Word-frequency effects in sound change as a consequence of perceptual asymmetries: An exemplar-based model

Todd, Pierrehumbert, Hay (2019)

In a nutshell

- RQ: How can high-frequency words change faster than low-frequency words in some cases, slower in other cases, and at the same rate in yet other cases?

- Kind-of-an-answer: This puzzle can be answered by giving substantial weight to the role of the listener.

- Content: The authors present an exemplar-based computational model of regular sound change, in which they show that high-frequency words can change:

  1. at the same rate as low-frequency words when a phoneme category moves without encroaching on the acoustic space of another
  2. faster than low-frequency words when it moves toward another
  3. slower than low-frequency words when it moves away from another

Background

- The authors argue that a possible sound change is shaped by the listener. Therefore, it is essential to acknowledge that listeners are also speakers and that any cognitive change to sound representation in the listener's perceptual system will be reflected in the speech of that listener-turned-speaker.
    - Ohala (1981) – sound change occurs as a result of the misperception of sounds in their phonological environment.
    - Harrington et al. (2018) – sound change is gradient in all words via biases involved in correct perception of spoken words.
    - The authors of this study (Todd et al. 2019) – sound change is a gradient transformation of a phoneme over time. Focus is on the rate at which words of different frequencies participate in regular sound change.

FAH

- Frequency Actuation Hypothesis (FAH) claims that the word-frequency effect will be different for different kinds of sound change (Phillips 1984).

Which words change faster?	Physiologically motivated changes (*matter* → *madder*)	Non-physiologically motivated changes (*tune* → *toon*)
High-frequency words	✓	
Low-frequency words		✓

FAH

- The most intuitive application of the FAH to regular sound change makes the assumption that gradient phonetic change results primarily from iterated biases in the speaker's phonetic implementation, and thus predicts that high frequency words should always change faster than low-frequency words.

- However, the FAH does not uniformly support the data.

Study	**Sound change**	**Result**	**FAH**
Bermudez-Otero et al. (2015)	[t] → [ʔ]	Low- and high-frequency words change *at the same time*	?
Hay and Foulkes (2015)	[t] → [ɾ] or [d]	High-frequency words change faster	✓
Hay et al. (2015)	[ɛ] → [ɪ]	Low-frequency words change faster	X

Hypotheses

Hypotheses	**Condition**
Low- and high-frequency words change *at the same time*	When a phoneme moves with no acoustic ambiguity
High-frequency words change faster	When a phoneme moves toward another, increasing acoustic ambiguity
Low-frequency words change faster	When a phoneme moves away from another, decreasing acoustic ambiguity

Goal

- To develop a computational model that tries to capture the results of the three studies.

Study	Sound change
Bermudez-Otero et al. (2015)	[t] → [ʔ] between vowels as in *mitten* which is produced as "*mi'en*"; words of all frequencies change at the same rate
Hay and Foulkes (2015)	[t] → [ɾ] or [d] between vowels as i produced as "*madder*"; high-frec faster
Hay et al. (2015)	[ɛ] → [ɪ] is a part of the *push chain* to [ɛ] low-freq. words chan

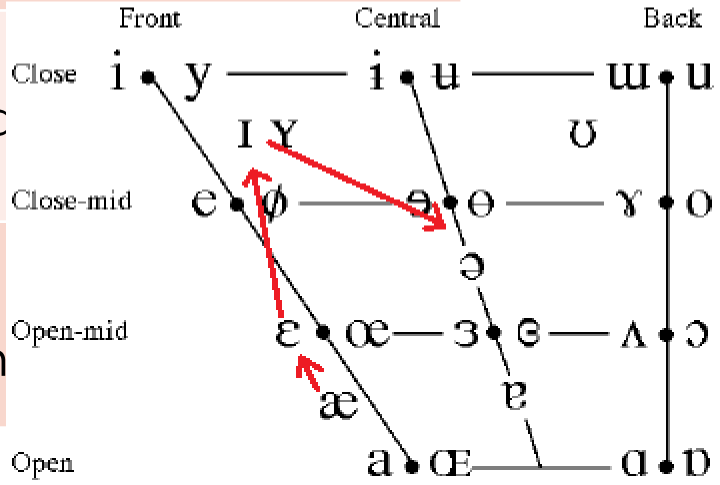

Modeling the Sound Change

Change	Method of modeling
/t/ - glottaling	A single isolated phoneme category subject to a consistent production bias.
/t/ - tapping	Two-phoneme category whereby one is biased toward the other, focusing on the category which is subject to the bias.
/ɛ/ - raising	Two-phoneme category whereby one is biased toward the other, focusing on the category which is not subject to the bias.

In /t/-tapping, /t/ is the Pusher, while in the /ɛ/-raising, /ɛ/ is the Pushee

Model Desiderata

- The model must produce certain frequency-independent properties that are related to the maintenance of structure over the course of change.

