

Siegfried K. Berninghaus and Bodo Vogt

A Stochastic Model of the Co-evolution of Networks and Strategies

Abstract:

We consider a theoretical model of co-evolution of networks and strategies whose components are exclusively supported by experimental observations. We can show that a particular kind of sophisticated behavior (anticipatory better reply) will result in stable population states which are most frequently visited in co-evolution experiments.

1. Introduction

The steadily increasing interest in the problem of the *co-evolution of networks and strategies* is inspired by two different strands of research. In traditional evolutionary game theory the focus was on the evolution of strategy configurations in populations of players interacting bilaterally with the whole population via simple two-person games. Moreover, coordination games were often utilized as instances of such two-person base games. The question then was whether population members would eventually coordinate on one particular equilibrium of the one-shot game when randomly matched¹ with each other sufficiently many times. The uniform and stable choice of an equilibrium that evolves over time in the symmetric two-person coordination game is often interpreted as the evolution of a convention (see e.g. Berninghaus 2003).

Models on the evolution of conventions in large populations were soon extended to include populations endowed with a social network structure (see Berninghaus and Schwalbe 1996a; 1996b; Berninghaus and Vogt 2006). It is a basic assumption in such models that population members will exclusively be matched with members of their reference (or neighborhood) group and not with any other members of the population. However, reference groups are usually assumed to be overlapping.² The main emphasis in this strand of research was on determining which conditions would guarantee either that one convention would eventually be uniformly spread across the population or that several conventions would coexist indefinitely.

¹ See e.g. Young 1993; Amir and Berninghaus 1996.

² Consequently, an agent may belong to several reference groups. Non-overlapping reference groups would not add a really new aspect because they can simply be regarded as sub-populations without a network structure.

A major drawback of this type of model is that the network structure is exogenously imposed (by the model maker) rather than allowed to evolve from individually made decisions. Since the results of strategy choice in networks crucially depend on the particular network structure, the choice of a network structure (like a ‘circular’ or ‘rectangular’ network structure) is by no means trivial. But who can say which network shape is the ‘correct’ one to use in our models?

In another strand of literature, network structures are introduced as endogenous variables resulting from a strategy choice of an appropriately formulated non-cooperative (network) game (see e.g. Bala and Goyal 2000; Haller and Saranghi 2003; Jackson and Wolinsky 1996). However, models of strategic network formation often lack of a satisfactory motivation for the individual payoffs. In this literature, network games are commonly interpreted as information exchange games in which the players’ payoffs are generated by extracting valuable information from the members with whom an agent is connected.

Obviously, both approaches are flawed. Hybrid models addressing the co-evolution of both strategies and networks have proven to be superior (see e.g. Skyrms and Pemantle 2001; Berninghaus and Vogt 2006; Bramoulier et al. 2004). In most of these models the co-evolution problem was posed as a static Nash problem in an appropriately defined non-cooperative game framework in which players simultaneously choose network partners and the strategies they will employ against these partners in a two-person game.

In a recent experimental study (Berninghaus et al. 2008) it was found that the static Nash concept does not play a major role in the co-evolution problem; during the course of the experiment Nash configurations were almost never observed. But subjects engaged in another type of regular behavior which placed them into another type of stable strategy configuration that differs significantly from Nash-type configurations. In this paper, we formulate a simple, stylized dynamic strategy adaptation process whose elements are exclusively taken from observed behavior. From a formal point of view this adaptation process can be interpreted as an instance of a Stochastic Learning System (General Random System with Complete Connections = GRSCC).

A unique feature of our design is that the 2×2 base game is an anti-coordination game (of the *Hawk-Dove* type). A ‘convention’ generated by this game is a ‘regularity in behavior’ whereby the players take opposite strategy choices when they are matched with each other. The overwhelming part of the literature on the co-evolution process is devoted to *pure coordination games* as base games. We justify our choice of anti-coordination games as base games in the co-evolution process elsewhere (Berninghaus et al. 2008).

2. The Model

2.1 The Base Game

We consider a set $I = \{1, \dots, n\}$ of n agents who are engaged in playing the same 2×2 game with each of their direct neighbors in the network.³ If two players i and j are linked, they play the symmetric 2×2 normal form game $G = (\Sigma, \pi(\cdot))$ with strategy set $\Sigma := \{H, D\}$ and payoff function $\pi(\cdot): \Sigma \times \Sigma \rightarrow \mathbb{R}$ characterized by the following payoff table with $a > b > c > d > 0$.

