

## Signaling conventions: Who learns what where and when in a social network?

**Previous Research.** Lewis (1969) invented signaling games to show that linguistic conventions can arise even without a pre-existing language simply from regularities in communicative behavior. Signaling games have since become a standard model for the pragmatic evolution of semantic meaning (cf. Steels, 1997; Nowak and Krakauer, 1999; Skyrms, 2010). In order to understand the applicability and conceptual adequacy of signaling game models, the most important theoretical question that needs to be addressed is under which circumstances stable signaling conventions can arise. This paper contributes to this question by showing how language evolution depends on the social structure of a population.

Following a general trend in evolutionary game theory, recent studies have started to probe into the simplifying assumption underlying classical evolutionary dynamics that populations of agents are homogeneous, i.e., barring of social structure. Dispensing with this artificial assumption, Zollman (2005), for instance, has demonstrated how coexistent language conventions can evolve if the population of language users is arranged on a lattice. Wagner (2009) studied the same dynamic on so-called  $\beta$ -graphs which exhibit more realistic *small-world properties*, namely a high clustering coefficient, paired with a low characteristic path length (Watts and Strogatz, 1998). Wagner’s simulations showed for  $\beta$ -graphs (defined below) that (i) the higher the clustering coefficients the larger the fractions of players that acquire a signaling convention, and (ii) the lower the characteristic path length the smaller the number of connected regions of agents that use the same signaling convention.

**Contribution.** This paper intends to extend this line of research into the relation between social structure and evolving patterns of language use. Following Wagner, we look at the evolution of signaling conventions on  $\beta$ -graphs. Adding to previous studies, however, we investigate the *local network properties* associated with (regions of) agents that have successfully learned a language or not. In a nutshell, our experiments show that languages form preferentially on locally highly connected subgraphs; borders between languages fall preferentially on regions “in between” highly connected subregions. Our results therefore not only show *that* different language conventions can coexist, but also *where* to expect uniformity and language contact.

Also, previous work focused on the deterministic imitate-the-best dynamics. This dynamics assumes that an imitating agent somehow can access the relative, time-averaged communicative success of all her neighbors. Instead, we assume that agents update their behavior based on their actual previous interactions with their neighbors. We consider both (i) reinforcement learning (RL) and (ii) a version of best response behavior (BR) based on beliefs acquired from previous encounters (fictitious play). These dynamics differ in the level of entailed *rationality* of learners. Interestingly, although varying the rationality of learners had effects on time course and overall success of learning, results about where languages form on the network appeared to be largely *independent* of the learning dynamics.

**Experimental Set-Up.** A  $\beta$ -graph is obtained by first considering a ring of nodes where each node is connected to its  $k$  nearest neighbors. Subsequently, for each node, its  $k/2$  left neighbors are rewired to a random vertex  $n$  with probability  $\beta$ . For our analysis, we created 10 such  $\beta$ -graphs with 300 nodes,  $k = 6$  and  $\beta \in \{.08, .09, .1\}$ . These parameter choices ensured the small-worldliness of our networks that we had to keep small for obtaining enough data points at manageable computation costs. For each network, we ran 20 trials with either only BR- or only RL-learners. Agents played standard Lewisean cooperative signaling games for state-action coordination with equiprobable states and no signaling costs. Communication happened randomly between neighbors on the network, and each agents’ behavior was updated separately after each round of communication the agent was involved in. We recorded strategies of agents in suitably chosen regular intervals. Each trial ran until at least 90% of agents had acquired a language, or each network connection had been used 3000 times in either direction. The latter was to ensure a compromise between a manageable running time and sufficient time for learning, but also because we were interested in the results of learning after a reasonable time-span, not necessarily in the limit behavior.