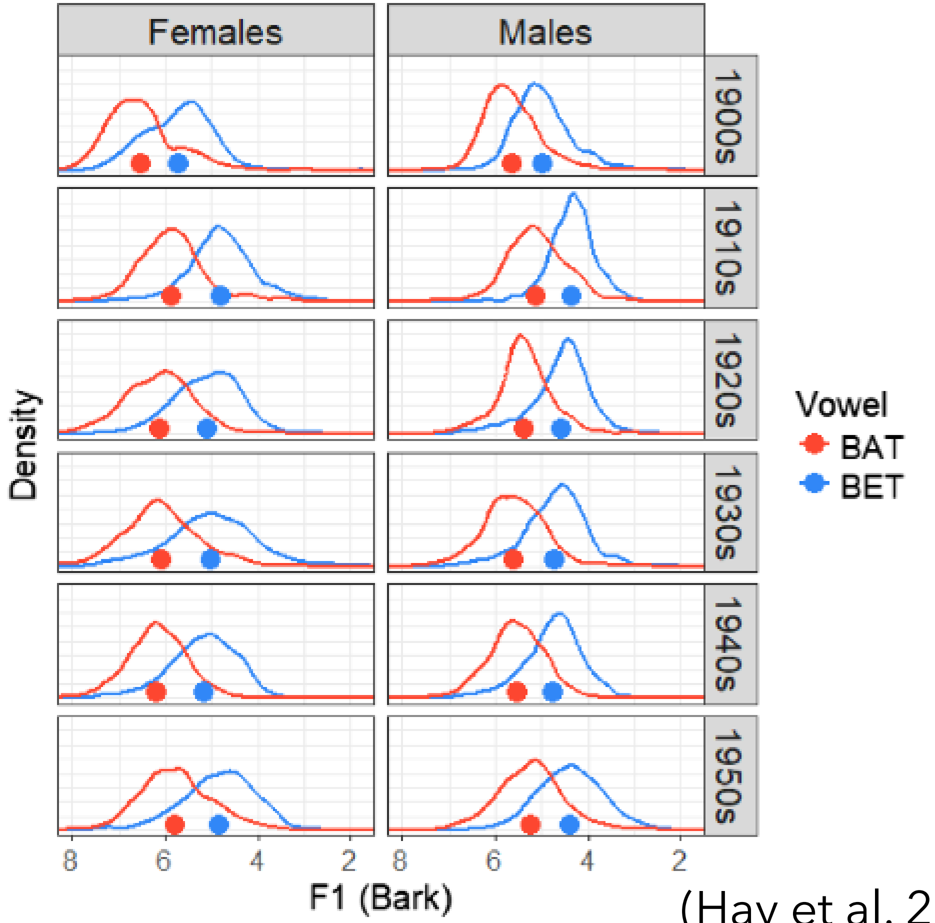

(Hay et al. 2015)

- The model should (**1**) generate the movement of each category, (**2**) maintain the shape (width and skewness) of each category; and in two-category interaction, it should maintain (**3**) the distance between and (**4**) overlap of the categories

Model Desiderata

- The model must balance a number of forces.

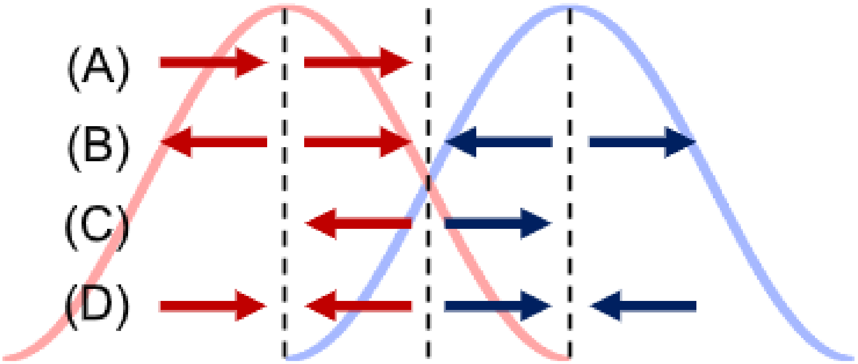

- Forces that act on the phoneme category distribution in order to meet the model desiderata
- (**A**) Intrusive force – Pusher → Pushee
- (**B**) Spreading force – Far side of the Pushee retreats
- (**C**) Repulsive force – Pushes 2 categories apart
- (**D**) Squeezing force – Categories not too wide or skewed

- A single-category model initiates and balances A, B, D.
- A two-category model initiates and balances A, B, C, D.

Exemplar Theory

- Categorization of a perceived stimulus entails the comparison with *exemplars* – episodic traces of experienced instances – of other stimuli, stored in memory (Nosofsky,1986). For speech perception, this is taken to mean that listeners store richly-detailed memories of spoken words as they experience them, which they use as a basis of comparison for categorizing other instances of spoken words.

