University of Groningen

Co-evolution of conventions and networks

Corten, R.; Buskens, V.W.

Published in:
Social Networks

DOI:
[10.1016/j.socnet.2009.04.002](https://doi.org/10.1016/j.socnet.2009.04.002)

IMPORTANT NOTE: You are advised to consult the publisher's version (publisher's PDF) if you wish to cite from it. Please check the document version below.

Document Version
Publisher's PDF, also known as Version of record

Publication date:
2010

[Link to publication in University of Groningen/UMCG research database](#)

Citation for published version (APA):

Corten, R., & Buskens, V. W. (2010). Co-evolution of conventions and networks: an experimental study. *Social Networks*, 32(1), 4-15. <https://doi.org/10.1016/j.socnet.2009.04.002>

Copyright

Other than for strictly personal use, it is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license (like Creative Commons).

The publication may also be distributed here under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license. More information can be found on the University of Groningen website: <https://www.rug.nl/library/open-access/self-archiving-pure/taverne-amendment>.

Take-down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Downloaded from the University of Groningen/UMCG research database (Pure): <http://www.rug.nl/research/portal>. For technical reasons the number of authors shown on this cover page is limited to 10 maximum.

Co-evolution of conventions and networks: An experimental study[☆]

Rense Corten^{*}, Vincent Buskens

Department of Sociology/ICS, Utrecht University, Heidelberglaan 2, 3584 CS Utrecht, The Netherlands

ARTICLE INFO

Keywords:

Network dynamics
Co-evolution
Experiment
Coordination
Simulation

ABSTRACT

We study the emergence of conventions in dynamic networks experimentally. Conventions are modeled in terms of coordination games in which actors can choose both their behavior and their interaction partners. We study how macro-level outcomes of the process in terms of Pareto-efficiency and heterogeneity depend on initial conditions. Moreover, we examine the underlying processes at the microlevel. Predictions are derived from a game-theoretic model which is applied to our experimental conditions by means of computer simulation. The results provide mixed support for the macro-level hypotheses, and indicate possible directions to improve the model at the micro level.

© 2009 Elsevier B.V. All rights reserved.

1. Introduction

In many social and economic interactions, people have an interest in aligning their behaviors with one another. We speak the same language to communicate, we agree to the same traffic rules in order to drive safely, and when writing an article together, it helps if computer programs are compatible. In economic interactions, if trade takes place in the marketplace, traders must at least manage to meet at the same time and place. In a more abstract sense, coordination problems are central to the problem of social order and the emergence of institutions (see Hume [1739–40], 1978; Hardin, 2007). For instance, Hobbes' *Leviathan* presupposes that citizens coordinate on a leader to solve the problem of social order.

This paper is concerned with the situations in which coordination is problematic. It studies the role of social networks in how actors handle coordination problems if the social network can be changed by the actors and co-evolves with the actors' behavior in coordination problems. Moreover, we study the effects of information availability in these networks on coordination. We use a laboratory experiment in which subjects play coordination games while choosing interaction partners. In this experiment, we test hypotheses derived from a game-theoretic model reflecting the

experiment. We apply analytical methods and computer simulation to derive our hypotheses.

1.1. Coordination, conventions, and networks

Often, coordination problems are resolved by conventions, i.e., behavioral patterns that are mutually expected and self-reinforcing (Lewis, 1969). Everyday conventions include traffic rules (driving left or right), technological standards (GSM frequencies), spelling standards, and rules for appropriate behavior, such as dress codes and table manners (Elias, 1969; Ullmann-Margalit, 1977; Coleman, 1990). These conventions share the feature that, once established, none of the actors involved has an incentive to deviate from the convention, provided that others do not deviate.

In some coordination problems, there is no reason to prefer one convention over another. The driving problem is a prominent example, but the situation probably also holds for many etiquette rules (e.g., does anyone really prefer "black tie" over "blue tie"?). In other cases, possible conventions are ranked according to their utility for all actors involved. For instance, we may choose between everyone being self-sufficient or everyone specializing in one type of labor, where the latter is more efficient (provided that others do the same). A further classification might be made according to the consequences of coordination failure. While some actions have higher or lower payoffs if they are also chosen by others, actions may have different consequences when they are not chosen by others. When two people fail to coordinate in trying to lift a heavy object, the one who does not lift is better off than the one who unsuccessfully lifts. Still, both would have preferred to lift the object together. Similar risks apply to many collective action problems (cf. Hardin, 1995). Situations where it is problematic to reach socially and individually optimal conventions are central to this paper. The coordination game in Fig. 1 represents such a situation for two actors.

[☆] We gratefully acknowledge useful comments and suggestions by Werner Raub, Stephanie Rosenkranz, Károly Takács, Karen Cook, and Jeroen Weesie. This study is part of the Polarization and Conflict Project CIT-2-CT-2004-506084 funded by the European Commission-DG Research Sixth Framework Programme. This paper reflects only the authors' views and the Community is not liable for any use that may be made of the information contained therein. Additional funding was provided by Utrecht University through the High Potentials 2004 subsidy for the research program "Dynamics of Cooperation, Networks, and Institutions."

^{*} Corresponding author.

E-mail addresses: r.corten@uu.nl (R. Corten), v.buskens@uu.nl (V. Buskens).

		Player 2				Player 2	
		LEFT	RIGHT			LEFT	RIGHT
Player 1	LEFT	d, d	b, c		LEFT	20, 20	0, 10
	RIGHT	c, b	a, a		RIGHT	10, 0	14, 14

$d > a > c > b$

Fig. 1. A coordination game and a numerical example.

