Unifying horizontal and vertical interactions in the Bayesian language game

Bas Cornelissen
Willem Zuidema
Institute for Logic, Language and Computation
University of Amsterdam

IN A NUSTELL We propose a Bayesian framework that connects Bayesian *iterated learning* (IL) to *naming games* (NG), the two main branches of agent-based models of cultural language evolution. Surprisingly, the two appear to be closely related: the extremes of continuum. We find that Bayesian IL trivializes the effect of cultural evolution, whereas the Bayesian NG yields non-trivial, lineage-specific languages.

TECHNICALITIES The BNG is inspired by Bayesian IL models[4] and extends the NG of De Vylder & Tuyls.[5] Agents essentially perform repeated Bayesian updating in a Dirichlet-categorical:

$$\text{SPEAKER} \quad \begin{cases} \theta_t \mid \alpha_{t-1} & \sim \text{Dirichlet}(\alpha_{t-1}) \\ x_i \mid \theta_t & \sim \text{Categorical}(\theta_t) \end{cases}$$

$$\text{HEARER} \quad \alpha_{t+1} := \alpha_t + c_t$$

where θ is a language, x a production, α the param. of Dirichlet, and c a vector of counts.

FRAMEWORK

1 Bayesian naming game

The speaker picks a language (a distribution over words), accounting for its innate biases towards certain languages, and produces a word. The hearer observes the word and updates the probabilities it assigns to all languages in a Bayesian fashion. On average, this makes the observed word more likely to be produced when the hearer later becomes a speaker. Alignment is thus increased.

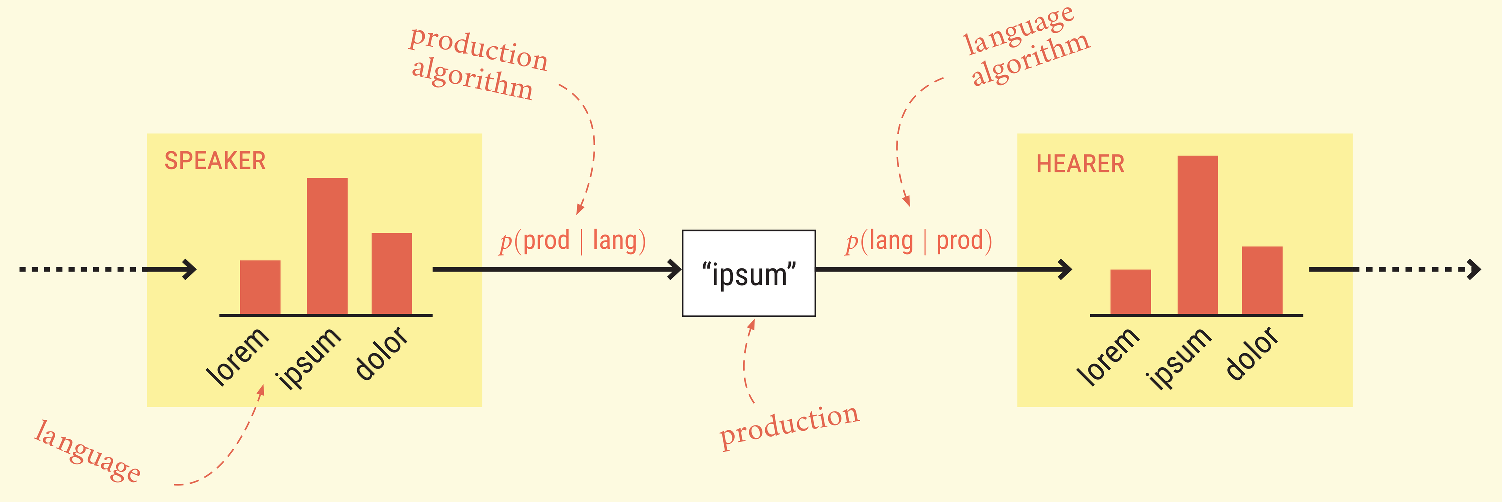

$$p(\text{lang} \mid \text{prod}) \propto p(\text{prod} \mid \text{lang}) \cdot p(\text{lang})$$

Alternative interpretation: a language as a distribution over *linguistic features* rather than words.[1]

2 Population model

The **life expectancy**, in a random walk, interpolates between IL ($\gamma = 1$) and NG ($\gamma = \infty$).