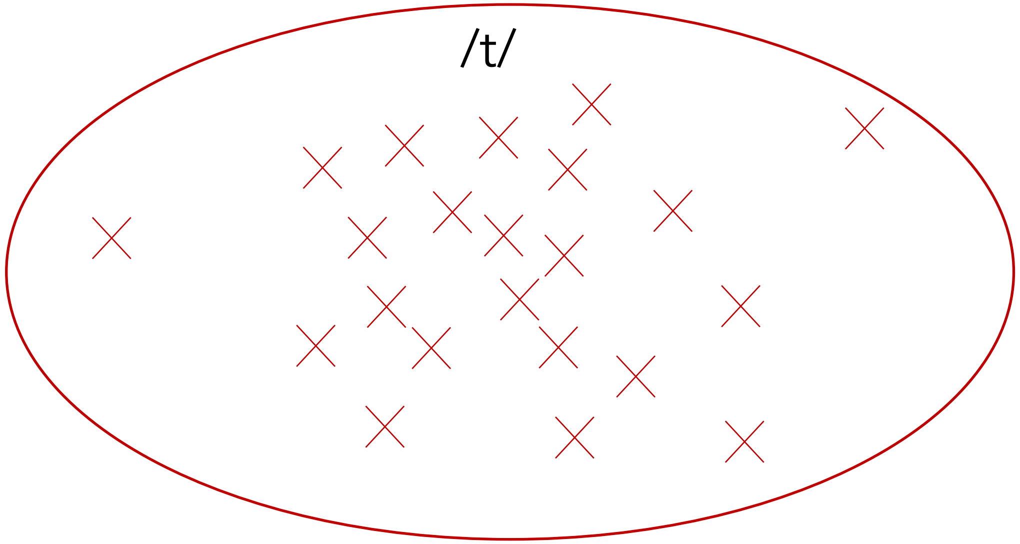

Model Description

Level of representation	Meaning
Category	A generalization over experienced instance of a phoneme (e.g. a vowel), stored in memory -- /æ/ from 'map', 'lab', 'sad', etc. 1-2 categories
Type	An abstract template of a word containing a particular phoneme that contains information about the *frame* (e.g. onset and coda /m/ and /p/ in the word 'map'), and the category (e.g. nucleus vowel /æ/); Around 92 types per category
Exemplar	Experienced perceptual-acoustic realizations of the category phoneme; 492 exemplars per category
Type frequency	The number of exemplars for a given type. Based on word log-frequencies in a large corpus; 1-12 range
Exemplar space	The distribution of exemplars across a granularized perceptual-acoustic dimension e.g. vowel F1). Assumed to be shared across perception and production; The space in this study is one-dimensional;

