No learning bias favoring phonetically grounded rules? Revisiting vowel harmony

Alexander Martin, Adriana Guevara-Rukoz, Thomas Schatz & Sharon Peperkamp

Introduction: Phonotactic patterns and phonological rules tend to be both structurally simple, i.e. they involve a small number of features, and phonetically-grounded, i.e. they reflect constraints on speech production and perception. While there is abundant evidence for a learning bias in favor of structurally simple patterns and rules, a learning bias in favor of phonetically grounded patterns and rules is absent or at best weak (Moreton & Pater, 2012ab). Here, we revisit the question of a learning bias in favor of phonetically grounded rules. We focus on vowel harmony, which is both typologically frequent and phonetically grounded, and compare it to the logical equivalent vowel disharmony, which is exceedingly rare and lacks phonetic grounding. Previous work has found no evidence for a learning bias favoring harmony (Pycha et al. 2003; Skoruppa & Peperkamp 2011). We provide evidence that there *is* a (small) learning bias in favor of harmony. Moreover, we show that this bias compounds over time with transmission from one individual to the next. Thus, this bias could contribute to the typological prevalence of harmony systems.

Study 1: Learning biases have been shown to be easier to detect 1/ with a production rather than a perception task (Skoruppa et al. 2006), and 2/ by presenting learners with exceptions to the rule (Baer-Henney et al. 2014); both manipulations increase cognitive demands. We designed an experiment in which auditory exposure consisted of 32 pairs of French nonce items that were said to consist of the singular and corresponding plural forms of nouns in an invented language, such as /vizɛ̃/ - /vizɛ̃tɛl/, /kovu/ - /kovutɔl/ (harmonic), or/vizɛ̃/ - /vizɛ̃tɔl/, /kovu/ - /kovutɛl/ (disharmonic). Thus, the plural suffix had two forms, /tɛl/ and /tɔl/. There were 5 exposure groups, in which the singular/plural pairs were 1/ all harmonic, 2/ all disharmonic, 3/ mostly harmonic (75% harmonic - 25% disharmonic), 4/ mostly disharmonic (25% harmonic - 75% disharmonic), or 5/ half-half (control condition: 50% harmonic - 50% disharmonic). French participants were tested on both exposure items and novel items. In each trial, a singular was presented auditorily, and they had to orally provide the plural form they believed to correspond to the stem. In a first experiment with 78 participants we obtained - besides an effect of Consistency - an effect of Rule (β = 0.98, *SE* = 0.43, $\chi^2(1)$ = 4.87, *p* = .03), with better performance on harmony than on disharmony. The effect of Rule was present in the inconsistent input conditions (β = 1.10, *SE* = 0.41, $\chi^2(1)$ = 6.55, *p* = 0.01), but not in the consistent ones (β = 0.91, *SE* = 0.90, $\chi^2(1) < 1$). In a second experiment with 32 speakers, focused on the inconsistent conditions, we did not replicate the effect of Rule. Collapsing all the data from the two experiments, however, the effect was trending (β = 0.70, *SE* = 0.36, $\chi^2(1)$ = 3.62, *p* = 0.057; see Figure 1), with numerically better performance on harmony than on disharmony.

Study 2: Learning biases might be easier to detect when non-native stimuli are used, which reduces the likelihood that participants encode the stimuli orthographically. We tested 173 English speakers on the same French stimuli as those in Study 1. There were only two exposure groups, with 100% harmony and 100% disharmony, respectively. Moreover, as the test was administered online, we used a perception task. In each trial, a singular was followed by two plurals, one with /tɛl/ the other with /tɔl/, and participants had to indicate which plural they thought was correct. A logistic mixed-effects model analysis revealed a robust effect of Rule, with better performance for harmony than for disharmony (β = 0.21, *SE* = 0.07, $\chi2(1)$ = 9.58, $p < .002$; see Figure 2).

Study 3: Both experiments using iterated learning and modelling studies have shown that weak biases may be amplified by transmission within a population (Reali & Griffiths 2009; Kirby 2001). We show this effect for vowel harmony in a hybrid approach: we use the data from Study 1 to initiate chains, whose evolution is simulated by probabilistically sampling from observed behavior in the different experimental input conditions. We started by using the distributions of responses from the different input conditions (i.e., 0, 25, 50, 75, 100% harmony) to estimate response distributions for all other possible values of input harmony. The simulation was based on individual "languages" (chains of transmission) which evolved from one "generation" to another. At each generation *n*, the output harmony was sampled from the estimated distribution associated with the input for this generation (i.e., output at generation $n-1$). The output value from generation *n* was then passed on as the input for generation $n+1$. For example, if the output harmony of generation $n-1$ was estimated at 76%, a value was randomly drawn from the estimated distribution associated with this value, say 82%. This value was then considered as the output for generation *n* and as the input for generation *n*+1, and so on. We found that after only a few generations, the number of chains with mean percentages of harmony situated below the 20th percentile sharply decreased, while the number of those with mean percentages of harmony situated above the 80th percentile increased ($t = 34$, $p < .0001$; see Figure 3). These results show that despite the initial bias being small, it compounded over time, such that chains that were strongly harmonic vastly outnumbered chains that were strongly disharmonic.