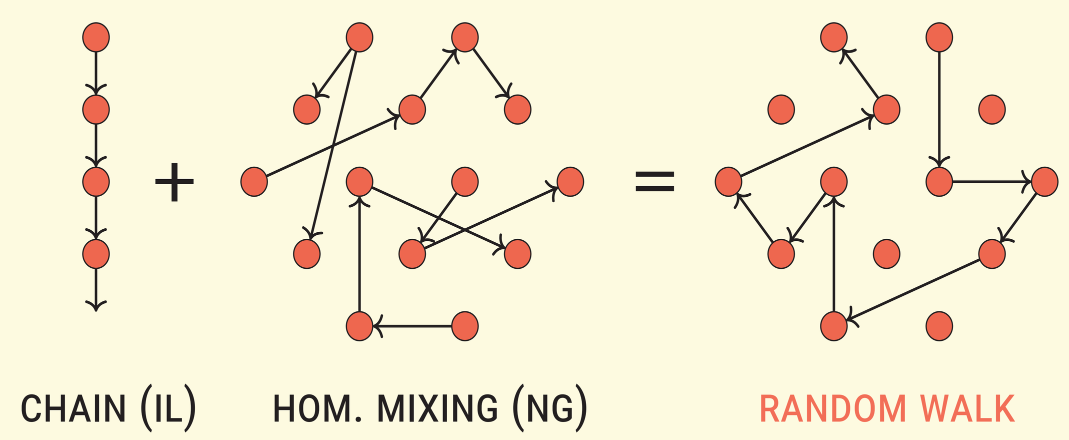

3 Strategies: MAP—sample

Sample languages and productions or use the ones with **maximum probability** (MAP)?

lang. strategy $\propto p(\text{lang} \mid \text{prod})^{\eta}$, prod. strategy $\propto p(\text{prod} \mid \text{lang})^{\zeta}$

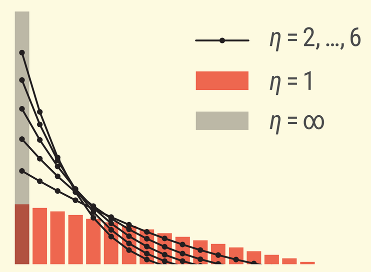

Exponentiating a distribution: sampling ($\eta = 1$) and MAP ($\eta = \infty$).[2]

= Bayesian language game

Unified view of Bayesian agent-based models of cultural language evolution: from Bayesian IL to BNG, with all strategies.[3]

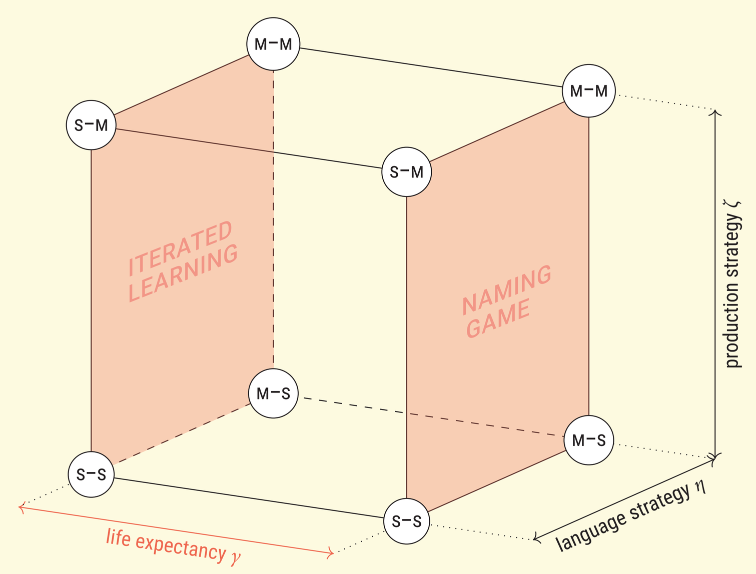

Parameter space. S–M = a sampling lang. strategy ($\eta = 1$) and maximizing prod. strategy ($\zeta = \infty$)

1 See e.g. Reali & Griffiths (2010), whose IL model is identical to the BLG with a life expectancy of 1 (i.e., the IL-variant)

2 Kirby, Dowman & Griffiths (2007) use exponentiation to interpolate between samplers and maximisers.

3 It should be noted that in this framework, agents are *not* perfect Bayesian reasoners when they use a maximising production strategy (i.e., $\zeta > 1$).