Model Description

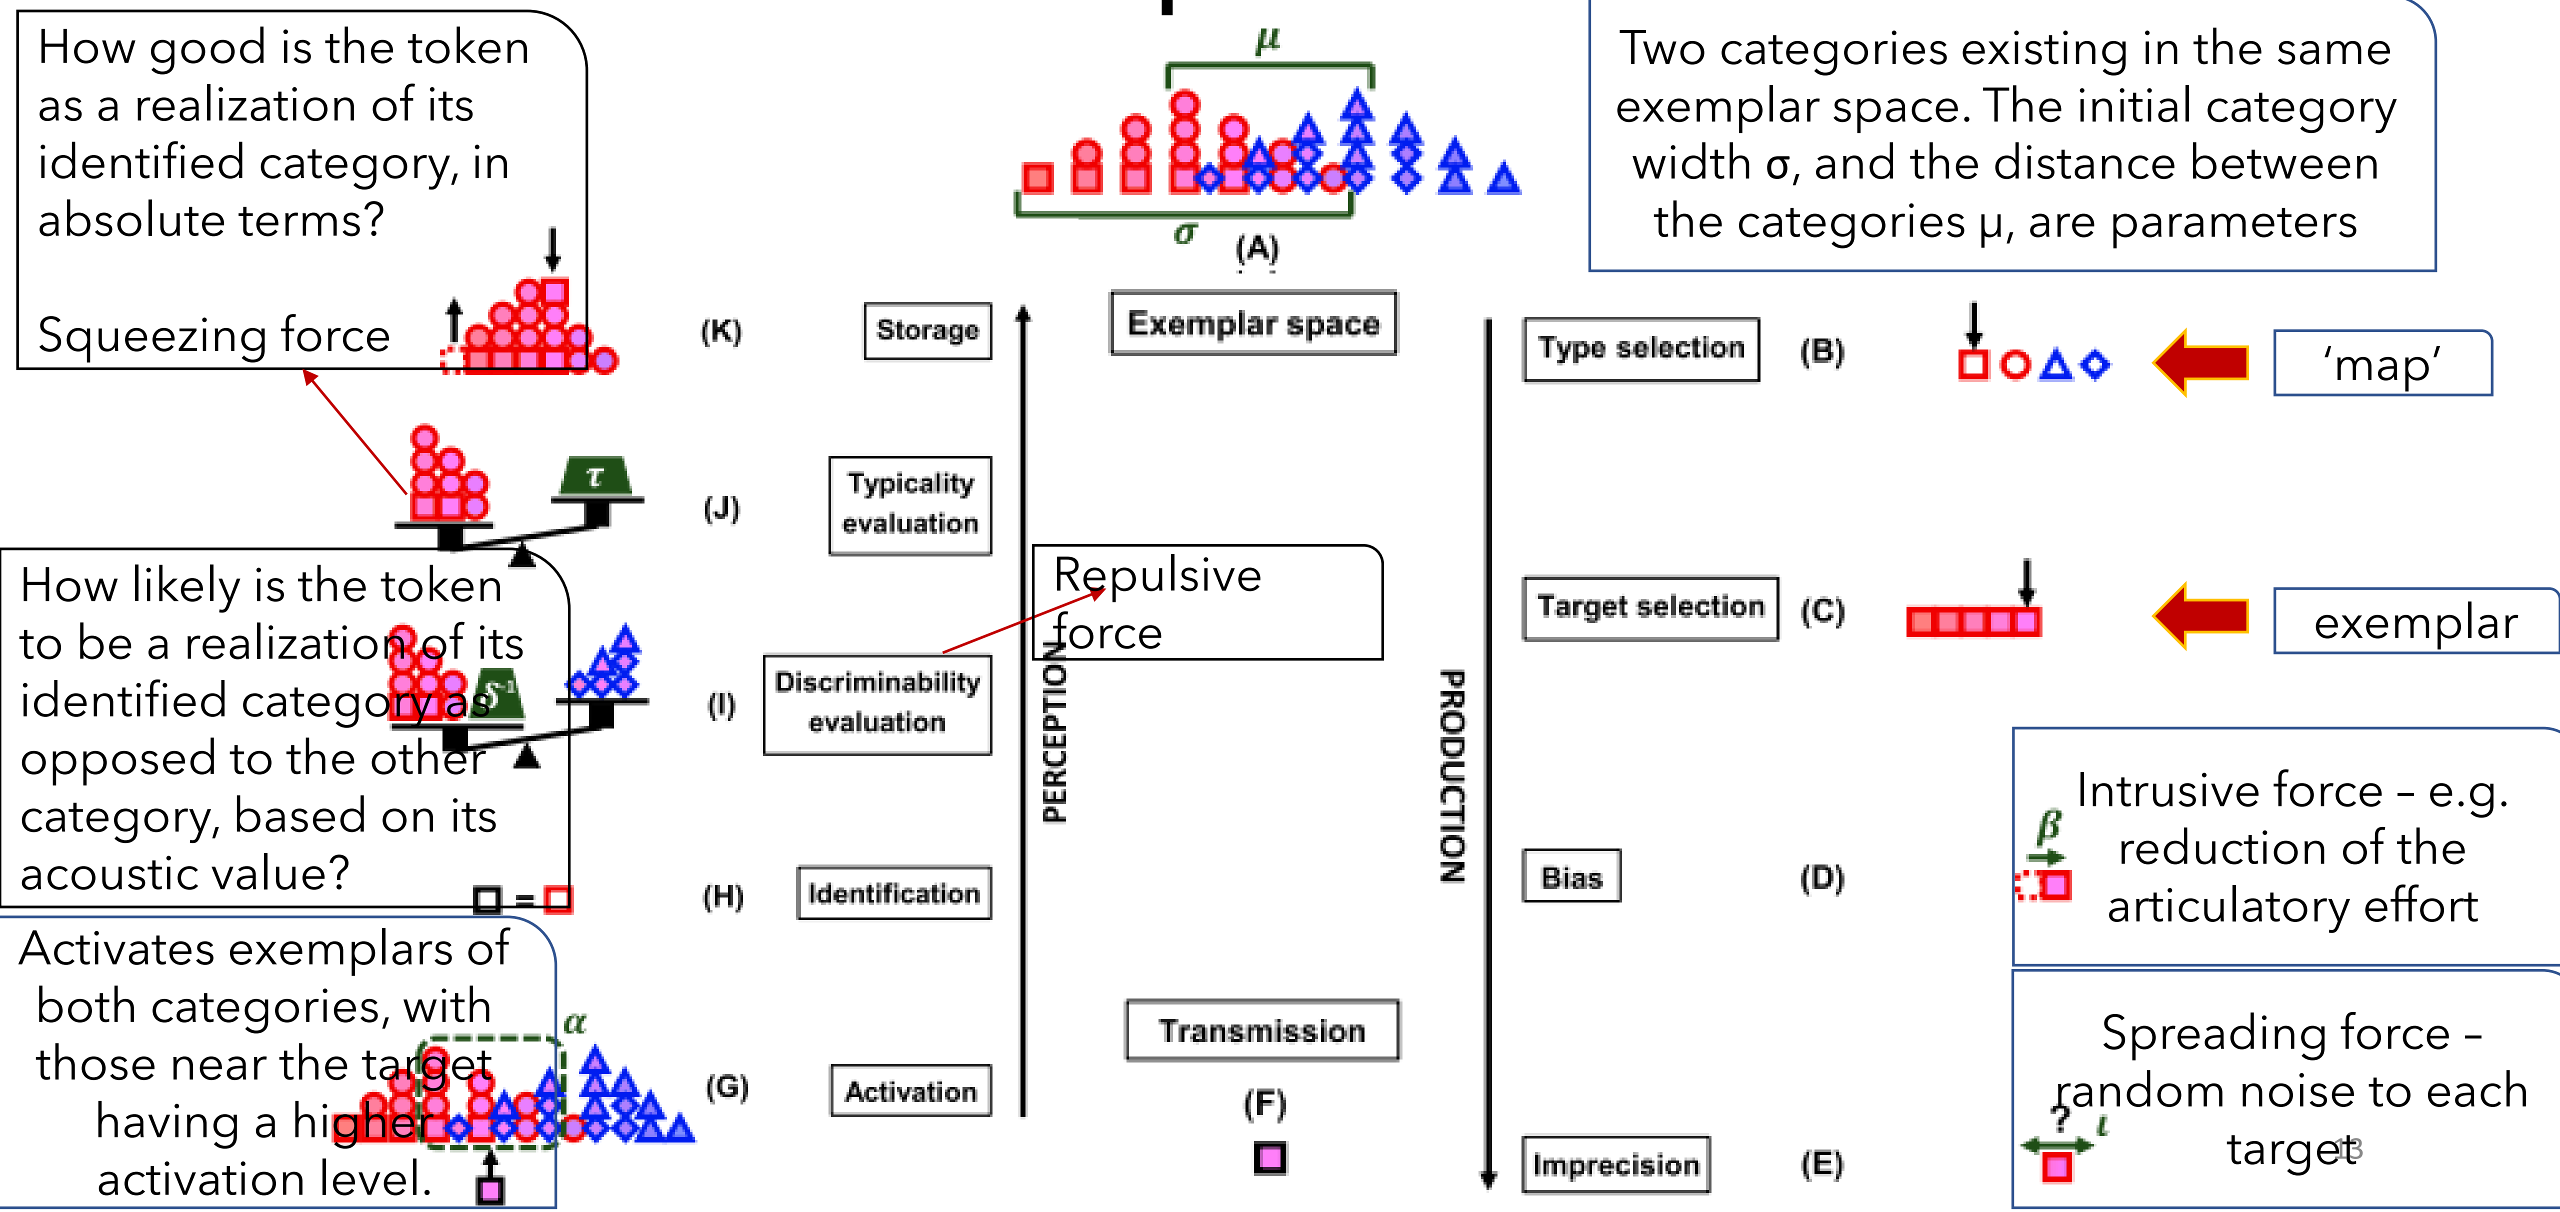

Two categories existing in the same exemplar space. The initial category width σ, and the distance between the categories μ, are parameters

How good is the token as a realization of its identified category, in absolute terms?

Squeezing force

How likely is the token to be a realization of its identified category as opposed to the other category, based on its acoustic value?

Activates exemplars of both categories, with those near the target having a higher activation level.

Repulsive force

'map'

exemplar

Intrusive force – e.g. reduction of the articulatory effort

