The Change of Signaling Conventions in Social Networks

Roland Mühlenbernd*

Eberhard Karls University Tübingen,
Department of Linguistics,
Wilhelmstr. 19, 72074 Tübingen, Germany
eMail: roland.muehlenbernd@uni-tuebingen.de

Abstract. To depict the mechanisms that have enabled the emergence of semantic conventions, philosophers and researchers particularly access a game-theoretic model: the *signaling game*. In this article I argue that this model is also quite appropriate to analyze not only the emergence of a semantic convention, but also its change. I delineate how the application of signaling games helps to reproduce and depict mechanisms of semantic change. For that purpose I present a model that combines a signaling game with *innovative reinforcement learning*; in simulation runs I conduct this game repeatedly within a multi-agent setup, where agents are arranged in social network structures. The results of these runs are contrasted with an attested theory from sociolinguistics: the *'weak tie'-theory*. Analyses of the produced data target a deeper understanding of the role of environmental variables for the promotion of i) semantic change or ii) solidity of semantic conventions.

Keywords: signaling game, reinforcement learning, multi-agent account, social networks structure, mechanisms of language change

1 Introduction

"What are the mechanisms that can explain the emergence of semantic meaning?" Philosophers have been concerned with this question for a long time. Russell (1921) once said: "[w]e can hardly suppose a parliament of hitherto speechless elders meeting together and agreeing to call a cow a cow and a wolf a wolf." With this proposition Russell wanted to make the following point: the assumptions that i) *semantic meaning is conventional* and ii) *semantic conventions are a result of verbal agreements* lead to a particular paradox of the evolution of human language: language (as a tool to make verbal agreements) is needed for language (in form of semantic meaning) to emerge.

Lewis (1969) found a very elegant solution for this paradox: he argued that semantic conventions can arise without previous agreements, but just by regularities in communicative behavior. He expounded his point with a game-theoretical

* Special Thanks to Michael Franke, Shane Steinert-Threlkeld, three anonymous reviewers and the ERC Advanced Grant Project Group 'Language Evolution - the Empirical Turn' (EVOLAEMP) for comments and discussions.

model: the *signaling game*. This game basically models a communicative situation between a speaker and a hearer, and by playing this game repeatedly and using simple update rules to adjust subsequent behavior, both participants might finally agree on a convention without making an overt verbal agreement in advance (cf. Skyrms, 2010). In other words: semantic conventions can arise 'unconsciously'[1] by repeated communication and simple adaption mechanisms; and a signaling game is an elegant way to formalize these dynamics.

Apparently, quite similar mechanism can be assumed for the change of semantic conventions, or to be more precise: the innovation, shift and loss of semantic conventions. Like we cannot assume that speechless elders made agreements to call a wolf a 'wolf', we furthermore cannot assume that the people in the 1970s made a public announcement to use the word 'groovy' when they wanted to express that something is really nice, and another announcement in the 1980s, that people shouldn't use this word anymore. Just as semantic meaning can emerge in an unconscious and automatic way, in the same way, expressions arise, change their meaning, or get lost. Thus it seems to be plausible that a signaling game is an appropriate model to explain general mechanisms of semantic change.

To study the dynamics of semantic change, I use repeated signaling games in combination with an *update mechanism* that depicts unconscious behavior of decision making: a modified variant of *Roth-Erev reinforcement learning* (Roth & Erev, 1995), which is one of the most popular learning rule in combination with repeated signaling games (cf. Barrett, 2009; Barrett & Zollman, 2009; Skyrms, 2010). The novelty of this study is the fact that I use this setup to analyze the change rather than – as done is former studies – the emergence of semantic conventions. Furthermore, I conduct simulation experiments of communicating agents in social network structures to evaluate the environmental factors that might or might not support change or stability of semantic conventions.

This article is divided in the following way: in Section 2 I present related work that concerns mostly computational studies dealling with similar research questions. In Section 3 I i) discus the advantages and disadvantages of my model in comparison with others, and ii) motivate the research question of this paper in view of real-world studies from sociolinguistics that hypothesize or give evidence about the way network structure plays a part in the mechanisms of language change. In Section 4 I introduce some basic notions of repeated signaling games, reinforcement learning dynamics and network theory. Furthermore, I discuss a noteworthy extension for reinforcement learning, called *innovation* (cf. Skyrms, 2010; Alexander, Skyrms, & Zabell, 2012). It can be shown that this additional feature realizes an interesting interplay between stabilizing and renewing effects (cf. Mühlenbernd & Nick, 2014); and I adopt it for my experiments, which are described and analyzed in Section 5. A final conclusion is presented in Section 6.

[1] In a game theoretic sense, 'unconsciousness' of agents signifies that they do not deduce a particular decision, but rather learn it by optimizing behavior.

2 Related work

This section outlines relevant literature: four research directions of computational studies that investigate the development of conventions in a population of multiple agents by analyzing repeated interaction.

2.1 Emergence of behavioral conventions

A popular direction that involves game-theoretical models is the study of *behavioral convention*[2], that deals with the analysis of *coordination games*, such as the standard coordination game, the *prisoners' dilemma* or the *intersection game* (cf. Axelrod, 1984; Airiau, Sen, & Villatoro, 2014). As one of the first studies in this field, Young (1993) analyzed a repeatedly played n-player coordination game combined with a non-deterministic variant of the reinforcement learning rule *fictitious play* (Brown, 1951). Young was able to show that players always select a pure Nash equilibrium of the game as a behavioral convention. In a subsequent study, Shoham and Tennenholtz (1997) relaxed Young's precondition that all players play the same game by analyzing $n - k$ *stochastic social games* (standard coordination game and prisoners' dilemma), where n agents play k-player games. They i) defined and used the reinforcement learning rule *highest cumulative reward* (HCR), and ii) showed that a convention is guaranteed to emerge with all agents using HCR.

Still, note that Shoham and Tennenholtz (1997) used a model where agents interact with randomly selected partners from the whole population, thus they did not reconsider a particular network topology. A number of subsequent studies incorporated more realistic topologies, such as small-world networks.[3] In this field, Airiau et al. (2014) studied the emergence of behavioral conventions in different network topologies, also in a scale-free small-world networks. Their model involved the coordination game and the intersection game with varying parameters. Furthermore, they applied next to Fictitious Play two further kinds of reinforcement learning rules. With their experiments in scale-free networks, they found out that multiple regions of local conventions can emerge, and that a stable get-together of adjacent regions can be explained by particular network features. Unfortunately, Airiau et al. (2014) didn't point out the concrete properties that support emergence and stability of multiple conventions.

2.2 Emergence of lexical conventions

Another relevant research direction involves the study of *lexical conventions* by analyzing so-called *language games*, such as the popular *naming game* (Steels,

[2] Note: in the literature under discussion this phenomenon is solely called 'convention' (or sometimes 'norm'). I label it 'behavioral convention' to distinguish this phenomenon from more communication-related types of conventions.

[3] A good overview of subsequent studies that concern different particular network typologies is given in Airiau et al. (2014).

2002): such a game i) models an interaction between a speaker and a hearer who try to find names for objects to understand each other, and ii) is generally used to investigate how a common lexicon is established in a society. A naming game's lexicon depicts a mapping between a set of concepts and a set of words. A lexical convention is optimal in a society of agents, if all use the same lexicon with a *maximum specificity*: a one-to-one mapping between concepts and words.