This game has two pure Nash equilibria: (LEFT, LEFT) and (RIGHT, RIGHT). The payoffs are higher for both players if both play LEFT; therefore (LEFT, LEFT) is the *payoff-dominant* or *efficient* equilibrium. Choosing LEFT is also risky: if the other player plays LEFT or RIGHT with equal probability, the expected payoff of playing LEFT is lower than the expected payoff of playing RIGHT. Therefore, (RIGHT, RIGHT) is the *risk-dominant* equilibrium (Harsanyi and Selten, 1988). We refer to the actions associated with these two equilibria as payoff-dominant and risk-dominant actions. The equilibria in the coordination game can be interpreted as conventions. Throughout this paper, we mostly use the term “convention” to refer to these equilibria.

Early studies on multi-person coordination focus on global interaction (i.e., every actor interacts with every other actor). Such theoretical models suggest that risk-dominant conventions are more likely to occur even if they are inefficient (Kandori et al., 1993; Young, 1993). This assertion is mirrored in experimental studies that show how subjects’ behaviors often converge to risk-dominant conventions (Cooper, 1990; Van Huyck et al., 1990). More recent models recognize that in larger populations, actors adjust their behavior not to everyone but rather to their local environment. One reason for this is that actors can only observe behavior within a limited portion of the population. A more important reason is that the very nature of the interaction does often not imply global interaction. For example, one speaks the language spoken by those one talks to, which is typically not the entire population. Still, Young (1998) predicts that eventually everyone will play the risk-dominant behavior even in coordination games played in a network structure. These results are derived from a stochastic model in which actors make random ‘mistakes.’ More deterministic models find that the network matters for the likelihood of reaching payoff-dominant behavior (Berninghaus and Schwalbe, 1996; Berninghaus and Ehrhardt, 1998; Buskens and Snijders, 2008).

Although explicit theoretical models of coordination in networks are scarce, the topic is not: a large strand of sociological literature studies processes of social influence in networks (see Marsden and Friedkin, 1993, for an overview). A closely related body of literature exists on *threshold models* of diffusion (Granovetter, 1978), including a number of studies looking at network effects on diffusion (Abrahamson and Rosenkopf, 1997; Watts, 1999; Ehrhardt et al., 2006; Centola and Macy, 2007). The network coordination game applied in this paper may be interpreted as a game-theoretic representation of social influence or diffusion.

All of these studies consider networks to be exogenous. It is increasingly recognized, however, that the commonly observed relation between networks and behavioral similarity can be attributed not only to influence, but also to selection. Actors prefer to interact with others who have similar characteristics or behave similarly (McPherson et al., 2001). This implies that social networks also change in the feedback process between influence and selection (Snijders, 2001; Knecht, 2008). Behavioral dynamics can be expected to differ when networks are dynamic. For instance, differences in behavior are more likely to persist if groups that use different conventions self-segregate. This paper contributes to a better understanding of such co-evolution processes. Thus, the

main question is: *how do conventions in coordination problems and networks co-evolve?*

A number of studies address theoretical perspectives on this question. Jackson and Watts (2002) propose a game-theoretic model in which actors play coordination games in an endogenous network and derive conditions under which constellations of networks and behavior are stochastically stable depending on the cost of maintaining ties. Their main finding is that, whereas various network structures are possible, behavioral choices in the coordination game become homogeneous. Berninghaus and Vogt (2006) analyze a similar (though deterministic) model, and find that networks can emerge consisting of multiple unconnected groups, while different conventions are maintained within various groups. (See also Skyrms and Pemantle, 2000; Goyal and Vega-Redondo, 2005, for related models.) These studies provide general characterizations of networks that might emerge, but many different constellations are still usually at least theoretically possible. To examine which stable structures are more likely depending on various starting conditions, Buskens et al. (2008) apply computer simulations. They find that the *density* of the initial network has a strong influence on the way behavior develops: the higher the density, the stronger the influence of the initial behavioral distribution on the behavior’s emergent distribution. Moreover, they find that if the initial network is more *segmented*, then there is higher likelihood that two groups with different behavior will emerge. More generally, it is found that in a majority of the cases, a single convention is reached.

Our first aim is to empirically study which outcomes of the co-evolution process are more or less likely given initial conditions, such as initial network structures. For this purpose, our experimental setup includes three different initial eight-person networks: the full network, the circle network, and the “two-squares” network (see Fig. 2). By choosing these three network structures, we vary both network density and segmentation.

All models discussed above assume that actors are fully informed of all other actors’ past behavior. This assumption seems highly unrealistic for many real-life contexts. In large populations, it is typically impossible to keep track of all others’ behavioral choices. But even in smaller populations this might be difficult if behavior can only be discovered through interacting with or obtaining personal information about another actor. This could be the case for both opinions and other types of behavior (e.g., choice of technology, language). Our second aim is to examine what effect limited information has on the co-evolution of networks and conventions. We compare the situation in which actors observe past behavior of all other actors in the network to the situation in which they only observe their neighbors.

1.2. An experimental approach

Despite the growing theoretical interest in the co-evolution of networks and behavior, empirical evidence testing these theories is scarce. This is understandable given that the demands to data suitable for testing these theories are very high. More specifically, one needs detailed longitudinal information on social relations and individual behavior. To test predictions at the network or macro level, one needs sufficient variance and many observations at the macro level. While collecting field data that meet these requirements is not impossible (e.g. Knecht, 2008), practical difficulties mean that one usually must compromise on the number of observations at the network level, the number of observed time points, or the “depth” of observation at the individual level.

