's Health Care of Atlanta.

Correspondence concerning this article should be addressed to Sarah C. Creel, Department of Cognitive Science, University of California, San Diego, 9500 Gilman Drive, Mail Code 0515, La Jolla, CA 92093-0515. E-mail: creel@cogsci.ucsd.edu

port, & Tanenhaus, 2008) and as a mechanism of statistical language learning (Gómez, 2002; Maye, Werker, & Gerken, 2002). In both cases, high variability is thought to lead to detection or production of regular structures (for a more general approach, see, e.g., Nosofsky, 1986). For instance, learners regularize morphosyntactic input variability (Hudson Kam & Newport, 2005, 2009), tending to selectively produce the most frequent variant. Of course, for highly frequent forms, learners are more likely to maintain irregular (i.e., lexically specific) morphological items, such as the irregular-but-frequent English verb *run/ran* (Wonnacott, 2011; Wonnacott et al., 2008). Infant learners extract invariant structure in syntactic rules (Gómez, 2002) and phonological rules (Chambers, Onishi, & Fisher, 2003).

Of course, adults might be expected to be much less flexible in extracting invariant structures that contradict their existing language categories. Nonetheless, some studies suggest that even adult learners may be able to "hear through" variability in phonological forms. That is, adult learners may pick up on consistent phonological patterns across a set of word forms. Newport and Aslin (2004) found that adult listeners extracted recurring phonological patterns from phonologically varying input, recognizing repeated consonant-to-consonant patterns with vowels intervening or vowel-to-vowel patterns with consonants intervening (see also Bonatti, Peña, Nespor, & Mehler, 2005; Chambers, Onishi, & Fisher, 2010). Considering word learning directly, several studies suggest that adults learning novel words perceive similarity between words with the same syllable-initial consonants—specifically, they are more likely to confuse (Creel, Aslin, & Tanenhaus, 2006) or falsely recognize (Creel & Dahan, 2010) such words for each other. For instance, Creel and Dahan found that learners were more likely to falsely recognize "choob" /tʃub/ as the label for a "choop" /tʃup/ picture than they were to verify "joop" /dʒup/ as the label for "choop" /tʃup/. They also found effects of gradient similarity: Confusions were greater when the changed element differed by fewer phonological features (e.g., /dʒup/ and /tʃup/ were confused more often than /sæf/ and /bæf/). Creel et al. (2006) found that listeners more often confused words matched in syllable-initial consonants (/pibo/ and /pɑbu/) than those that matched in vowels (/pibo/ and /diko/; see Cutler, Sebastián-Gallés, Soler-Vilageliu, & van Ooijen, 2000, and van Ooijen, 1996, for related evidence of flexible vowel interpretation from a word reconstruction task). Although none of these studies directly tested the ability to map similar phonological forms to the same meaning, they suggest that even adult learners are sensitive to matching phonological structures despite noticeable phonological variability. This may be especially true when vowels are variable. If listeners can detect such phonological patterns, can they then map these patterns to meanings? Put another way, do listeners' confusions of highly phonologically similar words have a positive trade-off, allowing listeners to more easily map highly similar word forms to a single meaning?

One hint that learners might map partial phonological similarity to meaning comes from studies of phonological systematicity. Word learners appear to benefit when new words are similar in form to words they already know (Stamer & Vitevitch, 2012; Storkel, Armbrüster, & Hogan, 2006). However, this does not consider what happens when learners are acquiring similar words *simultaneously*. Simultaneous learning of similar words has been examined by Monaghan, Christiansen, and Fitneva (2011; see also

Monaghan, Mattock, & Walker, 2012), who tested the capacity of human learners and computational models to map phonologically similar or phonologically arbitrary sets of words to novel referents. They found that learners benefited from moderate phonological similarity within a word type (e.g., all action words ended in voiceless fricatives, and all nouns ended in velar stops). When phonological similarity between words in the same semantic category was too high, similarity created within-category confusion that made learning more difficult. Importantly for current purposes, when learning was measured in terms of mapping to the correct category—rather than the exact word-referent mapping—there were benefits for high phonological similarity. This suggests that listeners may have extracted phonologically systematic elements and mapped those to category-level semantic characteristics. If listeners can extract phonological patterns and map them to diffuse semantic categories, it is possible that they can extract phonological patterns from among accent variability and map those patterns to specific meanings.

Encoding Variability

Though ignoring (or downweighting) phonological variability is one potential solution to accent inconsistency, it has potentially negative consequences in a multilingual or multiaccent situation. In particular, listeners who actually ignored sound contrasts that existed in one phonological system but not the other would lose certain word distinctions. For instance, if a listener began to ignore the differences between /i/ and /I/ entirely, due to phonological inconsistency, then many existing words would become homophones (*beat/bit, feel/fill, deep/dip*). Learners might instead deal with phonological variability without loss of information by encoding multiple variants of the same form (see, e.g., Goldinger, 1996; Pierrehumbert, 2002; Ranbom & Connine, 2007)—one variant per accent. At least two lines of research suggest that this takes place. Sebastián-Gallés, Echeverria, and Bosch (2005) tested Catalan-native speakers who also knew Spanish in a Catalan-language lexical decision task. These listeners readily perceive the /e/ versus /ɛ/ distinction that is present in Catalan but not in Spanish, which has a single sound closer to /e/ (Sebastián-Gallés, Vera-Costán, Larsson, Costa, & Deco, 2009). In the Catalan lexical decision task, words containing /ɛ/ were sometimes mispronounced with /e/ (mimicking Catalan words with a Spanish accent). Catalan-native listeners often provided "word" responses to these Spanish-accent-like mispronunciations. However, they did not make "word" responses to the reverse mispronunciation (/e/ mispronounced as /ɛ/), which would be uncharacteristic of a Spanish accent. Sebastián-Gallés and colleagues interpreted this to mean that these listeners have stored additional lexical representations of Catalan words with Spanish accents, derived from experience with Spanish-accented speakers of Catalan. These results are consistent with encoding of multiple forms for phonologically inconsistent words.

Relatedly, Sumner and Samuel (2009) found evidence that bidialectal listeners stored both *r*-dropped and *r*-containing variants of words (e.g., both *sister* and *sistah*), with greater form priming within a dialect (prime *sister*, target *sister*) than between dialects (prime *sistah*, target *sister*). These two lines of research together suggest that listeners with experience to multiple phonological

systems store accent variants as additional word forms, in contrast to a wholesale tuning out of accent variability.

The Current Study

The research summarized above is equivocal in its predictions for learning under accent variability. Some studies suggest that listeners readily extract phonologically consistent elements when other phonological information is detectably inconsistent (Chambers et al., 2010; Newport & Aslin, 2004) and that listeners may even map phonological patterns to diffuse semantic characteristics (Monaghan et al., 2011, 2012). However, other results (Sebastián-Gallés et al., 2005, 2009; Sumner & Samuel, 2009) suggest that listeners must encode multiple form–meaning mappings in situations of accent inconsistency—which would presumably make learning slower. Effects of accent variability on vocabulary acquisition are as yet unknown. Thus, our major question is whether phonological inconsistency causes word learning difficulty and, if so, why.

Another outstanding question, and a secondary question in the current study, is whether certain groups of listeners are better at learning with variability. Fluent bilingual listeners seem adept at disregarding phonetic variability that is not relevant in the language at hand (e.g., Gonzales & Lotto, 2013; Quam & Creel, 2013; though see counterevidence from Cutler, Mehler, Norris, & Segui, 1992, who suggest that seemingly balanced French–English bilinguals only utilize the segmentation strategy of their dominant language). Bilinguals are also highly likely to experience accent variability by hearing each of their languages filtered through the phonology of the other by certain speakers. Finally, bilinguals may possess certain cognitive advantages for selective attention (e.g., Bialystok, 1999; Bialystok & Martin, 2004), and they have been shown to perform better than monolinguals in a vocabulary-learning task (Kaushanskaya & Marian, 2009). Do these characteristics of bilinguals give them an advantage in processing phonological inconsistency?

In the current pair of experiments, we taught listeners labels for 16 novel objects. For some listeners, two speakers always used the same label (e.g., *veeg* /vig/) for the same object. For other listeners, each speaker labeled the two objects with nonidentical but similar words: For instance, the male speaker called a picture *veeg* (/vig/), and the female speaker called it a *vig* (/vɪg/). Thus, the listener had to map two different labels to the same visual referent. By linking each accent to one of two highly discriminable speakers, we provided listeners with a potential context cue (talker identity) to which accent would be used. Experiment 1 examined differences in learning a one-accent vocabulary versus a two-accent vocabulary, where accents differed only in their vowels. Experiment 2 replicated and extended Experiment 1 by including multiple two-accent conditions of increasing accent separation, as well as a different-word vocabulary set to assess whether partial phonological overlap in the two-accent conditions was beneficial relative to the absence of phonological overlap. Experiment 2 also addressed whether the results of Experiment 1 were in part due to increased phonological competition.

Experiment 1

In this experiment, we asked two questions. First, is it more difficult to learn new words when they are not pronounced con-

sistently? To address this question, listeners learned in either a one-accent (consistent) condition or a two-accent (inconsistent) condition. Given previous research suggesting that vowel variability is particularly easy to ignore (e.g., Creel et al., 2006; Cutler et al., 2000; Nazzi, 2005; Nazzi & Bertoncini, 2009; Nazzi, Floccia, Moquet, & Butler, 2009; van Ooijen, 1996), accents differed only in their vowels. Our second question was whether there is a benefit for bilinguals (over monolinguals) in encoding phonologically variable words. To answer this second question, we tested English monolinguals and Spanish–English bilinguals. Bilinguals are accustomed to switching between speakers with different languages and phonological systems, and they have been shown to exceed monolinguals in vocabulary learning (Kaushanskaya & Marian, 2009). Spanish speakers in particular are accustomed to processing more broadly distributed vowels than in English, because Spanish contains only three front-vowel distinctions (vs. six in English; Fox et al., 1995). This might make it easier for them to ignore small gradations between the English vowels in our word set.

