Fuzzy Categorical Perception as a Buffer: Simulating How Cognitive Inefficiency Enhances the Robustness of Emergent Vowel Systems

Li Lei (leili@science.ru.nl)[1] & Tao Gong (gtojty@gmail.com)[2]
[1]Centre for Language Studies, Radboud University, Nijmegen, 6525 HT, The Netherlands
[2]Travelers, Hartford, CT 06183, USA

Abstract

Human categorical speech perception (CP) is traditionally modelled as deterministic to maximise individual efficiency, yet empirical evidence suggests it is intrinsically probabilistic or fuzzy. This study investigates this apparent cognitive inefficiency using an agent-based model of emergent vowel systems. We simulated populations playing Imitation Games under varying noise levels to compare the robustness of rigid (strong CP) versus fuzzy (weak CP) perception mechanisms. Contrary to the assumption that precision yields optimality, results demonstrate that the perceptual uncertainty of weak CP acts as a crucial buffer against environmental noise. Weak CP populations not only achieved higher communicative success but also maintained categories with significantly longer lifespans and lower cognitive maintenance costs. We argue that individual-level perceptual inefficiency is an adaptive feature essential for robust social coordination, functioning as a mechanism of loose coupling that prevents systemic collapse.

Keywords: fuzzy categorical perception; loose coupling; cultural evolution; perceptual uncertainty; cognitive economy

Introduction

Human communication systems, particularly speech, are frequently conceptualised as striving for maximum efficiency (Kemp et al., 2018; Liljencrants & Lindblom, 1972; Zipf, 2016). In the specific case of vowel systems, this efficiency is shaped by a complex interplay of acoustic, physiological, and social constraints: balancing communicative distinctiveness with the minimisation of articulatory and cognitive effort (Lindblom et al., 2011). However, biological systems are rarely ideal; they are full with noise, ambiguity, and apparent sub-optimality: traits that are increasingly recognised as adaptive communicative strategies rather than flaws (Fröhlich et al., 2025). This study examines a fundamental paradox in speech perception: how individual-level perceptual uncertainty (an apparent inefficiency) may function as a necessary condition for robust social coordination.

In the domain of sound change, computational models (specifically those grounded in exemplar theory) have traditionally relied on a mechanism of strong categorical perception (CP; see Liberman et al., 1957; McMurray, 2022) to simulate how agents categorise sounds (Johnson, 2007; Pierrehumbert, 2001). Strong CP operates as a deterministic, winner-take-all step-function (Harnad, 1987): a signal slightly closer to category A than B is treated as a 100% match for A. While speech science has largely moved past this assumption (McMurray, 2022), it remains a pervasive computational simplification (see Johnson, 2007). Importantly, this simplification aligns with the framework of resource-rational cognition (Lieder & Griffiths, 2020). By discarding perceptual uncertainty (fuzziness), strong CP minimises individual computational load, serving as a highly efficient and seemingly rational heuristic for a single observer.

In contrast, empirical evidence shows that human perception is better characterised by weak (fuzzy) CP (McMurray, 2022), which manifests as a probabilistic S-shaped curve. As signals approach the fuzzy boundaries between categories (Kallens et al., 2018), perceptual certainty degrades, and perceivers select categories probabilistically. The discrepancy between this empirical fuzziness and computational rigidity provides a unique methodological opportunity. By contrasting strong and weak CP, we can intuitively assess the macro-level consequences of implementing a resource-rational simplification, while isolating the specific contribution of perceptual fuzziness to system dynamics.

From a classical engineering perspective, weak CP represents a cognitive inefficiency. It introduces noise and uncertainty into the decoding process, seemingly hindering the optimal extraction of information. Why, then, would such a flawed mechanism persist in human cognition? We propose that what appears as inefficiency at the micro-level of the individual might act as a mechanism of cognitive economy and robustness at the macro-level of the population[1].

To test this hypothesis, we employ an agent-based model (ABM; see Holland, 2000) of emergent vowel systems: namely Imitation Games (IGs; De Boer, 2000, 2001). Building on a preliminary simulation that investigated emergent vowel systems under a single noise condition (Lei & Gong, 2026), the current study provides a more comprehensive analysis by systematically varying noise levels and refining the perceptual mechanism. We simulate a population of agents playing modified IGs, allowing phonological categories to self-organise from scratch. By systematically manipulating environmental noise and contrasting the rigid strong CP against the seemingly inefficient weak CP, we demonstrate that the very fuzziness of weak CP provides a crucial buffer. This

[1]While some exemplar-based models have incorporated probabilistic categorisation (e.g., Wedel, 2004), their primary focus remains on specific trajectories of sound change rather than the comparative robustness of perceptual mechanisms in emergent systems as discussed here.