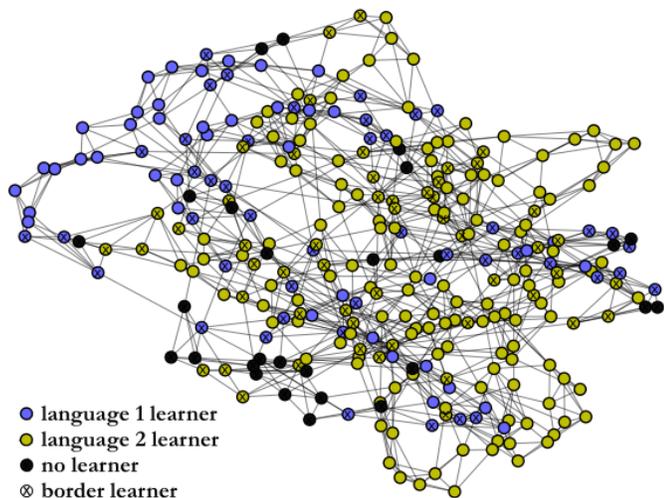
An agent’s behavior and manner of updating, of course, depended on the agent’s learner type. Crucially, agents strategies were modelled as behavioral strategies, i.e., independently for each choice point and independently for their roles as senders and receivers. The behavior of RL-players was modeled by a set of urns, one for each choice point. Each urn contains balls corresponding to each available action in proportion to the likelihood of choosing that action. Initially all urns were filled with 100 balls. Successful communication was reinforced by placing 14 additional balls of the successful type back into the urn. Additionally, we also assumed a mechanism of *lateral inhibition* (see Steels, 1997) by removing 4 balls of the successful type from *all* urns for the same player role. From previous studies we knew that in order to match the plasticity of RL-learners with that of BR-learners we should look at RL-learners with bounded memory and BR-learners with unbounded memory. We therefore had RL-learners “remember” only reinforcement effects of the previous 300 rounds (independently of role or choice point). The behavior of BR-learners was defined as usual. BR-learners had an unbounded memory of all previous encounters. Based on this, they derived a belief about the average strategy of their neighborhood (i.e., not keeping track of each individual neighbor) by maximum likelihood estimation. Behavior of BR-learners was then simply defined as a best response to this belief with random tie breaking. With an initially empty memory, BR-learners initially played entirely at random, just like their RL-cousins.

**Results.** Our main goal was to determine which *local* network properties best characterize where, on average, learning would be most likely successful. Towards this goal, we looked at what we would call *language regions*. A language region is a maximal subset of agents that have acquired the same language that forms a connected subgraphs. Despite the different learning dynamics, our data confirmed Wagner’s results that in small world networks like ours the number of language regions is small while the size of language regions is relatively big. Interestingly, for our relatively small networks of 300 agents, the observed frequencies  $f(n)$  of having  $n > 0$  language regions of a given type could be approximated by  $f(n) \approx a \cdot b^{-(n-1)}$  with  $a$  given by the observed frequency of a single language region of a given type and  $b \approx 3$ , independently of learning dynamics. On top of that, we also found that the set of all language regions of a given type, although not necessarily a connected subgraph itself, nonetheless was always a region of higher density, transitivity and clustering than the original network. This supports the conclusion that local cliquishness supports the evolution of a local language.

We also compared the specific local network properties of agents who had learned a language with those who had not. We were also interested to see who lives at the border of her language region and who lives in its interior. Let’s say that a *learner* is an agent who acquired the same signaling convention in both her sender and receiver role (for RL-learners this meant getting close enough to the pure strategy in question); a *border agent* is an agent whose neighborhood is not uniformly playing the same strategy as she herself (see graph below). Our results were by and large the same for both learning dynamics. Learners compare to non-learners by higher values of clustering (CL), degree centrality (DC), and eccentricity (EX), as well as by lower values of closeness centrality (CC) and betweenness centrality (BC) (see table below). Intuitively speaking, this means that in order to successfully learn a language in a social network an agent would have to be well embedded in a dense *local* structure. Globally well-connected agents, on the other hand, have difficulties learning a language in a heterogeneous network, because they might be torn between different locally firmly established conventions. This is also supported by looking at the average local network properties of those learners who were border agents, i.e., at the outside of their language region (second row of the table below). Border learners showed all properties of what are called “boundary spanners” in the network literature when compared to non-border learners. Most importantly, border learners had significantly higher betweenness centrality and significantly lower eccentricity and closeness centrality than inner region learners. Naturally, the difference between inner and border learners also showed in the time course of learning: inner learners acquired their language significantly faster than border learners.

A very surprising but welcome further result of our experiments was that the learning dynamics did not have significant impact on the local network properties that characterize regional learning success. Still, there were, of course, notable differences between learning dynamics. Firstly and obviously, RL-learners take much longer to learn than BR-learners. Moreover, learning dynamics differed in learning success. Let’s say here that a *pure learner* is an agent who learned the same strategy for both roles; a *mixed learner* has learned a different strategy for both roles; a *minimal learner* has learned a strategy for at least one role. We found that RL-learners generated significantly larger proportions of pure learners, but significantly smaller proportions of mixed and minimal learners than their more rational and more flexible BR-cousins.

**References.** Lewis (1969), *Convention*. ★ Novak & Krakauer (1999), “The Evolution of Language”, *PNAS*. ★ Skyrms (2010), *Signals*. ★ Steels (1997), “The Synthetic Modeling of Language Origins”, *Evolution of Communication*. ★ Wagner (2009), “Communication and Structured Correlation”, *Erkenntnis*. ★ Watts & Strogatz (1998), “Collective Dynamics of ‘Small-World’ Networks”, *Nature*. ★ Zollman (2005), “Talking to Neighbors”, *Philosophy of Science*.



	learner	border
CL	high	low
CC	low	high
BC	low	high
DC	high	high
EX	high	low

Table 1: Properties of learners compared to non-learners and of border learners compared to non-border learners.