There are at least two possible outcomes. On the one hand, learning words in two accents may be relatively easy because listeners can extract recurring phonological patterns despite vowel inconsistency (e.g., Newport & Aslin, 2004), due to a general tendency to downweight vowel information (Creel et al., 2006; Cutler et al., 2000) and to the low predictability of vowels in the current experiment. If this is the case, the vowel-shift condition should show accuracy as high and learning speed as rapid as in the one-accent condition. On the other hand, learning words in two accents may be harder than learning words in one accent because listeners are essentially learning two labels for each object. If so, listeners should be slower to learn and less accurate in the vowel-shift condition than the one-accent condition. Outcomes may also differ as a function of language background: Spanish–English bilinguals may be globally superior to monolinguals (see Kaushanskaya & Marian, 2009). They may also find learning in two accents easier than do monolinguals, both because they are accustomed to learning multiple labels for the same referent across languages and because they may be better able to ignore English front vowel variability due to their experience with Spanish front-vowel categories (Fox et al., 1995).

Method

Participants. Undergraduates (mean age = 20.9 years, *SD* = 3.0, range = 18–33 years; 6 ages not reported) who were either monolingual English speakers (*n* = 30) or bilingual Spanish–English speakers (*n* = 32) were recruited from the University of California, San Diego (UCSD) psychology participant pool and were compensated with course credit. An additional 16 participants were run but were not included in the final analysis due to failure to complete the experiment in the allotted 2-hr time slot (5), reporting native competence in languages besides Spanish and English (5), equipment problems (1), leaving early (1), extremely low test accuracy (1 outlier, who scored 66% correct; remaining participants' mean was 98%, *SD* = 2.6%, range = 87%–100%), or missing or unclear language background data (3). Monolingual participants used English almost exclusively, though some reported limited experience learning other languages, including Spanish, at ages 8 or later (*M* = 14.0, *SD* = 3.2). Those with Spanish experience (*n* = 11) reported Spanish fluency with a mean

of 1.67 ($SD = 0.87$), on a scale of 1 to 5, and Spanish use of 0–2 hr per week ($M = 0.20$, $SD = 0.63$). Bilingual participants reported fluency in both Spanish and English, with average reported age of Spanish acquisition at 0.8 years ($SD = 1.1$, range = 0–4) and English acquisition at 4.6 years ($SD = 3.6$, range = 0–13). For most bilingual participants (22 of 32, or 68.8%) we also obtained scores on Gollan, Weissberger, Runnqvist, Montoya, and Cera's (2012) Multilingual Naming Task (MINT). This picture-naming task assesses relative language dominance in English or Spanish by asking listeners to name a set of 68 pictures first in one language and then in the other. The dominance score is then calculated by subtracting the number of pictures correctly named in English from the number of pictures correctly named in Spanish (versions are being developed for other language pairs). Negative scores indicate English dominance. The bilingual participants scored −11.68 on average ($SD = 8.86$), which was significantly English dominant, $t(21) = 6.18$, $p < .0001$. Eighteen participants (82%) showed negative scores (English dominant), two (9%) showed scores of 0 (balanced), and two (9%) showed slightly positive scores (1 and 5, Spanish dominant).

We do not have a measure of monolingual language knowledge, but the preliminary norms for the MINT (Gollan et al., 2012) suggest that monolingual English speakers perform near ceiling (mean accuracy 95%). Gollan et al.'s self-described balanced bilingual Spanish–English speakers averaged a bit lower, 90.8%, on English and averaged 75.2% on Spanish. Our bilinguals, similarly, averaged 89.3% ($SD = 5.2$%) on English and 72.1% ($SD = 10.2$%) on Spanish. Note that the MINT seems to suggest greater English dominance than participants' self-reports, though, as Gollan et al. remarked, self-reports may take into account factors other than vocabulary. Nonetheless, our MINT results suggest that our bilinguals may score modestly lower on English vocabulary than our monolinguals would.

Stimuli. Auditory stimuli were recorded in a soundproof recording room in the Center for Research in Language at UCSD by a male and a female native English speaker who had both lived continuously in Southern California since the age of 9.

Auditory stimuli consisted of 16 consonant–vowel–consonant words (see Table 1) in two different artificial accents. The accent used was a *front vowel shift*, where English vowels produced with the tongue toward the front of the mouth (/i/, /ɪ/, /ɛ/, /æ/) were shifted "downward" in the mouth toward the next highest vowel (/ɪ/, /ɛ/, /æ/, /ɑ/). For instance, the word "dev" (/dɛv/) in one accent would be "dav" (/dæv/) in the other accent. Maye, Aslin, and Tanenhaus (2008; see also Houde & Jordan, 1998) previously applied this consistent front vowel shift with real English words to produce an accent. In the current study, unlike Maye et al., neither accent was familiar to listeners because the words were not familiar. That is, listeners did not know whether one talker was an accented variant of the other talker or vice versa. Words were designed such that certain pronunciations matched the initial consonant–vowel sequence of other words in the opposite accent (see Table 2). Participants in the one-accent condition were exposed only to nouns pronounced in one accent (A or B); participants in the vowel-shift condition were trained and tested on words in which one talker referred to objects in a different accent than the other talker. The accent difference chosen was not one that most listeners would have experienced in Southern California, though the mapping between the two accents bears some similarity to the relationship between U.S. English and accents found in New Zealand (Watson, Maclagan, & Harrington, 2000) and between Standard U.S. English and Northern California accents (e.g., Eckert, 2008). If anything, previous experience with such accents should make listeners' task easier, diluting accent-inconsistency effects.

The task required participants to learn the above-described novel names as labels for 16 unfamiliar object images. These images, originally used by Creel et al. (2006), consisted of a set of two-dimensional black-and-white nonsense shapes.

Procedure. Participants completed a questionnaire about their linguistic background immediately prior to participating in the experiment. Participants who described proficiencies in languages other than Spanish or English in this section were excluded from the analyzed data set.

The experiment was administered in a low-noise room containing only the testing equipment. Participants sat in a comfortable chair approximately two feet from a computer screen on which visual stimuli appeared. Sounds were heard through high-quality

Table 1
Stimuli Used in Experiments 1 and 2

One accent		Two-accent/vowel shift		Different vowels		Different codas		Different words	
dɪdʒ	dɪdʒ	dɪdʒ	dɪdʒ	dɪdʒ	dɛdʒ	dɪdʒ	dɪv	dɪdʒ	vɪg
dɪv	dɪv	dɪv	dɛv	dɪv	dɪv	dɪv	dɪdʒ	dɪv	vɛf
dɛg	dɛg	dɛg	dæg	dɛg	dæg	dɛg	dɛv	dɛg	væd
dæz	dæz	dæz	dɑz	Dæz	dɪz	dæz	dæg	dæz	vɑb
gɪb	gɪb	gɪb	gɪb	Gɪb	gɛb	gɪb	gɪd	gɪb	zɪv
gɪd	gɪd	gɪd	gɛd	gɪd	gɪd	gɪdʒ	gɪb	gɪd	zɛb
gɛz	gɛz	gɛz	gæz	gɛz	gæz	gɛz	gɛd	gɛz	zæs
gædʒ	gædʒ	gædʒ	gɑdʒ	gædʒ	gɪdʒ	gædʒ	gæz	gædʒ	zɑf
vɪg	vɪg	vɪg	vɪg	vɪg	vɛg	vɪg	vɪf	vɪg	dɪdʒ
vɪf	vɪf	vɪf	vɛf	vɪf	vɪf	vɪf	vɪg	vɪf	dɛv
vɛd	vɛd	vɛd	væd	vɛd	væd	vɛd	vɛf	vɛd	dæg
væb	væb	væb	vɑb	væb	vɪb	væb	væd	væb	dɑz
zɪv	zɪv	zɪv	zɪv	zɪv	zɛv	zɪv	zɪb	zɪv	gɪb
zɪb	zɪb	zɪb	zɛb	zɪb	zɪb	zɪb	zɪv	zɪb	gɛd
zɛs	zɛs	zɛs	zæs	zɛs	zæs	zɛs	zɛb	zɛs	gæz
zæf	zæf	zæf	zɑf	zæf	zɪf	zæf	zæs	zæf	gɑdʒ

Table 2

Examples of Noun Names for the Onscreen Target and Competitor Images for All Testing Trial Types in the Different-Onset Target and Same-Onset Target Categories

Target category	Onscreen target		Onscreen competitor	
	Accent A	Accent B	Accent A	Accent B
Different onset	vɪg	vɪg	dɛg	dæg
	vɪg	vɪg	zɪb	zɛb
Same onset	vɪg	vɪg	vɛd	væd
	vɪg	vɪg*	vɪf*	vɛf

Note. Participants in the one-accent condition learned accent A or accent B; participants in the two-accent condition learned both accent A and accent B. Asterisks denote cross-accent phonological competitors.

Sennheiser HD 280 Pro headphones. Participants were allowed to adjust the volume to a comfortable level. Participants interacted with the screen both with a standard keyboard and a mouse. Experimental stimuli were presented using MATLAB with the PsychToolBox 3 (Brainard, 1997; Pelli, 1997).