Salazar, Rodriguez-Aguilar, and Arcos (2010) investigate such a naming game in a scale-free network and small-world network topology (Watts & Strogatz, 1998) of 1000 agents. In their model, agents do not update their behavior by learning rules, but adopt lexical mappings of neighbors by a selection mechanism called *elitist selection*. In their experiments, agents have a lexical space of 10 concepts and 10 words. The simulation experiments showed that in each simulation run a global lexical convention with maximum or almost maximum specificity emerged, whereby conventions i) evolve much faster in a scale-free network, but ii) have a better average specificity in a small-world network. Furthermore, by integrating a *random innovation* mechanism that allows agents to invent new mappings, lexical conventions with maximal specificity emerge in every simulation run for both network topologys.

In an experimental approach, (Centola & Baronchelli, 2015) tested a similar game in online experiments: they found out that in dependence of the network typology either multiple local conventions (spatial and random network), or one global convention emerged (homogeneous mixed population).

2.3 Emergence of signaling conventions

Another line of research investigates the emergence of semantic meaning – in form of so-called *signaling conventions* – in structured populations, by analyzing the dynamics of repeated *signaling games*. A signaling game models a communication situation between a sender and a receiver, where the sender encodes an information state t with a message m, and the receiver decodes m with an interpretation state a. Simply put, a signaling convention depicts a compatible pair of patterns of communication between sender and receiver (see Section 4 for a detailed definition).

As a first study in this research direction, Zollman (2005) made experiments with a simple signaling game (two information/interpretation states and two messages) combined with the imitation rule *imitate-the-best* on a toroid lattice structure (100 $\times$ 100 agents). His results revealed the emergence of multiple regions using a local convention. Wagner (2009) adopted Zollman's model with a different network topology: a small-world network (Watts & Strogatz, 1998); and he did not only reproduce Zollman's result of the emergence of multiple regions, but also revealed the role of two properties of the network structure: i) the higher the *clustering coefficient*[4] of the whole network, the higher the probability that agents adopt signaling conventions at all, and ii) the higher the

[4] For the definition of these network properties I refer to Jackson's *Social and Economic Networks* (Jackson, 2008), Chapter 2.

Table 1. The different roles of agents in the emergence process of signaling conventions and their characteristic combinations of network properties.

	DC	CC	BC	CL
founding fathers	high	low	low	high
stabilizers	low	low	low	high
late-learners	high	high	high	low

average path length[4] of the network, the higher the number of regions with local signaling conventions. Both facts lead to the result that simulation runs on a small-world network generally result in the emergence of a small number of local signaling conventions.

Finally, Mühlenbernd and Franke (2012) analyzed agents playing the simple signaling game and update via Roth-Erev reinforcement learning (Roth & Erev, 1995) on small-world and scale-free network structures. As a basic result also here a small number of local language conventions emerged. A finer analysis revealed that an agent's individual network features have a quite high correlation with her role during the process by which a convention emerges: initiators of a language convention (founding fathers), immediate adopters of this initiation (stabilizers) and agents who tend to adopt late (late-learners) have quite different combinations of the particular network properties *degree centrality* (DC), *closeness centrality* (CC), *betweenness centrality* (BC) and *individual clustering* (CL), as shown in Table 1. Stochastic analysis confirmed that these correlations are significant.

2.4 The change of language use in social networks

Another line of research involves multi-agent simulations to analyze language change, without applying game-theoretic models. Nettle (1999) simulated the interactive behavior of members embedded in a grid structure where spatial distance represents social distance and each agent can only possess one of two competing variants of a *linguistic item.* In each step of a simulation run each agent can keep her current variant or can adopt the other one, in dependence on which one has the higher 'impact' value. This impact value is a combination of i) a social impact value that integrates the number, social status and social distance of other members using this variant, and ii) a functional bias of the variant. Nettle tested his system for a range of different parameter settings and came to the following results: i) a full substitution of one variant over the other can only take place when super-influential high-status agents are involved, and ii) a functional bias alone is never enough for a new variant to replace the old one, since there is always a high social impact value required.

Ke, Gong, and Wang (2008) adopted a light version of Nettle's impact equation, and integrated it in small-world networks. Their results revealed that a new variant can replace an old one even without super-influential agents, but it must have an enormously high functional bias in comparison to its competitor.

Table 2. Research directions that apply simulation studies to examine the nature of emergence/change of (mostly linguistic) conventions in multi-agent populations (no requirement for completeness).

	A	**B**	**C**	**D**
Research topic	Emergence of behavioral conventions	Emergence of lexical conventions	Emergence of signaling conventions	General mechanisms of language change
Object	behav. convention	lexical convention	signaling convention	linguistic item
Process	emergence	emergence	emergence	change
Adoption	learning	learning (innovation)	imitation/learning	imitation
Interaction	coordination games	language games	signaling games	influence flow
Selected Literature	Young (1993) Shoham & Tennenholtz (1997) Airiau et al. (2014)	Steels (2002) Salazar et al. (2010) Centola & Baronchelli (2015)	Zollman (2005) Wagner (2009) Mühlenbernd & Franke (2012)	Nettle (1999) Ke et al. (2008) Fagyal et al. (2010)

In sum, both Nettle (1999) and Ke et al. (2008) used network simulation models to investigate the propagation of a new variant, but both also integrated a functional bias – a network independent value – that plays an important role in their analyses. This additional complexity is adequate for the appropriate research question, but it masks the way how network features might influence the propagation processes in language change on their own.

The following simulation study can be considered as 'state of the art' in network simulation studies to investigate language change: Fagyal, Swarup, Escobar, Gasser, and Lakkaraju (2010) used 'scale-free' small-world networks with directed ties denoting the direction of influence, considering eight different competing variants. Members of the network i) have a status value proportional to their outgoing ties, ii) adopt a new variant of a neighbor (connected member) with a probability proportional to the neighbor's status, and iii) have only one variant at a time in their inventory. Note that Fagyal et al. – in contrast to Nettle and Ke et al. – i) did not consider any functional bias, and focused on the impact of social biases in terms of status, and ii) defined social bias only in terms of network structural features, since an agent's status is defined by her number of outgoing ties. This point advanced the social network approach by explaining language change in terms of network properties, and Fagyal et al. followed this direction by taking such properties into consideration exclusively. Their results showed first of all that the propagation of a variant is realized by 'central influential' members, which is in accordance with Nettle's result of super-influential agents being a necessary condition for society-wide spread of a variant. As a second result, they showed that 'peripheral low-connected' members – so-called loners – are the source for innovations.

Table 2 shows an overview of all four research directions **A**, **B**, **C** and **D**.

3 Discussion

In its most general sense, language change can be seen as a new linguistic variant replacing an old one across some set of contexts. Although each instance of linguistic change has its temporal and spatial inception, one of the greatest challenges in sociolinguistics is to determine which variables support the initiation and propagation of a new variant (cf. L. Milroy, 1980; Croft, 2000; Labov, 2010). To better understand language change, we should ask: Which social circumstances support linguistic innovation? And which social environment is a fertile ground for a new variant to spread?

An insightful theory in sociolinguistics about the role of 'social network structure' in language change is the 'weak tie'-theory (J. Milroy & Milroy, 1985): As a result of speaker innovation, a new variant i) emerges generally on so-called 'weak ties'[5], and ii) spreads via 'central' members[6] of the local community. A number of field studies indirectly support the 'weak tie'-theory (cf. Labov, 1973, 1991, 2001; Trudgill, 1988; L. Milroy & Milroy, 1992; J. Milroy, 1996; Llamas, 2000; L. Milroy & Gordon, 2003) .

