

# SIGNALING CONVENTIONS: WHO LEARNS WHAT WHERE AND WHEN IN A SOCIAL NETWORK?

ROLAND MÜHLENBERND

*Seminar für Sprachwissenschaft, University of Tübingen, Wilhelmstr. 19,  
Tübingen, 72074, Germany  
roland.muehlenbernd@uni-tuebingen.de*

MICHAEL FRANKE

*Seminar für Sprachwissenschaft, University of Tübingen, Wilhelmstr. 19,  
Tübingen, 72074, Germany  
michael.franke@uni-tuebingen.de*

Lewis (1969) invented signaling games to show that meaning convention can arise simply from regularities in communicative behavior. This paper contributes to the question how the formation of signaling conventions depends on the social structure of a population. Our results not only show that different language conventions can coexist, but also where to expect uniformity and language contact. We found that place and time of convention formation can be traced well to particular clusters of high/low values of suitable notions from formal network theory. Against prior expectations, we found that agent rationality is less important than network role in deciding how and when an agent adopts a convention.

## 1. Introduction

Lewisian signaling games have become a standard model for the pragmatic evolution of semantic meaning (cf. Steels, 1997; Nowak & 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. 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 for a so-called imitate-the-best dynamic 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 (defined below) which exhibit more realistic small-world properties, namely a high clustering coefficient, paired with a low characteristic path length (Watts & Strogatz,

1998). Wagner’s simulations showed that (i) the higher the clustering coefficients the larger the fractions of players that acquire a unique 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.

This paper probes deeper into the relation between social structure and language evolution in order to further our knowledge of synthetic evolutionary processes in structured populations and thereby to pave the way for a more thorough understanding of the sociological factors of linguistic variability. While previous related work has focused on studying which *global* network structures are especially conducive to innovation and its spread (Ke, Gong, & Wang, 2008; Fagyal, Swarup, Escobar, Gasser, & Lakkaraju, 2010), the present paper investigates more closely the *local* network properties associated with (regions of) agents that have successfully learned a language or not. In distinction to previous studies, we also focus not on imitation, but on used-based learning dynamics from evolutionary game theory. To study the effect of agent rationality on language evolutions, we considered best-response dynamics and reinforcement learning.

Our most striking results, in a nutshell, were these. Firstly, we found that languages form preferentially on locally highly connected subgraphs; borders between languages fall preferentially on regions “in between” highly connected subregions. Secondly, conventionalization depended crucially on local network properties, while the learning dynamics and the amount of agent rationality had hardly any noticeable effect. Thirdly, we compared the local properties of agents who had learned a language with those who had not, and of those who lived at the borders of language regions with those who lived in the interior. To characterize the differences we found, we made a distinction between *family men* and *globetrotters*, which we characterized by relative values of suitable clusters of properties from formal network theory. We found that learners and interior agents tend to be family men with tight local connections, while non-learners and border agents tend to be globetrotters with wide-ranging global connections. Finally, we found evidence that the first ones to adopt a convention were globetrotters, while it is mainly family men who stabilize a convention after its inception.

## 2. Signaling games & learning dynamics

A signaling game is a game played between a sender  $S$  and a receiver  $R$ . Initially, nature selects a state  $t \in T$  with prior probability  $\Pr(t) \in \Delta(T)$ , which  $S$  observes, but  $R$  doesn’t.  $S$  then selects a message  $m \in M$ , and  $R$  responds with a choice of action  $a \in A$ . For each round of play, players receive utilities depending on (in the cheap-talk case we consider here) the actual state  $t$  and the response action  $a$ . We will here be concerned only with a simple variant of this game, which we call *Lewis game*: there are only two states that are equiprobable, two messages and two actions that correspond one-to-one with the states, indicated by the same index. Players share an interest in successful communication, expressed by utility

function  $U(t_i, a_j) = 1$  if  $i = j$  and 0 otherwise.