The main experiment presented novel vocabulary items in both a *training phase* and a *testing phase*. Each participant heard a particular speaker with a single, consistent accent throughout both phases (in the vowel-shift case, the speakers were heard using different accents than each other). In each training trial, two pictures appeared, and 500 milliseconds (ms) later, a phrase "Choose the X" (where X is a novel word) was played. The phrase "Choose the," which did not contain any front vowels, was pronounced identically across all speakers and conditions. Thus, no accent-specific information was present until the object label was spoken. The participant selected one of the two images as the referent by mouse-clicking on that image. To provide accuracy feedback, only the correct image stayed on screen. Training phase targets were paired with competitors that did not share an initial consonant (*different-onset*; 75% of trials) or shared an initial consonant (*same-onset*; 25%). Of the latter type, half of trials pictured two objects that had a consonant–vowel match between accents, such as hearing "Choose the /vɪg/" and seeing a /vɪg/-**/vɪg/** object and a **/vɪf/**-/vɛf/ object depicted (see Table 2). Therefore, if participants were unable to ignore vowel differences across accents, those exposed to two accented labels for a single object were more likely to confuse labels across accents in these trials due to overlap in two segments (/vɪ/), and those who heard only one accent used by the two speakers would not experience as much confusion due to overlap in only one segment (/v/). Participants progressed to the test phase after reaching a 90% accuracy threshold across a training block (128 trials).

In the testing block (192 trials), participants again heard a word on each trial and selected a referent shape, but no feedback was given. Test trials contained pairs of pictures with different-onset (50% of trials) or same-onset labels (50% of trials). For a subset of same-onset trials in the vowel-shift condition, the competitor object's alternate-accent label used the same initial consonant–vowel as the target object's label, such as hearing "Choose the /vɪg/" and seeing a /vɪg/-**/vɪg/** object and a **/vɪf/**-/vɛf/ object depicted.

During training and testing, each object occurred equally often in each screen position as target and as nontarget. Thus, listeners could not use visual layout cues to learn the correct responses.

Across participants, there were three different picture-to-word assignments, three trial orders (quasi-random with the constraint that no two adjacent trials had the same target picture), and four accent assignments (both accent A; both accent B; male A, female B; female A, male B). These factors were approximately evenly distributed across the different conditions. Approximately equal numbers of monolinguals and bilinguals learned one accent (16 monolinguals, 15 bilinguals) or two accents (14 monolinguals, 17 bilinguals).

Results

Learning rate. Preliminary analyses suggested no effects or interactions of different picture-to-word assignments, talker-to-accent-mappings, or quasi-random orders, so analyses were collapsed over these factors. Learning was faster for one-accent training than for vowel-shift training (see Figure 1), but effects of language background were not evident. An analysis of variance (ANOVA) was conducted on learning rate (number of blocks needed to reach criterion) with condition (single-accent, two-vowel accent) and language (monolingual, bilingual) as between-participants variables. Effect sizes are reported as generalized eta-squared, which equates for differences in between- versus within-participants designs (Bakeman, 2005; Olejnik & Algina, 2003; η_G^2 is identical to partial η^2 for between-participants designs). An effect of condition, $F1(1, 58) = 10.14$, $p = .002$; $\eta_G^2 = .15$, indicated faster learning in the single-accent condition ($M = 2.13$ blocks, $SD = 0.56$) than in the vowel-shift condition ($M = 2.94$ blocks, $SD = 1.29$). There was no effect of language, $F1(1, 58) = 0.05$, $p = .83$; $\eta_G^2 = .00$, nor was there a Language × Condition interaction, $F1(1, 60) = 0.57$, $p = .46$; $\eta_G^2 = .01$.

Test accuracy. Figure 2 shows accuracy on test trials by language background, learned accent(s), and trial type. ANOVAs were conducted on log-odds of test trial accuracy with pair type (different-onset words, same-onset words) as a within-participants factor and language and condition as between-participants factors. All factors were repeated measures within items. Prior to calculation of log odds, scores of 0 and 1.0 were transformed as suggested in Macmillan and Creelman (1991), by adding $1/(2N)$ to scores of 0 and subtracting $1/(2N)$ from scores of 1.0 (where N is the number of items tested). There was an effect of condition, $F1(1, 58) =$


Figure 1. Experiment 1, learning rate (blocks to reach criterion, with standard errors) by language background and condition. Mono. = monolingual; Biling. = bilingual.

Figure 2. Experiment 1, test accuracy with standard errors. mono. = monolingual; biling. = bilingual. Error bars indicate standard errors.

17.12, $p = .0001$; $F2(1, 31) = 61.5$, $p < .0001$; $\eta_G^2 = .16$, reflecting that accuracy was higher overall for the one-accent learners than the vowel-shift learners. An effect of pair type, $F1(1, 58) = 62.43$, $p < .0001$; $F2(1, 31) = 115.84$, $p < .0001$; $\eta_G^2 = .27$, resulted from lower accuracy on same-onset test trials than different-onset test trials. Both effects were qualified by an interaction of Condition $\times$ Pair Type, $F1(1, 58) = 8.03$, $p = .006$; $F2(1, 31) = 15.65$, $p = .0004$; $\eta_G^2 = .05$. This interaction appeared to result from a robust one-accent condition advantage for same-onset trials, $F1(1, 60) = 20.72$, $p < .0001$; $F2(1, 31) = 55.09$, $p < .0001$; $\eta_G^2 = .26$, and a smaller, marginal advantage on different-onset trials, $F1(1, 60) = 3.69$, $p = .06$; $F2(1, 31) = 7.24$, $p = .01$; $\eta_G^2 = .06$. There was no effect of language, nor interactions involving the language factor (all Fs < 1, except for items Language $\times$ Condition, $F2(1, 31) = 1.10$, $p = .30$; and items Language $\times$ Condition $\times$ Trial Type, $F2(1, 31) = 2.62$, $p = .12$).

A large portion of this lower accuracy may be ascribed to trials where the competitor object's label (as spoken by the other talker) overlapped substantially with the target label, such as hearing /dɪv/ and seeing a /div/ - /dɪv/ and a /dɛdʒ/ - **/dɪdʒ/**. To verify this, we compared shared-vowel same-onset test trials (see Table 2; hearing /dɪv/ with a /dɛdʒ/ - **/dɪdʒ/** on screen) in the vowel-shift condition, different-vowel same-onset trials (hearing /dɪv/ and seeing a /dɛg/-/dæg/ on screen) in the vowel-shift condition, and same-onset trials in the one-accent condition (which were all different vowel). Because some items had only different-vowel same-onset trials and not shared-vowel ones, items were not treated as a repeated measure when the comparison involved shared-vowel same-onset items. Analyses on accuracy log odds showed that the one-accent same-onset trials (.981 ± .019) were more accurate than different-vowel same-onset trials in the vowel-shift condition (.955 ± .043; $t1(60) = 3.21$, $p = .002$; $t2(31) = 3.47$, $p = .002$). However, both were more accurate than shared-vowel same-onset trials in the vowel-shift condition (.891 ± .093; vs. one-accent: $t1(60) = 8.00$, $p < .0001$; Welch's $t2(35.55) = 7.27$, $p < .0001$; vs. vowel-shift different-vowel: paired $t1(30) = 6.58$, $p < .0001$; Welch's $t2(46.80) = 3.88$, $p = .0003$).

Discussion

We asked whether learners would be affected by phonological inconsistency in a word-learning task. Accent inconsistency

slowed word learning, with equivalent results for monolingual and bilingual listeners. Greater difficulty for vowel-shift learning is consistent with accounts that bidialectal listeners encode multiple forms of the same word (Sebastián-Gallés et al., 2009; Sumner & Samuel, 2009), because presumably it is more time consuming to encode two forms than to encode one form. It is somewhat surprising, though, given previous evidence that listeners readily extract invariant phonological structure (Chambers et al., 2010; Newport & Aslin, 2004). It is particularly interesting in that listeners are thought to be more accepting of vowel changes than consonant changes (Cutler et al., 2000; van Ooijen, 1996)—that is, different vowels should have been easy to ignore. This suggests that phonological inconsistency does slow vocabulary acquisition.

We also failed to find an effect of bilingual status, counter to Kaushanskaya and Marian's (2009) findings of a bilingual advantage in a word-learning task. In particular, bilinguals in our task did not learn faster than monolinguals, nor did they show a selective advantage for learning the vowel-shift vocabulary. There are several differences between the studies that might explain these divergent findings. One possibility is that our bilingual populations are somewhat different: Kaushanskaya and Marian's bilinguals were native English speakers who learned either Spanish or Mandarin as a very early second language (on average, age 5 [Spanish] or 2 [Mandarin]); our bilinguals had learned Spanish as a first language, concurrent with or followed by English. A second consideration is that Kaushanskaya and Marian asked listeners to associate novel words with familiar English "translations," whereas we asked listeners to associate novel words with unfamiliar shapes. A third consideration, returned to in the General Discussion, is that the words we used were phonotactically and phonetically Englishlike. Kaushanskaya and Marian do not provide examples of their novel words, but their descriptions suggest that words were consonant–vowel–consonant–vowel disyllables. Consonant–vowel syllable structure is much more common in other languages (including Spanish and Mandarin, the languages of Kaushanskaya and Marian's bilinguals) than it is in English. These cues to "Englishness" in our stimuli may have selectively activated the English phonological knowledge of bilingual listeners, causing them to encode the two accents as different words rather than collapse over vowel variation.

The account we propose based on Experiment 1's results is that within-word phonological variability slowed learning because listeners had to encode two phonological forms for each referent and because they experienced cross-accent phonological competition. However, an obvious counterexplanation is that the difficulty lay not in experiencing inconsistent pronunciations of words but in learning twice as many word forms. For example, memorizing a 32-word list would be more taxing than memorizing a 16-word list, apart from the additional task of mapping forms to meanings. This is a somewhat different explanation than saying that the difficulty stems from inconsistency in the particular words that get mapped to the same referent—that /vig/ and /vɪg/ are discernibly different. If the latter explanation is true, listeners should be gradiently affected by the degree of phonological overlap between the two labels for an object.

