Cognition 179 (2018) 89–101

Contents lists available at ScienceDirect

Cognition

journal homepage: www.elsevier.com/locate/cognit

Original Articles

Neutralization and homophony avoidance in phonological learning

Sora Heng Yin[a,*], James White[b]

[a] Department of Linguistics, New York University, 10 Washington Place, New York, NY, 10003, United States
[b] Department of Linguistics, University College London, Chandler House, 2 Wakefield Street, London WC1N 1PF, United Kingdom

ARTICLE INFO

Keywords:
Phonology
Language change
Learning bias
Functional load
Neutralization
Artificial language

ABSTRACT

Previous research has suggested that homophony avoidance plays a role in constraining language change; in particular, phonological contrasts are less likely to be neutralized if doing so would greatly increase the amount of homophony in the language. Most of the research on homophony avoidance has focused on the history of real languages, comparing attested and unattested (hypothetical) phonological changes. In this study, we take a novel approach by focusing on the language learner. Using an artificial language learning paradigm, we show that learners are less likely to acquire neutralizing phonological rules compared to non-neutralizing rules, but only if the neutralizing rules create homophony between lexical items encountered during learning. The results indicate that learners are biased against phonological patterns that create homophony, which could have an influence on language change. The results also suggest that lexical learning and phonological learning are highly integrated.

1. Introduction

A crucial function of speech sounds is that they allow speakers to contrast words. For instance, the English words *tab* and *dab* differ only in their initial sounds ([t] vs. [d]), suggesting that /t/ and /d/ represent two basic sound categories, or phonemes, in English (Hayes, 2009). However, the phonological rules of a language sometimes result in the neutralization of phonemic contrasts. For example, the flapping rule in American English affects both /t/ and /d/, changing them into an alveolar flap [ɾ] between vowels when the second vowel is unstressed; as a result, the distinct lexical items *pat* [pæ**t**] and *pad* [pæ**d**] have the same phonetic realization when the suffix *–ing* is added (i.e., *patting* and *padding* are both pronounced as [pæɾɪŋ]).[1]

We can distinguish two levels of neutralization. First, there can be neutralization at the lexical level, such as when flapping results in the same pronunciation of the words *patting* and *padding.* This level of neutralization could be called derived homophony (Silverman, 2012, p. 4). Second, there is neutralization within the phonological system. Neutralizing phonological rules, by definition, eliminate a contrast between two (or more) phoneme categories; for instance, the contrast between /t/ and /d/ is lost in American English in flapping contexts because both are realized as the same sound, [ɾ]. These two levels of neutralization (i.e., phonological and lexical) are clearly related, but are partially distinct. The application of a neutralizing phonological rule may result in homophony, or it may not. For instance, flapping occurs in the word *getting* [gɛɾɪŋ], but there is no lexical neutralization in this particular case because *ged* and *gedding* are not existing words of English. In sum, while neutralizing phonological rules have the *potential* to create homophony, the amount of *actual* homophony that they create can vary depending on the contents of the lexicon.[2]

Neutralization poses a challenge for our understanding of language change and typology. On one hand, neutralization creates ambiguity, which reduces the communicative efficiency of a language. On the other hand, neutralization is not uncommon in the world's languages (for an overview, see Silverman, 2012). Are there pressures against the development of neutralizing rules given that they increase ambiguity? If so, which mechanisms are responsible for such pressures, and which factors influence whether a neutralizing rule will eventually develop in a language?

According to the functional load hypothesis (Hockett, 1967; King, 1967; Martinet, 1952; Wedel, Kaplan, & Jackson, 2013; see also earlier work: Gilliéron, 1918; Jakobson, 1931; Mathesius, 1931), the likelihood of two phonemes being neutralized over the course of language change depends on the amount of information that they carry: pairs of

* Corresponding author.
E-mail addresses: sora.heng.yin@nyu.edu (S.H. Yin), j.c.white@ucl.ac.uk (J. White).

[1] Strictly speaking, flapping in American English is not completely neutralizing from a phonetic perspective. Subtle acoustic differences remain between /t/ and /d/ after flapping, including a small (but statistically reliable) difference in the duration of the preceding vowel. However, adult speakers are unable to reliably distinguish flapped /t/ and /d/ in perception (Herd, Jongman, & Sereno, 2010).

[2] We use the term 'rule' throughout the paper as a straightforward way of referring to a context-sensitive change from one sound to another that could be generalized to novel cases. We are not suggesting that the cognitive structure of the grammar is organized as in traditional rule-based theories of phonology (e.g., Chomsky & Halle, 1968).

https://doi.org/10.1016/j.cognition.2018.05.023
Received 26 June 2017; Received in revised form 26 March 2018; Accepted 28 May 2018
0010-0277/

phonemes that play a greater role in contrasting meaning in a language (i.e., have a higher functional load) are less likely to be neutralized over time. Applied to phonological rules, the theory predicts that a neutralizing rule resulting in a large amount of homophony is less likely to be adopted into a language than one resulting in little homophony, all else being equal.

A prediction that emerges from the functional load hypothesis is that the amount of homophony created amongst *existing* lexical items should affect the likelihood of a neutralizing rule being adopted or applied. As Kaplan (2011) points out, this outcome is not necessarily expected under all models of phonology. While models of phonology with connections to exemplar theory have embraced the idea that the lexicon and the phonological system are highly integrated (Bybee, 2002; Ernestus & Baayen, 2003; Hay, 2000; Pierrehumbert, 2001; Pierrehumbert, 2002; Ussishkin & Wedel, 2009; Wedel, 2004, 2006, 2007), traditional generative models of phonology have typically maintained a modular view of the lexicon and the phonology. In generative models, the domain of phonological analysis is considered to be the set of possible words in a language, not the set of existing words in the lexicon (e.g., see Halle, 1962; Ní Chiosáin & Padgett, 2009; Padgett, 2003, 2009).