	H	D
H	d, d	a, c
D	c, a	b, b

Table 1: Hawk-Dove game

The game in Table 1 is a generalized Hawk-Dove game, also known as a Chicken game. Action H is called a ‘hawk’ action while D is called a ‘dove’ action.⁴ Typically, in Hawk-Dove games dove players are better off when matched with other doves rather than with hawks, although playing dove against dove does not constitute an equilibrium of G .

As mentioned before, it is quite unusual to study the co-evolution problem with an anti-coordination game as the base game. We use this scenario in our theoretical model in order to stay as close as possible to the experimental setup.⁵

2.2 The Strategic Network Game

We do not impose a fixed interaction structure on the population of players but instead assume that networks can be built up from individual decisions. We assume that all players participate in a *network game*, which in this case is a simple normal form game. An individual strategy of player i in the network game is supposed to be a vector of ones and zeros, $v_i \in \{0, 1\}^n$. We say that player i establishes a link to player j if $v_{ij} = 1$; otherwise, it is equal to zero. A link between two players allows both players to play the Hawk-Dove game G . As a player cannot play the game alone, we assume $v_{ii} = 0$ for all $i \in I$. A connection between two players is assumed to exist bilaterally if at least *one* player expresses interest in opening one.⁶

Each strategy configuration $v = (v_1, \dots, v_n)$ in the network game generates a directed graph denoted by \mathcal{V}_v , where the vertices represent players and a directed edge between i and j , i.e., $v_{ij} = 1$, indicates that i opens a link with j . The

³ To avoid trivialities, we assume $n \geq 3$.

⁴ In the following, strategies in the 2×2 Hawk-Dove game will be called ‘actions’.

⁵ In these kinds of experiments, one cannot exclude the possibility that the subjects’ behavior would change completely if another base game was used instead. Consequently, our behavioral assumptions in section 3.1, which are taken from experiments with Hawk-Dove games, would soon become obsolete.

⁶ Based on Bala and Goyal 2000, this assumption simplifies modeling non-cooperative network formation games considerably.

set of i 's neighbors, given a network \mathcal{V}_v , is defined to be the set of all players with whom i opens a link (neighbors connected via an 'active link' of i , $v_{ij} = 1$) and the set of all players who open a link with i (neighbors to whom i is connected via a 'passive link' $v_{ji} = 1$). Therefore, we define the *set of all neighbors of i* by the set

$$N_i(\mathcal{V}_v) := \{j \mid \max\{v_{ij}, v_{ji}\} = 1\}.$$

The set of *neighbors reached via an active link* depends solely on i 's network strategy vector v_i . It is defined by

$$N_i^a(v_i) := \{j \mid v_{ij} = 1\}.$$

The cardinality of this set is denoted by the natural number $n_i^a(v_i) := |N_i^a(v_i)|$.

Opening a link with another player is not costless. For the sake of simplicity, we assume that connection costs k (> 0) are constant per connection and do not vary among the players. Player i 's total payoff in the network game consists of the aggregate payoff that she can extract from playing with her neighbors minus the costs of establishing her active links.

2.3 The Extended Network Game

Combining the two non-cooperative games described in the previous subsections 2.1 and 2.2, we model the strategic situation of a player as a *non-cooperative game* in which an individual strategy choice is composed of the *simultaneous choice of neighbors* via links $v_i \in \{0, 1\}^n$ in the network game and *actions* $\sigma_i \in \{H, D\}$ in the Hawk-Dove game. We denote this *extended network game* by

$$\tilde{G} = (S_1, \dots, S_n; P_1(\cdot), \dots, P_n(\cdot)),$$

with strategy sets $S_i := \{0, 1\}^n \times \{H, D\}$ and payoff functions $P_i : S_1 \times \dots \times S_n \rightarrow \mathbb{R}$ where

$$P_i(s_i, s_{-i}) := \begin{cases} 0, & \text{if } N_i(\mathcal{V}_v) = \emptyset \\ \pi(\sigma_i, H) \sum_{j \in N_i(\mathcal{V}_v)} \mathbf{1}_{\{\sigma_j = H\}} + \pi(\sigma_i, D) \sum_{j \in N_i(\mathcal{V}_v)} \mathbf{1}_{\{\sigma_j = D\}} \\ -k n_i^a(v_i), & \text{else.} \end{cases}$$

The vectors s_{-i} and σ_{-i} denote the strategies $s_j \in \{0, 1\}^n \times \{H, D\}$ and the actions $\sigma_j \in \{H, D\}$ chosen by each player $j \neq i$. \tilde{G} is a normal form game. A strategy configuration $s = (s_1, \dots, s_n) = ((\sigma_1, v_1), \dots, (\sigma_n, v_n))$ induces a social network represented by a directed graph \mathcal{V}_v and a vector of actions in the Hawk-Dove game's action set $\{H, D\}$. An important consequence of our particular payoff definition is that player i may benefit from a connection to j even though she does not have to pay for it (i.e., if $v_{ij} = 0$ but $v_{ji} = 1$ holds).