Further, Experiment 1 does not distinguish between difficulty due to within-word phonological variability and difficulty due to cross-accent phonological competition. Perhaps listeners were otherwise good at ignoring variability but were slowed during learning by phonological competition between pairs of pictured items, which was weaker in the one-accent condition. That is, listeners in the vowel-

shift condition were forced to distinguish between extremely similar items (like /vIg/ in one accent and /vIf/ in the other) on training trials, but listeners in the one-accent condition were not. Analyses of first-training-block accuracy bear this out: There was no difference in log-odds accuracy on different-onset trials for one-accent versus vowel-shift learners (.774 vs. .746; $t1(60) = 1.02$, $p = .31$; $t2(31) = 2.27$, $p = .03$) but a large difference on same-onset trials (.742 vs. .653, $t1(60) = 2.97$, $p = .004$; $t2(31) = 3.24$, $p = .003$), driven particularly by vowel-shift trials containing cross-accent competitors ($M = .615$; $t1(60) = 4.26$, $p < .0001$; Welch's $t2(37.89) = 4.20$, $p = .0002$). Having to distinguish similar-sounding words on training trials may have slowed learning on its own, rather than or in addition to actual difficulty encoding phonologically inconsistent labels. The roles of phonological competition versus phonological inconsistency could be tested by removing phonological competition but retaining the phonologically inconsistent labels.

A final issue not resolved by the current experiment is the role of talker-specific context. In the vowel-shift condition, each accent was produced by a single talker. This means that, during the test, talker identity could be used as a cue to which word would be spoken or which accent would be used. We did not test what happened when this pattern was broken. Thus, we do not know the extent to which listeners benefited from consistent talker–word or talker–accent relationships to recognize words.

The next experiment was a modified replication of Experiment 1, which addressed the issues we identified above: the effects of pho-nological competition, whether listeners are learning different variants versus learning different words, and the role of talker specificity in learning. Listeners again learned labels for pictures, but in addition to learning in one versus two accents, some listeners learned completely different word pairs for the same referent. To examine effects of phonological inconsistency alone, without effects of phonological competition, all trials presented pictures with labels that were phono-logically distant from one another. Finally, test trials included both original-talker trials and switched-talker trials, to assess whether lis-teners were encoding words talker-specifically.

Experiment 2

In this experiment, listeners learned labels for objects in one of five conditions. Two conditions were endpoints of label similarity: the one-accent stimuli from Experiment 1, with identical labels for the same picture (/vig/ = /vig/), and a different-word condition designed to maximize the distance between the two labels for a referent (/vig/ = /dIdʒ/). The remaining conditions were all two-accent con-ditions, constituting intermediate levels of label similarity: *vowel-shift* (/didʒ/ = /dIdʒ/, as in Experiment 1); *different vowels* created by a nonsystematic mapping of vowels to each other (/dIdʒ/ = /dɛdʒ/; see Table 1); and different word-final consonants (i.e., *different codas*; /dIdʒ/ = /div/; see Table 1). The words used in the different-word condition were identical to those in the vowel-shifted two-accent condition; importantly, similar words mapped to the same picture in the vowel-shift condition (/didʒ/ = /dIdʒ/ and /vig/ = /vIg/), and similar words mapped to different pictures in the different-word condition (/didʒ/ = /vIg/ and /vig/ = /dIdʒ/; see Table 1). Note that the one-accent and vowel-shift conditions mimicked Experiment 1.

To eliminate effects of phonological competition on learning speed, all training trials in all conditions presented pictures with labels that were phonologically distant from one another. In particular, all train-

ing trials presented pictures whose labels did not begin with the same consonant in any accent (see Table 3). Finally, to assess whether listeners were depending to some extent on talker-specific word knowledge, we included trials during the test phase in which each talker switched to using the other talker's accent.

We asked three questions. First, is it simply more difficult to learn 32 words than to learn 16 words, or is it specifically difficult to encode within-word variability? If results in Experiment 1 simply reflected difficulty encoding more word forms, then learning rate in the current experiment should be identical between two-accent conditions and the different-word condition, which should both take longer than the one-accent condition. However, if listeners specifically had difficulty encoding within-word variability, the number of blocks needed to reach the accuracy criterion should scale with the dissimilarity of the two labels applied to the same picture. Second, we asked whether phonological inconsistency effects in Experiment 1 were due to the presence of phonological competitors on training trials. If so, learning rate in the current experiment should now be more similar between one-accent and all two-accent conditions. Third, if listeners depended on talker identity to recognize words, accuracy should be higher on test trials where the talker used his or her own accent than on trials where the talker switched to the other talker's accent. If listeners did not depend on talker identity, accuracy should not differ between original-talker trials and switched-talker trials.

Method

Participants. $N = 120$ new participants (mean age = 21.1 years; $SD = 3.1$; range = 18–44 years) were recruited from the UCSD Human Participant Pool and received course credit for participation. Given the null effect of language background in Experiment 1, participation was not restricted by language, but language background data were obtained from all participants. Three additional participants were replaced due to missing ques-tionnaire data (2) or confusion in how the experiment worked (1).

As expected, most participants (104 of 120) reported English as the language they felt most comfortable with. Language back-ground data are summarized in Table 4. Twenty-six participants (21.7%) reported no knowledge of another language (4; 3% of total participants) or reported beginning language instruction after age 10 (22; 18% of total participants). Of individuals reporting fluency in multiple languages (94 of 120, or 73%), mean reported age of

Table 3
Example Target and Competitor Trials in Experiment 2

Condition	Target label	Distractor label
One accent	didʒ	gib
	væb	gId
Shifted vowel	didʒ/dIdʒ	gib/gIb
	væb/vɑb	gId/gɛd
Different vowel	didʒ/dɛdʒ	gib/gɛb
	væb/vib	gId/gid
Different coda	didʒ/div	gib/giv
	væb/væd	gIdʒ/gIb
Different word	didʒ/vIg	gib/zIv
	væb/dɑz	gId/zɛb

Note. Each target picture was paired only with other pictures that did not share an initial consonant.

Table 4
Language Background Data, Experiment 2

Background	N	English		Other language		BDS score	Choose English
		Started learning	Full-time school	Started learning	Full-time school		
Group							
Monolingual	26	1.2 (0.9)	15.1 (3.8)	13.4 (1.5)[a]	−32.8 (0.7)	1.5 (2)	1.00
English-dominant	31	2.7 (1.2)	15.4 (2.6)	2.7 (2.7)	−19.9 (3.8)	0.4 (1.4)	0.97
Moderate English-dominant	30	3.8 (1.5)	15.5 (1.9)	1.9 (1.3)	−11.2 (2.5)	1.9 (4)	1.00
Balanced	33	7.7 (4.2)	10.9 (3.5)	1.3 (0.9)	2.9 (7.1)	5.1 (4)	0.61
Language	% of N = 94[b]						
Spanish	20.4%						
Korean	19.4%						
Cantonese	10.2%						
Mandarin	8.2%						
Chinese (unspecified)	5.1%						
Japanese	5.1%						
Vietnamese	5.1%						
Arabic	4.1%						

Note. BDS = Bilingual Dominance Survey.
[a] Began schooling in a language other than English; some numbers reflect once-daily classes. Four participants in this group reported no exposure to a foreign language.　[b] Two participants each spoke Russian, Tagalog, or Chiu Chow/Teochew Chinese. One participant (1%) each spoke Armenian, Bosnian, Burmese, Farsi, German, Gujarati, Hindi, Hungarian, Indonesian, Italian, Portuguese, Punjabi, Slovak, Tamil, Telugu, or Urdu.

English acquisition was 4.8 years (SD = 3.5; range = 0–13 years, except for one participant who learned English at 22, which was 22 years before experimental participation). All but 10 of them (8.3% of total participants) were exposed to English before age 10. Language dominance was calculated with Dunn and Fox Tree's (2009) Bilingual Dominance Survey (BDS). Bilinguals tended to be English dominant, with a bilingual-dominance score (English dominant = negative) of −9.1 (SD = 10.7), which was significantly English biased, $t(93)$ = 8.22, p < .0001. A positive BDS score reflects equal or greater dominance in the individual's other language; only 19 individuals (20.2% of bilinguals) scored at 0 or greater. Thus, the sample included a number of bilinguals, but, as in Experiment 1, most were mildly to strongly dominant in English, and most had been exposed to English fairly early in life.

Stimuli. All spoken-word stimuli were previously recorded for Experiment 1 but were presented in different arrangements (see Table 1). All conditions shared most words with the other conditions, but a few words occurred only in some conditions in order to create full sets of stimuli. Pictures were the same as Experiment 1.

Design. The conditions we tested can be divided into three groups: one accent (as in Experiment 1); two accents (three different versions, described below); and different words (described below). The single-accent vocabulary was identical to that in Experiment 1. The first two-accent condition—the *vowel shift* condition—was also identical to the vowel-shift condition in Experiment 1. The remaining conditions were new to Experiment 2. The *different vowel* condition also contained word pairs that differed only in their vowels but with greater phonological distance between labels (1.5 phonological features on average) and less systematicity—vowels were exchanged with each other haphazardly rather than undergoing a relative shift. In the *different coda* two-accent condition, words differed only by their final (coda) consonants, again with greater phonological distance between paired labels relative to the vowel shift condition (2 features). The different-vowel and different-coda vocabularies were constructed from a superset of the stimuli in Experiment 1; a few words from the original two-accent condition were replaced in order to create a

balanced set of word pairs. The final learning condition was a *different-word* set: Each picture was labeled by two dissimilar words. Importantly, this set was constructed by rearranging the stimuli used in the vowel-shift condition. Therefore, if the main difficulty of learning in two accents is simply encoding 32 word forms instead of 16, this condition should appear identical to the vowel-shift condition. However, if a large part of the difficulty rests in encoding nonidentical labels for the same object, this condition should be harder than the vowel-shift condition. There were three word-to-picture assignments. In combination with five conditions, two talker-to-accent assignments, and two test block orders, the result was 60 unique conditions, each of which was run twice.