4 e.g. Griffiths & Kalish (2007). In the IL setting ($\gamma = 1$), the Bayesian language game reduces to Reali & Griffiths (2009).

5 De Vylder & Tuyls (2006) proved convergence in a deterministic variant of a naming game of which ours is a near-direct generalisation. They use queue-agents and languages are always relative frequencies, which form a discrete subset of the simplex. We use the full simplex and no queue.

Griffiths, T. L., & Kalish, M. L. (2007). Language Evolution by Iterated Learning With Bayesian Agents. *Cognitive Science, 31*(3), 441–480. http://doi.org/10.1080/15326900701326576

Kirby, S., Dowman, M., & Griffiths, T. L. (2007). Innateness and culture in the evolution of language. *Proceedings of the National Academy of Sciences of the United States of America, 104*(12), 5241–5245.

Reali, F., & Griffiths, T. L. (2010). Words as alleles: connecting language evolution with Bayesian learners to models of genetic drift. *Proceedings. Biological Sciences / The Royal Society, 277*(1680), 429–36.

De Vylder, B., & Tuyls, K. (2006). How to reach linguistic consensus: A proof of convergence for the naming game. *Journal of Theoretical Biology, 242*(4), 818–831.

FINDINGS

Behaviour of the Bayesian naming game

In the BNG, coherence is reached with a **lineage specific language** that **reflects the biases**, but is non-trivially shaped by the cultural process.

The BNG reproduces **wide, constrained variation**. The variation is determined by the strenght of the bias, and the average of different lineages gives the bias.