Conclusion: While the increases in cognitive demands implemented in Study 1 yielded weak results, the use of non-native stimuli in Study 2 proved useful to reveal a learning bias favoring harmony. Moreover, Study 3 showed that even a weak, non-significant, bias compounds over time. Due to vowel-to-vowel coarticulation, vowel harmony is more likely to emerge in any given language than vowel disharmony. We showed that that in addition, typology could be shaped by a small learning bias in favor of harmony.

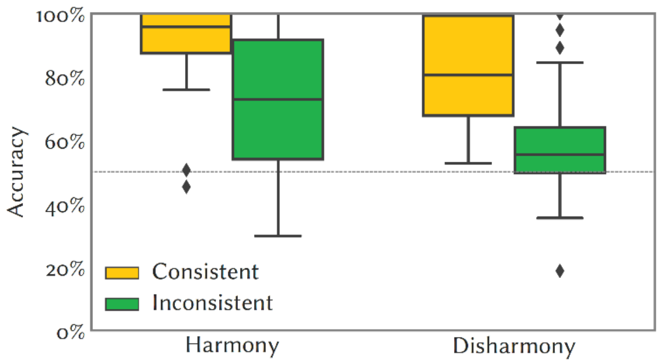

Figure 1: Accuracy scores in Study 1

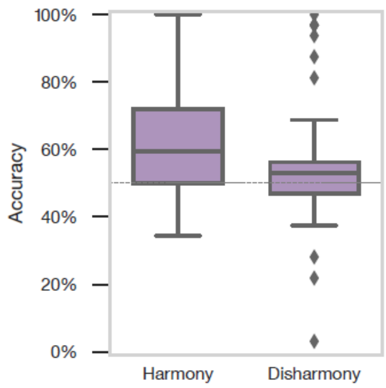

Figure 2: Accuracy scores in Study 2

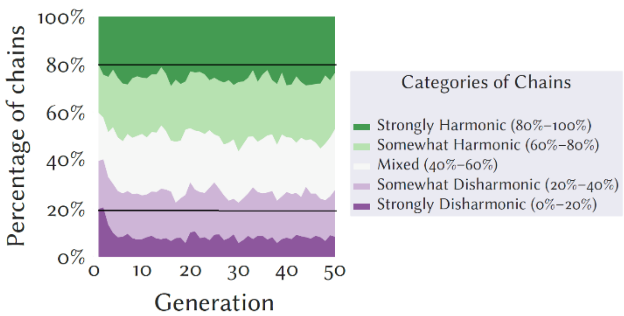

Figure 3: Results from the transmission simulation in Study 3.
The shaded areas represent the percentage of chains that fall within a given interval over time.

References

Baer-Henney, D., F. Kügler & R. van de Vijver (2014). The interaction of language-specific and universal factors during the acquisition of morphophonemic alternations with exceptions. *Cognitive Science* 39, 1537-69.

Kirby, S. (2001). Spontaneous evolution of linguistic structure. *IEEE Transactions on Evolutionary Computation* 5, 102-110.

Moreton, E. & J. Pater (2012ab). Structure and substance in artificial-phonology learning. Part I: Structure. Part II: Substance. *Language and Linguistics Compass* 6, 686-701 / 702-718.

Peperkamp, S., K. Skoruppa & E. Dupoux (2006). The role of phonetic naturalness in phonological rule acquisition. *Proceedings of BUCLD* 30, 464-475.

Pycha, A., P. Nowak, E. Shin & R. Shosted (2003). Phonological rule-learning and its implications for a theory of vowel harmony. *Proceedings of WCCFL* 22, 101-114.

Reali, F. & T. Griffiths (2009). The evolution of frequency distributions: Relating regularization to inductive biases through iterated learning. *Cognition* 111, 317-328.

Skoruppa, K. & S. Peperkamp (2011). Adaptation to novel accents: feature-based learning of context-sensitive phonological regularities. *Cognitive Science* 35, 348-366.