Journal of Language Evolution, 2017, 77–93
doi: 10.1093/jole/lzx009
Advance Access Publication Date: 18 May 2017
Research article

Category competition as a driver of category contrast

Andrew Wedel[1,*] and Ibrahim Fatkullin[2]

[1]Department of Linguistics, University of Arizona, Tucson, AZ, 85721, USA and [2]Department of Mathematics, University of Arizona, Tucson, AZ, 85721, USA

*Corresponding author: wedel@email.arizona.edu

Abstract

Some mental categories map to percepts which are products of human behaviors, such as linguistic signals. Because behavior is learned and updated by experience, biases in the way a behavior is perceived can influence how it is reproduced, allowing behaviorally based categories to evolve over time. Here we show that this perception–production feedback loop can itself promote preservation of contrast between categories. Using both simulation and analytical tools, we show that asymmetries in the mapping of perceptual variants to competing categories acts to sharpen category boundaries. Evidence from patterns of change in modern languages is consistent with this mechanism. Because the ability to maintain a large number of distinct signal/meaning categories is a prerequisite for complex language, this cognitively general mechanism may have contributed to the initial evolution of the language faculty.

Key words: categorization; phoneme merger; computational simulation; spandrel; contrast.

1. Introduction

The perception of difference is the basis of categorization. While many familiar categories are differentiated through stable, external features of the world (e.g. *rock*, *tree*), behavioral categories are at least partially grounded in learned, culturally produced and transmitted features. Categories based on these learned features differ from externally grounded categories in a functionally interesting way because any biases in the way such a behavior is perceived, learned, or reproduced can influence its form, allowing the behavior itself to change over time (Griffiths and Kalish 2007; Kirby et al. 2015). A classic example of the way a behavioral object can change in this way is the Telephone Game, in which one person whispers a phrase into the ear of another, who in turn whispers what he or she heard into the ear of the next person, and so on. After the phrase has passed through a number of people, it often is found to have changed into a quite different, but still intelligible phrase. Two factors contribute to this outcome: the use of a whisper adds noise that degrades transmission accuracy, and the bias on the part of the listener to map the degraded percept to some grammatical sequence of real words recreates an interpretable phrase. The two together create rapid change that remains nonetheless within a space constrained by the language. More generally, feedback over cycles of perception, categorization, and reproduction results in a self-referential system that can evolve through noise or bias anywhere in that cycle (Boyd and Richerson 1985; Henrich and Boyd 2002; Brighton, Smith and Kirby 2005; Kirby, Dowman and Griffiths 2007). Language is an example of this kind of categorial system (Andersen 1973; Lindblom et al. 1984; Ohala 1989; Labov 1994; Hurford 1999; Kirby 1999; Bybee 2001; Blevins 2004; Beckner et al. 2009).

A body of research across multiple domains shows that the act of perceiving can 'retune' perceptual categories, influencing future categorization behavior (reviewed in Goldstone 1998). In the domain of language, the act of mapping a speech percept to a category not only influences future categorization behaviour, but also production from that category (reviewed in Pisoni and Levi 2007). Although there appears to be contextual constraints on this phenomenon (see e.g. Nielsen 2011), it can be robustly identified in experimental studies. For example, if listeners are prompted to identify a pronunciation variant with a particular word, they are more likely to spontaneously map similar pronunciations to that word in the future (Norris et al. 2003; Maye et al. 2008). At the same time, if listeners are exposed to a variant pronunciation, they are more likely to produce that variant in the future (Goldinger 2000; Pardo 2006; Nielsen 2011; Levi 2015).

Early on, Pierrehumbert (2001) recognized that a consequence of this 'perception–production feedback loop' is that when two categories are similar enough that productions from those categories are potentially confusable, the category distinction itself should weaken and eventually collapse, resulting in the loss of one of the categories. In fact, distinct sound categories do often merge during the course of language change in a way that is consistent with Pierrehumbert's model predictions. For example, in many dialects of American English, the vowel category / ɔ / has merged with the category / ɑ /, such that the originally distinct pronunciation of *caught* ([kɔt]) is now indistinguishable from that of *cot* ([kɑt]; Labov, Ash and Boberg 2006).

However, despite the fact that all sound categories are to some extent confusable with perceptually adjacent categories, sound categories in a language never merge to the extent that word distinctions can no longer be efficiently encoded. What keeps this from happening? Some theories of language have posited that sound category contrast is directly monitored and preserved when necessary (e.g. Martinet 1955; Flemming 1995), but more recent explanatory forays have looked for mechanistically local, less teleological explanations for the maintenance of contrastive sound categories in language (e.g. Berrah et al. 1996; Browman and Goldstein 2000; de Boer 2001; Wedel 2004; Oudeyer 2006; Wedel 2012; Hall et al., manuscript under review). For example, previous work using the same computational model employed here has shown that category contrast can be maintained when the degree to which a percept updates the category system is inversely correlated with ambiguity, where more ambiguous percepts influence the category system more weakly through lower storage probability or weight (Wedel 2006, 2012; Winter and Wedel 2016).

However, previous work has suggested that competition between categories may itself drive maintenance of category distinctions in the absence of any more direct mechanism of selection for contrast (Blevins and Wedel 2009). Our goal here is to explore and explicate this phenomenon. To do this, we will provide both modeling and analytic arguments that the categorization/production cycle itself can promote persistence of signal distinctions between existing categories. First, we show that given the existence of feedback between perception and production, competition between perceptually adjacent categories in their boundary region promotes a decrease in category overlap, promoting maintenance of category distinctiveness over time. We then show two theoretically interesting conditions in which category distinctions are predicted to more easily collapse. First, if productions are nearly always mapped to the intended category by a listener despite perceptual ambiguity, as when context robustly disambiguates categories, there is little competition between the categories and category overlap is less inhibited. Greater overlap can thereby promote category loss if a future generation of learners acquires the overlapped distribution as one category. Second, we show that if production from a category is prompted primarily by factors within the categorial system, such as a category's prior frequency of successful competition for percepts, then distinctions between perceptually adjacent categories are more likely to collapse because one category can more rapidly grow at the expense of the other.

Neither the computational model nor the analytic approach presented here are intended as direct representations of human language with all its attendant properties. Rather, they should be understood as investigations of simple information transmission systems which share relevant properties with language. Our intention here is to show that contrast maintenance between competing categories can be an intrinsic by-product of production/perception feedback, without any direct or indirect selection against ambiguity in production or perception (see Wedel 2012 for additional discussion). We do not argue that this mechanism must in fact contribute significantly to contrast maintenance in present-day human language; nonetheless, in the discussion, we review patterns of language change that are consistent with this mechanism. Further, we note that given evidence that perception/production feedback is a basic property of cognition (reviewed in Oudeyer 2002), this property may have served as a spandrel (Gould and Lewontin 1979) for the initial evolution of the human language faculty by providing a pre-existing feature necessary for a stable communication system (Oudeyer 2006).

2. Category competition promotes category contrast

Mapping a percept to a category label requires comparing the percept in some fashion to the contents addressed by that category label. Many models of speech recognition characterize this process of comparison in terms of a competition between categories on the basis of their similarity of their contents to the object (reviewed in Frauenfelder and Floccia 1998; Gaskell and Marslen-Wilson 2002; Pierrehumbert 2006; McQueen 2007). The competition model accounts for the observation that ambiguous percepts are categorized both more slowly and more variably: when multiple categories are initially activated by a percept, it takes longer for the competition to produce a winner, and it is also more probable that a 'wrong' category will win (Luce and Pisoni 1998). Here we show how competition can also promote maintenance of contrast between the contents of competing categories within a production/perception feedback loop.

We will model category contents as a set of *exemplars*, that is, percepts previously mapped to that category (Hintzman 1986; Nosofsky 1988; Walsh et al. 2010). As a visual representation of how competition can promote sharper category boundaries if categories update with experience, imagine a set of exemplars divided between two category labels (Fig. 1). Each exemplar represents a percept that was previously mapped to the category label, and we imagine an initial condition in which there is some degree of category overlap. If the two categories compete for a new percept on the basis of the percept's similarity to their stored exemplars, that percept will be more likely to be assigned as a new exemplar to the category that contains a denser set of exemplars in the same region of perceptual space as the percept (Hintzman 1986; Nosofsky 1988; Goldinger 1996). The new exemplar resulting from mapping the percept to the winning category strengthens the connection of that category label to that region of the perceptual dimension, with the consequence that any new percept in that region will be yet *more* likely to be assigned to that category. Provided that there is some mechanism for information turnover, such as memory decay (e.g. Pierrehumbert 2001; Wedel 2006; Beckner and Wedel 2009, see also Spike et al. 2013), this competition-based positive feedback will result in a relative reduction in the degree of category overlap (see Section 2 and the Appendix 2 for discussion of factors that influence the abruptness of the overlap boundary).

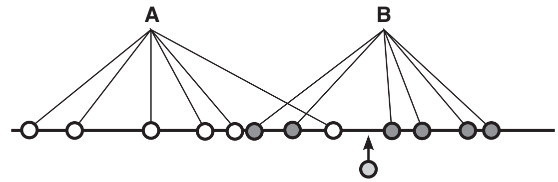

Figure 1. New percept more likely to be categorized as B.