A. Lineage specificity and reflection of the bias

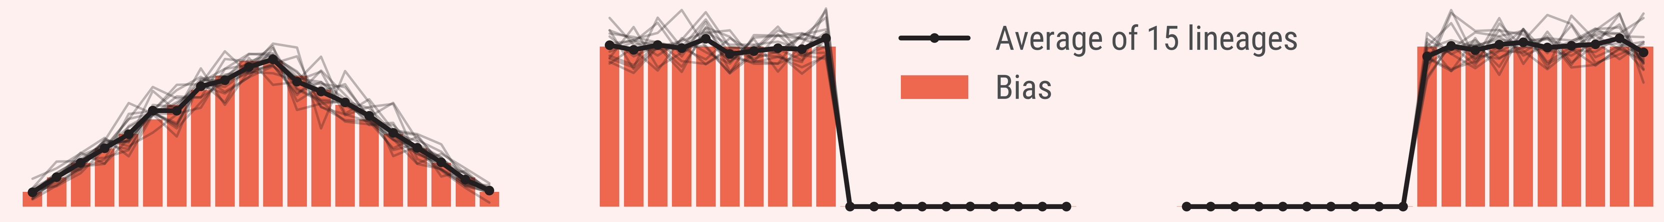

Characterising the Bayesian language game

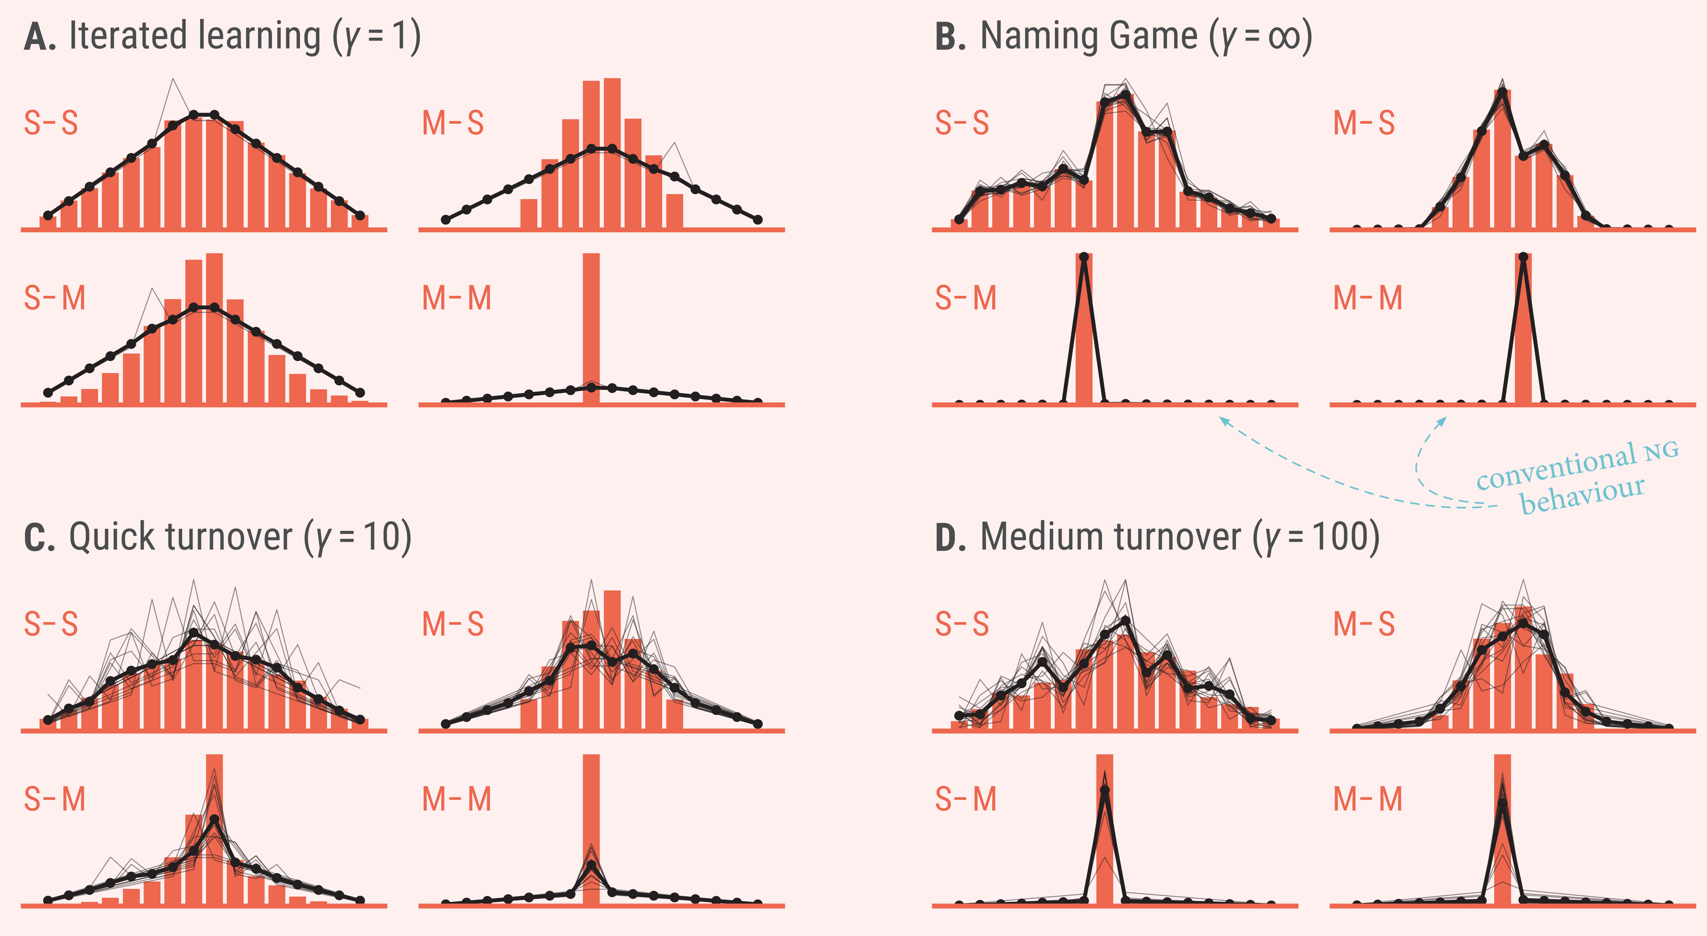

Outcomes of the BLG for IL, NG and two intermediate pop. turnover rates, for all 'extreme' strategies. The bias is a 'pyramid'.

Average language in population — Singe agent's language — Used language

The Bayesian IL model seems predictably determined by the biases (even for non-samplers). A higher life expectancy (lower turnover) results in stronger non-trivial cultural effects. The resulting language is shaped by both the bias and the cultural process.

CONCLUSIONS Iterated learning and naming games are naturally connected in the *Bayesian language game*. One extreme case, the Bayesian NG, reveals new behaviour: lineage-specific langauges that reflect innate biases of the learners. In comparison, Bayesian IL models seem to trivialize the cultural process. This highlights the importance of horizontal interactions.