Note that also already discussed computational studies support the 'weak tie'-theory in parts. E.g. Fagyal et al. (2010) showed that a new variant spreads via 'central influential' members. On the other hand, the study of Mühlenbernd and Franke (2012) revealed that the initiators of a new convention to evolve – the founding fathers – have a high degree centrality value, but rather low values in closeness and betweenness centrality. This shows that one must be careful with the interpretation of particular labels of individual network features such as centrality, since they can be defined in diverse ways. Furthermore, none of the presented computational studies has tested the impact of the strength of a tie on innovation. All in all, with the aim to check the 'weak tie'-theory within a computational model, the main research question of this study is as follows:

> *Which individual network features support i) the innovation process of a new variant, and ii) the spread of this variant to become a convention?*

Note that research direction **D** is the only one that uses computational studies for the *change* of a linguistic conventions, whereas all other directions deal with the *emergence* of conventions. Furthermore the three studies of research direction **D** have another thing in common: they depict individual language change simply as the mechanism of one linguistic variant replacing another one, determined by influence values. Therefore they abstract from an essential feature of language: communicating information from a speaker to a hearer (see Mühlenbernd & Quinley, 2013).

Research directions **A**, **B** and **C** give insights in the way conventions emerge in a structured population. As classical game-theoretic studies, directions **A** and **C** draw on relevant properties, such as the Nash equilibrium, which is a great advantage in analyses of results. Furthermore, directions **B** and **C** have an explicit model of a communication process between a speaker and a hearer. In this sense, the model of a signaling game allows for i) describing the communication

[5] Weak ties are links that have a low strength, often defined by frequency or multiplexity of the connection. Weak ties connect mostly detached communities.

[6] Note that centrality in a network can be defined in multiple ways, as introduced in Section 4.4.

process in an explicit way, and ii) performing a game-theoretic analysis at the same time.

Having the advantages of signaling games as well as the research question in mind, I chose to apply modeling techniques in accordance with direction **C** to reproduce computational dynamics in accordance with direction **D** in a more fine-grained way. All in all, this study is primarily a continuation of research direction **C**, and it is inspired by research direction **D** particularly in the sense, that the change rather than the emergence of a linguistic convention is the process under investigation. In the following section I introduce the relevant technical background for my computational study.

4 Signaling Games, Learning and Networks

This section gives a coarse technical and theoretical background to understand the important concepts of this article: the signaling game, reinforcement learning, the innovation mechanism, and some basic notions of network theory.

4.1 Signaling Games

A signaling game $SG = \langle\{S, R\}, T, M, A, Pr, U\rangle$ is a game played between a sender S and a receiver R. T is a set of information states, M is a set of messages and A is a set of interpretation states (or actions). $\Pr \in \Delta(T)$[7] is a probability distribution over T and describes the probability that an information state is topic of communication. $U : T \times A \to \mathbb{R}$ is a utility function that generally determines how well an interpretation state matches an information state.

Let us take a look at the simplest variant of the game where we have two information/interpretation states and two messages: $T = \{t_1, t_2\}$, $M = \{m_1, m_2\}$, $A = \{a_1, a_2\}$, a flat probability distribution: $Pr(t) = 1/|T|\ \forall t \in T$, and a simple utility function that gives a positive value iff the interpretation state a matches the information state t, marked by the same index: $U(t_i, a_j) = 1$ iff $i = j$, else 0. Figure 1 shows the *extensive form game* of this simple variant, which depicts the way this game is played: an information state $t \in T$ is chosen with prior probability $Pr(t)$ (here both states are equiprobable), which the sender wants to communicate to the receiver[8] by choosing a message $m \in M$. The receiver wants to decode this message by choosing an interpretation state $a \in A$. Communication is successful iff the information state matches the interpretation state, which is indicated by the utility value at the end of each leaf. In this study I only consider signaling games that are $n \times n'$-games as defined in Definition 1.

Definition 1 **($n \times n'$-signaling game).** *An $n \times n'$-signaling game is a signaling game $SG = \langle\{S, R\}, T, M, A, Pr, U\rangle$ with:*

$|T| = |A| = n$, $|M| = n'$, $\forall t \in T : Pr(t) = 1/|T|$ *and* $U(t_i, a_j) = \begin{cases} 1 \text{ if } i = j \\ 0 \text{ else} \end{cases}$

[7] $\Delta(X)$ denotes the set of all probability distributions over a random variable X.

[8] Informally spoken, the information state came to the sender's mind. In game theory we say that the state is chosen by an invisible participant, called nature N.

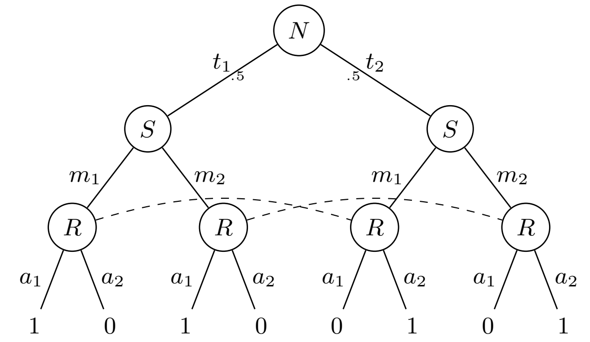

Fig. 1. Extensive form game for the 2×2- game. Note that the dashed lines denote situations the receiver cannot distinguish, since he does not know the information state t_i of the sender, but only the message m_j he received.

Fig. 2. The two signaling systems of the 2×2-game.

Note that messages are initially meaningless in this game, but meaningfulness can arise from regularities in behavior. Behavior is here 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.

Now, what circumstances can tell us that a message is attributed with a meaning? The answer is: this can be indicated by the combination of sender and receiver strategy, called strategy profile. A message has a meaning between a sender and a receiver, if both use pure strategies that constitute a specific *isomorphic strategy profile*. For the 2×2-game there are exactly 2 such strategy profiles, as depicted in Figure 2. Here in profile L_1 the message m_1 has the meaning of state t_1/a_1 and message m_2 has the meaning of state t_2/a_2. For profile L_2 it is exactly the other way around.

Lewis (1969) called such strategy profiles *signaling systems*. Such signaling systems have interesting properties. Not only is the meaning of a messages defined, but it can also be shown that signaling systems i) ensure perfect communication and maximal utility, ii) are Nash equilibria over expected utilities (Crawford & Sobel, 1982), and iii) are evolutionary stable states (cf. Wärneryd, 1993; Huttegger, 2007). Additionally, note that the number of signaling systems increases disproportionately with the number of states and messages: an $n \times n'$-game has ${}^{n'!}/_{(n'-n)!}$ different possible signaling systems.

At this point it is explained how semantic meaning can be expressed by participants' communicative behavior: a message has a meaning, if sender and receiver communicate according to a signaling system. But this does not explain at all, how participants come to such a signaling system in the first place, by expecting that messages are initially meaningless. To explore the paths that might lead from a meaningless to a meaningful message, it is necessary to explore the process that leads from participants' arbitrary communicative behavior to a behavior that constitutes a signaling system. Such a process can be simulated by repeated signaling games, where the participants' behavior is guided by *update dynamics*, such as imitation and learning rules. An increasingly used learning rule in combination with repeated signaling games is *Roth-Erev reinforcement learning* (cf. Barrett, 2009; Barrett & Zollman, 2009; Skyrms, 2010).

4.2 Reinforcement Learning

Reinforcement learning can be captured by a simple model based on urns, also known as *Pólya urns* (Roth & Erev, 1995). 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 choice. By adding or removing balls from an urn after each access, an agent's behavior is gradually adjusted. For signaling games, the sender has an urn $\mho_t$ for each state $t \in T$, which contains balls for different messages $m \in M$. The number of balls of type m in urn $\mho_t$ designated with $m(\mho_t)$, the overall number of balls in urn $\mho_t$ with $|\mho_t|$. If the sender is faced with a state t she draws a ball from urn $\mho_t$ and sends message m, if the ball is of type m. The same holds in the same way for the receiver. The resulting *sender response rule* σ and *receiver response rule* ρ is given in Equation 1 and 2, respectively.

$$\sigma(m|t) = \frac{m(\mho_t)}{|\mho_t|} \quad (1) \qquad\qquad \rho(a|m) = \frac{a(\mho_m)}{|\mho_m|} \quad (2)$$

The learning rule is realized by changing the urn content in dependence of the communicative success. The standard account works as follows: if communication via t, m and a is successful, the number of balls in the sender's urn $\mho_t$ is increased by $\alpha \in \mathbb{R}$ balls of type m (sender update), and the number of balls in urn $\mho_m$ is increased by $\alpha \in \mathbb{R}$ balls of type a (receiver update).[9] In this way successful communicative behavior is more probable to reappear in subsequent rounds.