Procedure. Participants signed consent paperwork and sat in a sound-treated room with a computer running MATLAB. MATLAB first presented two dialog boxes. One asked participants to report basic demographic data, and the other asked participants questions about their language background. These included the questions from Dunn and Fox Tree's (2009) BDS, an instrument designed to assess a bilingual's relative dominance in one language or another.[1] We changed from using the MINT to using the BDS because the BDS is more language general: Although the MINT (Gollan et al., 2012) has to be adjusted for each pair of languages tested, the BDS is designed to assess dominance in any two languages, allowing use of a broader subset of the participant pool at our language-diverse university. Previous work in our lab has shown strong correlations between the MINT and the BDS (Bregman & Creel, 2012; Quam & Creel, 2013), though note that the BDS seems to

---

[1] Note that it is possible that presenting the language background survey before the experiment itself primed participants to attend to language differences. One might suggest that heightened attention to language differences results in increased attention to the different labels in the experiment. However, it is not clear how attention to language differences across all participants would generate the different performance patterns in different experimental conditions.

estimate participants in our population as slightly less English dominant than the MINT does. Participants then proceeded to the main experiment.

As in Experiment 1, participants saw pairs of pictures and heard a spoken label, and they were asked to guess which picture was labeled. Unlike in Experiment 1, picture pairs appearing in the same training trials had labels phonologically dissimilar from each other in both accents (see Figure 3 and Table 3), so that listeners did not experience strong phonological competition during training. In fact, listeners in the one-accent and two-accent conditions could have performed perfectly in the experiment by associating a single initial consonant with each picture, as each picture occurred only with other pictures whose labels began with a different consonant (or consonant*s*, in the different-word condition).

Training trials continued in 64-trial blocks until participants reached 90% accuracy in a block, or until they completed 12 blocks. They then proceeded to two 64-trial blocks of nonreinforced test trials, which were identical in content to the training trials with one exception: One of the two test blocks (order counterbalanced across participants) presented the learned words with each talker using the other talker's accent, whereas the other block presented each talker using his or her original accent. Following the two test blocks was a 64-trial discrimination posttest. Each trial of the posttest presented two printed words in pseudo-phonetic English spelling (e.g., VEEG and VIG) corresponding to spoken words used in the experiment. On each trial, the participant heard one spoken label (/vig/) and was asked to select the printed word that matched it (in this case, VEEG). This posttest was designed to assess listeners' ability to identify the sounds that differed between the two labels for an object. For the two-accent conditions, paired words were always the two words that had referred to the same object. For the one-accent and different-word conditions, the pairs were the same as those used in the vowel-shift condition.

Results

Learning rate. Participants learned fastest in the single-accent condition and slowest in the different-word condition, with the two-accent conditions falling in between (see Figure 4a). An


Figure 3. Experiments 1 and 2. Gray bars: average phonological-feature distance between the spoken word on a training trial (e.g., /vig/) and the other picture's label (in two-accent cases, the closer of the two labels). White bars: phonological-feature distance between pairs of words that were accent equivalents of each other (e.g., /vig/ and /vIg/). Error bars indicate standard errors. diff. = different.


Figure 4. Experiment 2, learning rates in each condition, (a) overall and (b) with a median split on Bilingual Dominance Scale scores. Note that blocks in Experiment 2 were half as many trials as in Experiment 1. Error bars indicate standard errors. diff. = different; vwl = vowel.

analysis of covariance (ANCOVA) on learning rate (in blocks) was conducted with condition (one-accent, vowel-shift, different-vowel, different-coda, different-word) as a between-participants factor and BDS score as a continuous factor (mean-centered). BDS was significant, $F1(1, 110) = 6.62$, $p = .01$; $\eta_G^2 = .06$, reflecting increased learning time as dominance in a language other than English increased. That is, less English-dominant listeners took more time to reach criterion, rather than showing a bilingual advantage. An effect of condition, $F1(4, 110) = 17.95$, $p < .0001$; $\eta_G^2 = .39$, reflected different learning rates across conditions. Individually, Tukey's honestly significant difference (HSD) tests indicated that the different-word condition ($M = 9.79$, $SD = 2.72$) was significantly slower than all other conditions; in addition, the different-coda condition ($M = 6.63$, $SD = 3.56$) was significantly slower than the single-accent condition ($M = 3.71$, $SD = 2.26$) and the vowel-shift condition ($M = 4.04$, $SD = 3.71$). Unlike in Experiment 1, the single-accent and vowel-shift conditions did not differ. The interaction of BDS × Condition was marginal, $F1(4, 110) = 2.24$, $p = .07$; $\eta_G^2 = .08$. It appeared to stem from faster learning by more English-dominant participants in the different-

vowel condition ($r(22) = .57$, $p = .004$; see Figure 4b). It is unclear why this condition would differ as a function of English dominance. This result should be interpreted with caution due to the marginal nature of the interaction, but it is in the opposite direction of a bilingual advantage in word learning. Nonetheless, both sets of participants experienced increasing difficulty as an object's labels became less similar.

Accuracy in first block of training trials. Inspection of performance across blocks suggested that participants in the different-word condition were becoming frustrated or confused as the task continued, with some individuals showing random accuracy in later blocks. Although this may reflect the particular difficulty of this condition—in line with our hypothesis—it raises the question of whether the learning-rate analysis was contaminated by a high rate of participant mental fatigue. To obtain an additional measure of learning difficulty, we measured accuracy only in the first block of training trials (see Figure 5a). We computed an ANCOVA on log-odds-transformed accuracy, with condition and BDS score (participants analysis only) as factors.[2] The results were roughly similar to the learning rate analysis: Condition was significant, $F1(4, 110) = 9.23$,

$p < .0001$; $F2(4, 155) = 22.92$, $p < .0001$; $\eta_G^2 = .25$). There was neither an effect of BDS, $F1(1, 110) = 1.05$, $p = .31$; $\eta_G^2 = .01$ (see Figure 5b) nor a BDS × Condition interaction, $F1(4, 110) = 0.54$, $p = .70$; $\eta_G^2 = .02$. Tukey HSD tests indicated that the single-accent accuracy ($M = .701$, $SD = .104$) differed from different-vowel ($M = .624$, $SD = .080$), different-coda ($M = .616$, $SD = .075$), and different-word ($M = .556$, $SD = .071$) accuracy and that vowel-shift accuracy ($M = .661$, $SD = .098$) differed from different-word accuracy (these were also significant by items Tukey tests). As in the learning rate data, the single-accent and vowel-shift conditions did not differ; this was a different outcome than that observed in Experiment 1. These results mirror the learning-rate analyses in suggesting that increasing phonological distance between the two labels for a picture results in slower learning.

Test accuracy. To gauge whether listeners were depending at all on talker identity to recognize words, we compared accuracy on original-talker test trials to accuracy on switched-talker test trials (see Figure 6). Accuracy varied across conditions, which is unsurprising for two reasons. First, conditions with faster learning showed more rapid improvement during training, meaning that they would tend to reach a higher accuracy level than criterion by the end of a training block. Second, accuracy varied across conditions because some participants reached the test without reaching criterion accuracy. Overall, accuracy for original-talker trials was .908 ($SD = .148$) and that for switched-talker trials was .905 ($SD = .137$). An ANOVA was conducted on log-odds-transformed accuracy with BDS as a continuous factor (mean-centered; participants analysis only), condition and test order (original trials first, switched trials first) as between-participants factors, and talker (original, switched) as a within-participants factor. An effect of condition, $F1(4, 100) = 8.49$, $p < .0001$; $F2(4, 155) = 75.31$, $p < .0001$; $\eta_G^2 = .23$, suggested differences in accuracy depending on learning condition. An effect of BDS, $F1(1, 100) = 7.46$, $p = .007$; $\eta_G^2 = .06$, suggested that the increase in accuracy with increasing English dominance was statistically significant. The effect of talker was significant by participants, $F1(1, 110) = 4.37$, $p = .04$; $F2(1, 155) = 1.15$, $p = .28$; $\eta_G^2 = .00$, as was the Condition × Talker interaction, $F1(1, 110) = 3.10$, $p = .02$; $F2(4, 155) = 2.36$, $p = .06$; $\eta_G^2 = .01$). Individual paired t tests indicated that the interaction resulted from a significant difference (after Bonferroni correction) only for the vowel-shift group, $t(23) = 2.88$, $p = .009$; $t(31) = 2.49$, $p = .02$, with better performance in the same-talker condition (.979) than in the changed-talker condition (.960).

Posttest. Accuracy on the identification posttest was high (see Table 5), reflecting the general English dominance of our sample: More English-dominant listeners were better at matching a spoken word (/vig/) with its correct orthographic form (VEEG, not VIG). Log-odds-transformed posttest accuracy also correlated significantly with BDS scores ($r = .32$, $p < .0001$), such that lower English dominance resulted in lower posttest accuracy. This is consistent with less accurate perception of English speech sounds by less English-dominant speakers, though it may also reflect weaker knowledge of complex grapheme–phoneme correspondence rules in English.

Figure 5. Experiment 2, accuracy in the first block of training trials, (a) overall and (b) with a median split on Bilingual Dominance Scale scores. Error bars indicate standard errors. diff. = different; vwl = vowel.

---

[2] The items ANOVA did not contain the BDS score, as this was identical for items within a condition except for the two different accents used in the one-accent condition.

Discussion

Experiment 2 tested learning of words where labels for a particular picture were phonologically consistent, phonologically inconsistent, or extremely different. As in Experiment 1, participants were slower to learn words with phonological inconsistency than to learn phonologically consistent ones. However, the patterns were somewhat different than in Experiment 1. First, Experiment 2 repeated the one-accent and vowel-shift conditions in Experiment 1 but with phonological competition trials removed. Unlike in Experiment 1, the one-accent and vowel-shift conditions were equivalently rapidly learned. This suggests that the main driver of difficulty in the vowel-shift case in Experiment 1, with a subtle, single-feature vowel shift, was phonological competition rather than nonidentical word forms. One might suggest that the different samples across the two experiments generated this change; however, there were no effects of language background in these two conditions in either Experiment 1 or Experiment 2, implying that language background contributed relatively little to these effects. Further, the language-background difference that did appear in Experiment 2 was that more bilingual participants had greater difficulty learning similar labels for the same referent. Thus, a greater proportion of bilinguals in Experiment 2 should, if anything, heighten two-label learning difficulty relative to Experiment 1, not diminish it.