A natural benchmark solution of the extended network game \tilde{G} is the *Nash equilibrium*, which in our context is defined as follows:

Definition 1

The strategy configuration $s^* = ((\sigma_1^*, v_1^*), \dots, (\sigma_n^*, v_n^*))$ is a Nash equilibrium of \tilde{G} if

$$\text{for all } i \in I: P_i(s_i^*, s_{-i}^*) \geq P_i(s_i, s_{-i}^*) \text{ for all } s_i \in S_i.$$

In an equilibrium, no player has an incentive to unilaterally change either her link choice or her action choice $\sigma_i^* \in \{H, D\}$ or both. The Nash equilibria s^* for the static game \tilde{G} were characterized completely by Bramoulier et al. (2004) and by Berninghaus and Vogt (2006). A Nash equilibrium of the extended network game \tilde{G} is characterized both by a network structure and a distribution of H and D players denoted by n_H and n_D ($n_H + n_D = n$).

The existence and the characterization of Nash configurations s^* crucially depends on the base game's payoffs a, b, c, d and the connection costs k (for a detailed analysis see Berninghaus and Vogt 2006). We will not go into all details of characterizing the static Nash solution. In the sections below we will provide more information on this solution as needed.

Our theoretical analysis is based on the experimental results about the co-evolution of networks and strategies conducted by Berninghaus et al. (2008). We derive our theoretical assumptions directly from experimental observations on strategy choice. The experimental results show that Nash configurations do not play a significant role in the co-evolution process. Subjects instead seem to anticipate the remaining players' responses to their own strategy changes and take these responses into account when formulating their own strategy choices. This behavior can render some Nash constellations unstable or stabilize some non-Nash constellations. This behavior pushes the population close to strategy configurations which are called *one step ahead stable* states (see Berninghaus et al. 2008). We will expand upon this concept in the following sections.

3. Dynamics

Let us consider the extended network game \tilde{G} . Studying the co-evolution of networks and strategies is equivalent to studying the temporal evolution of $\{s_t\}_t$. From $\{s_t\}_t$ we will proceed to derive additional stochastic processes in the hopes of providing more insight into the co-evolution process.

3.1 Action and Link Adaption

Our dynamic model is mainly based on two important observations⁷ (in a stylized form) from our experimental results. We formulate these observations as an assumption:

⁷ Our experiments were conducted in *continuous time*, i.e., subjects could switch actions or network strategies at any moment during the experiment (30 minutes). Assumption 1 in some sense mimics this particular design.

Assumption 1

- (1) We subdivide the co-evolution process $\{s_t\}_t$ into two types of rounds:
- In rounds of type A one player is randomly selected to switch actions.
 - In rounds of type B one player is randomly selected to build or sever one link.
 - After a type A round there follow only type B rounds until a Nash network is reached in the network game (given the action configuration prevailing in the last type A round).
- (2) When adapting links (in type B rounds), agents employ a myopic better reply rule, i.e., they randomly choose a link which may result in a payoff increase; otherwise, they keep the status quo link.
- (3) When adapting actions (in type A rounds), agents employ a better reply rule by anticipating the network induced by their action change, i.e., they randomly switch their action, which may result in an anticipated payoff increase; otherwise, they keep the status quo action.

Typically, in our experiment, we observed many more link changes than action changes.⁸ After switching from hawk to dove or vice versa, the remaining players usually adapted the links that had become detrimental after the last action change.

In the following we make a simplifying technical assumption:⁹

Assumption 2

For the payoffs of the base game and the connection costs, the following inequality holds:

$$a > b > k > c > d.$$

Two consequences immediately follow from Assumption 2:

Result 1

- a) Given a particular action configuration $\sigma = (\sigma_1, \dots, \sigma_n)$, only the following networks are compatible with Nash networks: One-way connections¹⁰ from each hawk player to all dove players and one-way connections among all dove players.¹¹
- b) Given an arbitrary action configuration $\sigma \in \{H, D\}^n$ and a link configuration $v = (v_1, \dots, v_n)$, the result is a Nash network associated to σ . If one player switches actions and that results in a new action configuration $\sigma' \neq$

⁸ A tentative explanation of this feature goes as follows: In action changes all players connected with the deviating player are concerned, while link changes concern only one player. A player with a limited computation capacity may prefer link changes since she then has to deal with one reaction only.