The learning mechanism can be intensified by *lateral inhibition*: if communication via t, m and a is successful, not only will the number of ball type m in urn $\mho_t$ be increased, but also will the number of all other ball types $m' \in M \setminus \{m\}$ be decreased by $\gamma \in \mathbb{R}$. Similarly, for the receiver. Franke and Jäger (2012) introduced the concept of *lateral inhibition* for reinforcement learning in signaling games and showed that it leads the system more speedily towards pure strategies.

Furthermore, *negative reinforcement* can be used to punish unsuccessful behavior. It changes urn contents in the case of unsuccessful communication in the following way: if communication via t, m and a is unsuccessful, the number of

[9] Note that the *number* of balls is just a metaphor for better comprehensibility of the principle, and therefore the incremental value per urn can be $\in \mathbb{R}$.

balls in the sender's urn $\mho_t$ will be decreased by $\kappa \in \mathbb{R}$ balls of type m; and the number of balls in the receiver's urn $\mho_m$ will be decreased by κ balls of type a.

Note that Roth-Erev reinforcement learning has the property to slow down the learning effect: if the total number of balls in an urn increases over time, but the rewarding value α is a fixed value, then the learning effect mitigates. A way to prevent the learning effect from slowing down is to keep the overall number of balls $|\mho|$ on a fixed value Ω by scaling the urn content appropriately after each round of play. Such a setup is a variant of so-called Bush-Mosteller reinforcement (Bush & Mosteller, 1955). All in all, a reinforcement learning setup for a signaling game can be given as defined in Definition 2.

Definition 2 (Reinforcement Learning Setup). *A reinforcement learning setup for a* $n \times n'$*-signaling game* $SG = \langle \{S, R\}, T, M, A, Pr, U \rangle$ *is a tuple* $RL = \langle (\sigma, \rho), \alpha, \kappa, \gamma, \Omega, \phi \rangle$*, with:*

- *sender response rule* $\sigma : T \to \Delta(M)$
- *receiver response rule* $\rho : M \to \Delta(A)$
- *reward value* $\alpha \in \mathbb{R}$
- *punishment value* $\kappa \in \mathbb{R}$
- *lateral inhibition value* $\gamma \in \mathbb{R}$
- *urn size value* $\Omega \in \mathbb{R}$
- *function determining initial urn settings* ϕ

I would like to make three points clear: first of all, note that the response rules (Equations 1 and 2) of reinforcement learning describe the way agents make their decisions. It should be clear from the equations that such decisions are not based on any assumption for rationality. In other terms: an agent does neither make any use of assumptions about the other agent's behavior, nor has she direct access to the other agent's history. Agents also do not try to achieve a future goal. Agents simply choose more probably that strategy that has worked better in the course of their own history. This distinguishes reinforcement learning from other learning rules that incorporate a stronger notion of rationality, such as Fictitious Play (Brown, 1951).[10] Note that this lack of rationality in agents is desirable, since language change basically happens in an unconscious manner, thus without the presence of rational decisions, as discussed in Section 1.

As the second point, note that the notions of reinforcement learning as used here differ from classical notions in computer science. For example, Sutton and Barto (1998) identify three obligatory main subelements of a reinforcement learning setup: a *policy*, a *reward function*, and a *value function*. These subelements can also be found in the given setup: the policy is represented by the response rules (Equations 1 and 2), the reward function is represented by the update mechanism that I described informally (including the reward value α and punishment value β), and the value function can be represented by an expected util-

[10] For a discussion and comparison of reinforcement learning and Fictitious Play in signaling game playing agents, see e.g. Mühlenbernd (2011).

ity function over strategy pairs.[11] Furthermore, note that the classical definition of a reinforcement learning problem in computer science is "a straightforward framing of the problem of learning from interaction to achieve a goal" (Sutton & Barto, 1998, p. 51). As mentioned before, this also differs to the given account, where agents are not intrinsically goal-oriented, but simply driven by the need to communicate successfully.

Thirdly, note that this account of reinforcement leaning does not allow for much *exploration*, which is the possibility to test new strategies, even if the current strategy works well. Once an agent has learned a strategy that works perfectly with her communication partners – in other words: a signaling system – she has no incentive to change her strategy a bit.[12] This point is an unrealistic shortcoming of the given account, since if human language would work that way, it would not change once its semantic system is established for all relevant concepts of a society. But linguistic meaning is only stable from a myopic point of view, since it is incessantly changing in the long run. These changes are result of the fact that there is always a possibility for exploration, may it be conscious or unconscious, and this possibility can be realized by a mechanism that I will call *innovation*, which will be introduced in the next section.

4.3 The Mechanism of Innovation

With the goal to analyze dynamics concerning the change of signaling conventions, an essential additional feature for the reinforcement learning setup is *innovation.* The basic idea stems from Skyrms (2010) – he calls it *invention* – and can be described as follows: each sender urn contains, next to the balls for each message, an additional ball type, which Skyrms calls *black balls.* Whenever the sender draws a black ball from this urn, she sends a completely new message that has never been sent before. In other words, the sender invents a new message. Further experiments with this setup were made by Alexander et al. (2012) for 2-players games and by Mühlenbernd and Nick (2014) for spatial structures.

The second study (Mühlenbernd & Nick, 2014) used a reinforcement learning setup with negative reinforcement and lateral inhibition. In this setup the black balls of the agents' sender urns can increase and decrease in dependence of communicative success. By naming the total number of an agent's black balls her *force of innovation*, the study revealed an interesting relationship between society-wide force of innovation and communicative success: increasing communicative success leads to decreasing force of innovation, and vice versa.[13]

[11] For a definition of expected utilities over strategy pairs, see e.g. Mühlenbernd (2011). Note that signaling systems maximize expected utilities, therefore they are optimal according to such a value function.

[12] As mentioned earlier: signaling systems are Nash equilibria over expected utilities and evolutionary stable strategies. Furthermore, in combination with the current reinforcement learning setup, agents that have learned a signaling system stick with it with a zero probability to change.

[13] It was shown for experiments with 3-agent populations that the force of innovation and communicative success reveal a significant negative correlation.

Unfortunately, the fact that the number of possible messages is verbatim unlimited (since the innovation mechanism produces a new message every time a black ball is drawn) leads to the phenomenon that larger populations will probably never find an agreement (or at least need an unmanageable amount of runtime), but end up in a chaos of a never ending production of new messages. This phenomenon was shown even for a little community of 6 agents. But it has also been shown that by limiting the possible message set, this problematic nature of a never ending chaos can be avoided (Mühlenbernd & Nick, 2014).

In such a game with a limited message set, agents do not send a completely new message, when they draw a black ball, but send a randomly chosen message from a fixed message set M. In this way the game keeps its innovative nature (if $|M| \gg n$), but avoids runtime problems. Such a game is called an $n \times n'_m$-signaling game and defined as given in Definition 3.