In G. Boleda, I. Dautriche, M. Maldonado, C. Saldana and T. Torrent (Eds.), *Proceedings of the 48th Annual Conference of the Cognitive Science Society*

hidden economy allows the vowel system to absorb noise that would otherwise shatter a system built on rigid, optimal precision. Thus, we argue that in the evolution of communication, individual cognitive inefficiency may not be a bug, but a feature essential for social robustness, serving as a mechanism of loose coupling that prevents the system from shattering under its own rigidity. The remainder of this paper is structured as follows. Section 2 details the ABM, focusing on the implementation of weak CP within the IGs framework. Section 3 reports the simulation results, highlighting the comparative robustness of the two perceptual mechanisms under increasing noise. Finally, Section 4 discusses the theoretical implications of these findings for understanding the role of cognitive inefficiency in the emergence of communication.

Computational Experiments

To investigate the impact of perceptual fuzziness on system robustness, we implemented the IGs framework, a classic model for simulating the self-organisation of vowel systems. We modified the perceptual module of the agents to compare two distinct categorisation mechanisms: strong and weak CP.

Imitation Games

The simulation involves a population of agents equipped with a synthetic vocal tract and an auditory system. As illustrated in Figure 1, agents interact through a series of pairwise language games to establish a shared vowel repertoire from scratch (the agent without repertoire will randomly generate a category before interacting). Each cycle of the game consists of four distinct steps (De Boer, 2000):

1. **Initiation**: Two agents are randomly selected from the population. One is assigned the role of the initiator (speaker) and the other the imitator (listener). The initiator selects a vowel from their existing repertoire and articulates it. Crucially, acoustic noise is added to the produced signal to simulate environmental interference and articulatory variability.
2. **Imitation**: The imitator perceives the noisy signal and maps it to the closest vowel category in their memory (see the following subsection for the specific mechanisms). The imitator then attempts to reproduce the signal and send it back to initiator.
3. **Feedback**: If the signal matches the original category that initiator chose at first, the initiator attempts to send a success signal to imitator; otherwise, the game fails.
4. **Updating**: Both agents update their repertoires based on the interaction outcome.

During the last step, the agents track the frequency of use and communicative success for each vowel category. If imitation succeeds, both agents shift the prototype of the used category slightly towards the perceived signal. This helps facilitate convergence. If imitation fails, the imitator's response depends on the reliability of its chosen category. If the chosen category is well-established (having a high success-to-use ratio), the agent adds a new category at the position of the signal. If not, the agent shifts the existing prototype towards the signal to correct its representation. In addition, categories with low success rates are removed to maintain the quality of the repertoire. Crucially, agents also merge categories that are acoustically too close to prevent overcrowding. This merge threshold is linked to the noise level, ensuring categories remain distinct. As we will discuss, this threshold also defines the boundary of uncertainty in the weak CP condition. Specific update parameters follow De Boer (2000).

Weak CP Implementation and Validation

To investigate the impact of perceptual fuzziness, we modified the categorisation module of the agents. While the standard strong CP employs a deterministic winner-take-all rule-selecting the prototype P_i that minimises the perceptual distance to signal S: our weak CP implementation introduces probabilistic selection within a specific acoustic horizon.

We defined an acoustic threshold T_{aco} (functionally equivalent to the merge threshold used in repertoire updating). When the distance between a signal and a category falls within this range ($d_j \leq T_{aco}$), the agent enters a confusion stage, selecting categories probabilistically via a softmax function:

$$P(P_i|S) = \frac{e^{-\text{dist}(S,P_i)}}{\sum_j e^{-\text{dist}(S,P_j)}} \quad (1)$$

with the temperature parameter set to 1. This ensures that perception becomes stochastic only when signals are acoustically ambiguous. To validate and compare the performance of these mechanisms, we conducted a computational identification task using stimuli derived from the Finnish /y/-/i/ continuum (Aaltonen et al., 1997). In the original study, stimuli varied in second formant, F_2 (1520 Hz to 2966 Hz) while other formant values (F_1, F_3, and F_4) were fixed. We presented these stimuli to agents pre-stored with /y/ and /i/ prototypes. We sampled the F_2 continuum at 5 Hz intervals, presenting each stimulus 10,000 times with random acoustic noise. Particularly, the IGs framework operates in a non-linear F_1-F_2' perceptual space (effective second formant, calculating by F_2-F_4). As shown in Figure 2a, stimuli that are equidistant on the Mel scale form two distinct clusters in the internal space of the model, revealing a perceptual leap or acoustic fault line between F_2 = 1780 Hz and 1850 Hz. This confirms the existence of a psychophysical boundary within the model itself.