Although messages are initially meaningless in this game, meaningfulness arises from regularities in behavior. Behavior is defined in terms of strategies. A *behavioral sender strategy* is a function  $\sigma : T \rightarrow \Delta(M)$ , and a *behavioral receiver strategy* is a function  $\rho : M \rightarrow \Delta(A)$ . A behavioral strategy can be interpreted as a single agent's probabilistic choice or as a population average. For a Lewis game, exactly two isomorphic strategy profiles constitute evolutionary stable states (Huttegger, 2007). In these, strategies are pure (i.e., action choices have probabilities 1 or 0) and messages associate states and actions uniquely, like so:

$$L_1: \quad \begin{array}{l} t_1 \rightarrow m_1 \rightarrow a_1 \\ t_2 \rightarrow m_2 \rightarrow a_2 \end{array} \quad L_2: \quad \begin{array}{l} t_1 \quad m_1 \quad a_1 \\ \quad \times \quad \times \\ t_2 \quad m_2 \quad a_2 \end{array}$$

Classical evolutionary game theory assumes a homogenous population of agents and studies evolutionary processes on the aggregate population level. In this paper we focus instead on more fine-grained agent-based evolutionary dynamics. Agents repeatedly play a Lewis game with those agents they are connected with in their social network, and adapt their behavioral strategies based on learning from previous interactions. We consider two kinds of *learning dynamics* that differ with respect to how rational the learning agents are assumed to be: best-response dynamics (BR) and reinforcement learning (RL).

**Best-response dynamics.** The idea of best-response dynamics is simple: agents remember the past plays that they have been engaged in and derive from their memory a belief about their opponents' behavior; it is to that belief that they play a rational best response. We assume here that agents form a belief about the collective behavior of all of their neighbors, not keeping track of each agent separately. More concretely, a given agent's belief about his neighborhood's receiver (sender) behavior  $B_r(a|m)$  ( $B_s(t|m)$ ) is just a behavioral receiver (sender) strategy derived by keeping track of *all* of the agent's past interactions. The sender's expected utility for sending  $m$  in state  $t$  is  $EU_s(m|t) = \sum_{a \in A} B_r(a|m) \times U(t, a)$ . Accordingly, the receiver's expected utility is  $EU_r(a|m) = \sum_{t \in T} B_s(t|m) \times U(t, a)$ . A *best response* is an action choice that maximizes expected utility. This gives rise to the following *response rules for best-response dynamics*:

$$\sigma(m|t) = \begin{cases} \frac{1}{|\arg \max_m EU_s(m|t)|} & \text{if } t \in \arg \max_m EU_s(m|t) \\ 0 & \text{else} \end{cases} \quad (1)$$

$$\rho(a|m) = \begin{cases} \frac{1}{|\arg \max_a EU_r(a|m)|} & \text{if } a \in \arg \max_a EU_r(a|m) \\ 0 & \text{else} \end{cases} \quad (2)$$

**Reinforcement learning with limited memory.** Reinforcement learning can be captured by a simple model based on urns, also known as *Pólya urns* (cf. Roth &

Erev, 1995; Skyrms, 2010). An urn models a behavioral strategy, in the sense that the probability of making a particular decision is proportional to the number of balls in the urn that correspond to that action choice. By adding or removing balls from an urn after each encounter, an agent’s behavior is gradually adjusted. For signaling games, the sender has an urn  $\Omega_t$  for each state  $t \in T$ , which contains balls for different messages  $m \in M$ . The number of balls of type  $m$  in urn  $\Omega_t$  is designated with  $m(\Omega_t)$ , the overall number of balls in urn  $\Omega_t$  with  $|\Omega_t|$ . If the sender is faced with a state  $t$  she draws a ball from urn  $\Omega_t$  and sends message  $m$ , if the ball is of type  $m$ . The same holds *mutatis mutandis* for the receiver. The resulting *response rules for reinforcement learning* are:

$$\sigma(m|t) = \frac{m(\Omega_t)}{|\Omega_t|} \quad (3) \quad \rho(a|m) = \frac{a(\Omega_m)}{|\Omega_m|} \quad (4)$$

The learning dynamics are realized by changing the urn content dependent on the communicative success. In detail: if communication via  $t$ ,  $m$  and  $a$  is successful, the number of balls in urn  $\Omega_t$  is increased by  $\alpha \in \mathbb{N}$  balls of type  $m$  and reduced by  $\gamma \in \mathbb{N}$  balls of type  $m' \neq m$ . Similarly, for the receiver. In this way successful communicative behavior is more probable to reappear in subsequent rounds. In our experiments, all urns were initially filled with 100 balls and we set  $\alpha = 10$  and  $\gamma = 4$ . From previous work (Mühlenbernd, to appear) we knew that in order to match the plasticity of different learning dynamics, we should consider BR-learners with *unbounded* memory but RL-learners with *bounded* memory. For that reason, an RL-learners’ urns only reflected the impact of the last 300 interactions (irrespective of role) that the agent was engaged in. With an initially empty memory, BR-agents initially played entirely at random, just like their RL-cousins.