Definition 3 **($n \times n'_m$-signaling game).** *An $n \times n'_m$-signaling game is an $n \times n'$-signaling game with $m = |M|$ possible messages and an optimal number of n' messages.*[14]

By conducting experiments with $n \times n'_m$-signaling games and a reinforcement learning setup with innovation on a large population (e.g. a 100×100 toroid lattice structure), Mühlenbernd and Nick (2014) revealed two thing:

1. Once a population has learned one unique signaling convention and reaches perfect communication, the force of innovation has dropped to zero. In other words, the society has reached a stable state: with usage of a unique semantic convention and without any spirit for innovation.
2. If multiple local convention coexist, communication cannot be globally optimal, and therefore the force of innovation is never zero. Such a situation causes little changes in the interaction between connected agents of neighboring *local signaling conventions* (as defined in Definition 5) and might lead to the change of the whole signaling convention of an area over time.

The second point resembles exactly the dynamic picture expected from the phenomenon under investigation: *change induced by variation.* Since in this study I am interested in the evaluation of more realistic social network structure, so-called small-world networks, I introduce some basic notions of network theory in the following section.

4.4 Basic Notions of Network Theory

To ensure that a network structure resembles a realistic interaction structure of human populations, it should have *small-world* properties; c.f. Jackson (2008) found out that these properties show in the analysis of human friendship networks. According to this line of studies, the essential two properties of small-world networks are i) a short *characteristic path length*, and ii) a high *clustering*

[14] Note that *optimal number n'* defines the number of messages which are necessary to create a signaling system.

coefficient (Watts & Strogatz, 1998).[15] Additionally, most often human networks display a third property, namely to be *scale-free*: the frequency of agents with ever larger numbers of connections roughly follows a power-law distribution. In this sense I consider a special kind of a scale-free network, which is both scale-free and has small-world properties (Barabási & Albert, 1999). This network type can be constructed by a particular *preferential attachment* algorithm – called Holme-Kim algotithm – that takes two parameters: m that controls the *network density*, and p that controls the *clustering coefficient* (Holme & Kim, 2002).[4]

A main goal of this work is to investigate the relationship between the change of meaning conventions and the structural properties of the network and its members. My experiments showed that by analyzing the agents' behavior and dynamic features during a simulation run, there seems to be an explanatory value of network properties that express an agent's globally and/or locally connectivity and her embeddedness. Following Mühlenbernd and Franke (2012), in order to capture these properties more adequately, I investigated suitable notions from social network theory: *degree centrality* (DC) describes the local connectivity of an agent, *closeness centrality* (CC) her global centrality, *betweenness centrality* (BC) her global connectivity in terms of information flow, and *individual clustering* (CL) her embeddedness into the local structure.[4]

As I argued in Section 3, also the strength of ties between agents might play an important role in language change. Easley and Kleinberg (2010) showed that the strength of a tie between two agents has basically a strong linear correlation with the overlap of both agents' neighborhoods. To keep things easy I define the strength of a tie by this neighborhood overlap. Furthermore, since my analysis deals with agents rather than with ties between them, I calculate an agent's *ties strength* TS as the average strength value of all ties of this agent:

Definition 4 (Ties Strength). *For a given network $G = \langle N, E \rangle$, whereby N is a set of nodes (here also agents) and $E \subseteq N \times N$ is a set of edges, the* ties strength *$TS(n)$ of an agent $n \in N$ is defined as*

$$TS(n) = \frac{\sum\limits_{m \in N(n)} \frac{N(n) \cap N(m)}{N(n) \cup N(m)}}{|N(n)|} \tag{3}$$

whereby $N(n) = \{m \in N | \{n, m\} \in E\}$ is the set of neighbors of agent $n \in N$.

Note: the notions of DC, CC, BC, CL and TS describe *static* network properties of an agent, since they do not change during a simulation run and are determined by the network structure and the agent's position inside it.

As a final remark, since agents in populations generally agree on signaling systems as groups, I call such a group-wide signaling system a *signaling convention*, as given in Definition 5.

[15] For the definition of these network properties I refer to Jackson's *Social and Economic Networks* (Jackson, 2008), Chapter 2.

Definition 5 (Signaling Convention). *For a given network structure of agents that play the repeated signaling game with their connected neighbors, a signaling convention is a signaling system that is used by a group of connected agents.*

5 Simulating Language Change

In my experiments agents in a social network communicate via signaling games and update via innovative reinforcement learning. This leads to the effect that i) multiple local conventions emerge (cf. Zollman, 2005; Wagner, 2009; Mühlenbernd, 2011; Mühlenbernd & Franke, 2012), and ii) agents invent new messages from time to time, since communication is not perfectly successful in a society with multiple conventions and therefore the force of innovation stays on a non-zero level (cf. Mühlenbernd & Nick, 2014). As my experiments will show, while mostly invented messages disappear as fast as they appear, from time to time new variants can spread and establish new regional conventions. Therefore, I analyze if particular structural features support emergence and spread of innovation. I examine if the results support the 'weak tie'-theory.

5.1 Experimental Settings

I conducted simulation runs of agents that are placed in a social network structure. Per simulation step every agent communicates by playing a signaling game with each of her direct neighbors. Each agent updates her behavior by innovative reinforcement learning. The concrete settings of the experiments were as follows:

- *network structure*: a scale-free network with 500 agents (Holme-Kim algorithm (Holme & Kim, 2002) with $m = 2$ and $p = .8$)
- *signaling game*: a $3 \times 3_9$-signaling game
- *reinforcement learning setup* (see Definition 2): Bush-Mosteller reinforcement with negative reinforcement, lateral inhibition ($\alpha = 1$, $\kappa = 1$, $\gamma = 1$, $\Omega = 20$), innovation, and the following initial condition ϕ: each sender urn contains only black balls, and each receiver urn has an equiprobable distribution of each ball type
- *initial simulation condition*: for the first 100 simulation steps the network is divided in 10 connected components, and agents communicate exclusively with neighbors of the same component
- *stop condition*: reaching 100,000 simulation steps
- *number of simulation runs*: 10

Since I am interested in the mechanisms that show how and why semantic conventions change, not how they evolve from the scratch, the simulation runs were started with the given *initial simulation condition*, which ensures that 10 already established local signaling conventions are given from the beginning (see Figure 3). In the following the results of the simulation runs are presented.

Fig. 3. Exemplary scale-free network with 500 agents and 10 already established local signaling conventions, differenced by color.

5.2 Global Values

To get a global picture of how the population behaves I measured the following global values:

- *average communicative success*: the utility value averaged over all plays during a simulation step
- *average force of innovation*: the number of black balls averaged over all agents after a simulation step
- *global number of signaling conventions*: the total number of different signaling conventions existent in the population at the given simulation step

Figure 4 shows the global values over time for an exemplary simulation run. This example depicts the general nature of all simulation runs in this experiment: the average communicative success increases up to a value of around .85, and the number of signaling conventions decreases down to a value below 10. Furthermore, both values oscillate quite strongly: the communicative success oscillates between .8 and .9, the force of innovation oscillates between .4 and .6, and the number of signaling conventions oscillates between 6 and 10.