Figure 2b and Figure 2c compare the resulting identification curves. The strong CP mechanism produces a steep curve; although environmental noise prevents a perfect vertical transition, the boundary is sharp. In contrast, weak CP exhibits a sigmoid-like trend, yet the curve is not a perfectly smooth S-shaped. This irregularity arises because the non-linear F_1-F_2' space implicitly discretises the stimuli into clusters (as seen in Figure 2a). Specially, however, a distinct confusion stage is observable (particularly before F2' $\approx$ 1805 Hz), where categorisation is non-deterministic. This demonstrates that while the location of the boundary is constrained by the acoustic fault line of the model, the nature of the transition (specifically the fuzziness essential for our hypothesis) is successfully activated by the weak CP mechanism.

Evaluation Metrics

We observe and quantify these changes using the following four evaluation metrics.

1. **Success Rate**: This metric measures the communicative efficiency of the population. It is defined as the average percentage of all agent interactions (De Boer, 2000, p. 453) that are successful within each timestep.
2. **Category Variance**: This metric, from Belpaeme et al. (2007, p. 188), measures the similarity of categories within the population, namely, the degree of convergence on a shared vowel system. It is first defined for two agents (A_i and A_j) as the average minimum distance between their respective category sets (C_i and C_j). As shown in Equation 2, k_i and k_j are the number of categories for each agent ($k_i = |C_i|$), and $d(c_a, c_b)$ is the perceptual distance between categories c_a and c_b.

$$CV(A_i, A_j) = \frac{1}{k_i + k_j}\left(\sum_{c_a \in C_i} \min_{c_b \in C_j} d(c_a, c_b) + \sum_{c_b \in C_j} \min_{c_a \in C_i} d(c_a, c_b)\right) \tag{2}$$

The metric we use is the average category variance $\overline{CV}(P)$ for the entire population P (of N agents), as shown in Equation 3. A lower $\overline{CV}(P)$ value indicates a more unified and converged vowel system.

$$\overline{CV}(P) = \frac{2}{N(N-1)} \sum_{i=1}^{N-1} \sum_{j=i+1}^{N} CV(A_i, A_j) \tag{3}$$

3. **Number of Categories**: This metric is used to observe the stability of the vowel system (Belpaeme et al., 2007, p. 187). It is defined as the average number of vowel categories possessed by agents in the population at each timestep.
4. **Category Lifespan:** This is an original metric for this study, designed to measure the cognitive economy required for agents to maintain their vowel repertoires. Categories can be removed or merged during the Updating step. We assign a unique ID to each newly generated category; the duration for which this ID persists in a single run (i.e., the number of timesteps from its creation to its removal or merging) is considered the category's lifespan. This metric records the average lifespan of all categories that existed during a single run.

Simulation Settings and Analysis

We employed a 2 (CP mechanism) × 5 (noise level) factorial design. The population size was fixed at $N = 20$ agents (consistent with De Boer, 2000). The noise parameter varied across five levels (0.08, 0.09, 0.10, 0.11, 0.12), a range selected to span the critical perceptual confusion region identified in the previous section. For each of the 10 conditions, we conducted 30 independent runs. A single run consisted of 10,000 timesteps, where one timestep represents a complete interaction cycle for all agent pairs. To facilitate analysis, time-series data from each run were downsampled to 250 equidistant points.

Given the stochastic nature of the simulations, violations of homogeneity of variance were anticipated (and confirmed via Levene's test). Since our analysis assesses the converged macroscopic state of the system (calculated as the average of the final 10 data points) rather than micro-level trial data, we employed robust two-way ANOVAs based on trimmed means to address the observed heteroscedasticity. This method provides valid test statistics even when standard parametric assumptions are violated.

Results

In this section, we compare the performance of the strong and weak CP mechanisms across four key metrics to comprehensively evaluate the robustness of the emergent vowel systems under increasing environmental noise,