Experiment 2 further expands on Experiment 1 by suggesting that difficulty in mapping two forms to one referent is gradient: As the distance between two forms increased, listeners took longer to learn the mappings, even though phonological competition was minimized. Additionally, the sheer number of word forms to be learned (32 in the two-accent and two-word conditions, 16 in the single-accent condition) appeared to be less a factor than similarity within a label pair: The slowest learning occurred in the different-word condition, which employed exactly the same word set as that in the vowel-shift condition.

Interestingly, there was limited evidence of talker-specific storage: Except in the vowel-shift condition, listeners were equally accurate on test trials where talkers switched to the other talker's accent as when

Table 5
Posttest Accuracy in Each Condition

Condition	Pairs tested	Accuracy (SD)
One accent	Vowel shift	.940 (.080)
Vowel shift	Vowel shift	.941 (.061)
Different vowel	Different vowel	.970 (.040)
Different coda	Different coda	.961 (.058)
Different word	Vowel shift	.944 (.070)

they used their own accent. Previous work has demonstrated talker-specificity effects on word recognition (Creel, Aslin, & Tanenhaus, 2008; Creel & Tumlin, 2011; Goldinger, 1996, 1998; Palmeri, Goldinger, & Pisoni, 1993). However, in Creel and colleagues' previous word-learning studies showing talker-specificity effects, effects were seen most on trials where words were phonologically similar and when the dependent measure was visual fixations. In the current experiment, only accuracy was measured, and the contrasted words were phonologically distinct. Thus, the lack of a talker effect in most conditions may reflect that experimental conditions here were not conducive to detecting talker-specificity effects. Alternately, the effect in the vowel-shift condition might be taken as adaptation to talker-specific speech patterns (see, e.g., Eisner & McQueen, 2005; Kraljic & Samuel, 2005). That is, learners are able to slightly adjust their phoneme representations to each talker's accent, meaning that the most accent-like condition (vowel shift) is the one in which talker-specificity effects surfaced.

As in Experiment 1, we did not find a bilingual facilitation effect in any condition. In fact, the only effect of lower English dominance in our sample was a negative one. This raises questions about the role of bilingualism in vocabulary acquisition: Why would bilinguals experience a negative impact of phonological inconsistency if they are more accustomed to it? We return to this in the General Discussion.

One might also ask how much our results were affected by listeners' interpretations of the two artificial accents they heard. Participants in Experiment 2 completed a post-experiment questionnaire (8 of 120, or 7%, of participants did not complete the questionnaire due to a programming error). Participants were asked, among other things, what they thought the purpose of the experiment was; whether they noticed that pictures had two names; and whether they noticed anything interesting about the voices. Most listeners (80% or greater in each condition, except for the single-accent condition;[3] see data in Table 6) reported noticing multiple names for multiple pictures. Interestingly, many listeners in the two-accent conditions made reference to different accents or different pronunciations or commented on the similar names for the same picture (close-vowel: 70%; two-vowel: 50%; two-coda: 48%). Some participants, when asked if they had noticed two names for one picture, even explicitly wrote that "I noticed them being pronounced differently, but I did not hear a completely


Figure 6. Experiment 2, test accuracy as a function of condition and talker match. Error bars indicate standard errors. diff. = different; vwl = vowel.

---

[3] Single-accent learners often illusorily reported hearing multiple names for the same picture, though at lower rates than other groups. We verified that this did not actually occur in their exposure. These illusory multiple-name reports are similar to other studies we have run where listeners report having heard the same word in multiple voices, when they actually had not (illusory conjunctions; Creel & Tumlin, 2011).

Table 6

Post-Experiment Questionnaire Data Reflecting Awareness of Cross-Accent Similarities, Experiment 2

Condition	N	Noticed 2 forms	Referred to accent	Noted different pronunciation	Noted label similarity	Accent, pronunciation, or similarity
One accent	23	65%	22%	30%	4%	52%
Vowel shift	23	83%	43%	30%	9%	70%
Two vowel	22	91%	18%	23%	23%	50%
Two coda	21	90%	10%	38%	10%	48%
Two word	23	100%	4%	4%	0%	9%

different word for the same picture." (This was counted as a "noticed different names for same shape" response.) Another participant wrote that the two voices "pronounced the same names a little differently." These data suggest that listeners in the two-accent conditions did parse the two forms of the same word as related. However, listeners in the two-word condition rarely (2 of 23, or 9%) described the words as differing in pronunciation or accent.

Mentions of label similarity did not correspond to strong performance differences among the two-accent learners. There were no differences in learning rate or in first-block accuracy. There was a significant difference in test accuracy ($p = .02$), which disappeared ($p = .40$) when listeners who never reached criterion were removed (2 [5%] of those mentioning similarity, 6 [20%] of those not mentioning similarity). The most that can be said is that listeners who demonstrated awareness of label similarity were numerically more likely to reach the accuracy criterion.

Note that the vowel-shift condition showed the highest proportion of mentions of name similarity. Based on this, it is plausible that listeners in the two-accent (vowel-shift) condition in Experiment 1 also processed the two labels for each picture as differences in accent or pronunciation, rather than as wholly different word forms. Of course, to the extent that listeners did parse the two labels for a picture as wholly different word forms in Experiment 2, they appeared to benefit from phonological overlap nonetheless.

General Discussion

At the outset, we asked whether word learning is affected by phonological inconsistency. In both Experiment 1 and Experiment 2, listeners were faster to learn and were more accurate when they learned the words in a single accent than when they learned words simultaneously in two accents. Listeners were also faster to learn in two accents than when they learned in two "languages" (two very different labels for each word; Experiment 2). This implies that, even for listeners who have mature speech-sound representations, listeners capitalize on the gradient phonological similarity between the two forms that map to a referent. In fact, when phonological competition is eliminated as a source of difficulty in learning (Experiment 2), listeners in the highest similarity two-accent condition—the vowel-shift condition—learned as quickly as single-accent learners. This suggests that perceptible but very small-magnitude inconsistencies in word forms—here, a single phonological feature of the vowel—may not generate much difficulty when one is learning

words in multiple accents, apart from increases in phonological competition.

We had also asked whether bilinguals, who are more accustomed than monolinguals to processing multiple phonological systems, might have an advantage in learning under accent variability. Counter to expectations, bilingual status did not confer a notable advantage in either experiment and was slightly disadvantageous in Experiment 2. This contrasts with Kaushanskaya and Marian's (2009) earlier findings of a bilingual advantage in vocabulary learning.

The novel contribution of this simulation of multiaccent learning is that high within-category variability may slow the learning process, due both to increased phonological competition and to within-word inconsistency. Although numerous studies have demonstrated that learning new minimal-pair vocabulary items (e.g., *bih* and *dih*) is more difficult than learning dissimilar items (e.g., Creel et al., 2006; Stager & Werker, 1997), this study is the first to our knowledge to examine the reverse case (though see Mattock, Polka, Rvachew, & Krehm, 2010): how learners perform when highly similar but distinguishable word forms must be regarded as the same. In combination with previous research, our findings suggest that similar word forms are easier to map to the same referent. This means both that highly similar forms are difficult to map to different referents (Creel et al., 2006; Stager & Werker, 1997) and that highly similar forms are relatively easier to map to the same referent, provided that phonological competition is minimized (current study). Note that highly similar does not need to mean identical: Even if two forms do not correspond perfectly, they nonetheless facilitate form–meaning mapping more than less similar forms (see also Monaghan et al., 2011, 2012).

This work has implications for language acquisition as well. Accounts of infants' progress toward language-specific speech perception commonly invoke the notion that attention to within-category variability is deleterious to word learning (e.g., Werker & Tees, 1984); yet, very few studies have explored whether this is empirically true. We show that adults are indeed impacted by "knowing too much"—perceiving within-category variability. Of course, it is possible that phonological variability affects the adult learners we consider here more than it would affect young language learners. Young children have less entrenched native phonological systems (e.g., Oh et al., 2011), which might make it easier for them to merge across native contrasts. On the other hand, research from our lab (Creel, in press) suggests that monolingual preschool-aged children in a simplified word-learning task also have greater difficulty learning words with phonological inconsis-

tency than without. In the following sections, we discuss the implications of these results for learning under accent inconsistency and why bilingual listeners do not exceed monolinguals on our vocabulary learning task.

Extracting Regularities or Encoding Variability?

We initially posited two ways in which learners might deal with phonological inconsistency when learning new words in multiple accents. On the one hand, they might ignore the variable elements and extract the consistent ones; on the other, they might store each form as a separate, distinct word. Our two experiments suggest that neither of these accounts is completely correct. Listeners do not seem to extract regularities so much as they seem to be clustering across learned elements. If they were just extracting regularities, the similarity between inconsistent elements—for instance, how phonologically or acoustically close to each other the two alternating vowels were—should not matter. However, Experiment 2 suggested that learners were sensitive to the degree of dissimilarity of the inconsistent elements. This implies that learners are encoding both forms, but as they encode both forms, they can cluster them gradiently, with the overall phonological distance between the two labels affecting ease of clustering. As the two forms become farther apart in phonological space, clustering becomes difficult (or impossible). This suggestion of clustering fits more broadly with a class of speech- and word-recognition theories that suggest that learners include rich acoustic-phonetic detail in their speech representations rather than discarding it (Creel et al., 2008; Creel & Tumlin, 2011; Goldinger, 1996). This rich detail then allows coherent mappings to emerge. The data are also consistent with accounts of more abstracted lexical organization in which listeners do not store every detail but do store more than one form for a word (see Ranbom & Connine, 2007, who provide evidence that the frequency of phonologically reduced variants of particular words influences recognition). In any event, it is particularly interesting that multiword mappings can emerge in adults, who readily distinguish the label pairs used. Nonetheless, this presents a consistent mirror image pattern to adults' difficulty mapping gradiently phonologically similar forms to different meanings (e.g., Creel & Dahan, 2010): Words that are harder to map to different meanings are easier to map to the same meaning, even if they are discernibly nonidentical.