Especially the oscillation of the number of signaling conventions is an indicator for local reactivity. To get a better understanding of what is actually happening, Figure 5 shows a sequence of the first 10,000 simulation steps for the

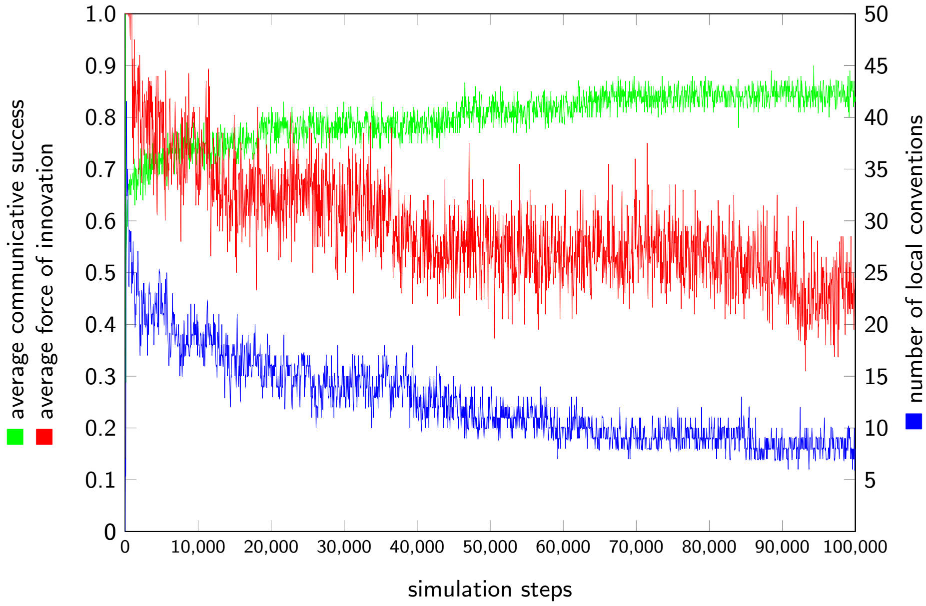

Fig. 4. Simulation run of a $3 \times 3_9$-signaling game in a population of 500 agents placed in a small-world network: average communicative success, force of innovation and the number of society-wide signaling convention over 100,000 simulation steps.

number of learners for 6 different signaling conventions: here regions of new conventions emerge, grow to a specific amount and (mostly) get eventually extinct. This pattern shows quite nicely how semantic change is realized: an innovation is made at one point in time and place, then it spreads and its number of speakers increases to a specific amount and constitutes a region of a new signaling convention. Such a region might decrease in members and finally get extinct at one point. The research question of this study – as discussed in Section 3 – involves the relationship of an agent's static network properties and her contribution to the process of semantic change.[16] Note that such static network properties were already defined in Section 4.4. The next step is now to define agent's properties that characterize her contribution to stability, innovation or spread. For that purpose I define *dynamic properties* that depict the behavioristic characteristics of an agent during a whole simulation run.

5.3 Agent Features

Considering the dynamic picture of language change in the simulation runs, I was interested in detecting specific roles of agents that might support language

[16] E.g. the 'weak tie'-theory assumes a high negative correlation of an agent's tie strength (TS, see Definition 4) and her contribution to innovation, i.o.w. to start new regions of signaling conventions.

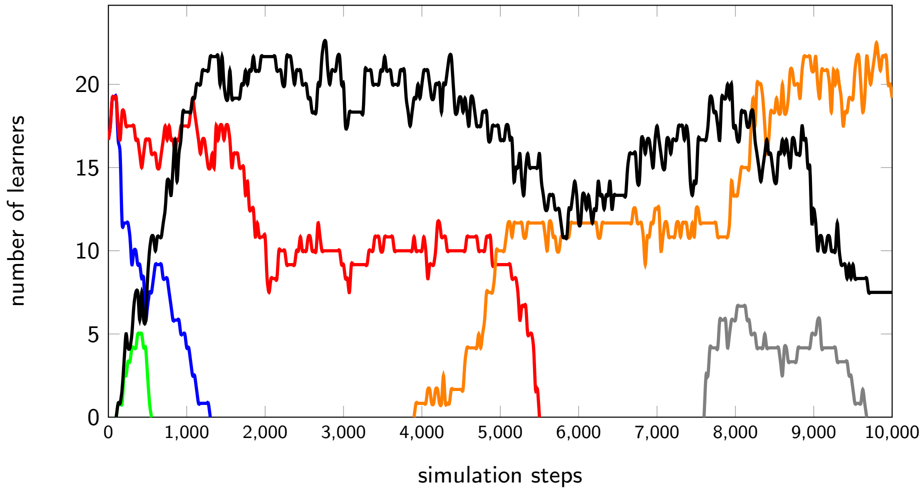

Fig. 5. Simulation run of a 3×3^9-game in a population of 500 agents placed in a social network: the number of learners for 6 specific different signaling conventions (distinguished by different colors) for the first 10,000 simulation steps.

change or strengthen local conventions. Following the study of Mühlenbernd and Franke (2012), I was particularly interested in the way an agent's *static* structural features and *dynamic* behavioral features might correlate. Static features are given by an agent's network properties *ties strength* TS, *degree centrality* DC, *closeness centrality* CC, *betweenness centrality* BC and *clustering coefficient* CL, as introduced in Section 4.4.

Dynamic features of an agent can be defined by her behavior in relation i) to her former behavior or ii) to other agents' behavior during a simulation run. Since I was interested in the way agents were involved in the innovation and spread of a new variant, I defined the dynamic features *innovation skill* and *impact*. To compare these values with a number of further dynamic features, I also defined features called *loyalty*, *majority preference*, *interiority*, *fraternity*, *mutual intelligibility* and *adaptivity*. For an agent n these features are defined as follows:

- *loyalty* $LOY(n)$: the proportion of simulation steps agent n played her favorite strategy (most often played strategy)
- *majority preference* $MAJ(n)$: the proportion of agents using the same strategy that agent n uses at a given simulation step, averaged over all simulation steps
- *interiority* $INT(n)$: the proportion of simulation steps for which agent n has exclusively neighbors using the same strategy
- *fraternity* $FRA(n)$: the proportion of agent n's neighbors using the same strategy that agent n uses at a given simulation step, averaged over all simulation steps

- *mutual intelligibility* $MI(n)$: the average MI[17] value of agent n to her neighborhood at a given simulation step, averaged over all simulation steps
- *adaptivity* $AD(n)$: the proportion of simulation steps at which agent n switched to a neighbor's strategy
- *impact* $IMP(n)$: the proportion of simulation steps at which a neighbor of agent n switched to agent n's strategy
- *innovation skill* $INV(n)$: the proportion of simulation steps at which agent n switched to a new strategy, which no neighbor is actually using

5.4 Feature Analysis

In my analysis I measured the correlation of all 5000 data points[18] and for each possible combination of features. The resulting plot is shown in Figure 6: correlations are depicted as circles (lower-left half) and correlation values (upper-right half), whereby i) the size of the circles as well as the color saturation represents the strength of the Pearson-correlation, and ii) the color represents the direction of the relationship (positive: blue, negative, red).[19]

The results show first of all: the data support the weak-tie theory, since i) INV has a high negative correlation with TS, and ii) IMP reveals a high positive correlation with all three centrality properties DC, CC and BC. Therefore, agents that have contributed much to the innovation of new conventions (high INV value) tend to have very weak ties on average (low TS value), and vice versa. Furthermore, agents who have had a strong impact on neighbors (high IMP value) are – in terms of all three centrality values – centrally located in the network. All in all, agents with weak ties were mostly responsible for the innovation of new conventions, which were afterwards propagated mostly by central members.

There are further high correlations involving innovation skills: first of all, INV has a very high negative correlation with LOY. Intuitively, this relationship can go in two directions: first of all, agents who stick mostly to one convention (high LOY value) do not contribute much to the innovation of new conventions (low INV value). And second, agents who contribute much to the innovation of new conventions (high INV value) do not stick mostly to one convention (low LOY value). Both directions are quite reasonable, and, taken together, indicate that conservative agents (high LOY value) and innovative agents (high INV value) constitute separate groups. Additionally, INV has a very high negative correlation with MI. This makes sense, since very innovate agents (high INV value) produce conventions that are not know to the neighboring agents, and their mutual intelligibility is – at least initially – quite low (low MI value). All together, the high correlation of INV with LOY and with MI both is not a very

[17] The mutual intelligibility value MI reproduces the expected utility for two different strategy pairs. For the definition see Mühlenbernd and Nick (2014), Definition 3.