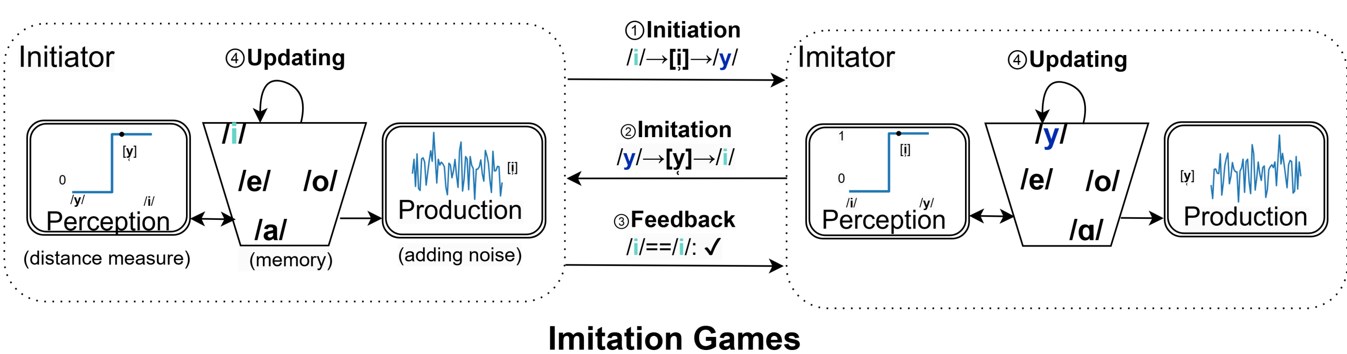

Figure 1: A flowchart of the core mechanism of the IGs. This diagram, adapted from De Boer (2000, 2001), illustrates the four steps of interaction between an initiator and an imitator (i.e., a minimal pair)

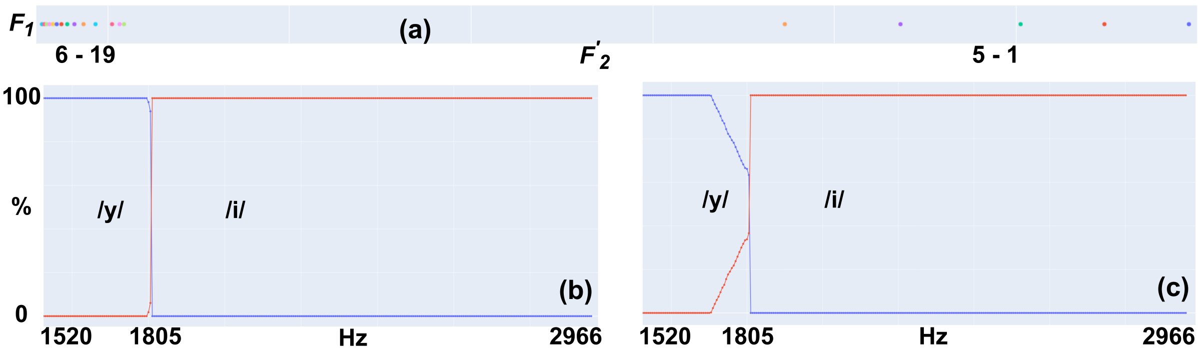

Figure 2: The non-linear mapping of the 19 stimuli from the Aaltonen et al. (1997) /y/-/i/ continuum within the perceptual space (F_1-F_2'; De Boer, 2000), and performance of the two perceptual decision mechanisms on the /y/-/i/ continuum identification task. In panel a, although the stimuli are equidistant on the Mel scale (30-mel intervals), they show clear clustering and a perceptual leap in the F_1-F_2' space. In the panel b, identification curve under the strong CP mechanism (original IGs). In the panel b, identification curve under the weak CP mechanism (temperature parameter is 1; noise parameter is 0.09). The X-axis represents the stimuli along the F_2 continuum, and the Y-axis represents the percentage of stimuli categorised as /i/.

Success Rate

We first evaluated the macroscopic performance of the emerged vowel systems by analysing the success rate of the IGs (Figure 3). A robust factorial ANOVA confirmed a significant interaction between the CP mechanism and noise level regarding communicative success. To deconstruct this interaction, we examined the post-hoc pairwise comparisons (Tukey-adjusted) specifically for the population size of $N = 20$.

Contrary to the expectation that high-fidelity perception is optimal in low-noise settings, the results indicated that the weak CP mechanism provided a statistically significant advantage across the entire noise spectrum tested. Even at the lowest noise level (0.08), weak CP systems achieved a significantly higher mean success rate (after convergence) than strong CP systems (Estimate = -0.014, $SE = 0.002$, $p < .001$).

This performance gap persisted and stabilized as environmental noise increased. At the critical noise threshold of 0.10, the disparity remained highly significant (Estimate = -0.020, $p < .001$), a trend that continued through to the highest noise level of 0.12 (Estimate = -0.015, $p < .001$).

Category Variance

We next examined the internal coherence using category variance (Figure 3), where lower values indicate a higher degree of similarity among agents' vowel repertoires. Interestingly, for the small population size of $N = 20$, the clear advantage of weak CP observed in communicative efficiency was not mirrored in system convergence. Post-hoc comparisons revealed no statistically significant difference in the mean category variance between strong CP and weak CP mechanisms at any noise level.