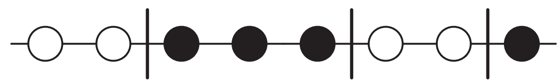

Figure 2. In a one-dimensional model, the category label of each element can be reconstructed if we know the boundaries between domains of like category labels, marked in this example with vertical lines. As a consequence, the evolution of the system can be described in terms of movement of these domain boundaries rather than change in the individual element category labels.

For the purposes of this argument, we assume that each percept, no matter its form, is mapped to and updates one category. Previous work using the same general computational model (Wedel 2006, 2012; Blevins and Wedel 2009; Winter and Wedel 2016) has explored a more direct mechanism for contrast maintenance in which the degree to which a percept influences the category system is correlated with its perceptual ambiguity, where more ambiguous percepts influence the category system more weakly through lower storage probability and thereby lower influence on the trajectory of change in the system. Here, we show that even when the perceptual properties of a percept do not influence its ability to update the category system *per se*, category competition contributes to reduction in category overlap.

The division of the dimensional space between categories that compete in this way is an example of a more general phenomenon known as *coarsening*, which occurs in systems with positive feedback between neighboring elements (Ratke and Voorhes 2002). We can illustrate coarsening in more detail in a simple one-dimensional lattice made up of black and white elements where in each time step an element may change its color (Fig. 2). We are interested in cases in which a given system element interacts with some other elements (its 'neighbors') to a greater degree than others. In the simple one-dimensional case above, we define the neighbors of any given element as the immediately adjacent elements. If we say that elements only have some positive probability of changing their color if they fail to match in color with a neighbor, regions of consistently black or white elements (which we may refer to as *domains*) will

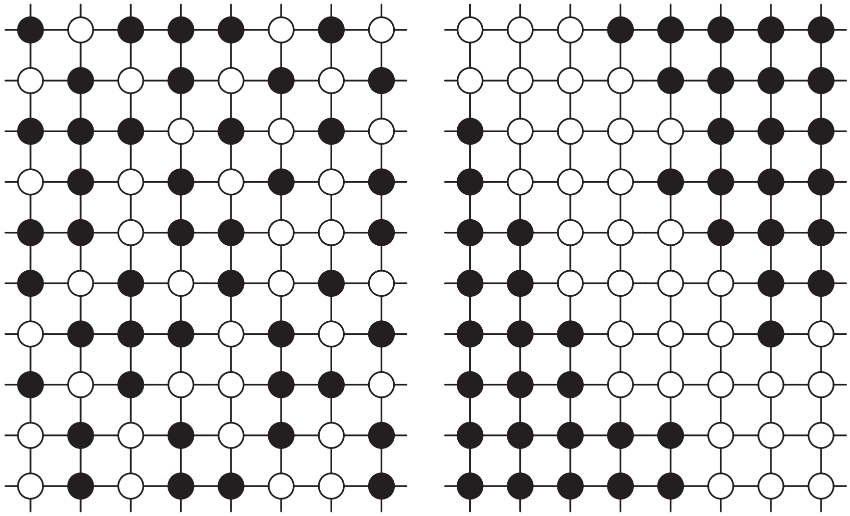

Figure 3. The coloring in the left lattice is highly irregular and neighboring elements are likely to have different colors. We say that such configurations are fine. Likewise, the lattice on the right is coarser, in that neighboring elements often share their color. For many types of lattices the average domain size is a suitable measure of the coarseness of the configuration.

grow at the expense of mixed regions (Wu 1982; Ben-Naim et al. 1996; Derrida and Zeitak 1996; Fatkullin and Vanden-Eijnden 2003). This occurs because color change only occurs at points in the lattice where black and white elements are next to one another, with the result that change is frequent in mixed regions and absent within consistent regions. Whenever some domain of a given color disappears, so does its boundary, thus, the overall system steadily evolves to minimize the number of boundaries between different domains. In a finite lattice, eventually a single domain of one color will occupy the entire space (for a computational example, see Wedel 2012).

This approach generalizes to category spaces of higher dimensionality and to arbitrarily fine-grained representations of category contents. 'Coarse' versus 'fine' structures are compared in a two-dimensional lattice in Fig. 3. Observe that the distinguishing feature of the coarse configuration is the smaller total boundary length between domains with different values.[1] In general, any local interactions that promote a local reduction in the degree of contact between unlike elements will promote global coarsening, and vice versa. See Appendix 3 for additional discussion. We will argue below that within a production/perception feedback loop, competition-based coarsening provides a plausible account for the maintenance of boundaries between competing categories.

2.1 Noise versus entrenchment

As reviewed above, experienced detail influences future production and categorization behavior. This has an interesting ramification for the evolution of categories under the influence of noise. Given that some level of experienced detail is retained in memory and influences future behavior, random noise in production and perception processes will continually introduce new variants into memory, promoting broadening of the category over time (see Pierrehumbert (2001) and Wedel

1 *Conservative* interactions, that is, those that favor change toward that of neighbors, promote growth of domains of the same value, whereas *rebellious* interactions, in which change is more probable away from neighbors, promote mixing and a finer structure. The relative contribution of conservative versus rebellious interactions in a system determines the average size of a domain of like neighbors in a system. A typical domain size exists at a given probability of conservative and rebellious interactions because the number of opportunities for conservative change declines as a pattern becomes coarser whereas the number of opportunities for rebellious changes increases (and vice versa as a pattern gets finer). The average domain scale at which there are approximately equal numbers of conservative and rebellious interactions is its equilibrium domain scale. When such a system is taken out of equilibrium by some short-term disturbance it will return to equilibrium, restoring its typical domain structure.

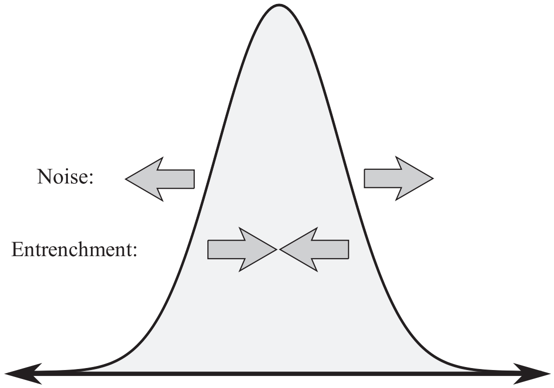

Figure 4. Noise versus entrenchment in perception and production.

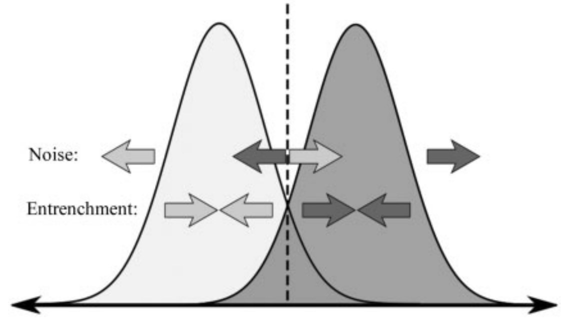

Figure 5. Variant trading between adjacent, competing categories. When two categories are adjacent, the balance of noise and entrenchment is no longer equivalent throughout the category distributions. While entrenchment continues to promote symmetrical tightening of each category, noise no longer has the same broadening tendency at the boundary region between the categories because some variants are assigned to the adjacent category. This *variant trading* diminishes the balancing effect of noise on the boundary-adjacent sides of each category, with the result that the net category center movement over time is away from the boundary.

(2006) for additional discussion). However, categories do not in fact steadily broaden over time. This is consistent with a broad range of evidence that mechanisms of production and perception can themselves promote tightening, or *entrenchment* of individual category distributions (reviewed in Oudeyer 2002), counteracting the broadening effect of noise.

Noise (upper arrows in Fig. 4) creates variation in the course of production or perception. Given some degree of retention of detail in perception and storage, noise promotes broadening of the range of variants stored in a category, abstractly represented here as a distribution in one dimension. The effect of any process promoting entrenchment (lower arrows in Fig. 4) in production or perception processes promotes tightening of the distribution over time. In the absence of other factors, the variance of a category should reflect an equilibrium between noise and entrenchment processes. If the sources of noise and entrenchment are unbiased throughout the category distribution, drift of the category center along the perceptual dimension will not be biased in any particular direction.