[18] Data points are the agents' features; for 10 simulation runs with 500 agents each.

[19] Due to the fact that some numbers are hard to spot, all values of these *Pearson-correlations* are also given in Table 3 (Appendix).

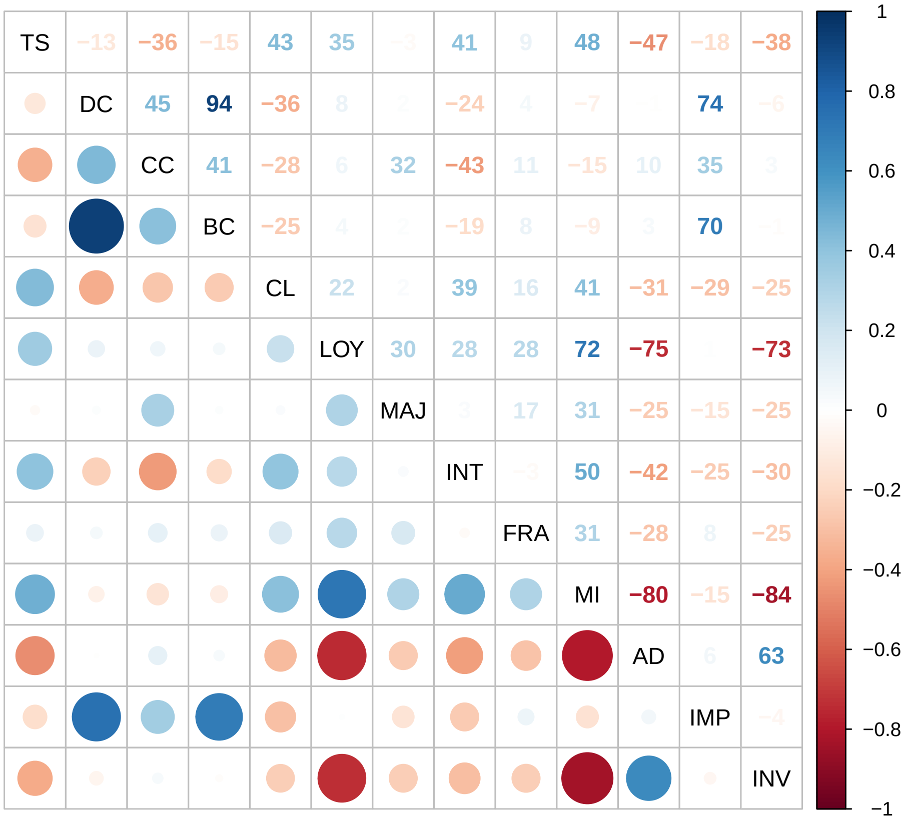

Fig. 6. The correlations for all different pairs of features: the static network properties *ties strength* TS, *degree centrality* DC, *closeness centrality* CC, *betweenness centrality* BC and *clustering coefficient* CL; and the dynamic behavioral features *loyalty* LOY, *majority preference* MAJ, *interiority* INT, *fraternity* FRA, *mutual intelligibility* MI, *adaptivity* AD, *impact* IMP and *innovation skill* INV.

surprising result, but is shows that the resulting correlations are comprehensible and not counterintuitive by any chance.

Finally, another dynamic feature that significantly correlates with static features is INT. Agents with a high INT value spend most of the time inside a language region, thus having only neighbors which use the same convention. As their static features reveal, they are strongly embedded in their local environment (positive correlations with CL and TS), but rater peripherally positioned (negative correlations with DC, CC and BC). These agents are probably the ones who sustain local conventions, since they show the feature combination of *stabilizers* of the study by Mühlenbernd and Franke (2012) (see Table 1).

Note that all correlations discussed here are statistically highly significant, each with a p-value below 0.01.

6 Conclusion & Outlook

To understand the nature of languages change, one of the greatest challenges in sociolinguistics is determining which social variables support the initiation and propagation of a new linguistic variant (cf. Labov, 2010). An insightful theory about the role of 'social network structure' in initiation and propagation processes of language change is the 'weak tie'-theory (J. Milroy & Milroy, 1985). This theory was directly and indirectly supported by a number of field studies (cf. L. Milroy, 1980; Trudgill, 1988; Labov, 2001). But those empirical studies concentrate mostly on so-called 'egocentris personal networks' (L. Milroy & Gordon, 2003) and fail to capture the entirety of a wholes population social network structure (for a discussion, see e.g. Mühlenbernd and Quinley (2017)). With this study, I present an alternative approach for testing theories from sociolinguistics/social sciences by applying a multi-agent account that allows for a precise definition and modeling of social network structure and features.

The central model for communication is a game-theoretic one: the signaling game. To analyze the dynamics of language change with this model might be – on the first view – an ambitious challenge by considering that signaling games are designed in a way that players are generally attracted to convention and stability. For all that, I was particularly interested in the way environmental variables in terms of network structure might describe characteristics that promote or mitigate semantic change. For that purpose I made simulation experiments on social network structures of agents that play the signaling game repeatedly with connected neighbors and update their behavior by a simple dynamics: Roth-Erev reinforcement learning. I extended this learning account by an additional feature – innovation – that supports the changing nature of the population's dynamics.

In my analysis I compared different features of agents – static network properties and dynamic behavioral properties of agents – to extract the characteristics of different roles participating in language change. All in all, the results support the 'weak tie'-theory, under the assumption that the strength of a tie is defined by its neighborhood overlap, as given in Definition 4.[20]

Since this study gives only a first impression where to look for network induced forces of language change, there are at least two directions necessary to reveal deeper results. First of all, the current data should be further analyzed by i) detecting causal dependencies between features that deliver more informative data, or ii) using regression models to find out if there are non-trivial interactions – e.g. non-linear dependencies – between network properties and dynamic features of agents. Second, my current results indicate to analyze further i) static properties, such as information flow measures (Jackson, 2008) or closeness vitality (Koschützki et al., 2005); and ii) dynamic features like individual force of innovation, number of known messages, or the growth magnitudes of an agent's newly innovated signaling system. These two additional steps are

[20] An open issue here is to test the 'weak tie'-theory when the strength of a tie is defined in other ways (cf. Mühlenbernd & Quinley, 2017).

currently investigated and can hopefully enrich subsequent work by delivering deeper insights into the role of innovation in dynamics of semantic change.

References

Airiau, S., Sen, S., & Villatoro, D. (2014). Emergence of conventions through social learning: Heterogeneous learners in complex networks. *Autonomous Agents and Multi-Agent Systems*, *28*(5), 779-804.

Alexander, J., Skyrms, B., & Zabell, S. (2012). Inventing new signals. *Dynamic Games and Applications*, *2*(1), 129-145.

Axelrod, R. (1984). *The evolution of cooperation*. New York: Basic Books.

Barabási, A.-L., & Albert, R. (1999). Emergence of scaling in random networks. *Science*, *286*(5439), 509–512.

Barrett, J. A. (2009). The evolution of coding in signaling games. *Theory and Decision*, *67*, 223–237.

Barrett, J. A., & Zollman, K. J. S. (2009). The role of forgetting in the evolution and learning of language. *Journal of Experimental and Theoretical Artificial Intelligence*, *21*(4), 293–309.

Brown, G. (1951). Iterative solution of games by fictitious play. In T. Koopmans (Ed.), *Activity analysis of production and allocation* (p. 375). New York: Wiley.

Bush, R., & Mosteller, F. (1955). *Stochastic models of learning*. New York: John Wiley & Sons.

Centola, D., & Baronchelli, A. (2015). The spontaneous emergence of conventions: An experimental study of cultural evolution. *Proceedings of the National Academy of Sciences*, *112*(7), 1989–1994.