These results also fit with Monaghan et al.'s (2011, 2012) studies on systematicity of form–meaning mappings. Monaghan et al. found that systematic phonological–semantic mappings make semantic category identification easier. Our results, Experiment 2 particularly, could be construed as phonological–semantic mappings easing semantic category identification, except that in our studies, each semantic category was a single object. When the labels for each "semantic category" (picture) were phonologically similar to each other, learning was easier than when the labels were phonologically arbitrary. Our results suggest that these regularities occur on the scale of single words.

Partial phonological similarity has implications for acquisition of two languages with high numbers of cognates. In fact, Albareda-Castellot, Pons, and Sebastián-Gallés (2011), who were studying bilingual language acquisition, argued that two-accent learning could be considered an extreme case of two-language learning with high cognate overlap. Our results suggest that as two cognate forms are closer and closer together, word–meaning mapping becomes easier and easier. Thus, the more cognates two languages have and the more phonologically similar those cognates are, the more the two languages will facilitate each other. Although language acquisition is typically thought of in dichotomous terms—one language, two languages—it might be construed more productively as a continuum from fully consistent single-language input on one end to input in two highly dissimilar languages on the other end. As Albareda-Castellot et al. pointed out, two-dialect acquisition and acquisition of two languages with high proportions of cognates might represent more intermediate points along that continuum. This suggests that research on bilingual acquisition and research on bidialectal acquisition might benefit each other.

Bilingualism and Word Learning

One somewhat surprising outcome (see Kaushanskaya & Marian, 2009) was the lack of a bilingual advantage in word learning. Previous studies have ascribed advantages to bilingual listeners in vocabulary learning due to better ability to inhibit competing representations (Kaushanskaya & Marian, 2009) or better phonological knowledge (Mattock et al., 2010; Quam & Creel, 2013). Why were bilinguals not at an advantage here? One possibility is that our bilinguals were too English dominant to benefit fully from a cognitive advantage in word learning, though Kaushanskaya and Marian reported a similarly English-dominant sample. A second possibility is that our bilinguals and monolinguals differed on other, unmeasured characteristics that militated against a bilingual advantage. These unmeasured characteristics might include generally better working memory in more monolingual participants. They might also include greater familiarity with multiple-choice testing among listeners with longer lengths of residence in the United States (i.e., the English monolinguals and English-dominant bilinguals vs. less English-dominant bilinguals). A third possibility is that, because the words we used were phonetically and phonotactically American English-like (unlike Kaushanskaya and Marian's), our words might activate English-based phonological knowledge (see Gonzales & Lotto, 2013; Quam & Creel, 2013) and lead listeners to encode English-relevant sound characteristics. This would confer an advantage to listeners with more English phonological knowledge.

Finally, the word-learning task of Kaushanskaya and Marian (2009) was different than our word-learning task: They asked participants to map phonologically consistent novel words to English translations, whereas we asked participants to map phonologically inconsistent novel words to pictures. Interestingly, Kaushanskaya and Marian suggest that their bilingual benefit may stem from greater ability to suppress other English translations and that novel-referent mapping might yield a different outcome, foreshadowing our current set of results. It is up to future research to discern whether our results differ from Kaushanskaya and Marian's due to task differences, population differences, or subtle phonetic and phonological cues in the words themselves.

Alternative Interpretations of the Current Findings

Although we have discussed our results as indicating patterns of learning in two accents, one might question whether listeners were

processing the variability between speakers as something other than accent variation. For instance, they might have regarded the situation as one where one speaker consistently produces speech errors relative to the other speaker, or they might have regarded one speaker as reliable, and the other as unreliable. We believe that listeners' processing of the variability as speech error rather than accent is actually consistent with our account of accent processing: If a particular speaker habitually produces a set of speech errors, one can regard it as an idiolect (essentially an accent that is limited to one person). The parallels between idiolect adaptation and accent adaptation have been noted by multiple researchers (e.g., Eisner & McQueen, 2005; Kraljic & Samuel, 2005). That is, processes that allow a listener to process idiolectal variability should also allow a listener to process accent variability. However, we do not think that our listeners regarded one speaker as reliable and one as unreliable or untrustworthy. The "unreliability" account makes a specific prediction: that listeners will (at least initially during learning) show chance performance for one speaker but above chance performance for the other. This would lead to a negative accuracy correlation between the two voices, as people who are more accurate on one voice (above chance) should be less accurate (chance) on the other and vice versa. However, this is not the case in our data: Considering just the first block of training trials in each experiment, accuracy in the male voice and the female voice were positively correlated in all two-accent conditions. This suggests that listeners were not tuning out one of the speakers as unreliable but were encoding word forms from both speakers.

Still, there might be differences in performance as a function of the degree of listeners' awareness of accentedness. One difference between the current study and some others (e.g., Maye et al., 2008; White & Aslin, 2011) is that we did not preadapt listeners to an accent; listeners had to infer accent differences—if they did infer them—based on their slowly increasing word knowledge. This leaves open a particularly interesting question: Would preadaptation to two accents alleviate difficulty in two-accent learning? That is, would listeners learn more rapidly from dual-accent input if they had already stored the relationships between the two accents? If so, this would be consistent with the hypothesis that listeners encode not only dual forms but the contexts in which they occur and the relationships between them. Alternately, would they choose to ignore or downweight input from the talker who has a nonstandard accent, which would presumably slow down learning? This remains an interesting question for future research.

Conclusion

The current study demonstrated that phonological inconsistency slowed vocabulary learning in adults. However, there was a relative learning advantage for nonidentical but similar word forms over phonologically distant word forms. For forms that were inconsistent in only a single feature, the major difficulty in learning seemed to be increased phonological competition rather than inconsistency. Results imply that phonological inconsistency, such as that encountered while learning words in different accents, may negatively impact vocabulary acquisition. How such effects play out over the longer term, during child language acquisition and during diachronic language change, remains to be seen.

References

Albareda-Castellot, B., Pons, F., & Sebastián-Gallés, N. (2011). The acquisition of phonetic categories in bilingual infants: New data from an anticipatory eye movement paradigm. *Developmental Science, 14,* 395–401.

Bakeman, R. (2005). Recommended effect size statistics for repeated measures designs. *Behavior Research Methods, 37,* 379–384. doi: 10.3758/BF03192707

Bialystok, E. (1999). Cognitive complexity and attentional control in the bilingual mind. *Child Development, 70,* 636–644. doi:10.1111/1467-8624.00046

Bialystok, E., & Martin, M. M. (2004). Attention and inhibition in bilingual children: Evidence from the dimensional change card sort task. *Developmental Science, 7,* 325–339. doi:10.1111/j.1467-7687.2004.00351.x

Bonatti, L. L., Peña, M., Nespor, M., & Mehler, J. (2005). Linguistic constraints on statistical computations. *Psychological Science, 16,* 451–459.

Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision, 10,* 433–436. doi:10.1163/156856897X00357

Bregman, M. R., & Creel, S. C. (2012). Learning to recognize unfamiliar voices: The role of language familiarity and music experience. In N. Miyake, D. Peebles, & R. P. Cooper (Eds.), *Proceedings of the 34th Annual Conference of the Cognitive Science Society* (pp. 144–149). Austin, TX: Cognitive Science Society.

Chambers, K. E., Onishi, K. H., & Fisher, C. (2003). Infants learn phonotactic regularities from brief auditory experience. *Cognition, 87,* B69–B77. doi:10.1016/s0010-0277(02)00233-0

Chambers, K. E., Onishi, K. H., & Fisher, C. (2010). A vowel is a vowel: Generalizing newly learned phonotactic constraints to new contexts. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 36,* 821–828. doi:10.1037/a0018991

Creel, S. C. (in press). Impossible to ignore: Word-form inconsistency slows preschool children's word learning. *Language Learning and Development.*

Creel, S. C., Aslin, R. N., & Tanenhaus, M. K. (2006). Acquiring an artificial lexicon: Segment type and order information in early lexical entries. *Journal of Memory and Language, 54,* 1–19. doi:10.1016/j.jml.2005.09.003

Creel, S. C., Aslin, R. N., & Tanenhaus, M. K. (2008). Heeding the voice of experience: The role of talker variation in lexical access. *Cognition, 106,* 633–664. doi:10.1016/j.cognition.2007.03.013

Creel, S. C., & Dahan, D. (2010). The effect of the temporal structure of spoken words on paired-associate learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 36,* 110–122. doi:10.1037/a0017527

Creel, S. C., & Tumlin, M. A. (2011). On-line acoustic and semantic interpretation of talker information. *Journal of Memory and Language, 65,* 264–285. doi:10.1016/j.jml.2011.06.005

Cutler, A., Mehler, J., Norris, D., & Segui, J. (1992). The monolingual nature of speech segmentation by bilinguals. *Cognitive Psychology, 24,* 381–410. doi:10.1016/0010-0285(92)90012-Q

Cutler, A., Sebastián-Gallés, N., Soler-Vilageliu, O., & van Ooijen, B. (2000). Constraints of vowels and consonants on lexical selection: Cross-linguistic comparisons. *Memory & Cognition, 28,* 746–755. doi: 10.3758/BF03198409