For instance, at the lowest noise level (0.08), the systems

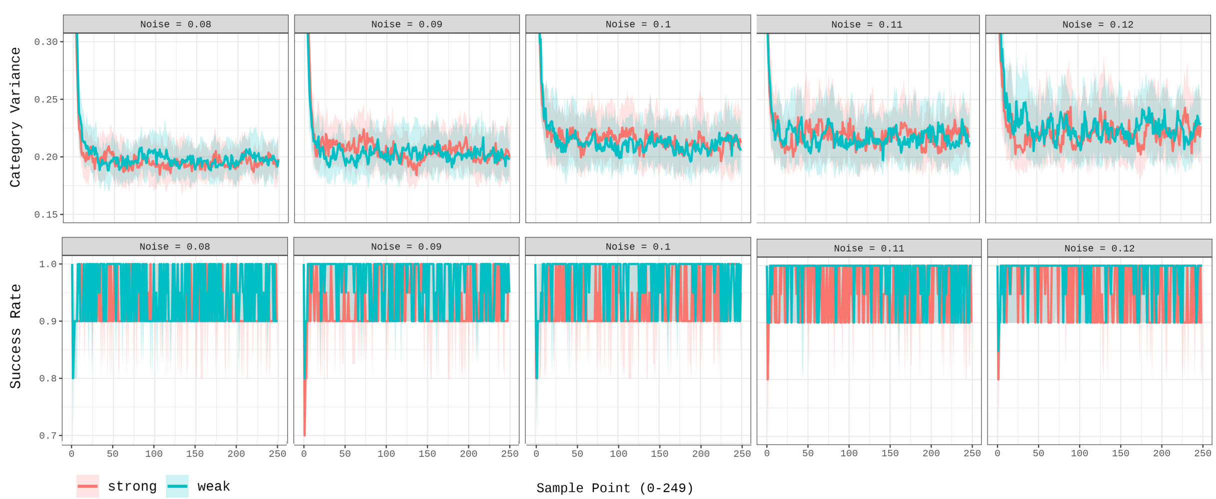

Figure 3: Time-series trajectory for success rate and category variance.

showed comparable variance ($p = .263$). Crucially, even at the highest noise level (0.12), where strong CP systems typically degrade in larger populations, the difference in category variance for the $N = 20$ group remained non-significant (Estimate = -0.001, $SE = 0.002$, $p = .539$).

Number of Categories

We further investigated the structural complexity by analysing the number of distinct vowel categories maintained by the population (Figure 4). Post-hoc comparisons for the $N = 20$ population revealed a complex, non-linear interaction between noise and perceptual mechanisms. In the low noise condition (0.08), strong CP systems maintained a significantly higher mean number of categories than weak CP systems (Estimate = 0.298, $SE = 0.072$, $p < .001$).

At intermediate noise levels (0.09 and 0.10), the difference between the two mechanisms became statistically non-significant ($p > .05$), indicating a convergence in system size. In particular, as noise escalated to 0.11, the trend reversed: weak CP systems successfully maintained a significantly larger inventory than strong CP systems (Estimate = -0.177, $p = .014$). At the highest noise level (0.12), the difference once again became non-significant ($p = .406$).

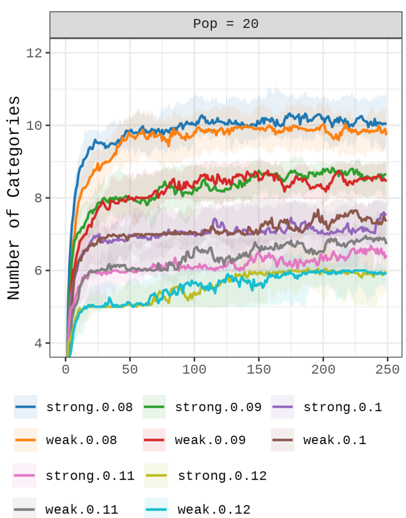

Figure 4: Time-series trajectory for the number of categories.

Category Lifespan

Finally, we quantified the cognitive economy by measuring category Lifespan.

Table 1 presents the descriptive statistics for the $N = 20$ population. The data suggests that categories under the weak CP mechanism persisted significantly longer than those under strong CP across all noise conditions.

Statistical analysis confirmed this observation. The robust ANOVA indicated a significant main effect of the CP mechanism. As detailed in the post-hoc pairwise comparisons, strong CP categories had significantly shorter lifespans than weak CP categories at every noise level tested (all estimates negative, $p < .001$).

Table 1: Mean category lifespan.

Noise	Weak CP		Strong CP	
	M	*SD*	*M*	*SD*
0.08	282.20	41.62	245.76	30.03
0.09	268.82	48.30	215.70	32.95
0.10	247.52	44.46	188.97	28.67
0.11	224.81	34.88	181.73	31.20
0.12	201.34	33.79	161.56	24.96

For example, at the critical noise threshold of 0.10, weak CP categories survived on average nearly 60 timesteps longer than their strong CP counterparts (Estimate = -58.5, $p < .001$).