⁹ In the experiment we assumed three different levels of connection costs k to analyze the impact of connection costs on linking behavior. To keep this exposition short, we focus on one particular constellation, i.e. that between payoffs and the level of connection costs.

¹⁰ Since connections are not costless and both partners benefit from a connection, two-way connections cannot be part of a Nash configuration.

¹¹ Note that the connections among dove players are not determined uniquely. In an equilibrium network, it is even possible that only one dove player will end up keeping all active connections with the remaining doves.

σ , then a possibly newly established Nash network v' reached from v is unique.

Part a) follows easily from Assumption 2. Hawks earn $(a - k) > 0$ from an active connection with a dove, while doves extract negative payoffs $(c - k) < 0$ from active connections with hawks. Because $(b - c) > 0$, active connections between doves are profitable. Therefore, to each action configuration we can associate several Nash networks differing only with respect to the distributions of active dove connections.

Part b) shows that the multiplicity of Nash networks disappears if the population starts from a given Nash network in reaction to a changed action configuration. For if a player changes her actions, then some links may become unprofitable or it may become profitable to open new links. Irrespective of the different sequences of players involved in the link adaptation, the new network v' is unique.

3.2 A Stochastic Strategy Adaptation Model

Having made some preparatory remarks about our stylized dynamic strategy adaptation process, we will now discuss the formal structure of the stochastic co-evolution model in more detail. We formulate a general stochastic process containing the essentials of the strategy adaptation process as described in Assumption 1. We interpret this process as an instance of a more general *Stochastic Learning Model* called a *Generalized Random System with Complete Connections* (GRSCC).

Definition 2

A GRSCC is characterized by a tuple $((W, \mathcal{W}), (X, \mathcal{X}), \Gamma, P)$ where

(W, \mathcal{W}) and (X, \mathcal{X}) are measurable spaces,

$\Gamma : W \times X \rightsquigarrow W$ is a transition probability,

$P : W \rightsquigarrow X$ is a transition probability.

A GRSCC describes a stochastic learning system with state space W and parameter space X . A learning system which starts in state w generates a conditional probability distribution $P(\cdot|w)$ on (X, \mathcal{X}) . State variable and realized parameter constellations x determine the probability of the next period state via the conditional probability $\Gamma(\cdot|w, x)$. The drawing in Figure 1 may illustrate how the elements of a GRSCC interact.




Figure 1: Sketch of a GRSCC

In learning applications of a GRSCC, the set of states W describes the internal states of a learning system which that will be adapted from one period to the next depending on the occurrence of some external events interpreted to be elements of X . In our co-evolution model we simply define $W := \Sigma$ and $X := V$. The GRSCC in this paper thus describes a stochastic process of the evolution of action configurations $\{\sigma_t\}_t$ and networks $\{v_t\}_t$ in a highly stylized way.

If we start from a (historically) given action configuration σ , then the stochastic network adaptation process in type B periods (see Assumption 1) will almost surely end up in an equilibrium network v with probability $P(v|\sigma)$. Given v in the following type A round, a randomly selected player will eventually choose a different action such that a new action configuration σ' will occur with probability $\Gamma(\sigma'|\sigma, v)$. It can easily be shown that the resulting state process of the GRSCC $\{\sigma_t\}$ is a Markov process on Σ with transition probabilities

$$Q(A | \sigma),$$

for each measurable $A \subset \Sigma$ (see Herkenrath 1976 for details). By composing a type A round action configuration σ_τ with a resulting equilibrium network, denoted by v_τ , the stochastic process $\{\sigma_\tau, v_\tau\}_\tau$ is a Markov chain.¹²

The network adaptation process: Given action and network configuration (σ, v) in type A period t , in each following type B period $t+k$ a randomly selected agent randomly chooses one link.¹³ If this link improves her payoff, the status quo link will be substituted; otherwise, the status quo choice will be kept for another period (**better reply**). Obviously, this behavior generates a Markov process $\{v_{t+k}\}_k$ of networks. The long-run behavior of this process is described by a lemma in Friedman and Mezetti (2000). To apply this result we only have to show that the separate network game (given a particular action configuration) fulfills the so-called *weak finite improvement property* (weak FIP).