As an alternative, we suggest laboratory experiments. Experiments have a number of well-known advantages that make them the preferred research method for behavioral approaches: the experimenter is in considerable control of incentive structures, information availability, and other ingredients of game-theoretic

Fig. 2. Initial networks used in the experiment.

models that are hard to measure in real-life situations (Crawford, 1997; Camerer, 2003). Moreover, behavior can be unambiguously observed in the laboratory. Accurate information on relations and behaviors can be recorded at every time-point, and one can relatively easily observe multiple networks and then examine the effects of various conditions at the network level.

Experiments therefore allow for an explicit micro–macro perspective. We vary conditions at the macro level (such as initial network structure and information availability), and study the effects of these conditions on macro-level properties (i.e., emergent conventions and the network structure). Because we also observe all individual behavior, we can place the process “under the microscope” and study *how* exactly macro-level conditions lead – via individual behavior – to macro-level outcomes (cf. Coleman, 1990). When the theoretical model fails to correctly predict outcomes at the macro level, it is possible to examine which aspects of the micro-foundations are responsible for deviations. Moreover, understanding the individual level processes that drive macro-level processes’ dynamics might also help to predict which of many stable states are more or less likely to occur. Studying these micro-level processes is our third aim.

A detailed examination at both the micro and macro levels is useful, given current network evolution models. These models make specific assumptions about, for instance, individual behavior and information use. In real-life settings, such model assumptions are hard to measure, which makes it difficult to assess which aspects of a model are most empirically problematic. Experiments are useful for developing and fine-tuning theoretical models before they are tested more broadly using real-life data. We do not advocate experiments, however, as the *only* way to study network evolution. Studying network evolution through experimental methods also has disadvantages. As always with laboratory experiments, the external validity of findings obtained under abstract laboratory conditions is lower than for real-life data. Another problem is that practical considerations usually prohibit using groups that approximate group sizes considered typical in real-life human interaction. Therefore, we consider experiments as a merely useful intermediate step between developing network evolution models and “messy” field research on real-world phenomena.

We aim to take maximal advantage of experimental methods’ benefits by explicitly making the experimental design and theoretical model as similar as possible. We use computer simulations to generate predictions tailor-made to our experimental conditions. In this way, we hope to minimize the misfit between the model and experimental conditions, and thereby obtain strong tests of our hypotheses. Corbae and Duffy (2008) conducted one other experiment on coordination games in dynamic 4-person networks. They find that, in the presence of shocks, only networks consisting of pairs are stable. By comparison, we use 8-person networks and focus on the efficiency of emerging behavior, the influence of initial network structures, and information availability.

Our study reflects the micro–macro approach sketched above. First, we specify a formal model of the co-evolution of coordination and networks taking into account the arguments for limited information availability. We analytically characterize this model’s stable (macro-)states. Second, we conduct computer simulations to generate more precise macro-level predictions of the experimental conditions. We formulate hypotheses at both the micro and macro levels. Third, we report the results of an experiment that tests our hypotheses.

2. Model and simulation

2.1. The model

First, we define the underlying game: a coordination game played in a network. Actors interact if there exists a tie between them. The actors with whom an actor interacts are called *neighbors*. Actors play a repeated multi-person coordination game as shown in Fig. 1 with their neighbors, choosing only one of two actions against all neighbors. In each period, they receive payoffs from all interactions with their neighbors.

We assume that actors update their actions according to myopic best-reply behavior (cf. Kandori et al., 1993); that is, they optimize their payoff in the current period, assuming that the other actors act as they did in the previous period. When actors are indifferent between two actions, they do not change their behavior. It is easy to verify that actors play the payoff-dominant action (LEFT) if and only if the proportion of neighbors playing LEFT is at least $(a - b)/(a - b - c + d)$. We refer to this quantity as the *risk threshold*.

Maintaining ties is costly. In each period, actors pay for each tie. In real life, people face constraints on time and effort in the maintenance of social relationships. In related models, this is often translated into the assumption that there is a fixed upper limit to the number of ties one can maintain. We generalize this assumption using a convex cost function, such that the marginal tie cost increases with the number of ties. Specifically, the cost of t ties in period p to actor i is given by $k(t)_{ip} = \alpha k_{ip} + \beta k_{ip}^2$, with $\alpha = 6$ and $\beta = 1$. An alternative interpretation is that interactions have decreasing marginal returns: the net benefits of interactions decrease with every additional relation.

We introduce network dynamics through the following assumptions. In every period, all actors can propose *one* new tie to another actor, or they can remove one existing tie. Ties are created by mutual consent: a new relation is formed only if both parties agree to it. Existing ties can be dissolved unilaterally. This assumption is a consequence of ties representing *interactions*, which by nature require the consent of both parties involved. Actors are assumed to choose a change in ties (if any) that yields the highest expected payoff, given the actions of other actors in the previous period. If several changes yield the same payoff, a random choice is made between them. Both in the simulation and the experiment, the dynamic process

described above runs in three phases per period:

- (1) Each actor initiates at most one change in ties. Either a new tie is proposed, an existing tie is severed, or nothing is changed. Bilateral proposals immediately result in ties, and removals are immediately implemented.
- (2) Actors choose to accept or reject incoming (unilateral) proposals from phase 1.
- (3) Actors choose an action in the multi-person coordination game in the network that results from phases 1 and 2. They receive their payoffs for this period.

While earlier models assumed that actors observed the behavior of all actors in the network, we introduce limited information. We assume that the extent to which actors observe others' behavior in the previous period depends on the network. We study two information regimes:¹

- (1) *Local information*: actors only observe the behavior of their neighbors;
- (2) *Global information*: actors observe the behavior of all actors.