2.2 Variant trading between competing categories

However, the equilibrium between noise and entrenchment can be disturbed when two categories approach one another along some shared dimension. In this case, variation at the outer edges of the shared dimension still contributes to broadening to the same degree, but variation in the boundary region does not. This happens because extreme variants of a category that approach another category may be perceived as members of that opposing category (Guy 1996). This process of *variant trading* between the two categories disrupts the equilibrium between noise and entrenchment, resulting in a bias in movement of both category centers away from the zone of greatest competition (Fig. 5). The tendency for categories to drift away from one another through variant trading is equivalent to a mathematical random walk with a wall. In a random walk, the walker's most likely position remains centered at the starting point even though its positional probability cloud broadens with time. However, if a wall is introduced that prevents the walker from continuing beyond some point in one direction, the walker's most likely position drifts steadily away from the wall. This drift is driven by the fact that when the walker is adjacent to the wall, turns toward the wall are blocked, while turns away from the wall are not. The result is a drift in average position away from the wall, where the rate of drift diminishes as the wall recedes into the distance. In the case under discussion here, the boundary region between the two categories is equivalent to the wall in the random walk: all extreme variants can move the category center away from the boundary, but not all extreme variants can move it toward the boundary, because some of those variants are assigned to the opposing category. If we assume for illustration that the same number of variants are traded in both directions across the boundary, these traded variants are equivalent to random walkers who 'bounce' off a wall. This variant trading results in steepening of the category distributions in the boundary region. The resulting distributional skew allows entrenchment within the category to pull the centers of the categories away from the region of category competition (see Winter and Wedel (2016) for related discussion of category skewing through competition and its effect on category system evolution).

The effect of variant trading on category contrast can be illustrated with a simple computational simulation of the evolution of two categories. Each category comprises a set of mappings to exemplars represented as points with a one-dimensional space with arbitrary values from 0 to 100, where each exemplar corresponds to a percept that has been previously identified with that category (note that both categories can contain exemplars at the same point along the perceptual dimension). In each round of the simulation, an output is produced from each category. This output is derived on a randomly chosen target exemplar from the category (Pierrehumbert 2001), which is biased toward a weighted-average based on the target's distance to other exemplars for that category, serving as a model of entrenchment (e.g. Hintzman 1986; see Appendix 1 for details). In addition, a small amount of additional random noise is added to the output to introduce novel variation. The simulation then treats the output as a percept, which is probabilistically categorized on the basis of similarity to existing exemplars in each category. This percept is then stored as a new exemplar in the winning category. In this model, the interaction of three factors promotes maintenance of category contrast (1) the stochastic addition of new, occasionally more extreme variants through noise, (2) the countervailing reduction of within-category variation through entrenchment processes, and (3) variant trading over the region of category competition. See the Appendix 1 for a more detailed description of the model architecture; for additional examples of the use of this general model architecture to explore linkage between low-level phonetic variation and larger scale sound change through feedback between perception and production (Wedel 2006; 2012; Blevins and Wedel 2009; Winter and Wedel 2016).

As we will see below, in this simple model of the effects of noise versus entrenchment within a production/perception loop, variant trading promotes coarsening, such that the two categories evolve to reduce their degree of overlap. A central purpose of this article is to show that this model architecture produces stably contrastive categories when the probability of production from categories is influenced by external factors, such as a drive to communicate. Below, we show the behavior of the model when each category produces an output in each round. A meaningfully different case, in which probability of production is internally controlled by frequency of previous category use will be discussed in Section 2.4.

Figure 6 shows a starting point for a run of the simulation in which two categories (A, B) both start with

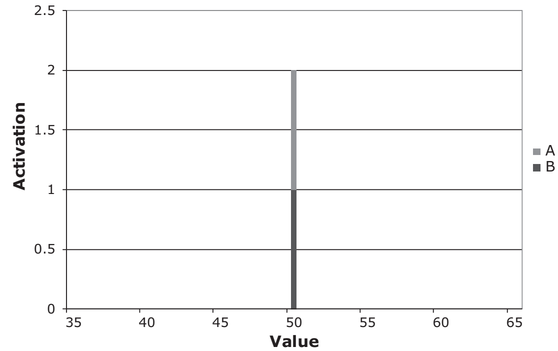

Figure 6. Initial weights for two categories (A, B) at point 50.

exemplars at the same point in the center of the perceptual space, that is, at a value of 50. Recall that in each cycle, an output from a category is produced from its set of exemplars with the addition of noise, and then assigned to the most similar category. In the first cycle in this example, each category will produce outputs with values near 50. Because each category is identical at this point however, assignment of this first set of outputs is equiprobable to category label A or B. However, as soon as some difference in the values of exemplars in each category is introduced through random noise, their ability to compete for percepts also becomes different. For example, if 49 is produced and assigned to category A in the first round, then the values in category A are now lower on average than those for category B. Subsequent productions from either category less than 50 are now more likely to be assigned to A, and productions that are greater than 50 are more likely to be assigned to B, and so on.

As noise in production introduces variation, the distribution of mappings to points in the perceptual space grows, but the feedback between previous and future categorization behavior ensures that the two categories reduce their degree of overlap (Fig. 7). Given the tendency toward entrenchment in each category, however, the categories tend to do more than simply spread out while splitting the dimension: each tends to consolidate, producing a bimodal distribution over the perceptual space (Fig. 8). (The probability with which categories remain adjacent versus drift apart depends on the relative strength of the noise versus entrenchment factors; see the Appendix 2 for discussion of the boundary properties of the model.)

We can demonstrate that this behavior is due to variant trading at the category boundary by modifying the simulation to eliminate category competition for percepts by always assigning outputs as new exemplars in

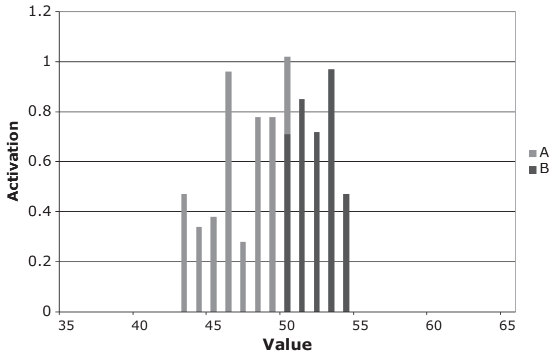

Figure 7. Distribution of weights after 20 rounds with category competition.

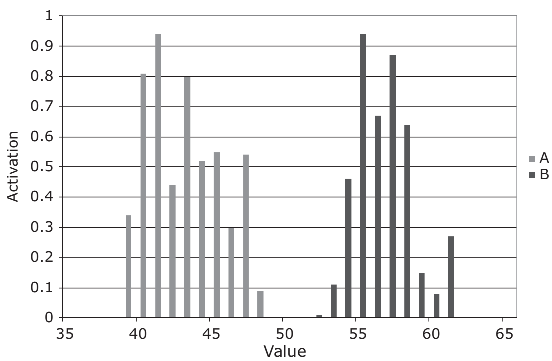

Figure 8. Distribution of weights after 1,000 rounds with category competition.

their parent category, no matter what their value. This eliminates variant trading: an extreme output that will still be stored in its parent category even if it happens to be more similar to exemplars in the alternative category. In this case, the two categories evolve independently of one another. Figure 9 shows the state of an example simulation without category competition at round 1,000. In the absence of variant trading, there is no mechanism promoting coarsening, and so there is no tendency for the two categories to split the distribution. (Note that within the model, this is distinct from category merger per se because the categories remain functionally independent.)

For further illustration, Figs. 10 and 11 compare the longitudinal trajectories of the two categories over 2,000 rounds of simulations in which categories compete for outputs, or do not compete, respectively.[2] In the

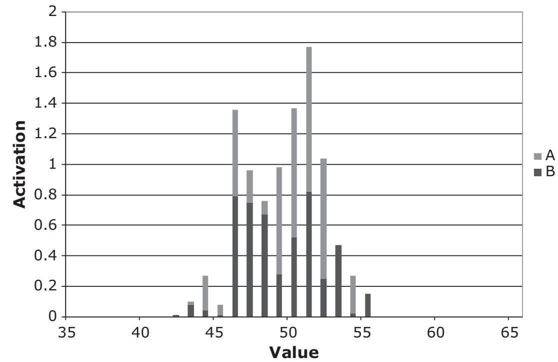

Figure 9. Distribution of weights after 1000 rounds without category competition.

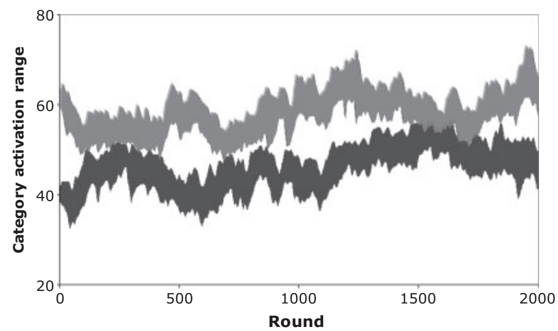

Figure 10. Evolution of two categories with category competition for outputs.

simulation shown in Fig. 10, variant trading prevents the two categories from occupying the same space, even though the categories frequently approach one another.

In the simulation shown in Fig. 11 on the other hand, outputs are always assigned to the intended category, that is, their category of origin, just as in the simulation output at a single point shown in Fig. 9. Variant trading does not occur and so the categories in the computational model can evolve to map to similar regions of the space.

2.3 External grounding of competing categories promotes stability

A wide range of experimental work shows that activation of a category through use, whether in production or perception, makes subsequent similar use more likely. In production, for example, many studies have shown that if speakers perceive a particular syntactic structure,

2 Both simulations are initialized with activations at 40 for category A and 60 for category B. The respective bandwidths represent the length along the dimension containing 90% of the total exemplar distribution in each category. Band edges were smoothed by averaging over a window of ±5 rounds from each point.