¹² Note our different usage of the time index: We use τ as an index of action configurations associated with the resulting equilibrium networks, while t denotes the real time time period. Therefore, $s_t = (\sigma_t, v_t)$ denotes the action and network configurations actually prevailing in each period t .

¹³ All alternatives are supposed to be selected with equal probability.

Definition 3

The network game fulfills the weak FIP if from each network profile $v \in V$ there exists a finite sequence of single player payoff improvements that end in a pure Nash equilibrium, where a configuration (σ, v_{-i}, v'_i) is a single player payoff improvement (over (σ, v_{-i}, v_i)) iff $P_i(\sigma, v_{-i}, v'_i) > P_i(\sigma, v)$.

The following lemma shows that our separate network adaptation model satisfies the weak FIP.

Lemma 1

Given a particular fixed action configuration σ in which agents can only switch links, the result is a separate network game. This separate network game satisfies the weak FIP.

According to Assumption 1, in a Nash network v^* all hawks have active connections to all doves and each dove is connected with the remaining doves by means of either an active or a passive connection.¹⁴ A network is not Nash if at least one link in v^* is severed. It is easy to see that starting from an arbitrary link configuration $v \neq v^*$ for a Nash network v^* , one can build a finite sequence of networks v ending up in v^* , which demonstrates the weak FIP.

It follows from a lemma by Friedman and Mezetti (Lemma 1, 2000) that the Markov process will almost surely end in a Nash network configuration v^* as an absorbing state and that the set of absorbing states in the Markov process will coincide with the set of Nash configurations in the separate network game. This defines our conditional probability $P(v|\sigma)$ of the GRSCC. From part b) of Result 1 we derive further properties of the probability measure $P(\cdot | \cdot)$:

Result 2

- a) Given an action configuration σ_1 , the support of $P(\cdot | \sigma_1)$ is the set of Nash networks.
- b) For $\tau \geq 2$ $P(\cdot | \sigma_\tau)$ is a degenerate probability measure concentrated at the unique Nash network reachable from the previously obtained Nash network.

Part a) immediately follows from Lemma 1 in Friedman and Mezzetti. Part b) follows from Result 1 b) and part a) since in period $\tau = 1$ an equilibrium Nash network is reached. From that time on the result of the network formation process is uniquely determined by the prevailing action configurations. Although the particular equilibrium network v_τ (for $\tau \geq 2$) reached is determined uniquely, the number of type B periods until the equilibrium is reached is not determined.

The action adaptation process: We consider an equilibrium network $v_{\tau-1}^*$ and an action configuration $\sigma_{\tau-1}$ which (without loss of generality) is assumed to have generated $v_{\tau-1}^*$ uniquely. Then a new type A period starts, i.e. a randomly selected player checks whether selecting a randomly generated action will make it profitable for her to switch actions. Let us denote by σ'_τ the action

¹⁴The direction of the links is ambiguous for dove players. Therefore, there exist multiple Nash networks.

configuration that results if the randomly selected player deviates. How will this player decide what to do? According to Assumption 1, each player is able to anticipate the support of the distribution $P(\cdot|\sigma'_\tau)$, i.e. the Nash network induced by σ'_τ . The selected player i will only change her action if her (anticipated) payoff is improved, i.e., if

$$P_i(\sigma'_\tau, f(\sigma'_\tau, v_{\tau-1})) > P_i(\sigma_{\tau-1}, v_{\tau-1}^*), \quad (1)$$

where f denotes a mapping which associates to each Nash network $v_{\tau-1}$ and action configuration σ'_τ the resulting Nash network $f(\sigma'_\tau, v_{\tau-1}) \in V$.

The probability $\Gamma(\sigma'_\tau|\sigma_{\tau-1}, v_{\tau-1})$ is equal to the probability that a randomly selected player's randomly selected action will exhibit an **anticipated better reply** as defined in inequality 1. From the theory of GRSCC we know that the process $\{\sigma_\tau\}_\tau$ is a Markov chain. Now we can argue that the weak FIP of the action adaptation process is satisfied in the same way as before. By considering type A periods only it is easy to see that from each action configuration σ a strategy configuration s^* can be reached in a finite number of steps by a single-player **anticipated payoff improvement**. We call such states *anticipatory Nash states*.