The information on behavior by actors other than neighbors is relevant only when actors make decisions about creating new ties. When actors update their behavior, they react only to their neighbors, so information about other actors does not play a role. Severing ties obviously only occurs between neighbors, so the distinction between the two information regimes is also irrelevant for this situation. Because network structure beyond direct adjacency does not enter into actors' considerations, other informational assumptions, e.g., what actors know about the network structure, does not affect our predictions.

We need one additional assumption to model how actors decide when choosing new neighbors if they only have information on their neighbors' behavior. We assume that actors use their neighbors' behavior to predict the behavior of others who they do not observe. If a proportion p of those that they observe (neighbors plus themselves) play action Y , they will assume that anyone else plays Y with probability p . Given this probability p , an actor proposes a tie to someone they cannot observe only if the expected benefits are larger than the tie costs.

2.2. Analytic results

We define a constellation of network ties and behavior as stable if two conditions are satisfied. First, no actor has an incentive to change his or her behavior, given the behavior of neighbors. Second, the network is *pairwise stable*, as defined by Jackson and Wolinsky (1996): no actor has an incentive to sever a tie, and no tie can be added without the consent of both parties.

Stable states can be formally characterized using the following definitions and theorems:

Definition 1. A (sub)network is t -full if and only if none of the actors have more than t ties and either (a) the addition of a tie causes at least one actor to have more than t ties, or (b) no ties can be added to the (sub)network.

Definition 2. $\bar{t}_Z$ is the maximum number of ties an actor wants to have if she chooses action Z , where $Z \in \{L(EFT), R(IGHT)\}$.

Theorem 1. Consider the co-evolution process as specified above. Assume that tie costs are equal to $k(t) = \alpha t + \beta t^2$, where $\alpha > b$ and $\beta > 0$. Under global information, networks are stable if and only if they satisfy one of the following conditions:

- (1) All actors choose the same action Z , where $Z \in \{L, R\}$, and the network is $\bar{t}_Z$ -full.
- (2) The network consists of two subnetworks of actors playing LEFT and RIGHT, and these subnetworks are $\bar{t}_L$ -full and $\bar{t}_R$ -full, respectively, and there are no ties between the two subnetworks.

Theorem 2. Consider the co-evolution process as specified above. Assume that tie costs are equal to $k(t) = \alpha t + \beta t^2$, where $\alpha > b$ and $\beta > 0$. Under local information, networks are stable if and only if they satisfy one of the following conditions:

- (1) All actors choose the same action Z , where $Z \in \{L, R\}$, and the network is $\bar{t}_Z$ -full.
- (2) The network consists of two subnetworks of actors playing LEFT and RIGHT, and these subnetworks are $\bar{t}_L$ -full and $\bar{t}_R$ -full, respectively. There are no ties between the two subnetworks, and for at most one action $Z \in L, R$, there exists an actor who chooses Z and has less than t_Z ties.

Theorem 1 states that, in any stable state, the network consists of one or more such subnetworks, within which all actors play the same action. The number of ties within each component is the maximum that the actors can afford given the payoffs to their actions and the tie costs. Within these boundaries, however, many different constellations are possible, such as only one component in which a single convention is played, several components playing the same convention, or several components playing different conventions (cf. Jackson and Watts, 2002. For the proof of Theorem 1, we refer to Buskens et al. (2008).

Under local information, our decision rule implies that actors in a homogeneous component guess that actors in other components also behave the same as themselves. The reason for this is that actors only observe their own behavior in the component. Therefore, we need the extra condition in Theorem 2. The argumentation for this extra condition is rather obvious given that, when this condition is not fulfilled, at least two actors want to establish a new tie. After the actors discover that they play different actions, two things might happen. Either the tie is severed again, or one of the actors changes behavior and becomes part of the other subnetwork. In the former situation, the two actors continue to create and sever the tie.

For our experimental conditions, we characterize the general results on stable states more precisely. The experiment is conducted with groups of eight players, who play one of the two games shown in Fig. 3. The tie costs are defined as $k(t_i) = 6t_i + t_i^2$, such that, for both games, $t_L = 7$ and $t_R = 5$. That is, given this cost function and the payoffs in the games, actors playing LEFT can profitably maintain, at most, seven ties (with other actors also playing LEFT), while actors playing RIGHT can profitably maintain, at most, five ties (with other actors also playing RIGHT).

		Player 2				Player 2	
		LEFT	RIGHT			LEFT	RIGHT
Player 1	LEFT	20,20	0,10	Player 1	LEFT	20,20	0,14
	RIGHT	10,0	14,14		RIGHT	14,0	16,16
Game I (low risk)				Game II (high risk)			

Fig. 3. The coordination games as used in the experiment.

¹ These two extreme scenarios are special cases of a more general model in which an actor can observe neighbors only at a specific distance. Additional simulation results on intermediate cases did not imply new substantive hypotheses that were interesting for further experimental testing.

According to [Theorem 1](#), this means that in our eight-actor setup the following types of constellations are stable under global information:

- All players playing LEFT, with all ties present;
- All players playing RIGHT, with all players with a degree less than five connected to one another. This may be a single component with all players having five ties, a single component with two players having two ties each, and six players with five ties, or two components: one full component of six players and a dyad.
- Heterogeneous constellations in which some players play RIGHT and others play LEFT. In these cases, the network will consist of two components with all ties present within the components.