Spreading force – random noise to each target

Single-Category Modeling

- Basic desiderata for such a model are that it: (**i**) generates movement of the category; and (**ii**) maintains the shape (width & skewness) of the category.

- Results of the model

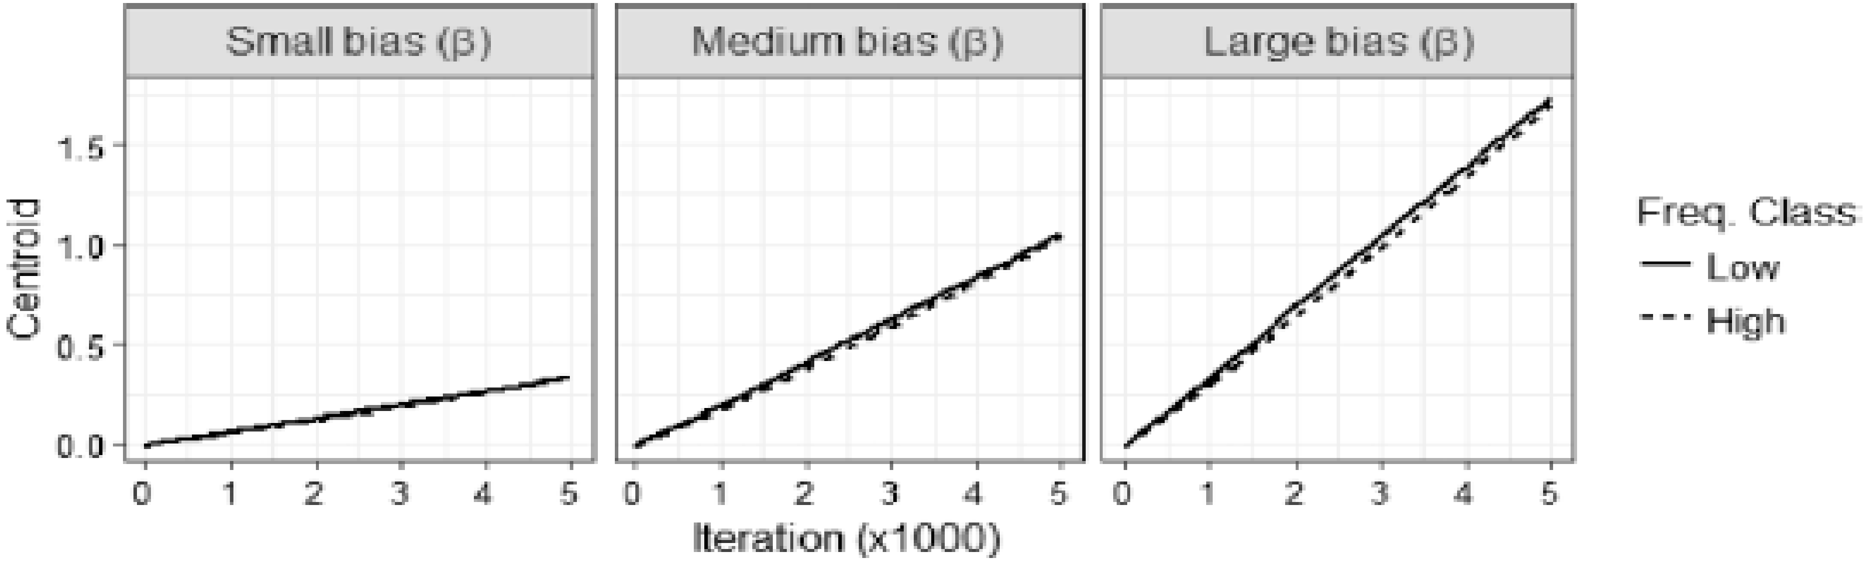

Fig. 5. Results of simulations with a single category ($\sigma = 0.8$) subject to varying degrees of bias, illustrating differences between low-frequency (solid) and high-frequency (dashed) types. For all degrees of bias, the centroid of the category distribution advances at the same rate for both low- and high-frequency types.