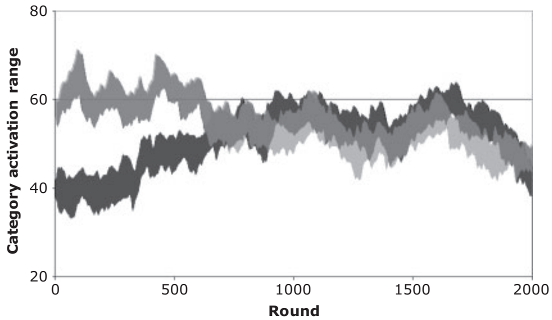

Figure 11. Evolution of two categories without category competition for outputs.

they are more likely to use that structure themselves in their subsequent productions (e.g. Smith and Wheeldon 2001). Priming is evident as well in production and perception of fine phonetic details: mapping a perceived phonetic variant to a given category by a listener makes it both more likely that the listener will map a similar variant to that category in the future (e.g., Norris et al. 2003; Kraljic and Samuel 2006; Maye et al. 2008), and also that spoken production from that category will be biased toward that variant (e.g. Goldinger 2000; Pardo 2006).

Within a categorization/production feedback loop model, priming introduces instability into the system, because activation of a category through use increases its probability of use in the future, both by increasing probability of production from that category, and through increasing the category's competitive success in perception. When categories are in competition in either production or perception, this positive feedback can promote the runaway increase in activation of one of the categories, resulting in the eventual allocation all memory to one category to the exclusion of all others that compete with it. As an example, in Pierrehumbert's (2001, 2002) exemplar-based model of lenition and merger of sound categories, priming effects are implemented both in production and categorization. Specifically, both the production frequency from a category, and its strength in competing for percepts was determined by the aggregate recency of exemplars in that category. As a result, a category that is recently successful in competing for percepts is both more likely to produce an output in the future and also to be more successful in competition in categorization. Both of these general features of this type of model are consistent with priming studies as noted above. As a consequence of this feedback, once categories approach one another and begin to significantly compete for an overlapping set of category outputs, one category tends to gain the upper hand and eventually extinguish the other. As we will argue below, this may be an appropriate model for the sound category mergers that often occur in language change. As noted above, the phoneme category / ɔ / has merged into the category / ɑ / in western varieties of North American English, such that the originally distinct pronunciations for *cot* ([kɑt]) and *caught* ([kɔt]) are now the same.

Competing categories do not always drift together and merge, however, and sometimes even seem to actively preserve mutual contrast (reviewed in Wedel 2012). The model developed here provides a potential account for this resistance, suggesting that preservation of contrast between competing categories is correlated both with variant-trading, and the degree to which use of those categories is influenced by external events rather than their own activation, for example, through their history of use. Recall that in the simulations presented in Figs. 8 and 9 above, an output was produced from each category in every round, without reference to how often percepts had been mapped to that category in the past.

When we modify our simulation to allow categories to compete in production as well as perception by making production probability contingent on past competitive success in categorization, we find a markedly different result. Just as in Pierrehumbert's (2001, 2002) simulations, when two categories approach one another within the perceptual space, one eventually gains the upper hand and extinguishes the other. In the simulation shown in Fig. 12 below, two outputs were produced and categorized in each round as before, but the category under production was randomly chosen in proportion to its activation, operationalized as the aggregate recency of its stored exemplars. In this case, when the two categories approach one another, whichever category happens to be initially more successful at competing for ambiguous percepts, gains activation at the expense of the other. Any initial imbalance allows the more active category to increasingly poach exemplars that would otherwise have been mapped to the other category, eventually resulting in drift to extinction of the less active category (see Pierrehumbert 2001; Oudeyer 2006: 80). As discussed above, the extinction of one phoneme category in favor of a perceptually adjacent one is a general phenomenon in language change.

However, anything that renders the probability of production from a category at least partially independent of the activation of the category inhibits this process of extinction-through-competition (cf. Figs 8 and 9 above). Lexical categories clearly have this property:

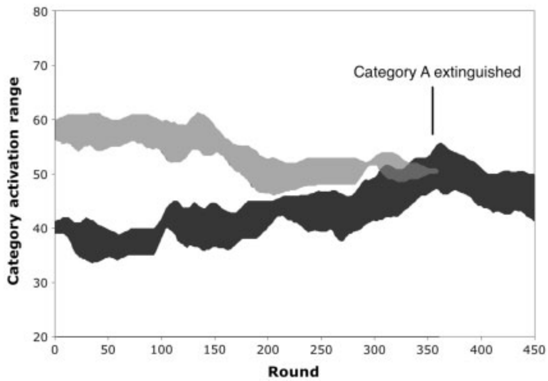

Figure 12. Category loss through feedback and competition.

although the probability that a word will be used may be influenced by the frequency of its prior use, it is also influenced by the current communicative context. In this way, external factors can interrupt the feedback loop between past and future production by providing context-driven boosts to low-frequency categories. The contribution of communicative function to the frequency of use should provide a measure of protection to lexical categories from the kind of feedback-driven loss of categories that we see in an activation-driven production model. In the following section, we provide analytical support for this conclusion.

2.4 An analytic approach to comparing externally-driven versus activation-driven production

While the model we have used in Sections 2.3 and 2.4 above is well-motivated psycholinguistically, its multiple interacting feedback loops make its behavior difficult to describe analytically. For this reason, it is useful to consider a simpler system which captures the essential properties of the model system in a more analytically tractable fashion. In this section, we develop a partial model of the system that allows us to compare the relative stabilities of two competing categories when production rate from each category is (1) externally determined or (2) determined by category activation. A key to making our model analytically tractable is the removal of category competition in perception by storing every output back into the category it came from. This is parallel to the controls run without category competition in the full model above (cf. Figs 9 and 11). In this case, the only avenue for competition between two categories is in the rate of production. At the end of this section we will discuss the contribution of perceptual category competition in the context of the analytical results developed here.

We start with two categories comprising category labels mapped to exemplars in memory. Each category is characterized by a relative activation level which is given simply as the number of exemplars currently mapped to that category label divided by the total number of exemplars in memory N. If n is the number of exemplars mapped to the first category label, the activation of the first category can be characterized by a number $x = n/N$, $0 \leq x \leq 1$, where when $x = 0$ the first category label is mapped to zero exemplars and is completely inactive and the second category is fully active, and vice versa. In each time step an output is produced from one of the two categories and stored back into the source category as a new exemplar, and then a random exemplar is deleted from the system (i.e. from a random category) to keep the total number of exemplars constant. Production and recategorization of an output in the context of random exemplar loss allows the number of exemplar mappings to category labels to change over time. Consider two extreme scenarios for output production probabilities:

- *Activation-driven production.* The probability of producing an output from each category is given by its current activation level, that is, x for the first category and $1 - x$ for the second.
- *Externally-driven production.* The probability of producing an output from each category is independent of its activation level and is set here to be equal to 1/2.

Evolution of our system is described by a random walk where in each time step, the quantity x may remain the same, increase by $\Delta x = 1/N$, or decrease by the same value of Δx with the following probabilities:

- *Activation-driven production*

$$\begin{aligned} \mathrm{P}\{x \mapsto x + \Delta x\} &= x(1 - x), \\ \mathrm{P}\{x \mapsto x - \Delta x\} &= x(1 - x). \end{aligned} \tag{1}$$

- *Externally-driven production*

$$\mathrm{P}\{x \mapsto x + \Delta x\} = \frac{1 - x}{2}, \qquad \mathrm{P}\{x \mapsto x - \Delta x\} = \frac{x}{2}. \tag{2}$$

Here, the symbol $\mathrm{P}\{\cdot\}$ denotes the probability of respective change. In the first case, the probability that the number of exemplars in the first category increases is the product of probabilities that (i) a new exemplar is categorized as belonging to the first category (probability is $x = n/N$) and (2) the deleted exemplar belongs to the second category (probability is $1 - x = 1 - n/N$). Note that this scenario is known in genetic and evolutionary biology as the Moran model (Moran 1958). In the second case, the probability that the number of exemplars in the first category increases is the product of probabilities that (1) a new exemplar is categorized as belonging to the first category (probability is 1/2) and (2) the deleted exemplar belongs to the second category (probability is same as the first case, i.e. $1 - x = 1 - n/N$). For both cases, over the course of time the value of x may reach 0 or 1. If x hits either of these boundaries, one of the categories becomes inactive. The time at which this happens is known as the exit time (or the fixation time), and we stop the evolution of the model at that point. Our goal now is to compare the expected exit times corresponding to these two scenarios. In the following, we will see that the expected exit time for the externally-driven scenario is much longer than that for the activation-driven scenario. In other words, a constant probability of production from a category promotes the stable coexistence of categories that share memory resources as in this model.

The exit times for these models (or their superposition) may be computed exactly using the Markov chain methods; however, the general expressions are rather complicated and not particularly enlightening. In order to get a clearer analytical insight, it is convenient to work with the scaling limits of the random walk for x, that is, to treat x as a continuous quantity (rather than a quantity changing by discrete steps of $\Delta x = 1/N$). This corresponds to an assumption that the total number of exemplars, N, is large, so that Δx is much less than 1 (in practice, $\Delta x = 0.1$ is already sufficiently small). We then also say that each time step takes Δt of 'real time'. Finally, let us use $x(t)$ to denote the activation level of the first category after $t/\Delta t$ time steps. Now let us consider the scaling limits for each scenario in more detail.