Under local information, both homogeneous constellations are still stable, but most of the heterogeneous constellations do not fulfill [Theorem 2](#). Consider the case where there are two fully connected components: one component of three players playing LEFT, and one of five players playing RIGHT. Because behavior within the subnetworks is homogeneous, these actors conclude that the actors in the other subnetwork also play the same action, in which case it would be profitable to form ties. Thus, ties are formed. Subsequently, depending on the specific payoffs in the game, LEFT-players might switch to RIGHT. Alternatively, the tie is severed again in the next period. As a result, the only heterogeneous stable constellations under local information consist of six RIGHT-players in a full component and a dyad of LEFT-players.

2.3. Simulation

Given the characterization of possible stable states in our experimental conditions, an open question is which of these stable states are *more or less likely* to arise given specific initial conditions. To derive such predictions, we perform computer simulations of our model with experimental conditions as parameters. The aim of the simulations is to derive sharp predictions for this parameter space, rather than explore the behavior of the model under many conditions (see [Buskens et al., 2008](#), for a broader investigation).

We study the following conditions:

- Two payoff sets for the coordination game as shown in [Fig. 3](#) that only vary in the risk involved in playing LEFT. Because the risk is lower in Game I than in Game II, we label them “low risk” and “high risk,” respectively; in both games, playing RIGHT is risk-dominant.
- Two information regimes as described above, in which actors can either observe the actions of all other actors or their neighbors’ actions only.
- Three initial networks with eight actors: the full network, the circle network, and the two-squares network (see [Fig. 2](#)).

- The propensity to play LEFT in the first period, using 2/8, 3/8, 4/8, 5/8, and 6/8.

The combination of these parameters leads to 60 different combinations. We simulate the network formation process until we obtain a stable situation for each of these combinations, and repeat the process 400 times. Altogether, this results in a simulated dataset of 24,000 cases.

On average, it took 40.37 tie changes and 2.65 behavioral changes to reach a stable state. Of all the simulations, 76.5% converged to a homogeneous constellation in which all actors played the same action in the coordination game. In the remaining 23.5%, heterogeneous constellations emerged. We define *efficiency* as the proportion of actors in the network playing the payoff-dominant action LEFT. Average efficiency over all simulation runs was .31. Overall, most situations converged to a homogeneous inefficient convention, and both efficient homogeneous outcomes and heterogeneous outcomes occurred less frequently.

For a more detailed examination of our results, we examined average efficiency levels and *heterogeneity* in stable states per information regime and initial network, separately for the two risk levels. To quantify the extent of heterogeneity, let π_L indicate the proportion of actors playing LEFT. Heterogeneity is defined as $h = 4\pi_L(1 - \pi_L)$. The measure varies between 0 and 1, and equals 1 when π_L is exactly 1/2 (indicating maximal heterogeneity), and equals 0 when π_L is either 0 or 1 (indicating homogeneity). The results are shown in [Figs. 4 and 5](#).

The amount of risk involved in playing LEFT has a strong effect on the resulting efficiency in stable states: efficiency is clearly lower when risk is higher (on average, the difference is .19). Effects of the initial network are clearly visible only under the high-risk condition, where efficiency appears to be lowest when starting from the two-squares network and highest when starting from the full network. These differences, however, are small. Effects of the information regime are virtually absent, except under the initial circle network with high risk, where efficiency is clearly lower under local information than under global information.

Heterogeneity is lower if risk is higher (–.21 on average). Thus, if we combine this finding with the results on efficiency in the previous paragraph, it appears that the higher efficiency in the low-risk condition is largely caused by more heterogeneous cases (rather than homogeneous efficient cases). We can also see an effect of the initial network. Especially in the low-risk condition, heterogeneity appears to be lower when starting with a full network. In the high-risk condition, this difference is also present, but only under global information. Unlike efficiency, heterogeneity clearly differs between the two information regimes. In all cases, except with the full initial network, heterogeneity is *lower* under *local* information. This seems somewhat counter-intuitive. The explanation is that with more information, it is easier to avoid actors who play a

Fig. 4. Average efficiency by risk level, initial network, and information regime (simulation results).

Fig. 5. Average heterogeneity by risk level, initial network, and information regime (simulation results).

different action. This is also consistent with the analytic result that most heterogeneous stable states under global information become unstable under local information.

Buskens et al. (2008), who analyze a similar model, report a strong effect of the initial behavioral distribution, as well as interaction effects of other parameters with this initial distribution. We examine this issue by running regression analyses with efficiency in the stable state as the dependent variable, and the initial proportion of actors playing LEFT (p_{LEFT}) and the density of the initial network (FULL) as independent variables. Because we look at only three different types of initial networks, density is, in practice, a dummy variable. As we are interested in the effect of initial density on the effect of the initial behavioral distribution, we also include an interaction effect between the two.

To estimate the effects of the independent variables on the proportion of actors playing LEFT, we treat each case (i.e., a group of eight actors) as a number of successes (actors playing LEFT), and apply logistic regression to predict the likelihood of success at the actor level. Because this approach inflates the number of observations, standard errors are adjusted accordingly.² We conduct the analysis for the two risk levels and the two information regimes separately (Table 1).

In both information regimes, the initial behavioral distribution has a very strong effect on efficiency. The main effect of the initial network's density (referring to the effect of density when the initial distribution of behavior is .5) is small and negative. There is also a rather strong positive interaction effect between density and the initial behavioral distribution. This means that the effect of the initial behavioral distribution is especially strong when the initial network is full, and smaller when the initial network is sparse. So, if networks are initially dense, having a majority perform a certain action leads to a stronger "pull" on the rest of the population.