Two-Category Modeling

- Basic desiderata for such a model are that it: (**i**) generates movement of one category in response to other; (**ii**) maintains the distance between the categories, (**iii**) maintains the shape (width and skewness) of the categories; and (**iv**) maintains the overlap of the categories.

- The authors develop the *basic model,* which proves to be successful. Why?

  1. By not storing tokens that fail the discriminability evaluation, the model avoids skewness-inducing overpopulation of the overlapping region between categories.
  2. By including the novel process of typicality evaluation, our model generates a squeezing force that keeps skewness in check while facilitating overlap.

Two-Category Modeling

- The model emphasizes the importance of the listener, because whenever the speaker is ambiguous, the listener is unlikely to store the token, and is thus unlikely to use it as a basis for future productions. The listener thus drives category interaction in our model by creating category repulsion via the discriminability force, with the speaker's constant bias serving to ensure that interaction persists in the face of this repulsion.

- Enriching the basic model with an empirically-grounded perceptual asymmetry allows for the generation of word frequency effects on rate of change that resemble those seen in New Zealand English /t/-tapping and /ɛ/-raising.

- All else being equal, the perceptual system is biased toward the recognition of high-frequency words, especially in the case of acoustically ambiguous tokens (lots of evidence for this provided).

Enhanced Model

- Approach: The authors encode this perceptual asymmetry in the model by varying the discriminability threshold, with *type frequency*. Tokens of high-frequency types receive lower δ than tokens of low-frequency types, making them more discriminable, i.e. more likely to pass the discriminability evaluation and be stored when encountered.

- Simulations with frequency-sensitive δ, keeping all other parameters fixed at the previously-tuned values.15 frequency-sensitive functions are constructed.