Dunn, A. L., & Fox Tree, J. E. (2009). A quick, gradient Bilingual Dominance Scale. *Bilingualism: Language and Cognition, 12,* 273–289. doi:10.1017/S1366728909990113

Eckert, P. (2008). Where do ethnolects stop? *International Journal of Bilingualism, 12,* 25–42. doi:10.1177/13670069080120010301

Eisner, F., & McQueen, J. M. (2005). The specificity of perceptual learning in speech processing. *Perception & Psychophysics, 67,* 224–238. doi: 10.3758/BF03206487

Fox, R. A., Flege, J. E., & Munro, M. J. (1995). The perception of English and Spanish vowels by native English and Spanish listeners: A multidimensional scaling analysis. *Journal of the Acoustical Society of America, 97,* 2540–2551. doi:10.1121/1.411974

Goldinger, S. D. (1996). Words and voices: Episodic traces in spoken word identification and recognition memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 22,* 1166–1183. doi:10.1037/0278-7393.22.5.1166

Goldinger, S. D. (1998). Echoes of echoes? An episodic theory of lexical access. *Psychological Review, 105,* 251–279. doi:10.1037/0033-295X.105.2.251

Gollan, T. H., Weissberger, G. H., Runnqvist, E., Montoya, R. I., & Cera, C. M. (2012). Self-ratings of spoken language dominance: A Multilingual Naming Test (MINT) and preliminary norms for young and aging Spanish–English bilinguals. *Bilingualism: Language and Cognition, 15,* 594–615. doi:10.1017/S1366728911000332

Gómez, R. L. (2002). Variability and detection of invariant structure. *Psychological Science, 13,* 431–436. doi:10.1111/1467-9280.00476

Gonzales, K., & Lotto, A. J. A (2013). bafri, *un* pafri: *Bilinguals' pseudoword identifications support language-specific phonetic systems.* Manuscript submitted for publication.

Houde, J. F., & Jordan, M. I. (1998, February 20). Sensorimotor adaptation in speech production. *Science, 279,* 1213–1216. doi:10.1126/science.279.5354.1213

Hudson Kam, C. L., & Newport, E. L. (2005). Regularizing unpredictable variation: The roles of adult and child learners in language formation and change. *Language Learning and Development, 1,* 151–195. doi:10.1080/15475441.2005.9684215

Hudson Kam, C. L., & Newport, E. L. (2009). Getting it right by getting it wrong: When learners change languages. *Cognitive Psychology, 59,* 30–66. doi:10.1016/j.cogpsych.2009.01.001

Kaushanskaya, M., & Marian, V. (2009). The bilingual advantage in novel word learning. *Psychonomic Bulletin & Review, 16,* 705–710. doi:10.3758/PBR.16.4.705

Kirby, S., Cornish, H., & Smith, K. (2008). Cumulative cultural evolution in the laboratory: An experimental approach to the origins of structure in human language. *Proceedings of the National Academy of Sciences, USA, 105,* 10681–10686. doi:10.1073/pnas.0707835105

Kraljic, T., & Samuel, A. G. (2005). Perceptual learning for speech: Is there a return to normal? *Cognitive Psychology, 51,* 141–178. doi:10.1016/j.cogpsych.2005.05.001

Macmillan, N. A., & Creelman, C. D. (1991). *Detection theory: A user's guide.* New York, NY: Cambridge University Press.

Mattock, K., Polka, L., Rvachew, S., & Krehm, M. (2010). The first steps in word learning are easier when the shoes fit: Comparing monolingual and bilingual infants. *Developmental Science, 13,* 229–243. doi:10.1111/j.1467-7687.2009.00891.x

Maye, J., Aslin, R. N., & Tanenhaus, M. K. (2008). The weckud wetch of the wast: Lexical adaptation to a novel accent. *Cognitive Science, 32,* 543–562. doi:10.1080/03640210802035357

Maye, J., Werker, J. F., & Gerken, L. (2002). Infant sensitivity to distributional information can affect phonetic discrimination. *Cognition, 82,* B101–B111. doi:10.1016/S0010-0277(01)00157-3

Monaghan, P., Christiansen, M. H., & Fitneva, S. A. (2011). The arbitrariness of the sign: Learning advantages from the structure of the vocabulary. *Journal of Experimental Psychology: General, 140,* 325–347. doi:10.1037/a0022924

Monaghan, P., Mattock, K., & Walker, P. (2012). The role of sound symbolism in language learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 38,* 1152–1164. doi:10.1037/a0027747

Nazzi, T. (2005). Use of phonetic specificity during the acquisition of new words: Differences between consonants and vowels. *Cognition, 98,* 13–30. doi:10.1016/j.cognition.2004.10.005

Nazzi, T., & Bertoncini, J. (2009). Phonetic specificity in early lexical acquisition: New evidence from consonants in coda positions. *Language and Speech, 52,* 463–480. doi:10.1177/0023830909336584

Nazzi, T., Floccia, C., Moquet, B., & Butler, J. (2009). Bias for consonantal information over vocalic information in 30-month-olds: Cross-linguistic evidence from French and English. *Journal of Experimental Child Psychology, 102,* 522–537. doi:10.1016/j.jecp.2008.05.003

Newport, E. L., & Aslin, R. N. (2004). Learning at a distance: I. Statistical learning of non-adjacent dependencies. *Cognitive Psychology, 48,* 127–162. doi:10.1016/S0010-0285(03)00128-2

Nosofsky, R. M. (1986). Attention, similarity, and the identification–categorization relationship. *Journal of Experimental Psychology: General, 115,* 39–61. doi:10.1037/0096-3445.115.1.39

Oh, G. E., Guion-Anderson, S., Aoyama, K., Flege, J. E., Akahane-Yamada, R., & Yamada, T. (2011). A one-year longitudinal study of English and Japanese vowel production by Japanese adults and children in an English-speaking setting. *Journal of Phonetics, 39,* 156–167. doi:10.1016/j.wocn.2011.01.002

Olejnik, S., & Algina, J. (2003). Generalized eta and omega squared statistics: Measures of effect size for some common research designs. *Psychological Methods, 8,* 434–447. doi:10.1037/1082-989X.8.4.434

Palmeri, T. J., Goldinger, S. D., & Pisoni, D. B. (1993). Episodic encoding of voice attributes and recognition memory for spoken words. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 19,* 309–328. doi:10.1037/0278-7393.19.2.309

Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision, 10,* 437–442. doi:10.1163/156856897X00366

Pierrehumbert, J. (2002). Word-specific phonetics. In C. Gussenhoven & N. Warner (Eds.), *Laboratory phonology VII* (pp. 101–139). Berlin, Germany: Mouton de Gruyter.

Quam, C., & Creel, S. C. (2013). *Both lifelong and immediate language context affect interpretation of lexical tone.* Manuscript submitted for publication.

Ranbom, L., & Connine, C. (2007). Lexical representation of phonological variation in spoken word recognition. *Journal of Memory and Language, 57,* 273–298. doi:10.1016/j.jml.2007.04.001

Sebastián-Gallés, N., Echeverria, S., & Bosch, L. (2005). The influence of initial exposure on lexical representation: Comparing early and simultaneous bilinguals. *Journal of Memory and Language, 52,* 240–255. doi:10.1016/j.jml.2004.11.001

Sebastián-Gallés, N., Vera-Constán, F., Larsson, J. P., Costa, A., & Deco, G. (2009). Lexical plasticity in early bilinguals does not alter phoneme categories: II. Experimental evidence. *Journal of Cognitive Neuroscience, 21,* 2343–2357. doi:10.1162/jocn.2008.21152

Smith, K., & Wonnacott, E. (2010). Eliminating unpredictable variation through iterated learning. *Cognition, 116,* 444–449. doi:10.1016/j.cognition.2010.06.004

Stager, C. L., & Werker, J. F. (1997, July 24). Infants listen for more phonetic detail in speech perception than in word-learning tasks. *Nature, 388,* 381–382. doi:10.1038/41102

Stamer, M. K., & Vitevitch, M. S. (2012). Phonological similarity influences word learning in adults learning Spanish as a foreign language. *Bilingualism: Language and Cognition, 15,* 490–502. doi:10.1017/S1366728911000216

Storkel, H. L., Armbrüster, J., & Hogan, T. P. (2006). Differentiating phonotactic probability and neighborhood density in adult word learning. *Journal of Speech, Language, and Hearing Research, 49,* 1175–1192. doi:10.1044/1092-4388(2006/085)

Sumner, M., & Samuel, A. G. (2009). The effect of experience on the perception and representation of dialect variants. *Journal of Memory and Language, 60,* 487–501. doi:10.1016/j.jml.2009.01.001

van Ooijen, B. (1996). Vowel mutability and lexical selection in English: Evidence from a word reconstruction task. *Memory & Cognition, 24,* 573–583. doi:10.3758/BF03201084

Watson, C. I., Maclagan, M., & Harrington, J. (2000). Acoustic evidence for vowel change in New Zealand English. *Language Variation and Change, 12,* 51–68. doi:10.1017/S0954394500121039

Werker, J. F., & Tees, R. (1984). Cross-language speech perception: Evidence for perceptual reorganization during the first year of life. *Infant Behavior & Development, 7,* 49–63. doi:10.1016/S0163-6383(84)80022-3

White, K. S., & Aslin, R. N. (2011). Adaptation to novel accents by toddlers. *Developmental Science, 14,* 372–384.

Wonnacott, E. (2011). Balancing generalization and lexical conservatism: An artificial language study with child learners. *Journal of Memory and Language, 65,* 1–14. doi:10.1016/j.jml.2011.03.001

Wonnacott, E., Newport, E. L., & Tanenhaus, M. K. (2008). Acquiring and processing verb argument structure: Distributional learning in a miniature language. *Cognitive Psychology, 56,* 165–209. doi:10.1016/j.cogpsych.2007.04.002