2.4. Overview of micro- and macro-level hypotheses

To conclude the theoretical portion of the analysis, we formulate the micro- and macro-level hypotheses that we will test. The *micro-level* hypotheses refer to the individual actor's behavior, as assumed in the model. The *macro-level* hypotheses are based on the simulation results.

The first micro-level hypothesis describes how actors decide on their actions when playing the coordination game. The model assumes that actors play according to a best-reply logic: they adapt their behavior to what their current neighbors played during the previous period. As stated in Section 2.1, actors only play

the payoff-dominant action LEFT if the proportion of their neighbors also playing LEFT is at least as large as the risk threshold. These assumptions translate directly into the following hypothesis:

Hypothesis 1.1. Actors play LEFT if and only if the proportion of their neighbors who played LEFT in the previous period exceeds the risk threshold.

The next two micro-level hypotheses relate to how actors decide to create or delete ties in the network. In the experiment, tie costs are chosen such that the cost of a tie between two actors playing different actions is always larger than the payoff. This leads to the following hypothesis:

Hypothesis 1.2. Actors sever ties with neighbors who played an action different from their own in the previous period.

Hypothesis 1.2 applies to both global and local information, because when deleting a tie, actors are already aware of a specific actor's previous behavior. When creating new ties, the situation is different. Actors can only observe the behavior of potential neighbors under global information, leading to the following hypothesis:

Hypothesis 1.3. Under global information, actors create new ties with other actors who played the same action as their own action in the previous period.

Under local information, actors cannot observe the behavior of potential neighbors, and are assumed to "guess" this behavior using the average behavior of their current neighbors. Given that actors only want to create ties with actors who they expect to play the same action as themselves (see above), this implies that, under local information, an actor only wants to create a tie with an unobserved other actor if enough of her current neighbors also play "her" action. We could specify an exact proportion of neighbors who must play the same action for that actor to be willing to establish a new tie with an unobserved other actor. Instead, we formulate an implication of this assumption in more general terms, and only predict the *direction* of the neighbors' behaviors' effect on the likelihood that new ties will be formed.

Hypothesis 1.4. Under local information, the higher the proportion of an actor's neighbors who play the same action as this actor, the higher the likelihood that the actor proposes a new tie.

The macro-level hypotheses are based on the results of the simulation. The first macro-level hypothesis concerns the effect of risk on efficiency, and follows from Fig. 4 and Table 1.

Hypothesis 2.1. The higher the risk involved in playing an efficient action, the lower the efficiency in stable states.

² To estimate this model, we use the "bologit" procedure in Stata 9 (StataCorp, 2005).

Table 1

Logistic regression for grouped data on efficiency, per information regime and risk level (simulation results).

	Low risk				High risk			
	Information regime							
	Local		Global		Local		Global	
	Coeff.	S.e.	Coeff.	S.e.	Coeff.	S.e.	Coeff.	S.e.
PLEFT	11.30	0.16	10.65	0.14	17.64	0.56	12.74	0.27
FULL	−0.36	0.08	−0.71	0.08	−2.61	0.37	−3.39	0.32
FULL × PLEFT	17.34	1.03	18.49	0.95	22.72	2.04	22.64	1.73
Constant	−6.79	0.10	−6.04	0.08	−13.46	0.41	−9.44	0.19
Number of groups	6000		6000		6000		6000	
Log pseudolikelihood	−13,461.51		−14,032.85		−7097.06		−9985.47	
McFadden's pseudo R^2	0.58		0.57		0.70		0.61	

A second set of hypotheses about the effects of the initial network and information regime on efficiency follows from Fig. 4 and Table 1.

Hypothesis 2.2. The higher the initial efficiency, the higher the efficiency in stable states.

Hypothesis 2.3. The higher the initial network's density, the lower the efficiency in stable states.

Hypothesis 2.4. The higher the initial network's density, the stronger the effect of the initial behavioral distribution on the emerging distribution of behavior.

Lastly, we derive a set of hypotheses that are concerned with effects on heterogeneity following from the results on heterogeneity in Fig. 5.

Hypothesis 2.5. Higher initial network density leads to lower heterogeneity in stable states.

Hypothesis 2.6. In initially sparse networks, more information leads to higher heterogeneity in stable states.

3. Experimental design

We test these hypotheses in a computer-aided experiment, designed to reflect both the assumptions of the theoretical model and its implementation in the simulation model as closely as possible. Subjects played the repeated coordination games used in the simulation, choosing one strategy with all neighbors, while also having the opportunity to choose with whom they interacted. The experimental conditions included the three initial networks used in the simulation (see Fig. 2), the two information regimes, and the two risk levels. The games used in the experiment are shown in Fig. 3. Groups of eight subjects played one of the games for 15 subsequent periods. In each period, they faced the following decisions (in this order):

- (1) Decide whether to change one relation: that is, propose one new link to another subject or unilaterally sever one existing link;
- (2) Decide whether to accept incoming proposals from other subjects;
- (3) Choose their behavior (LEFT or RIGHT).

Because subjects could also accept incoming proposals, more than one tie per subject may change in a given period. This setup is identical to the procedure used in the simulation. Only in the first period could subjects not change their network, because this was imposed as an experimental condition.

As in the simulation, the cost function for ties was $k(t_i) = 6t_i + t_i^2$, with t_i the number of ties for subject i . In practice, the cost function was not presented to the subjects literally; instead, the instruction

included a table showing the total costs for each possible number of ties.

Under global information, subjects were shown their current neighbors at the beginning of each period, as well as group members' behavior in the previous period. Under local information, subjects saw only the behavior of their *neighbors* from the previous period. At the end of each period, the resulting payoffs were shown, as were current neighbors and the behavior of either all subjects or only neighbors (again, depending on the information regime).