Enhanced Model

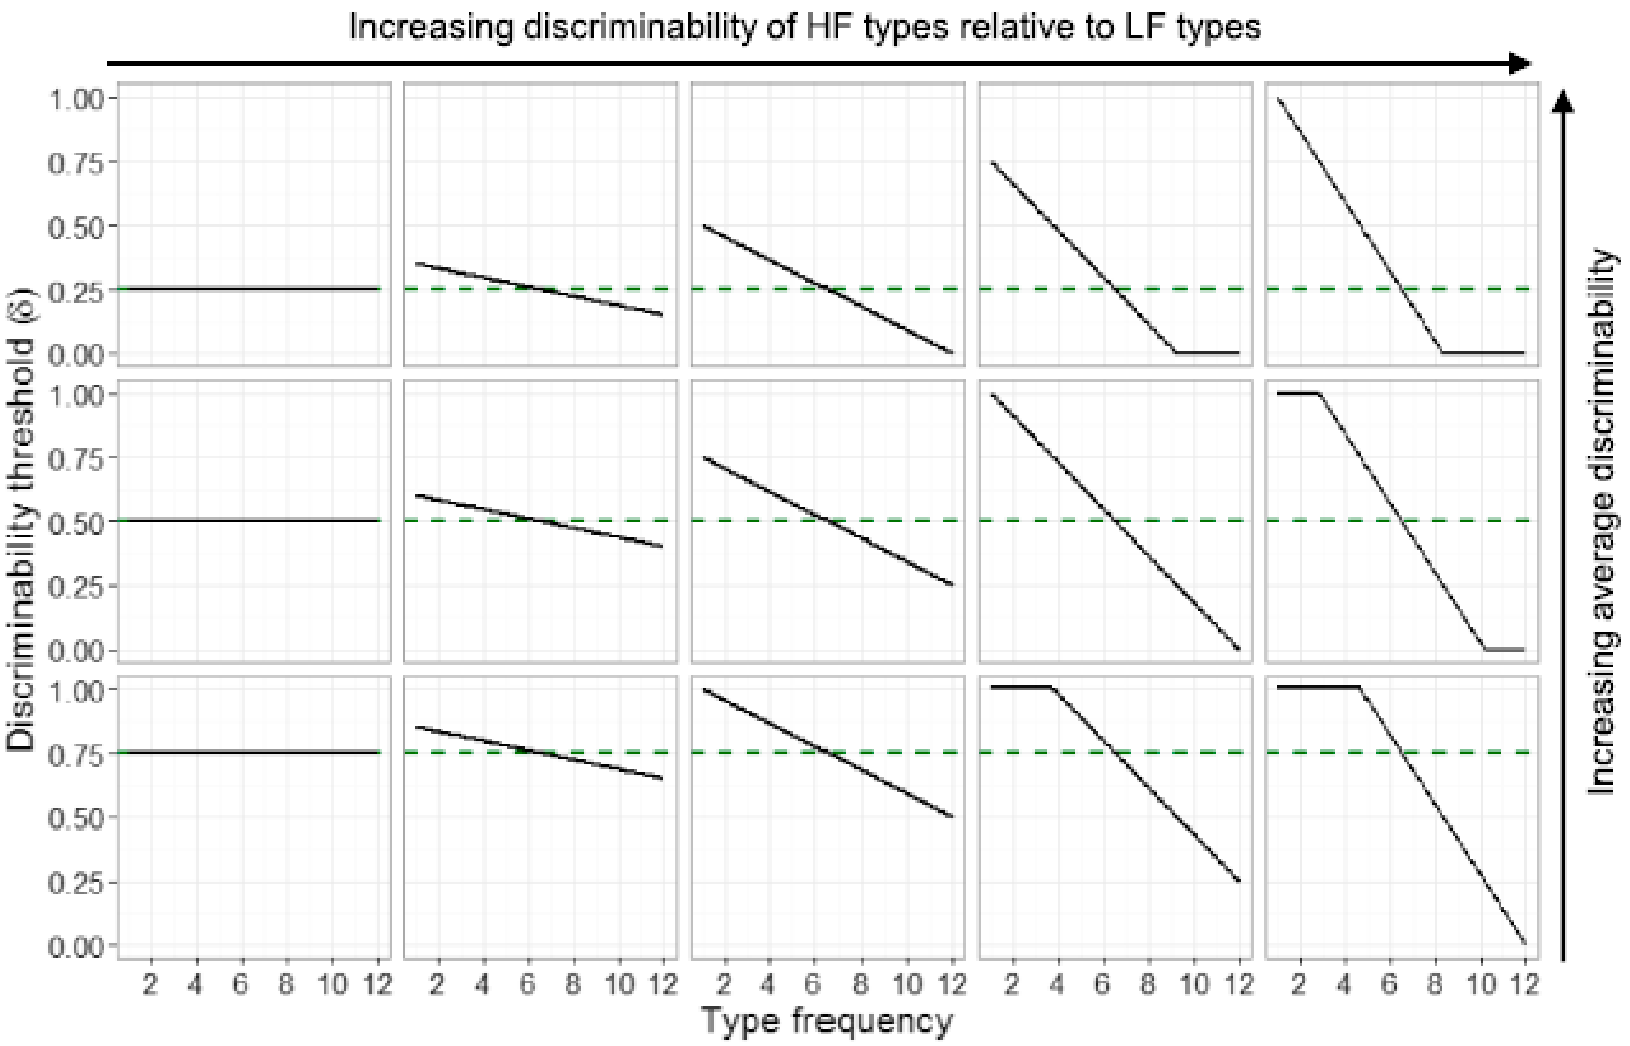

- The greater the type frequency, the lower the discriminability threshold, and therefore, the greater discriminability is generated.