Summary of Results

In summary, the simulation results demonstrate a consistent advantage for the weak CP mechanism, particularly in small populations ($N = 20$). While strong CP agents struggled to maintain communicative efficiency and structural stability as environmental noise increased, weak CP agents leveraged their perceptual uncertainty to achieve higher success rates, maintain larger vowel inventories (in distinct noise bands), and reduce the cognitive cost of repertoire updates.

General Discussion

This study examined a fundamental paradox in cognitive science: how individual-level perceptual uncertainty contributes to the robustness of population-level social coordination. By integrating a probabilistic weak CP mechanism into an ABM of vowel emergence, we demonstrated that the fuzziness of perceptual boundaries acts as a crucial buffer against environmental noise. While both strong and weak CP mechanisms could support the emergence of vowel systems under ideal conditions, our results revealed a critical divergence in how they achieve and maintain this order. The weak CP mechanism not only sustained higher communicative success in noisy environments but, more importantly, did so with significantly lower cognitive overhead.

The Hidden Cost of Cognitive Rigidity

Our results highlight a trade-off between communicative performance and cognitive economy. While the success rate for strong CP was comparable to or only slightly lower than that for weak CP in certain conditions, its category lifespan was dramatically lower. This disparity reveals a hidden cost of cognitive rigidity.

The strong CP mechanism, akin to the original IGs, achieves high success rates by relying on a powerful yet expensive error-correction strategy. When a deterministic categorisation fails due to noise, the agent must rapidly adapt by merging or deleting the faulty prototype. Our lifespan data (see Table 1) confirm that this correction comes at an extremely high cognitive overhead. aligning with broad empirical evidence that actively removing outdated representations and updating categorisation rules incurs substantial executive costs. Agents in the strong CP condition are forced to frequently destroy their existing categories and rebuild them around new signals to repair the hard failures caused by the brittleness of the system (e.g., Ashby and Maddox, 2005; Ecker et al., 2014).

This difference in category lifespan is directly rooted in the Updating logic of the IGs. Specifically, agents add new categories when interactions fail despite the high reliability of the chosen vowel. In the strong CP condition, rigid boundaries cause environmental noise to trigger hard failures even for well-established categories. These failures prompt the unnecessary addition of redundant categories, leading to acoustic overcrowding and subsequent forced removals. This constant cycle of creation and removal explains the high cognitive overhead and reduced category lifespan observed in our results. In contrast, weak CP avoids this overhead. The probabilistic nature of the softmax function allows agents to tolerate minor acoustic deviations without triggering a full-scale structural update.

This suggests that weak CP is not only more robust at the population level but also more cognitively plausible at the individual level. It strikes a better balance between the need for categorisation (to structure the world) and perceptual uncertainty (to tolerate its variability). Classical CP theory is often viewed as a trade-off between accuracy and cost; strong CP attempts to minimise cost through rigid categories, but paradoxically incurs high repair costs during interaction. Therefore, the perceptual instability introduced by weak CP functions effectively as a mechanism for cognitive economy, allowing the system to absorb noise rather than shatter under it.

Optimal Inefficiency in Evolutionary Dynamics

These findings can be interpreted through the theoretical framework of Bayesian iterated learning, which models how individual perceptual biases reshape language systems across generations (Kirby et al., 2014). In this framework, the strong CP mechanism functions as a maximiser strategy. As noted by Griffiths and Kalish (2007), a maximiser deterministically selects the most likely category based on the input. While this appears optimal for a single decision, it leads to distinct and often brittle evolutionary pathways. For instance, a maximiser achieves higher immediate accuracy than a probability matcher in binary choice tasks (Ferdinand, 2015). Our results align with this prediction: the strong CP agents, acting as maximisers, eliminated variation too rapidly in high-noise conditions. This rigid strategy, however, imposed a significantly higher maintenance cost (requiring frequent Updating operations) to sustain the stability of the vowel system.

Conversely, the weak CP mechanism functions as a probability matcher. It selects categories based on a probability distribution, similar to the softmax strategy observed in human decision-making under uncertainty (Lee et al., 2012). Although a probability matcher is technically suboptimal for individual classification accuracy (Ferdinand, 2015), our simulations demonstrate that this individual-level inefficiency translates into population-level robustness. This phenomenon can be explained by the non-linear dynamics of group interaction. As De Boer (2015) argues, even weak biases at the individual level can be amplified into dominant structures at the macro level. In our model, the probabilistic noise introduced by weak CP prevented the system from becoming trapped in local optima or collapsing under error. Therefore, the seemingly irrational behaviour of weak CP is not a flaw but an adaptive feature that exploits the non-linearity of self-organisation to maintain system stability.