An appropriate scaling limit for activation-driven production is to send Δx and Δt to 0 while keeping the ratio $D = \Delta x^2/\Delta t$ fixed (this ratio is called the diffusion coefficient).[3] The expected exit time may be found using the methods described, for example, in Gardiner (1985) and is given by:

$$T \approx -N^2\,(x_0 \ln x_0 + (1 - x_0)\ln(1 - x_0)). \qquad (3)$$

where x_0 is the initial value for $x(t)$. This formula relates the expected exit time T (measured in the number of time steps) to the total number of exemplars, N, and the initial activation level of the first category, x_0.[4]

Externally-driven production must be treated a bit differently. The principal distinction is that now there exists a constant average force that pushes the category activation level toward its equilibrium state, $x = 1/2$. This force is deterministic and its presence requires a different scaling limit, when a different ratio, $V = \Delta x/\Delta t$ is fixed. This implies that the diffusion coefficient, D, vanishes in the limit since $D = V\Delta x \to 0$ as $\Delta x \to 0$. In particular, this implies that in the scaling limit the system will never hit the boundary and will simply relax to the state when both categories are equally active. In practice, however, we formally keep the scaling limit D as a small quantity rather than setting it to zero.[5] The general formula is somewhat cumbersome, so we only present a particular case when $x_0 = 1/2$ (i.e. initially both categories are equally active):

$$T \approx \sqrt{\pi N/2}\, e^{N/2}, \quad \text{for} \quad x_0 = 1/2. \qquad (4)$$

The exit time here is again given in terms of the number of steps in the discrete random walk model. The essential observation is that the exit time is exponentially long with respect to the number of elements in the system, N. Since in the activation-driven production scenario, the exit time is only quadratic in N, the exit times for the externally-driven scenario are much longer. The graphs of the expected exit time T for $N = 10$ with different starting x_0 are displayed in Fig. 13.

2.4.1 Externally-driven production promotes stability in the context of category competition

We saw above that competing categories coexist for a longer time when their production is externally-driven rather than activation-driven. In a more general setting where both scenarios occur with some probability, the final result is an interpolation between the two extreme cases considered above. However, it is important to realize that even if the relative contribution of externally-driven production is small, its effects are overwhelming for sufficiently large N, that is, the exit time will immediately become much longer as soon as externally-driven

3 In the scaling limit described above, the function $x(t)$ is the solution of the stochastic differential equation, $dx(t) = \sqrt{2Dx(1-x)}\, dW_t$, where W_t is a standard Brownian motion.

4 The actual computation is easier to carry out in the prescribed scaling limit, which results in $T(x_0) = -(x_0 \ln x_0 + (1 - x_0)\ln(1 - x_0))/D$. Formula (3) is obtained by scaling back to the discrete model.

5 In this case the limiting process is a solution of $dx(t) = V(1/2 - x)\,dt + \sqrt{D/2}\,dW_t$. It may be found explicitly: $x(t) = \frac{1}{2} + e^{-Vt}(x_0 - 1/2) + \sqrt{D/2}\int_0^t e^{-V(t-s)}\,dW_s$. Observe that here we can see that if $D = 0$, the second, stochastic term vanishes and the evolution is purely deterministic with $x(t) \to 1/2$ as $t \to \infty$.

production contributes to even a small degree. This effect is due to the fact that externally-driven production introduces a deterministic drift towards the equilibrium $x = 1/2$ state, causing an exponential growth in exit times. This is an important point because externally-driven and activation-driven production represent two extremes, where in fact the probability of production of a lexical category is influenced both by external context and by internal activation (cf. the literature on production priming discussed above). This analysis suggests that provided the activation strength of a category does not fall below some threshold (conceptualized here as a low exemplar count), occasional boosts in activation provided by externally-prompted production may allow an infrequently used category to survive despite a more active neighboring category.

The fuller model of category competition featured in the simulations in Section 2.3 differs from this analytic model in the inclusion of a categorization decision, allowing outputs produced from one category to be mapped to the other in a process of variant trading. Two additional competing effects that influence exit times come into play in this richer model. First, a percept is more likely to be mapped to a more active category, all else being equal. As a result, when two categories are close enough to begin variant trading, the more active category will tend to outcompete the less active category in the boundary region, further reinforcing the asymmetry in activation. This positive feedback loop accelerates the extinction of the less active category beyond that predicted by the analytic model above.

Acting against this process, however, is the fact that the outputs from a less active category that are most dissimilar to the more active category are most likely to be correctly identified by a listener. This bias toward correct categorization of more contrastive outputs acts to shift the distribution of exemplars in the less active category away from the boundary over time. As long as production from the less active category is maintained at some significant level through reference to external properties, this ongoing 'selection' for more contrastive outputs mitigates the acceleration of category extinction under competition. As a consequence, to the extent that production is externally-driven as well as activation-driven, a less active category is more likely to shift away from a more active category rather than be swallowed up. A clear prediction of this model is therefore that category merger will be inhibited in relation to both the degree of competition between categories and the degree to which production from those categories is externally prompted. Semantically related words such as those in morphological paradigms represent a clear example from language that meets both these criteria: (1) given that they are semantically very similar, word categories within morphological paradigms compete strongly with one another on the basis of sound contrast; and (2) the probability of production from word categories is externally-driven to some degree, that is, influenced by external communicative context (Blevins and Wedel 2009; Hall et al., manuscript under review).

3. Discussion

Given category updating through production/perception feedback, we have argued that boundaries between competing categories will be preferentially preserved to the extent that the categories are not disambiguated by context:

- Given competing categories with externally-driven production, variant trading promotes maintenance of the boundary between the categories (Section 2.2, Figs 8 and 10).
- If variant trading between competing categories is prevented, as when perceptually similar outputs are disambiguated by context, erosion of the category boundary is not inhibited (Section 2.2, Figs 9 and 11).
- If production from competing categories is solely activation-driven, as in the priming model described in Pierrehumbert (2001), variant trading promotes extinction of the less active category by allowing growth of one category at the expense of the other (Section 2.3, Fig. 13).

Evidence from patterns of phoneme category merger are consistent with these predictions. Wedel et al. (2013a,b) provided statistical evidence that in a diverse set of languages, the larger the number of words that are distinguished by two phoneme categories, the more likely that the contrast between the phoneme categories will be preserved over the course of language change. How are these findings consistent with the variant trading mechanism of category contrast preservation? We know that ambiguous phoneme category percepts are likely to be identified by a listener as the intended phoneme if they are in an unambiguous lexical context (Norris et al. 2003; Maye et al. 2008)). For example, a fricative sound that is ambiguous between / s / and / f / will tend to be interpreted by a listener as an / s / if it occurs in the context '__ilk', because *silk* is a word, but *filk* is not. Conversely, the same ambiguous sound will tend to be interpreted as an / f / if it occurs in the context '__ive', because *five* is a word, but *sive* is not. The disambiguation of ambiguous phoneme category

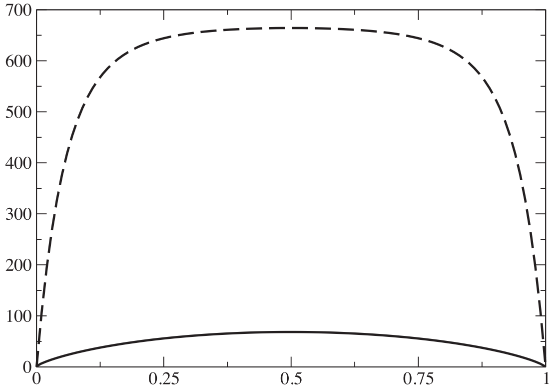

Figure 13. Expected exit times $T(x_0)$ for the activation-driven output production scenario (solid line), and for the externally-driven scenario (dashed line). T is measured in the number of steps in the random walk model, and x_0 is the initial relative activation level of the first category. The total number of exemplars is $N = 10$. Note that as the value of N grows, the exit times for the externally-driven scenario become increasingly long relative to those for the activation-driven scenario.

productions through local context corresponds to the *no-variant trading* case, in which category merger is not inhibited (Section 2.2, Figs 9 and 11). Conversely, we predict that variant trading is more likely to occur when an ambiguous percept for a phoneme is not distinguished by the immediate lexical context, such as when that phoneme distinction is the primary cue distinguishing the intended lexical item from another. For example, an ambiguous fricative between / s / and / f / is not distinguished by the context '__it', because both *sit* and *fit* are existing words of English.