### 3. Network games: design and basic notions

We modeled a structured population as a  $\beta$ -graph. A  $\beta$ -graph is obtained by first considering a ring of nodes where each node is connected to its  $k$  nearest neighbors and subsequently, for each node, rewiring its  $k/2$  left neighbors to a random vertex  $n$  with probability  $\beta$  (Watts & Strogatz, 1998). 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 each with either only BR- or only RL-agents. Agents played the standard Lewis game. Communication happened randomly between neighbors on the network, and each agent’s 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 short running time and sufficient time for learning, but also because we were interested

in the results of learning after a realistic time-span, not in limit behavior.

Our main goal was to investigate the relationship between meaning evolution and social network structure. The theoretical challenge here lies in adequately characterizing local network roles in terms of formal notions of network connectivity, which can never be crisp, but must necessarily be of a probabilistic nature. For our present purposes, however, a rather straightforward cross-classification based on whether an agent is globally and/or locally well-connected turned out to have high explanatory value. Using suggestive terminology, we will be mainly concerned with two types of agents, *family men* and *globetrotters*. The former have tight local connections, with less global connections; the latter show the opposite pattern plus a high degree of connectivity.

In order to capture these notions more adequately, we look at suitable notions from social network theory (Jackson, 2008): *betweenness centrality* (BC), *closeness centrality* (CC), *degree centrality* (DC), *individual clustering* (CL) and *eccentricity* (EX).<sup>a</sup> High values for BC and CC, as well as low values of EX, characterize agents that are globally at a central position in the network. So, for a measure of *global connectedness* we looked at relative values of these properties. On the other hand, a high value for CL should be considered a measure for the agent’s *local connectedness*. A high value for DC depicts a high degree of connectivity. Family men and globetrotters are thus characterized as follows:

	BC	CC	DC	EX	CL
family man	low	low	-	high	high
globetrotter	high	high	high	low	low

#### 4. Results

**Language regions.** In order to determine which local network properties best characterize where, on average, learning would be most likely successful, we looked at what we will 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 (2009) results that in small world networks like ours the number of language regions is small while the size of language regions is relatively big. Most of the time, two big language regions formed, one for each signaling convention. BR-dynamics, due to its slightly higher flexibility, was prone to produce a little more regional variability. On top of that, we also found that each connected language region of a given type had always a higher *average clustering* and *transitivity* value<sup>b</sup>

<sup>a</sup>For the definition of BC, CC, DC and CL we refer to (Jackson, 2008), chapter 2. EX of a node  $v$  is the maximum distance from  $v$  to all other nodes in the graph.

<sup>b</sup>Given the (sub)-graph  $G$ : average clustering depicts the average CL value over all nodes in  $G$  and transitivity depicts the fraction of all possible triangles in  $G$  that are in fact in  $G$ .

than the expected average value for a connected subgraph with the same size  $n$  (= number of nodes), as depicted in Figure 1. We may conclude from this that local cliquishness supports the evolution of a local language.




Figure 1. Comparing observed clustering and transitivity of language regions with expected values from randomly chosen connected subgraphs (solid lines, subgraph size along the  $x$ -axis).

**Learners and border agents.** Based on their learning success and network position, we classified agents into (i) *learners vs. non-learners* and (ii) *border agents vs. interior agents*. A learner is an agent who, by the end of a simulation run, has acquired the same signaling convention in both her sender and receiver role (for RL-agents this meant getting close enough to the pure strategy in question). Interior agents have only neighbors who learned the same language as they themselves, while border agents are agents whose neighborhood is not uniformly behaving in the same way that they do.

Our results were by and large the same for both learning dynamics: learners tend to be family men, border agents tend to be globetrotters (see Figure 2). 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. (Naturally, the difference between interior and border agents also showed in the time course of learning: interior agents acquired their language significantly faster than border agents.)

**Rationality.** A certainly surprising result of our experiments was that the learning dynamics did not have much impact on the local network properties that characterize regional learning success. Phrased more strikingly, we could conclude that an agent's location in the network was more influential to his behavioral adaptation than his rationality. Still, there were, of course, notable differences between

BR						
	learners		non-learners	border	interior	
CL	0.4524100	>	0.4266016	0.4041671	<	0.4877772
CC	0.2048951	<	0.2063953	0.2095079	>	0.2014160
BC	0.01307722	<	0.01414102	0.017328919	>	0.009838327
DC	0.02006919	≈	0.02002832	0.02052166	>	0.01970899
EX	7.698488	>	7.672088	7.584469	<	7.785561
RL						
	learners		non-learners	border	interior	
CL	0.4529768	>	0.4227674	0.4053346	<	0.4872364
CC	0.2048325	<	0.2070318	0.2096410	>	0.2012726
BC	0.01302537	<	0.01469637	0.01729162	>	0.00983325
DC	0.02006223	≈	0.02013186	0.02052492	>	0.01970263
EX	7.700152	>	7.653051	7.573122	<	7.795518

Figure 2. Average local network properties of learners vs. non-learners, and of border vs. interior agents by different learning dynamics. Symbols <, >, ≈ indicate whether differences in means are considered significant by a t-Test.