Crawford, V. P., & Sobel, J. (1982). Strategic information transmission. *Econometrica*, *50*, 1431–1451.

Croft, W. (2000). *Explaining language change: an evolutionary approach*. Harlow, Essex: Longman.

Easley, D., & Kleinberg, J. (2010). *Networks, crowds, and markets: Reasoning about a highly connecetd world*. Cambridge: Cambridge University Press.

Fagyal, Z., Swarup, S., Escobar, A. M., Gasser, L., & Lakkaraju, K. (2010). Centers and peripheries: Network roles in language change. *Lingua*, *120*, 2061–2079.

Franke, M., & Jäger, G. (2012). Bidirectional optimization from reasoning and learning in games. *Journal of Logic, Language and Information*, *21*(1), 117–139.

Holme, P., & Kim, B. J. (2002). Growing scale-free networks with tunable clustering. *Physical Review E*, *65*(2), 026107-1–026107-4.

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.

Koschützki, D., Lehmann, K., Peeters, L., Richter, S., Tenfelde-Podehl, S., & Zlotowski, O. (2005). Centrality indices. In D. Koschützki, K. Lehmann, L. Peeters, S. Richter, S. Tenfelde-Podehl, & O. Zlotowski (Eds.), *Network analysis* (Vol. 3418, pp. 16–61). Heidelberg/New York: Springer.

Labov, W. (1973). The linguistic consequences of being a lame. *Language in Society*, *2*(1), 81–115.

Labov, W. (1991). The three dialects of english. In P. Eckert (Ed.), *New ways of analyzing sound change* (pp. 1–44). New York: Academic Press.

Labov, W. (2001). *Principles of linguistic change, volume 2: Social factors.* Malden, MA: Blackwell.

Labov, W. (2010). *Principles of linguistic change, volume 3: Cognitive and cultural factors.* John Wiley & Sons.

Lewis, D. (1969). *Convention.* Cambridge: Harvard University Press.

Llamas, C. (2000). *Variation in the north-east of england.* Michigan State University. (Paper presented at NWAV 29)

Milroy, J. (1996). A current change in british english: Variation in (th) in derby. *Newcastle and Durham Papers in Linguistics*, *4*, 213–222.

Milroy, J., & Milroy, L. (1985). Linguistic change, social network and speaker innovation. *Journal of Linguistics*, *21*(2), 339–384.

Milroy, L. (1980). *Language and social networks.* Oxford: Blackwell.

Milroy, L., & Gordon, M. (2003). *Sociolinguistics: Method and interpretation.* John Wiley & Sons.

Milroy, L., & Milroy, J. (1992). Social network and social class: Towards an integrated sociolinguistic model. *Language in Society*, *21*, 1–26.

Mühlenbernd, R. (2011). Learning with neighbours: Emergence of convention in a society of learning agents. *Synthese*, *183*(S1), 87-109.

Mühlenbernd, R., & Franke, M. (2012). Signaling conventions: Who learns what where and when in a social network? In T. Scott-Phillips, M. Tamariz, E. A. Cartmill, & J. Hurford (Eds.), *Proceedings of the 9th international conference on the evolution of language (*Evolang IX*)* (p. 242-249).

Mühlenbernd, R., & Nick, J. (2014). Language change and the force of innovation. In S. Katrenko & K. Rendsvig (Eds.), *Pristine perspectives on logic, language, and computation* (Vol. 8607, pp. 194–213). Heidelberg/New York: Springer.

Mühlenbernd, R., & Quinley, J. (2013). Signaling and simulations in sociolinguistics. In K. Shwayder (Ed.), *University of pennsylvania working papers in linguistics* (Vols. 19, Article 16).

Mühlenbernd, R., & Quinley, J. (2017). Language change and network games. *Language and Linguistics Compass*, *11*(2), e12235.

Nettle, D. (1999). Using social impact theory to simulate language change. *Lingua*, *108*(2–3), 95–117.

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.
Russell, B. (1921). *The analysis of mind.* Unwin Brothers Ltd.
Salazar, N., Rodriguez-Aguilar, J. A., & Arcos, J. L. (2010). Robust coordination in large convention spaces. *AI Communications*, *23*, 357–372.
Shoham, Y., & Tennenholtz, M. (1997). On the emergence of social conventions: modeling, analysis, and simulations. *Artificial Intelligence*, *94*(1-2), 139-166.
Skyrms, B. (2010). *Signals: Evolution, learning & information.* Oxford: Oxford University Press.
Steels, L. (2002). Grounding symbols through evolutionary language games. In A. Cangelosi & D. Parisi (Eds.), *Simulating the evolution of language* (p. 211-226). Springer.
Sutton, R., & Barto, A. (1998). *Reinforcement learning: an introduction.* The MIT Press.
Trudgill, P. (1988). Norwich revisited: Recent changes in an english urban dialect. *English World-Wide*, *9*(1), 33–49.
Wagner, E. (2009). Communication and structured correlation. *Erkenntnis*, *71*(3), 377–393.
Wärneryd, K. (1993). Cheap talk, coordination, and evolutionary stability. *Games and Economic Behavior*, *5*(4), 532–546.
Watts, D. J., & Strogatz, S. H. (1998). Collective dynamics of small-world networks. *Nature*, *393*, 440–442.
Young, H. P. (1993). The evolution of conventions. *Econometrica*, *2*(1), 129-145.
Zollman, K. J. S. (2005). Talking to neighbors: The evolution of regional meaning. *Philosophy of Science*, *72*(1), 69–85.

Appendix: Correlation Values

Table 3. The *Pearson-correlation* values (accurate to two decimal places) over 5000 data points (10 simulation runs $\times$ 500 agents) for all different pairs of features: the static network properties *tie strength* TS, *degree centrality* DC, *closeness centrality* CC, *betweenness centrality* BC and *clustering coefficient* CL; and the dynamic behavioral features *loyalty* LOY, *majority preference* MAJ, *interiority* INT, *fraternity* FRA, *mutual intelligibility* MI, *adaptivity* AD, *impact* IMP and *innovation skill* INV.

	TS	DC	CC	BC	CL	LOY	MAJ	INT	FRA	MI	AD	IMP	INV
TS	1.0	-.13	-.36	-.15	.43	.35	-.03	.41	.09	.48	-.47	-.18	-.38
DC	-.13	1.0	.45	.94	-.36	.08	.02	-.24	.04	-.07	-.01	.74	-.06
CC	-.36	.45	1.0	.41	-.28	.06	.32	-.43	.11	-.15	.10	.35	.03
BC	-.15	.94	.41	1.0	-.25	.04	.02	-.19	.08	-.09	.03	.70	-.01
CL	.43	-.36	-.28	-.25	1.0	.22	.02	.39	.16	.41	-.31	-.29	-.25
LOY	.35	.08	.06	.04	.22	1.0	.30	.28	.28	.72	-.75	.01	-.73
MAJ	-.03	.02	.32	.02	.02	.30	1.0	.03	.17	.31	-.25	-.15	-.25
INT	.41	-.24	-.43	-.19	.39	.28	.03	1.0	-.03	.50	-.42	-.25	-.30
FRA	.09	.04	.11	.08	.16	.28	.17	-.03	1.0	.31	-.28	.08	-.25
MI	.48	-.07	-.5	-.09	.41	.72	.31	.50	.31	1.0	-.08	-.15	-.84
AD	-.47	-.01	.10	.03	-.31	-.75	-.25	-.42	-.28	-.80	1.0	.06	.63
IMP	-.18	.74	.35	.70	-.29	.01	-.15	-.25	.08	-.15	.06	1.0	-.04
INV	-.38	-.06	.03	-.01	.25	-.73	-.25	-.30	-.25	-.84	.63	-.04	1.0