Consistent with a role for local context, Wedel et al. (2013b) found that the number of minimal pairs which share word category (e.g. minimal pairs that are both nouns, both verbs, etc.) are significantly *more* predictive of merger probability than the number of minimal pairs that are of different word category. Because immediate sentence context tends to disambiguate word category membership, words in different categories are less confusable than words in the same category. Continuing our example, the minimal pair *sat˜fat* is less confusable in sentence context than *sit ˜ fit*, because *sat* is a verb, and *fat* is an adjective, while both *sit* and *fit* can be verbs, and can therefore be more confusable in context. The correlation between a greater number of same-category minimal pairs distinguished by some phoneme pair and a lower probability of phoneme category merger is consistent with the mechanism proposed here, because same-category minimal pairs provide more opportunity for categorization error and concomitant variant trading between competing phoneme categories. Above, we showed that given the possibility of variant trading, even a relatively small degree of external grounding for production of categories may work to inhibit category merger. This is a welcome result, since most instances of phoneme categories in lexicons do not distinguish minimal pairs (e.g. the / b / in 'badge' does not contrast with / p / to distinguish this word from the non-word 'padge'). In the crosslinguistic data analyzed by Wedel et al (2013b), mergers were not found between phoneme pairs which distinguished more than about thirty same-category minimal pairs. This finding may contribute some empirical basis for future hypothesis building about the way that category competition can contribute to contrast maintenance.

In contrast, in the absence of variant trading, frequency of use should potentiate merger by increasing the effective rate of change in the system. Wedel et al. (2013a) reported that merger probability for phoneme pairs which did *not* distinguish minimal pairs was positively correlated with greater frequency of the words they appeared in. This finding is consistent with any model of language change, including that proposed here, in which evolutionary change proceeds through variation introduced by usage (see e.g. Bybee 2006).

Although we have argued here that the effects of variant trading in a production/perception feedback loop on category evolution are consistent with patterns of category contrast maintenance in language change, we do not argue that it is the only mechanism for contrast maintenance, or even that it significantly contributes to contrast maintenance in current human languages. A number of other mutually compatible mechanisms for category contrast maintenance have been proposed (reviewed in Baese-Berk and Goldrick 2009; Buz and Jaeger 2016), and selection against ambiguous percepts as a possible mechanism to promote contrast maintenance has been explored in previous simulations using the same computational architecture (Wedel 2006, 2012; see also Hay et al. 2015). Rather, our intention has been to argue that the category-sharpening effect of variant trading is a basic property that comes for free given the production/perception feedback loops within competing categories. Under the assumption that basic social cognition and learning mechanisms predate language, this raises the possibility that the ability to develop and maintain a large system of distinct linguistic categories was already substantially in place at the time that the language faculty itself was evolving. In this way, the variant trading inherent in category competition may have served as a spandrel (Gould and Lewontin 1979) allowing the growth of a large set of distinct linguistic signal/meaning categories (for further discussion of this point, see Oudeyer 2006). This pre-existing faculty may have in turn influenced the evolutionary pathway taken by the modern language faculty (cf. Christiansen and Chater (2008) for arguments that language is adapted to the brain).

The existence of a sublexical, phonemic level of organization appears to be a property of all modern languages (Lindblom 2000; Studdert-Kennedy and Goldstein 2003). Given human limitations in producing and perceiving arbitrarily fine-grained articulations, this multilevel organization of signals is arguably a functional prerequisite for the very large vocabulary sizes typical of human languages (Hockett 1960; reviewed in Ladd 2013). We reviewed evidence above that in existing languages, variant trading of ambiguous word pronunciations may contribute to maintenance of phoneme category distinctions. In the early stages of language evolution, this mechanism may likewise have provided a pathway for development of multilevel linguistic signal systems in which meaningful words are built up from a set of contrastive, but individually meaningless phoneme categories.

Conflict of interest statement. None declared.

Appendix 1

Simulation architecture

This model architecture is conceptually based on that described in Pierrehumbert (2001, 2002) and further elaborated in Wedel (2006, 2012), Blevins and Wedel (2009), Winter and Wedel (2016). The dimensional space within the model is a set of points along a continuum with values from 0 to 100. Each category comprises a set of mappings to points corresponding to percepts that have been previously identified with that category; in the model runs shown in this article, there are always two categories. The mappings have weights varying between 0 and 1, and each mapping weight is exponentially decremented in each round of the simulation by division by a constant (here, 5) in a model of memory decay (Pierrehumbert 2001). In this case, for example, an exemplar in memory that is 100 cycles-old has an activation that is approximately two orders of magnitude lower than that of a newly stored exemplar. When an exemplar's activation falls below 0.01, it is deleted from the memory.

Production from a category begins by probabilistically picking one reference point within the space in relation to the mapping weights from that category. A population vector is then produced in relation to this point over all points mapped to that category (Guenther and Gjaja 1996; for an example of this architecture in which categories range over two dimensions rather than one, see Blevins and Wedel 2009). This vector is essentially a weighted average of all points mapped to the category, where both the Euclidean distance from the reference point and mapping strength influence each point's contribution (cf. Nosofsky 1988; see Hintzman 1986; Pierrehumbert 2002; Wedel 2006 for examples of production based on regions of the exemplar cloud as a model of entrenchment). The formula for the vector is given below, where p is the output population vector, y is each position within the perceptual space mapped to the category under production, w_y is the weight of the mapping to that point, x is the reference point chosen as the basis for production, and k is a scaling factor influencing the fall off of the contribution to the population vector of the point y relative to x:

$$p = \frac{\sum_y y w_y \, \mathrm{e}^{-k|x-y|}}{\sum_y w_y \, \mathrm{e}^{-k|x-y|}}. \tag{A.1}$$

The value of k used in the simulations shown here is 0.5. A Gaussian random variable with a SD 2 is added to the output; this serves to introduce variation into the system as it evolves, allowing categories to shift position along the dimension over time. This variation is biased slightly toward the center of the dimensional space, creating a fixed attractor in the system. Note that the bias toward the center of the dimension is not required for the simulation or the arguments based on it: it is included solely to promote category approach and overlap, so that the *avoidance* of category overlap in the presence of variant trading is clearer. This bias is calculated using a parabolic response curve given below, where b is the bias added to the output population vector, p is the output population vector, M is the number of points in the space, and G is a constant; b is subtracted from outputs greater than $M/2$ (here, 50) and added to those below it.

$$b = \frac{(p - M/2)^2}{G}$$

The value of G used in these simulations was 5000, giving a bias toward the center of 0.5 at the edges of the continuum. All else being equal, this bias shifts the distributions of both categories toward the center of the dimension over time, that is, toward 50; see Fig. 11 as an example.

After calculating an output from each category, the simulation turns around and treats the outputs as percepts to be assigned to one of the two categories on the basis of similarity. (Use of multiple agents rather than one does not qualitatively change the results of this simulation with respect to the arguments we make here, so for clarity we just model one agent in conversation with itself.) In these simulations, we use a variant of the Generalized Context Model for categorization (Nosofsky 1988) amended to take mapping weights into account. This is conceptually parallel to the population vector method used to produce outputs. The similarity S_i of a percept to a category i is calculated in terms of the distances from this percept (located at x) to each point y mapped to the category label i weighted by the mapping strength w_y, given in Equation (A.2):

$$S_i = \sum_y w_y e^{-k|x-y|}. \qquad \text{(A.2)}$$

The value of the scaling factor k we used in categorization is 2. A percept is then assigned to the most similar category using the Luce choice rule (Luce 1959) over the similarity scores for each category, where the probability of being assigned to a given category is equal to the similarity to that category divided by the summed similarity to all categories. When a new percept is identified as a member of a category, the weight of the mapping from the category label to the corresponding point is reset to 1. Note that both category labels can contain mappings to the same point along the perceptual dimension.

Appendix 2

Sensitivity of the simulation results to parameter values

The general requirements for the model are as follows:

1. External grounding of categories, prompting category production in a manner that is at least partially independent of category activation.
2. Introduction of variation, here accomplished through noise in production.
3. An updating record of variation in category contents.
4. Category competition on the basis of percept characteristics.
5. Entrenchment.

Otherwise varying the architecture of the model, for example, by increasing the number of dimensions (see Blevins and Wedel 2009), by producing outputs as averages over a square window (Pierrehumbert 2002: 65), or by categorizing percepts simply into the category with the closest mean does not qualitatively alter the model behavior.

The functions governing noise, categorization, and entrenchment in these simulations all have characteristic ranges over which they act determined by the constants used in the formulae described in the main text, and the maintenance of contrast between competing categories is driven by their interaction. When the ranges over which these functions act are quite divergent, the simulations no longer show the behavior illustrated here.

- The average width of categories is determined by the balance between noise and reversion to the mean. If the noise distribution is significantly broader relative to the function governing reversion to the mean, multiple independent peaks can develop even within one category and as a consequence contrast can arise trivially between categories.
- If the range of the similarity function (i.e. the population vector) in categorization is broad relative to the average width of the categories themselves (i.e. if the effective boundary zone between categories is relatively indistinct), significant variant trading between categories occurs even when their means are well separated.
- Because any percept assigned to a category will influence the range of future outputs from that category, a broad range of variable category assignment tends to pull categories together. Such a broad categorization range can be compensated by increasing the ranges of the functions for noise and entrenchment.

In the results presented above, each category output was calculated as a population vector over points in the space mapped to that category. This restriction is not required for the function of the model. Instead, the population vector can be calculated over the entire space, allowing points that are mapped to one category to influence production from another. This cross talk between categories in production promotes coalescence between categories that are tightly adjacent, but can be compensated by increasing the variance in categories relative to the scaling factor k in the population vector calculation shown in Equation (A.1) above.