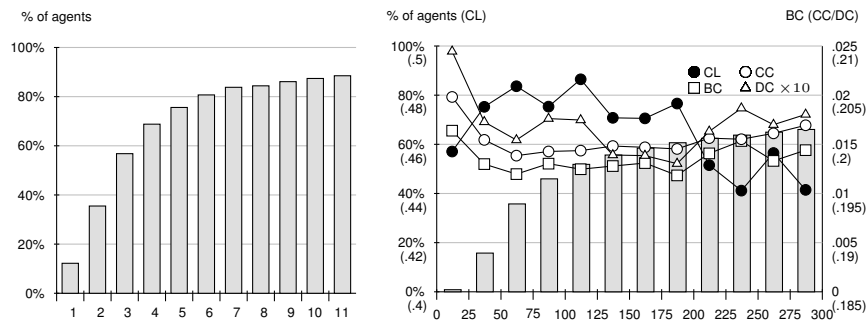


Figure 3. Temporal development of the proportion of agents having settled into their final language for BR-dynamics (left) and RL-dynamics (right). The right picture also plots the average values for CL, BC, CC and DC for those RL-learners who have settled into their final language during the specified interval of rounds. The graph shows that *founding fathers* (rounds 0-25) have properties similar to globetrotters, whereas *stabilizers* (rounds 25-200) have properties more similar to family men.

learning dynamics. The most obvious difference is that BR-learners settle into conventions much faster than RL-learners (see Figure 3).

**Founding fathers and stabilizers.** The slower RL-dynamics moreover showed a very interesting connection between the temporal development of meaning formation and network structure (see Figure 3 on the right): there seem to be three phases of conventionalization which affect different network roles. In phase 1 (0-25) the first agents to adopt a convention (*founding fathers*) have relatively

high centrality values (BC, CC and DC) and a low CL and therefore properties of globetrotters. In phase 2 (25-200) conventions were adopted by more locally well-connected family men (*stabilizers*) with lower centrality values and a high CL. The last agents to adopt a convention, (after ca. 200 rounds) again show the mark of globetrotters. This suggests the interpretation that a convention is usually sparked by globetrotters who are free approach a pure strategy, while it takes a locally well-connected set of family men to fix a meaning convention, so that it can also affect the (remaining) globetrotters.

## References

- Fagyal, Z., Swarup, S., Escobar, A. M., Gasser, L., & Lakkaraju, K. (2010). Centers and peripheries: Network roles in language change. *Lingua*, 120, 2061–2079.
- Huttegger, S. M. (2007). Evolution and the explanation of meaning. *Philosophy of Science*, 74, 1–27.
- Jackson, M. O. (2008). *Social and economic networks*. Princeton: Princeton University Press.
- Ke, J., Gong, T., & Wang, W. S.-Y. (2008). Language change and social networks. *Communications in Computational Physics*, 3(4), 935–949.
- Lewis, D. (1969). *Convention*. Cambridge: Harvard University Press.
- Mühlenbernd, R. (to appear). Learning with neighbours. *Synthese*, 1-23. (10.1007/s11229-011-9980-y)
- Nowak, M. A., & Krakauer, D. (1999). The evolution of language. *PNAS*, 96(14), 8028–8033.
- Roth, A., & Erev, I. (1995). Learning in extensive-form games: experimental data and simple dynamic models in the intermediate term. *Games and Economic Behaviour*, 8, 164–212.
- Skyrms, B. (2010). *Signals: Evolution, learning & information*. Oxford: Oxford University Press.
- Steels, L. (1997). The synthetic modeling of language origins. *Evolution of Communication*, 1, 1–34.
- Wagner, E. (2009). Communication and structured correlation. *Erkenntnis*, 71, 377–393.
- Watts, D. J., & Strogatz, S. H. (1998). Collective dynamics of small-world networks. *Nature*, 393, 440–442.
- Zollman, K. J. S. (2005). Talking to neighbors: The evolution of regional meaning. *Philosophy of Science*, 72(1), 69–85.