Perceptual Fuzziness as Loose Coupling

The macro-level robustness provided by weak CP can be theoretically understood through the lens of loose coupling (Orton & Weick, 1990). Unlike the tight integration of engineered systems (represented by our strong CP condition), human cultural systems thrive on looseness, which provides elasticity against environmental shocks. In the context of social interaction and cultural evolution (e.g., social tinkering, see Chater and Christiansen, 2025), the perceptual fuzziness of weak CP functions as a critical coordination buffer. It manages the variance of heterogeneous individuals, treating acoustically diverse and noisy signals as socially equivalent within a tolerant margin (Winter, 2014). Our simulations empirically demonstrate that this loose coupling is not merely a byproduct of flawed cognition, but a structural necessity. It is precisely this apparent inefficiency that resolves the tension between myopic individual errors and global stability, preventing the emergent vowel system from collapsing under the weight of rigid precision.

Limitations and Future Directions

First, as an exploratory ABM (see discussion in Bradfield, 2017), this study identified qualitative patterns rather than validating outcomes against historical linguistic corpora. Future work should therefore calibrate model parameters against specific cases of vowel shifting or merger (e.g., Chirkova and Gong, 2014) to test the predictive power of the weak CP mechanism. Second, while our simplified softmax function captures the essential property of probabilistic selection, it abstracts the complex psychophysical curves found in human perception (e.g., Baronchelli et al., 2010). Future iterations should thus incorporate empirically derived identification curves from actual categorisation experiments to validate whether the specific shape of the confusion zone affects the trajectory of sound change. Despite these limitations, this study provides a computational proof of concept that cognitive inefficiency, in the form of perceptual fuzziness, plays a functional role in the evolution of robust communication systems.

Acknowledgments

We would like to thank anonymous reviewers for their insightful suggestions and comments.

References

Aaltonen, O., Eerola, O., Hellström, Å., Uusipaikka, E., & Lang, A. H. (1997). Perceptual magnet effect in the light of behavioral and psychophysiological data. *The Journal of the Acoustical Society of America*, *101*(2), 1090–1105. https://doi.org/10.1121/1.418031

Ashby, F. G., & Maddox, W. T. (2005). Human category learning. *Annu. Rev. Psychol.*, *56*(1), 149–178. https://doi.org/10.1146/annurev.psych.56.091103.070217

Baronchelli, A., Gong, T., Puglisi, A., & Loreto, V. (2010). Modeling the emergence of universality in color naming patterns. *Proceedings of the National Academy of Sciences*, *107*(6), 2403–2407. https://doi.org/10.1073/pnas.0802485105

Belpaeme, T., De Boer, B., & Jansen, B. (2007). The dynamic emergence of categories through imitation. In C. L. Nehaniv & K. Dautenhahn (Eds.), *Imitation and social learning in robots, humans and animals: Behavioural, social and communicative dimensions* (pp. 179–194). Cambridge University Press. https://doi.org/10.1017/CBO9780511489808.013

Bradfield, J. (2017). The sound of a spherical cow. *Phonology*, *34*(2), 347–362. https://doi.org/10.1017/S0952675717000173

Chater, N., & Christiansen, M. H. (2025). Social tinkering: The social foundations of cultural complexity. *Behavioral and Brain Sciences*, 1–62. https://doi.org/10.1017/S0140525X25103981

Chirkova, K., & Gong, T. (2014). Simulating vowel chain shift in Xumi. *Lingua*, *152*, 65–80. https://doi.org/10.1016/j.lingua.2014.09.009

De Boer, B. (2000). Self-organization in vowel systems. *Journal of Phonetics*, *28*(4), 441–465. https://doi.org/10.1006/jpho.2000.0125

De Boer, B. (2001). *The origins of vowel systems*. Oxford University Press. https://doi.org/10.1093/oso/9780198299653.001.0001

De Boer, B. (2015). Biology, culture, evolution and the cognitive nature of sound systems. *Journal of Phonetics*, *53*, 79–87. https://doi.org/10.1016/j.wocn.2015.07.001

Ecker, U. K., Lewandowsky, S., & Oberauer, K. (2014). Removal of information from working memory: A specific updating process. *Journal of Memory and Language*, *74*, 77–90. https://doi.org/10.1016/j.jml.2013.09.003

Ferdinand, V. A. (2015). *Inductive evolution: Cognition, culture, and regularity in language* [Doctoral dissertation, University of Edinburgh]. http://hdl.handle.net/1842/11741

Fröhlich, M., Jäger, G., & Achimova, A. (2025). Rethinking ambiguity across species. *Neuroscience & Biobehavioral Reviews*, 106401. https://doi.org/10.1016/j.neubiorev.2025.106401

Griffiths, T. L., & Kalish, M. L. (2007). Language evolution by iterated learning with bayesian agents. *Cognitive Science*, *31*(3), 441–480. https://doi.org/10.1080/15326900701326576

Harnad, S. (1987). Psychophysical and cognitive aspects of categorical perception: A critical overview. In S. Harnad (Ed.), *Categorical perception: The groundwork of cognition* (pp. 1–25). Cambridge University Press.