Finally, the results of these simulations are not dependent on the horizontal architecture used here in which the cycle of production, perception, and updating is continuous within a single set of speakers. Similar results obtain when the cycle of production and perception is organized in unidirectional chains of transmission events across generations (for discussion, see Beckner and Wedel 2009).

Appendix 3

Scaling limits of lattice models

Lattice models for coarsening phenomenon in higher dimensions exhibit similar behavior to the one-dimensional

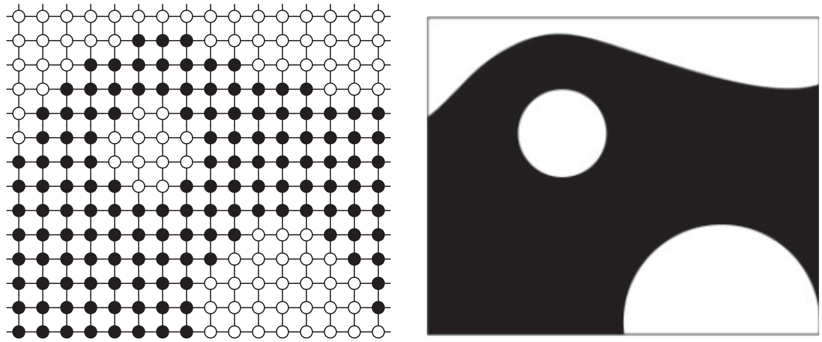

Figure C.1 A discrete lattice model (left panel) may be associated with a continuous model (right panel). In continuous models the central objects are the entire category domains (rather than the individual elements). Temporal evolution is then prescribed via the motion of the domain boundaries.

model described earlier. It is convenient to consider an appropriate *scaling limit*, a mathematical procedure by means of which a discrete model is replaced by a continuous one. For example, on a microscopic level, fluid dynamics is described by the motion of individual molecules. A scaling limit can be applied to arrive at a macroscopic description allowing these dynamics to be described in terms of density and currents (Kipnis and Landim 1998). In our particular context, the scaling limit allows us to consider the entire domains of similar elements rather than the individual elements themselves (Fig. C.1). The temporal evolution of the model may then be described in terms of the motion of the boundaries between the domains. In this way, the domain boundaries can be treated as curves (in two-dimensional models), or surfaces (in higher dimensional models). The motion of these surfaces may be described by their normal velocity (i.e. instantaneous velocity in the direction normal to the surface (Katsoulakis and Souganidis 1995)).

The precise law of this motion will depend on the details of interaction among the elements underlying the movement of the boundary. However, the principal criterion for coarsening remains the same. Namely, if the local law of motion tends to diminish the boundary, then globally the system will coarsen. On the other hand, if the local law of motion tends to increase the boundary, then on a global scale, the system will mix and become finer. For example, the interaction when an element switches to the category label of one of its neighbors (as described above), leads to so-called motion by mean curvature (Katsoulakis and Souganidis 1995), that is, where the normal velocity of the boundary is proportional to its mean curvature. Motion by mean curvature leads to coarsening in a system of an arbitrary dimension. In fact, one can show that in any dimensional model, a spherical domain of radius R will shrink and disappear during the time $T = R^2/2$ (assuming that the normal velocity is precisely equal to the mean curvature of the boundary). This provides a rough explanation for the phenomenon of coarsening: the 'lifetime' of the structures whose typical spatial scale is R is proportional to R^2, thus, the finer structures tend to disappear faster, allowing the larger ones to persist.

Funding

Fatkullin was partially funded by National Science Foundation (NSF) grants DMS-0807332 and DMS-1056471

References

Andersen, H. (1973) 'Abductive and Deductive Change', *Language*, 49: 765–93.

Baese-Berk, M., and Goldrick, M. (2009) 'Mechanisms of Interaction in Speech Production', *Language and Cognitive Processes*, 24: 527–54.

Beckner, C. et al. (2009) 'Language is a Complex Adaptive System', *Language Learning*, 59: 1–26.

——, and Wedel, A. (2009) 'The Roles of Acquisition and Usage in Morphological Change', *In Annual Meeting of the Berkeley Linguistics Society*, vol. 35, pp. 1–12. Berkeley, CA, USA: Berkeley Linguistics Society

Ben-Naim, E., Frachebourg, L., and Krapivsky, P. L. (1996) 'Coarsening and Persistence in the Voter Model', *Physical Review E*, 53: 3078–87.

Berrah, R. et al. (1996) 'From Form to Formation of Phonetic Structures: An Evolutionary Computing Perspective', in *Proceedings of the International Conference on Machine Learning*, Workshop on Evolutionary Computing and Machine Learning, ICML'96, Bari, Italy: Proceedings of Machine Learning Research

Blevins, J., and Andrew, W. (2009) 'Inhibited sound change: An evolutionary approach to lexical competition', *Diachronica*, 26.2: 143–183

——, and Wedel, A. (2009) *Inhibited Sound Change: An Evolutionary Approach to Lexical Competition*. Diachronica.

Boyd, R., and Richerson, P. J. (1985) *Culture and the Evolutionary Process*. Chicago: Chicago University Press.

Brighton, H., Smith, K., and Kirby, S. (2005) 'Language as an Evolutionary System', *Physics of Life Reviews*, 2: 177–226.

Browman, C. P., and Goldstein, L. (2000) 'Competing Constraints on Intergestural Coordination and Self-organization of Phonological Structures', *Bulletin De La Communication Parlee*, 5: 25–34.

Bybee, Joan L. (2006) 'From usage to grammar: The mind's response to repetition', *Language*, 82.4: 711–733.

Bybee, J. L. (2001) *Phonology and Language Use*. Cambridge: Cambridge University Press, 31: 489–509.

Christiansen, M., and Chater, N. (2008) Language as Shaped by the Brain. *Behavioral and Brain Sciences*.

de Boer, B. (2001) *The Origins of Vowel Systems*. Oxford: Oxford University Press.

Derrida, B., and Zeitak, R. (1996) 'Distribution of Domain Sizes in the Zero Temperature Glauber Dynamics of the One-dimensional Potts Model', *Physical Review E*, 54: 2513–25.

Fatkullin, I., and Vanden-Eijnden, E. (2003) 'Statistical Description of Interacting Brownian Walkers on the Line', *Journal of Statistical Physics*, 112: 155–63.

Flemming, E. (1995) 'Auditory Representations in Phonology', PhD Dissertation, UCLA.

Frauenfelder, U. H., and Floccia, C. (1998) 'The Recognition of Spoken Words', in FriedericiA. (ed.) *Language Comprehension: A Biological Perspective*, pp. 1–40. Berlin: Springer.

Gardiner, C. W. (1985) *Handbook of Stochastic Methods for Physics, Chemistry and the Natural Sciences*, 2nd edn. Berlin: Springer.

Gaskell, M. G., and Marslen-Wilson, W. D. (2002) 'Representation and Competition in the Perception of Spoken Words', *Cognitive Psychology*, 45: 220–66.

Georgopoulos, A. P. et al. (1983) 'Spatial Coding of Movement: A Hypothesis Concerning Coding of Movement Direction by Motor Cortical Populations', *Experimental Brain Research*, Suppl. 7: 327–36.

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: 1166–82.

—— (2000) 'The role of perceptual episodes in lexical processing', in Cutler A., McQueen J. M. and Zondervan R. (eds) *Proceedings of SWAP Spoken Word Access Processes*, pp. 155–9. Nijmegen: Max-Planck-Institute for Psycholinguistics.

Goldstone, R. L. (1998) 'Perceptual Learning', *Annual Review of Psychology*, 49: 585–612.

Gould, S. J., and Lewontin, R. C. (1979) 'The Spandrels of San Marco and the Panglossian Paradigm: A Critique of the Adaptationist Programme', *Proceedings of the Royal Society of London (Series B)*, 205: 581–98.

Griffiths, T. L., and Kalish, M. L. (2007) 'Language Evolution by Iterated Learning with Bayesian Agents', *Cognitive Science*, 31: 441–80.

Guenther, F. H., and Gjaja, M. N. (1996) 'The Perceptual Magnet Effect as an Emergent Property of Neural Map Formation', *Journal of the Acoustical Society of America*, 100: 1111–21.

Guy, G. R. (1996) 'Form and Function in Linguistic Variation', in Guy G. R. et al. (eds) *Towards a Social Science of Language: Papers in Honor of William Labov. Volume 1: Variation and Change in Language and Society*, pp. 221–52. Amsterdam: John Benjamins.

Hall, K. C. et al. (manuscript under review). The Message Shapes Phonology.

Hay, J. B. et al. (2015) 'Tracking Word Frequency Effects through 130 Years of Sound Change', *Cognition*, 139: 83–91.

Henrich, J., and Boyd, R. (2002) 'Culture and Cognition: Why Cultural Evolution Does Not Require Replication of Representations', *Culture and Cognition*, 2: 87–112.

