

What makes a phonological pattern learnable? A simple, unnatural rule in English

John Harris (UCL), Nick Neasom (UCL) and Kevin Tang (Yale)

There is a well-established collection of speaker-independent methods for discovering phonotactic patterns in languages, e.g. comparative reconstruction, phonological analysis and computational learning. There is also an increasingly varied collection of experimental methods for ascertaining how much of this patterning is actually internalised by speaker-hearers. In seeking to determine what makes a phonotactic pattern learnable or not, researchers have focused on a variety of factors, including phonological regularity, productivity, naturalness, and formal simplicity. Experimental studies have investigated various permutations of these factors, with results that are more or less surprising. For example, speakers have been shown to have internalised and to be able to productively apply (a) patterns that are regular, simple and natural (e.g. wug tests of English **-s**) but also (b) patterns that are irregular, relatively complex and not synchronically natural, such as English velar softening (e.g. Pierrehumbert 2006).

In this paper, we examine the English phonotactic pattern where consonants following /aw/ are restricted to coronals; hence **tout**, but not */tawk/, */tawp/ (e.g. Halle & Clements 1983). The pattern ('awT') is pretty regular, more so than velar softening. It is general, in that it affects a large swath of the lexicon. It is formally quite simple, more so than the **-s** pattern. And it is not natural. It is the synchronically accidental outcome of a series of largely unrelated sound changes; each of the changes might be natural, but their cumulative effect is not.

We report the results of two non-word judgement experiments designed to test the extent to which native speakers of English have tacit knowledge of the awT pattern. In both experiments, listeners were presented with C₁₋₃VC non-word stimuli containing the diphthongs /aw/, /ow/, /ij/, followed by a range of consonants, and were asked to judge how English-like they sounded. The selection of the non-words was controlled for lexical neighbourhood density, weighted by frequency. In the first test, listeners made forced choices between paired words distinguished solely by whether the vowel was followed by a coronal versus a non-coronal consonant. In the second, listeners rated individually presented stimuli on a scale of Englishness, drawn from a sample of around 1200 nonwords.

The question of whether speakers have implicit knowledge of a given phonotactic pattern can be approached in two stages: (a) do they have any tacit awareness of the pattern at all and, if so, (b) is the awareness commensurate with the pattern being stored as a grammatical rule? Broadly speaking, the results of both experiments show weak evidence of an awareness of awT but little or no evidence that this reflects grammaticalised knowledge. That is, to the extent that speakers have any tacit inkling of the pattern at all, it is probably not encapsulated in anything like a phonologist's rule or constraint. Where a coronal preference is observable, it does not generalise across different manners of articulation, as would be expected if there were a rule-driven bias towards formal simplicity. Also, the preference is influenced by onset size and lexical neighbourhood factors, which suggests subjects were making on-the-fly judgements of how much the non-words resemble real words.

We conclude that awT is a case where phonologists know more about a phonotactic pattern than speakers know. In the light of our results, we consider whether this should be attributed to the fact that awT is not natural (cf. Hayes & White 2013) or to other factors, such as that it is not involved in alternations.