While the functional load hypothesis is intuitively appealing, testing whether homophony avoidance actually affects the likelihood of neutralization has proven difficult. Many factors play a role during language change. Though we often know which phonological changes have occurred in a language, it is difficult to identify which other hypothetical changes could have occurred and determine why they did not occur (e.g., see discussion in Kaplan, 2011). In the current study, we take a novel approach to this issue by focusing directly on the language learner. On the basis of two artificial language experiments, we show that learners are less likely to acquire neutralizing rules than non-neutralizing rules, but only when the neutralizing rules create homophony. The results provide support for the functional load hypothesis and suggest that a learning bias is one mechanism through which homophony avoidance could shape phonological systems.

In the following sections, we review previous literature on homophony avoidance, provide a brief overview of the artificial language learning paradigm, and then introduce the experiments of the current study.

1.1. Homophony avoidance in language change

One approach to testing whether homophony avoidance plays a role in language change is to compare the amount of homophony created by attested phonological changes and hypothetical changes that did not occur. Relevant phonological changes include complete phoneme mergers (neutralization in all contexts) and the adoption of new phonological rules which neutralize contrasts in certain contexts. Several studies taking this approach have concluded that homophony avoidance plays a significant role in constraining phonological change. For instance, Silverman (2010) argued that even though Korean has several neutralizing rules, these rules create less homophony than other comparable (but non-occurring) rules would. Expanding on this research, Kaplan (2011) used Monte Carlo simulations to show that the attested neutralizing rules in Korean produce fewer homophones than we would expect based on the distribution of homophony created by randomly generated sets of comparable unattested rules.[3] Wedel et al. (2013) conducted a statistical analysis of neutralizing diachronic changes (mergers and new neutralizing rules) in nine languages. They found that the number of minimal pairs distinguished by a pair of phonemes was a significant predictor of whether the pair would be neutralized, even taking into account other factors such as overall phoneme probability.

[3] This was true for nouns. The results for verbs were more equivocal.

Other researchers, however, have raised doubts about the claim that functional load, or homophony avoidance, plays a significant role in phonological change. King (1967) analyzed the history of phoneme mergers in several Germanic languages and concluded that functional load is, at best, one of the least important factors in sound change. Sampson (2013) likewise argued that the history of phonological changes from Middle Chinese to Modern Chinese shows an increase in the amount of homophony between morphemes, which is inconsistent with homophony avoidance being a major factor in language change (see Kaplan, 2015 for a reply). Thus, the role of homophony avoidance in sound change remains an unsettled issue.

These disagreements underscore some major challenges facing this line of research. It is difficult to determine precisely how much homophony should be considered inconsistent with the functional load hypothesis. Even if we interpret the theory as stating that attested changes should produce less homophony than is expected by chance, it is difficult to identify which set of hypothetical changes should be used for calculating the 'expected' amount of homophony. Many interacting factors affect the likelihood that a phonological change will occur. Moreover, some phonological changes are inherently more likely to occur than others, independent from any effect of homophony avoidance. Defining a set of comparable hypothetical rules is far from a trivial endeavor (e.g., see discussion in Kaplan, 2011).

To show definitively that a hypothetical change failed to occur in a language due to homophony avoidance, one would need to demonstrate that the hypothetical change would have occurred if all conditions were identical except that it created less homophony (Kaplan, 2011). Blevins and Wedel (2009) cover cases in the history of natural languages where this was hypothesized to have occurred, but confirming that homophony avoidance was the cause in any particular case is challenging, particularly given the potential influence of other factors such as analogical change. In the history of natural languages, it is difficult to determine with confidence what speakers would have done if circumstances had been different.

However, using a tightly controlled artificial language learning experiment, we can indeed probe what learners will do under minimally different circumstances. Specifically, we can provide learners with the same set of rules to learn, varying only the amount of homophony created by those rules. This is the approach we took in the current study, allowing us to directly test whether homophony avoidance makes learners less likely to adopt neutralizing phonological rules.

1.2. Exploring learning biases using artificial language learning

The language learner plays an important role in shaping language change. Over many generations, even subtle learning biases can shift languages in certain directions (Culbertson, Smolensky, & Legendre, 2012; Kalish, Griffiths, & Lewandowsky, 2007; Kirby, Smith, & Brighton, 2004; Reali & Griffiths, 2009; White, 2017). Understanding the biases that learners bring to the language learning process is therefore a key component of explaining language change and typology. Artificial language learning experiments have emerged as a useful framework for probing learning biases with adults, children, and infants (Baer-Henney & van de Vijver, 2012; Carpenter, 2010; Chambers, Onishi, & Fisher, 2003; Cristià & Seidl, 2008; Culbertson & Newport, 2015; Culbertson et al., 2012; Fehér, Wonnacott, & Smith, 2016; Finley, 2011; Finley, 2017; Finley & Badecker, 2009; Finley & Badecker, 2012; Moreton, 2008; Newport & Aslin, 2004; Onishi, Chambers, & Fisher, 2002; Peperkamp, Skoruppa, & Dupoux, 2006; Saffran & Thiessen, 2003; Seidl & Buckley, 2005; Skoruppa & Peperkamp, 2011; Smith & Wonnacott, 2010; White, 2014; White & Sundara, 2014; Wilson, 2006; Wonnacott, 2011; Wonnacott, Brown, & Nation, 2017; for reviews see Gómez & Gerken, 2000; Moreton & Pater, 2012a; Moreton & Pater, 2012b). These studies have demonstrated that participants are able to learn novel linguistic patterns (including phonological rules) after brief exposure to an artificial language in the lab.