Hintzman, D. L. (1986) 'Schema Abstraction in a Multiple-Trace Memory Model', *Psychological Review*, 93: 411–28.

Hockett, C. F. (1960) 'The Origin of Speech', *Scientific American*, 203: 8996.

Hurford, J. (1999) 'The Evolution of Language and Languages', in Dunbar R., Knight C., and Power C. (eds) *The Evolution of Culture*, pp. 173–93. Edinburgh: Edinburgh University Press.

Jaeger, T. F., and Buz, E. (2016) 'Signal Reduction and Linguistic Encoding'. *Handbook of Psycholinguistics*. Hoboken, New Jersey, US: Wiley-Blackwell.

Johnson, K. (1997) 'Speech Perception Without Speaker Normalization', in Johnson K., and Mullennix J. W. (eds) *Talker Variability in Speech Processing*. San Diego, CA: Academic Press.

Ju, M., and Luce, P. A. (2006) 'Representational Specificity of Within-category Phonetic Variation in the Long-term Mental Lexicon', *Journal of Experimental Psychology: Human Perception and Performance*, 32: 120–38.

Katsoulakis, M. A., and Souganidis, P. E. (1995) 'Generalized Motion by Mean Curvature as a Macroscopic Limit of Stochastic Ising Models with Long Range Interactions and Glauber Dynamics', *Communications in Mathematical Physics*, 169: 61–97.

Kipnis, C., and Landim, C. (1998) *Scaling Limits of Interacting Particle Systems*. Grundlehren der mathematischen Wissenschaften, v.320. Berlin: Springer.

Kirby, S. (1999) *Function, Selection and Innateness: The Emergence of Language Universals*. Oxford: Oxford University Press.

——, Dowman, M., and Griffiths, T. L. (2007) 'Innateness and Culture in the Evolution of Language', *PNAS*, 104: 5241–5.

—— et al. (2015) 'Culture and Communication in the Cultural Evolution of Linguistic Structure', *Cognition*, 141: 87–102.

Kraljic, T., and Samuel, A. (2006) 'Generalization in Perceptual Learning for Speech', *Psychonomic Bulletin & Review*, 13: 262–8.

Labov, W. (1994) *Principles of Linguistic Change. Vol. 1: Internal Factors. (Language in Society; 20)*. Oxford, UK and Cambridge, MA: Blackwell.

——, Ash, S., and Boberg, C. (2006) *Atlas of North American English: Phonetics, Phonology, and Sound Change*. Berlin: Mouton de Gruyter.

Ladd, D. Robert (2012) 'What is duality of patterning, anyway?'. *Language and cognition*, 4.4: 261–273.

Levi, S. V. (2015) 'Generalization of Phonetic Detail: Cross-segmental, Within-category Priming of VOT', *Language and Speech*, 58/4: 549–62.

Lindblom, B. (1983) 'Economy of Speech Gestures', in MacNeilage P. (ed.) *The Production of Speech*, pp. 217–45. New York: Springer-Verlag.

—— (2000) 'Developmental Origins of Adult Phonology: The Interplay between Phonetic Emergents and the Evolutionary Adaptations of Sound Patterns', *Phonetica*, 57: 297–314

——, MacNeilage, P., and Studdert-Kennedy, M. (1984) 'Self-organizing Processes and the Explanation of Language Universals', in Butterworth B., Bernard C., and Dahl O. (eds) *Explanations for Language Universals*, pp. 181–203. Berlin: Walter de Gruyter.

Luce, P. A., and Pisoni, D. B. (1998) 'Recognizing Spoken Words: The Neighborhood Activation Model', *Ear & Hearing*, 19: 1–36.

Luce, R. D. (1959) *Individual Choice Behavior*. New York: Wiley.

Martinet, A. (1955) *Economie des changements phonetiques*. Bern: Francke.

Maye, J., Aslin, R. N., and Tanenhaus, M. K. (2008) 'The Weckud Wetch of the Wast: Lexical Adaptation to a Novel Accent', *Cognitive Science*, 32: 543–62.

McQueen, J. (2007) 'Eight Questions about Spoken Word Recognition', in Gaskell G. M. (ed.) *The Oxford Handbook of Psycholinguistics*, pp. 37–53. Oxford: Oxford University Press.

Moran, P. A. P. (1958) 'Random Processes in Genetics', *Mathematical Proceedings of the Cambridge Philosophical Society*, 54/1: 60–71.

Nielsen, K. (2011) 'Specificity and Abstractness of VOT Imitation', *Journal of Phonetics*, 39: 132–42.

Norris, D., McQueen, J. M., and Cutler, A. (2003) 'Perceptual Learning in Speech', *Cognitive Psychology*, 47: 204–38.

Nosofsky, R. (1988) 'Similarity, Frequency, and Category Representations', *Journal of Experimental Psychology, Learning, Memory and Cognition*, 14: 54–65.

Ohala, J. (1989) 'Sound Change is Drawn from a Pool of Synchronic Variation', in Breivik L. E. and Jahr E. H. (eds) *Language Change, Contributions to the study of its causes. [Series, Trends in Linguistics, Studies and Monographs No. 43]*, pp. 173–98. Berlin: Mouton.

Oudeyer, P.-Y. (2002) 'Phonemic Coding Might Be a Result of Sensory-Motor Coupling Dynamics', in Hallam B. et al. (eds) *Proceedings of the 7th International Conference on the Simulation of Adaptive Behavior*, pp. 406–16. Cambridge: MIT Press.

—— (2006) *Self-organization in the Evolution of Speech*. Oxford: Oxford University Press.

Pardo, J. S. (2006) 'On Phonetic Convergence during Conversational Interaction', *Journal of the Acoustical Society of America*, 119: 2382–93.

Phillips, B. S. (1984) 'Word Frequency and the Actuation of Sound Change', *Language*, 60: 320–42.

Pierrehumbert, J. (2001) 'Exemplar Dynamics, Word Frequency, Lenition, and Contrast', in Bybee J., and Hopper P. (eds) *Frequency Effects and the Emergence of Linguistic Structure*, pp. 137–57. Amsterdam: John Benjamins.

—— (2002) 'Word-Specific Phonetics', in Gussenhoven C., Rietvelt T., and Warner N. (eds) *Laboratory Phonology VII*, pp. 101–39. Berlin, New York: Mouton de Gruyter.

—— (2006) 'The Next Toolkit', *Journal of Phonetics*, 34: 516–30.

Pisoni, D. B., and Levi, S. V. (2007) 'Some Observations on Representations and Representational Specificity in Speech Perception and Spoken Word Recognition', in Gaskell G. (ed.) *The Oxford Handbook of Psycholingusitics*, pp. 3–18. Oxford University Press.

Ratke, L., and Voorhees, P. W. (2002) *Growth and Coarsening: Ripening in Material Processing*. p. 17. Berlin: Springer.

Recanzone, G. H., Schreiner, C. E., and Merzenich, M. M. (1993) 'Plasticity in the Frequency Representation of Primary Auditory Cortex Following Discrimination Training in Adult Owl Monkeys', *Journal of Neuroscience*, 13: 87–103.

Smith, M., and Wheeldon, L. (2001) 'Syntactic Priming in Spoken Sentence Production — An Online Study', *Cognition*, 78: 123–64.

Spike, M. et al. (2013) 'Learning, Feedback and Information in Self-Organizing Communication Systems', in *Proceedings of the 35th Annual Conference of the Cognitive Science Society*, pp. 3442–7.

Studdert-Kennedy, M., and Goldstein, L. (2003) 'Launching Language: The Gestural Origin of Discrete Infinity', in Christiansen M. H. and Kirby S. (eds) *Language Evolution: The States of the Art*, pp. 235–54. Oxford, UK: Oxford University Press.

Walsh, M. et al. (2010) 'Multilevel Exemplar Theory', *Cognitive Science*, 34: 537–82.

Wedel, A. (2004) 'Category Competition Drives Contrast Maintenance Within an Exemplar-based Production-perception Loop', in *Proceedings of the seventh meeting of the ACL special interest group in computational phonology*, pp. 1–10. Barcelona, Spain: Association for Computational Linguistics, Barcelona, Spain, July 2004.

—— (2006) 'Exemplar Models, Evolution and Language Change', *The Lingusitic Review*, 23: 247–74.

—— (2012) 'Lexical contrast maintenance and the development of sublexical contrast systems'. *Language and Cognition*, 4: 319–355.

——, Kaplan, A., and Jackson, S. (2013a) 'High Functional Load Inhibits Phonological Contrast Loss: A Corpus Study', *Cognition*, 128/2: 179–86.

——, Jackson, S., and Kaplan, A. (2013b) 'Functional Load and the Lexicon: Evidence that Syntactic Category and Frequency Relationships in Minimal Lemma Pairs Predict the Loss of Phoneme Contrasts in Language Change', *Language and Speech*, 56: 395–417.

Winter, B., and Wedel, A. (2016) 'The Co-evolution of Speech and the Lexicon: The Interaction of Functional Pressures, Redundancy, and Category Variation', *Topics in Cognitive Science*, 8: 503–13.

Wu, F. Y. (1982) 'The Potts Model', *Reviews of Modern Physics*, 54: 235–68.