Appendix A, available as an electronic supplement, provides some translated screen-shots from the computer interface, as well as the translation of the complete instructions. The experiment was in Dutch. The interface did not provide any information about the structure of the network, besides showing who the subject's direct neighbors were. This means that the two-squares network and the circle network could not be distinguished by subjects. That is, on screen, the two-squares network and the circle network looked exactly the same. Therefore, differences in outcomes between these two conditions can only be due to the dynamic process. At the bottom of the screen, subjects could review the complete history of their interactions. In this way, we meant to reduce unobserved differences between subjects in their ability to memorize previous events. No communication between subjects was allowed. The experiment was programmed and conducted with z-Tree software (Fischbacher, 2007). After the experiment, subjects were paid €0.01 for every point earned during the experiment. All sessions took place at the Experimental Laboratory for Sociology and Economics (ELSE) at Utrecht University.

Subjects were recruited from amongst undergraduate students at Utrecht University using the internet recruitment system ORSEE (Greiner, 2004). Using this system, we recruited subjects year-round and then invited them for experiments as needed.

For each session, 16 subjects were invited. Subjects were randomly assigned to workstations, and received instructions on paper after a short verbal introduction. After reading the instructions, the subjects played three practice periods, in which they played against simulated subjects instead of against one another. This facet was explicitly communicated. During the entire experiment, subjects were never deceived and were always allowed to ask for assistance from the experimenters (which happened rarely).

After the practice periods, the subjects played the experimental game for 15 periods in two groups of eight. After these 15 periods, the subjects were re-matched into two different groups of eight, and again played the game for 15 periods (though under a different condition). The new condition always involved a different initial network, a different information regime, and the same payoffs.

4. Results

We first present some general macro-level results, and then proceed to test the macro-level hypotheses (Section 4.1). In Section 4.2,

Fig. 6. Average proportion LEFT-choices in six conditions.

we test the micro-level hypotheses on individual decisions in the coordination game. In Section 4.3, we test the micro-level hypotheses on linking decisions.

4.1. Macro-level results

The experiment involved 12 sessions with 192 subjects, over 90% of whom were students, mostly freshmen. Students came from over 30 different fields, the most numerous of these being sociology, economics, and psychology. Out of all subjects, 61% were female and the average age was 22.1 years. Each session took about 75 min to complete. Subjects earned €12.30 on average. Because subjects played under two conditions, we had a total of 48 groups. One of the sessions (with two groups) did not completely run until the 15th period in the second condition due to a technical problem, such that we only have data from 46 groups for some of the analyses.

Fig. 6 presents the average proportion of people choosing LEFT per initial network per period for each information regime. From this figure it is immediately clear that subjects had a tendency to choose LEFT. A majority of subjects chose LEFT in all but one con-

dition from the first period. This is largely consistent with earlier experimental research on coordination games (cf. Straub, 1995). In the final period, 74% of all subjects played LEFT. Of all groups, fifteen converged on playing LEFT and only four converged on playing RIGHT. Eight of the fifteen LEFT-playing groups also reached a stable state in tie choices; that is, they had established a full network after 15 periods. None of the RIGHT-playing groups managed to reach a stable state in tie choices, although their behavior was stable. One group reached a stable state with heterogeneous behavior under global information: after 15 periods, the network consisted of a full component of five subjects playing LEFT, and a full component of three subjects playing RIGHT, with no links between the two components. A second group converged on this constellation under local information, although this was not stable given model assumptions.

To investigate the macro-level outcomes of the co-evolution process, we take all observations from the 15th period. Table 2 shows the efficiency and heterogeneity for each experimental condition, and also shows totals for the three initial networks and two information regimes. Figs. 7 and 8 present a graphical impression of these results in a similar fashion as that for the simulation (see Figs. 4 and 5).

To test Hypothesis 2.1, we compare the average efficiency between risk levels. Clearly, efficiency is lower in the high-risk condition (.92 vs. .67), which is significant using a Mann–Whitney test ($p = 0.006$) and supports Hypothesis 2.1.

In accordance with Hypothesis 2.3, efficiency is lowest for the full initial network (rightmost column of Table 2). This difference, however, is not significant. Moreover, closer inspection reveals that the difference only occurs under global information; under local information, the difference is either zero or reversed.

To test Hypotheses 6–8 on efficiency, we run a regression analysis similar to that in the simulated data (see Table 1): logistic regression for grouped data on the number of LEFT-choices in period 15 as the dependent variable. The unit of observation in this analysis is the individual decision; standard errors are adjusted to

Table 2

Efficiency and heterogeneity per experimental condition ($N = 46$).

Initial network		Risk level and information regime								
		Low risk			High risk			Total		
		Local	Global	Total	Local	Global	Total	Local	Global	Total
Two squares	Efficiency	0.97	0.75	0.90	0.66	1.00	0.83	0.81	0.92	0.86
	Heterogeneity	0.11	0.50	0.24	0.30	0.00	0.15	0.20	0.17	0.19
Circle	Efficiency	1.00	0.97	0.98	0.41	0.81	0.61	0.70	0.89	0.80
	Heterogeneity	0.00	0.11	0.05	0.45	0.47	0.46	0.23	0.29	0.26
Full	Efficiency	1.00	0.75	0.88	0.63	0.53	0.58	0.81	0.64	0.73
	Heterogeneity	0.00	0.00	0.00	0.25	0.33	0.29	0.13	0.16	0.14
Total	Efficiency	0.99	0.84	0.92	0.56	0.78	0.67	0.78	0.81	0.79
	Heterogeneity	0.04	0.14	0.09	0.33	0.27	0.30	0.18	0.21	0.20

Efficiency = efficiency, Heterogeneity = heterogeneity.