Enhanced Model

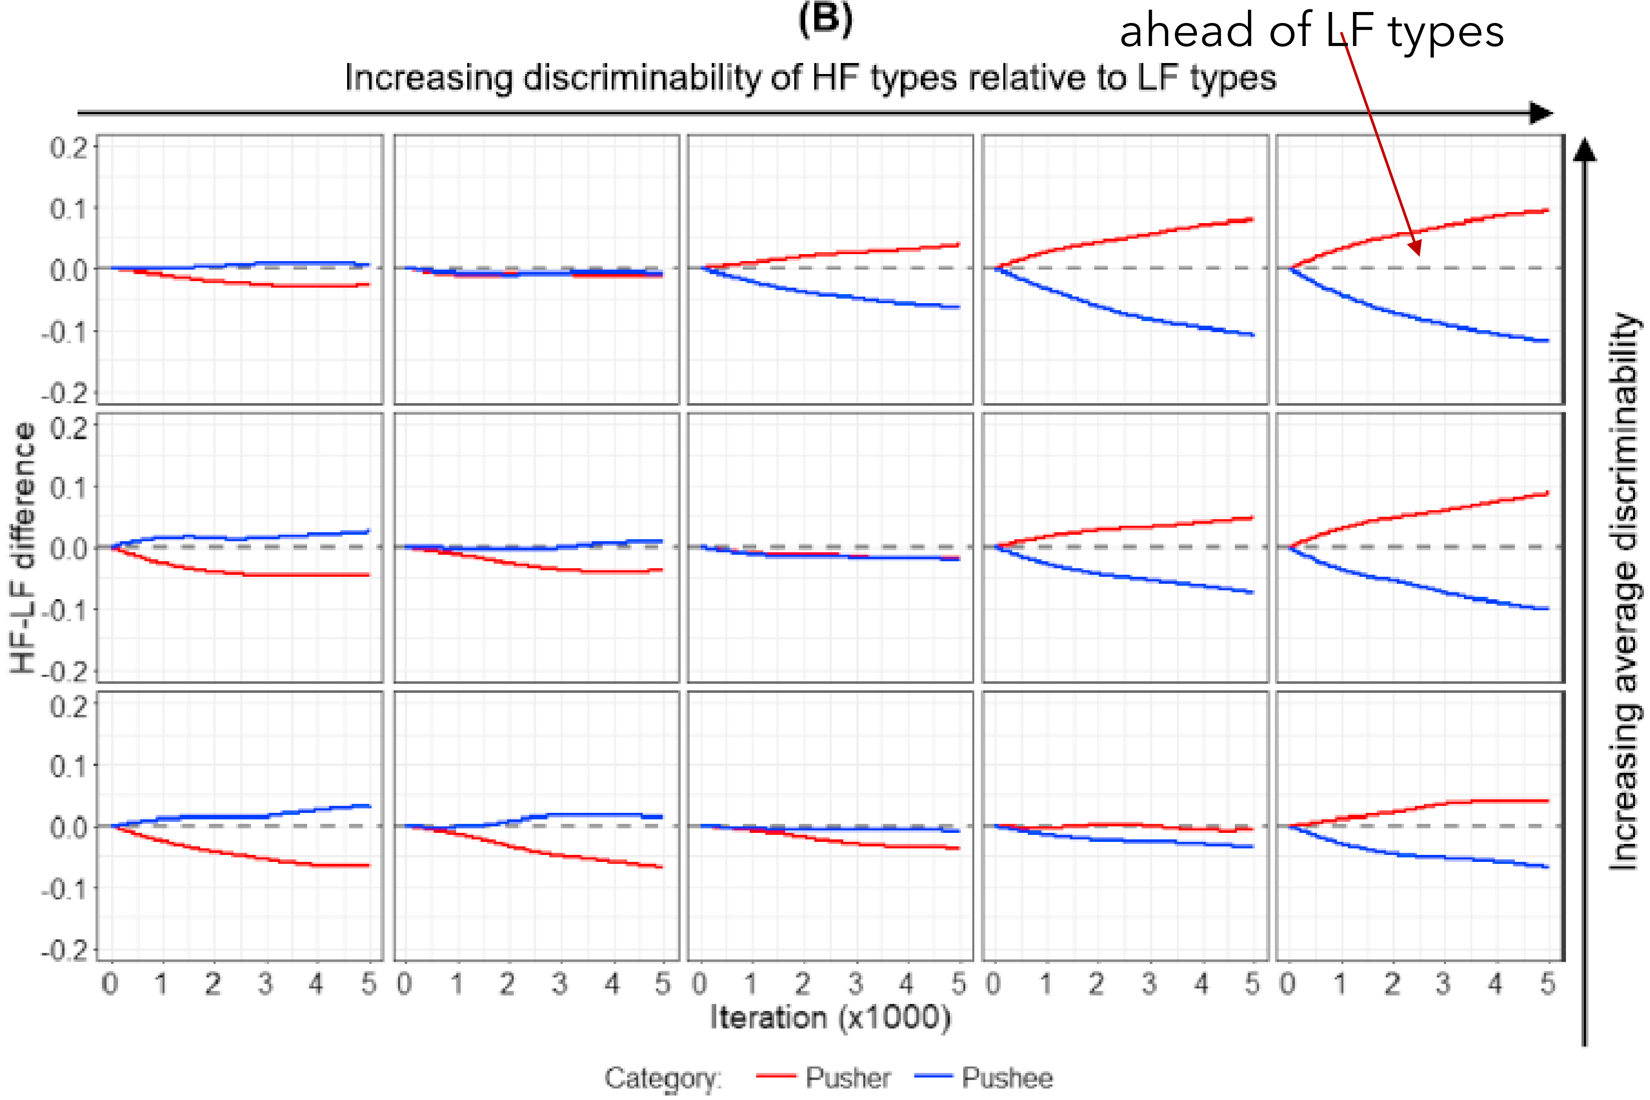

To what extent HF types are ahead of LF types

- The higher the discriminability, the greater the rate of change of HF words in the Pusher, and the greater the rate of the LF types in the Pushee.
- This means that the model manages to describe the behaviour of the /t/-tapping (Pusher-red) and /ɛ/-raising (Pushee-blue)

Discussion

- Unlike previous exemplar-based models of word frequency effects in sound change (Pierrehumbert, 2001, 2002), this model successfully generates different kinds of word frequency effects in different kinds of changes. This success follows from the conception of sound change not merely as the iteration of articulatory biases in the speaker, but rather as the result of balancing emergent forces that stem from both the speaker and the listener, where words of different frequencies are crucially assumed to be differentially sensitive to the perceptual forces in the listener.

- Unlike the FAH model, this model supports the data.

Discussion

- The case studies captured by the model can be united under consideration of how each category's change affects its discriminability.

Change	**Discriminability**
/t/ - glottaling	When the change has no impact on discriminability
/t/ - tapping	When the change acts to decrease discriminability for a category
/ɛ/ - raising	When the change acts to increase discriminability for a category

Conclusion

- The study presents an exemplar-based computational model of regular sound change that generates appropriate single-category movement and two-category interactions, and reflects key (frequency-independent) properties of real sound changes.

- The listener plays a central role in sound change allows it to predict different effects on word frequency on rate of change in different kinds of sound change, which match all of the empirical results that exist at the time of writing.

- The model thus shows how word frequency-based asymmetries in perception can generate word frequency effects on rate of sound change, without similar asymmetries in production, which differ systematically for different kinds of sound change.