Definition 4

A strategy configuration s^* is called an *anticipatory Nash state* if for each player i the following inequality holds:

$$P_i(\sigma^*, v^*) \geq P_i(\sigma', f(\sigma', v^*))$$

for each $\sigma' = (\sigma_1^*, \dots, \sigma_{i-1}^*, \sigma'_i, \sigma_{i+1}^*, \dots, \sigma_n^*)$.

Essentially, anticipatory Nash states are equivalent to so-called **one step ahead stable states**, which were introduced in Berninghaus et al. (2008). The experimental results show that the population of players mostly remains close to one step ahead stable states. Although anticipatory Nash states seem to have a formal structure similar to pure Nash states, there exist non-negligible differences. There are states which are Nash but not anticipatory Nash and vice versa (for more details see Berninghaus et al. 2008).

Result 3

Given Assumptions 1 and 2, the co-evolution process will almost surely end up in an anticipatory Nash state.

Our arguments show that σ_τ satisfy the weak FIP (extended by the *anticipated payoff improvement* concept). Applying Lemma 1 from Friedman and Mezzetti again, we conclude that the only absorbing states of this Markov process are action configurations that are immune to anticipated payoff improvements.

4. Concluding Remarks

The co-evolution of networks and strategies involves individual strategic decision problems that are not easy to model. At the same time, it is certainly not

reasonable to model the co-evolution process without referring to any empirical results. We therefore modeled our co-evolution process on our experimental findings.

However, our results have to be interpreted quite carefully. We believe that the particular experimental results which guided our theoretical model were mainly driven by our particular base game (Hawk-Dove) along with the ‘continuous time’ design. It is therefore not clear if our theoretical model would be equally successful in co-evolution experiments utilizing different base games or other design schemas. To test how robust our theoretical model really is, we certainly need to conduct additional experiments with varying designs.

References

- Amir, M. and S. K. Berninghaus (1996), “Another Approach to Learning and Mutation in Games”, *Games and Economic Behavior* 14, 19–43.
- Bala, V. and S. Goyal (2000), “A Non-cooperative Model of Network Formation”, *Econometrica* 68, 1181–1231.
- Berninghaus, S. K. (2003), “The Evolution of Conventions—Some Theoretical and Experimental Aspects”, in: M. J. Holler et al. (eds.), *Jahrbuch für Neue politische Ökonomie* 21, 101–132, Tübingen: Mohr Siebeck.
- (2006), “Evolution of Conventions in Populations with Local Interaction Structures”, in: M. Morlock, C. Schwindt, N. Trautmann and J. Zimmermann (eds.), *Perspectives on Operations Research: Essays in Honor of Klaus Neumann*, Wiesbaden: Dt. Univ.-Verl., 353–378.
- and U. Schwalbe (1996a), “Conventions, Local Interaction, and Automata Networks”, *Journal of Evolutionary Economics* 6, 297–312.
- and — (1996b), “Evolution, Interaction and Nash-Equilibria”, *Journal of Economic Behavior and Organization* 29, 57–85.
- and B. Vogt (2006), “Network Formation in Symmetric 2x2 Games”, *Homo Oeconomicus* 23(3/4), 421–466.
- , K. M. Ehrhart and M. Ott (2008), “Myopically Forward-Looking Agents in a Network Formation Game: Theory and Experimental Evidence”, SFB 504, Disc. paper No. 08/02, University of Mannheim.
- Bramoulier Y., D. Lopez-Pintado, S. Goyal and F. Vega-Redondo (2004), “Network Formation and Anti-coordination Games”, *International Journal of Game Theory* 33, 1–19.
- Friedman, J. W. and C. Mezetti (2001), “Learning in Games by Random Sampling”, *Journal of Economic Theory* 54, 261–292.
- Haller, H. and S. Sanraghi (2003), “Nash Networks with Heterogeneous Agents”, DIW Discussion paper, No. 337.
- Herkenrath, U. (1976), *Der assoziierte Markov-Prozess verallgemeinerter zufälliger Systeme mit vollständigen Bindungen*, Dissertation, Universität Bonn, Mathematisch-Naturwissenschaftl. Fakultät.

- Jackson, M. O. and A. Wolinsky (1996), "A Strategic Model of Social and Economic Networks", *Journal of Economic Theory* 71, 44–74.
- Skyrms, B. and R. Pemantle (2000), "A Dynamic Model of Social Network Formation", *Proceedings of the National Academy of Sciences of the USA* 97, 9340–9346.
- Young, H. P. (1993), "The Evolution of Conventions", *Econometrica* 61(1), 57–84.