Holland, J. H. (2000). *Emergence: From chaos to order*. Oxford University Press.

Johnson, K. (2007). Decisions and mechanisms in exemplar-based phonology. In M.-J. Solé, P. S. Beddor, & M. Ohala (Eds.), *Experimental approaches to phonology* (pp. 25–40). Oxford University Press. https://doi.org/10.1093/oso/9780199296675.003.0003

Kallens, P. A. C., Dale, R., & Smaldino, P. E. (2018). Cultural evolution of categorization. *Cognitive Systems Research*, *52*, 765–774. https://doi.org/10.1016/j.cogsys.2018.08.026

Kemp, C., Xu, Y., & Regier, T. (2018). Semantic typology and efficient communication. *Annual Review of Linguistics*, *4*, 109–128. https://doi.org/10.1146/annurev-linguistics-011817-045406

Kirby, S., Griffiths, T., & Smith, K. (2014). Iterated learning and the evolution of language. *Current opinion in neurobiology*, *28*, 108–114. https://doi.org/10.1016/j.conb.2014.07.014

Lee, D., Seo, H., & Jung, M. W. (2012). Neural basis of reinforcement learning and decision making. *Annual Review of Neuroscience*, *35*(1), 287–308. https://doi.org/10.1146/annurev-neuro-062111-150512

Lei, L., & Gong, T. (2026). Rethinking categorical perception as an evolutionary prerequisite for speech. In S. Hartmann, M. Sibierska, M. Fröhlich, Y. Jadoul, M. Josserand, T. Matzinger, K. Mudd, J. Nölle, M. Pleyer, S. Wacewicz, & P. Żywiczyński (Eds.), *The Evolution of Language: Proceedings of the 16th international conference (evolang xvi)* (pp. 252–262). The Evolution of Language Conferences. https://doi.org/10.17617/2.3696655

Liberman, A. M., Harris, K. S., Hoffman, H. S., & Griffith, B. C. (1957). The discrimination of speech sounds within and across phoneme boundaries. *Journal of Experimental Psychology*, *54*(5), 358. https://doi.org/10.1037/h0044417

Lieder, F., & Griffiths, T. L. (2020). Resource-rational analysis: Understanding human cognition as the optimal use of limited computational resources. *Behavioral and brain sciences*, *43*, e1. https://doi.org/10.1017/S0140525X1900061X

Liljencrants, J., & Lindblom, B. (1972). Numerical simulation of vowel quality systems: The role of perceptual contrast. *Language*, *48*(4), 839–862. https://doi.org/10.2307/411991

Lindblom, B., Diehl, R., Park, S.-H., & Salvi, G. (2011). Sound systems are shaped by their users: The recombination of phonetic substance. In N. Clements G. & R. Ridouane (Eds.), *Where do phonological features come from? Cognitive, physical and developmental bases of distinctive speech* *categories* (pp. 67–98). Academic Press. https://doi.org/10.1075/lfab.6.04lin

McMurray, B. (2022). The myth of categorical perception. *The Journal of the Acoustical Society of America*, *152*(6), 3819–3842. https://doi.org/10.1121/10.0016614

Orton, J. D., & Weick, K. E. (1990). Loosely coupled systems: A reconceptualization. *Academy of Management Review*, *15*(2), 203–223. https://doi.org/10.5465/amr.1990.4308154

Pierrehumbert, J. B. (2001). Exemplar dynamics: Word frequency, lenition and contrast. In J. L. Bybee & P. J. Hopper (Eds.), *Frequency and the emergence of linguistic structure* (pp. 137–158). John Benjamins Publishing Company. https://doi.org/10.1075/tsl.45.08pie

Wedel, A. (2004). Category competition drives contrast maintenance within an exemplar-based production/perception loop. In J. Goldsmith & R. Wicentowski (Eds.), *Proceedings of the 7th meeting of the acl special interest group in computational phonology: Current themes in computational phonology and morphology* (pp. 1–10). Association for Computational Linguistics. https://doi.org/10.5555/1622153.1622154

Winter, B. (2014). Spoken language achieves robustness and evolvability by exploiting degeneracy and neutrality. *BioEssays*, *36*(10), 960–967. https://doi.org/10.1002/bies.201400028

Zipf, G. K. (2016). *Human behavior and the principle of least effort: An introduction to human ecology*. Ravenio books.