Fig. 7. Average efficiency in the last period by risk level, initial network, and information regime (experimental results).

Fig. 8. Average heterogeneity in the last period by risk level, initial network, and information regime (experimental results).

account for the fact that individuals are clustered in 46 groups. To improve statistical power, we pool observations from different conditions and analyze them simultaneously, using control variables for the various conditions. As independent variables, we use the distribution of behavior in period 1 (PLEFT), a dummy variable indicating whether the initial network was the full network (FULL), and the interaction between the two. The main effect of the network dummy refers to the situation where efficiency in period 1 is .5, because this variable is centered at .5. Moreover, we include a dummy for local information (LOCAL) and a dummy for high-risk level (HIGH RISK). Table 3 presents the results of the analysis. There is a strong positive effect of initial behavior on behavior in the last period, which confirms Hypothesis 2.2. Also, there is a smaller but significant negative effect of starting in the full network, which supports Hypothesis 2.3.

There is a positive but insignificant interaction effect between the initial proportion playing LEFT and the full network. Hypothesis 2.4 can therefore not be confirmed.

Unlike in Table 2, there is no significant effect of HIGH RISK after controlling for the behavior in the first period. This suggests that differences between risk conditions in the final period are caused by differences in subjects' decisions in the first period, when they are not yet reacting to other subjects, and not by differences in the co-evolution process.

Table 2 shows average heterogeneity in the last period by experimental condition. Hypothesis 2.5 predicts that a higher density of the initial network leads to lower heterogeneity. Heterogeneity is indeed lower in the full network condition than in the other two conditions combined. This difference, however, is not significant using a Mann–Whitney test. Thus, Hypothesis 2.5 cannot be supported.

Hypothesis 2.6 predicted that more information leads to greater heterogeneity, especially in low-density networks. Heterogeneity is lower under local information in the circle network (.23 vs. .29), but not in the two-squares network (.20 vs. .17). Moreover, these

Fig. 9. Proportion of subjects playing LEFT by proportion of their neighbors playing LEFT in the previous period.

differences are not significant. Also, if we compare heterogeneity between local and global information over all networks, there is no significant difference.³

4.2. Individual behavior I: decisions in the coordination game

We analyze individual behavior to assess the extent to which the model reflects actual decision-making by subjects. Hypothesis 1.1 states that subjects should play LEFT only if the proportion of their neighbors who played LEFT in the previous period exceeds the risk threshold. First, we plot average efficiency against the distribution of neighbors' behavior in the previous period for both risk levels separately. Under the assumption that people exclusively play a best reply against what their neighbors did in the previous period, subjects are expected to play RIGHT under low risk as long as less than 58% of their neighbors play LEFT. Under high risk, this percentage is 73%.

Fig. 9 shows a somewhat more complicated picture. Clearly, subjects switch to LEFT at lower proportions than the .58 and .73 thresholds. Already, the proportion of LEFT-choices increases considerably in both cases at levels above .35. This indicates that subjects tend to take on more risk than the simple myopic best-reply heuristic of the model assumes. However, subjects' behavior is strongly associated with their neighbors' behavior in the previous period. For the majority of subjects, the threshold for switching lies between

Table 3

Logistic regression for grouped data of the proportion playing LEFT (experimental results).

	Coeff.	S.e.	p
PLEFT	8.77	2.71	0.00
FULL	−2.25	0.87	0.01
FULL × PLEFT	8.81	7.07	0.21
LOCAL	−0.76	0.61	0.22
HIGH RISK	−0.32	0.71	0.67
Constant	−2.38	2.32	0.30
Number of obs.	368		
Log pseudolikelihood	−100.09		
McFadden's pseudo R ²	0.46		

³ If we use regression analysis to predict heterogeneity with multiple predictors simultaneously, we find the same results.

Table 4

Logistic random intercept regression on playing LEFT, per information condition.

	Information regime					
	Global			Local		
	Coeff.	S.e.	p	Coeff.	S.e.	p
NEIGHTHRES	1.01	0.45	0.03	−0.06	0.44	0.89
NEIGHLEFT	3.78	0.67	0.00	5.98	0.77	0.00
EGOLEFT	1.52	0.29	0.00	2.88	0.26	0.00
GROUPELEFT	2.11	0.77	0.01	–	–	–
HIGH RISK	−0.16	0.37	0.67	−1.10	0.38	0.00
PERIOD	0.01	0.03	0.69	0.06	0.03	0.07
NUMTIES	0.06	0.08	0.49	0.07	0.08	0.37
PART	0.76	0.39	0.05	0.25	0.30	0.40
Constant	−4.07	0.55	0.00	−4.05	0.57	0.00
Var. ind. level	1.21	0.53		0.85	0.45	
Var. group level	0.00	0.00		0.00	0.00	
Number of obs.	2520			2677		
Log likelihood	−361.25			−324.22		

.4 and .6. Although the payoff functions are different in both risk conditions, subjects' reactions to their neighbors are very similar.

Fig. 9 shows only the bivariate relation between neighbors' behavior and the subjects' own behavior. To isolate the effect of neighbors' behavior from those of other effects, we conduct a logistic regression analysis with behavior as the binary dependent variable and subjects' decisions in every period as the unit of observation. Because observations within and between subjects are not independent, we use a model with a random intercept at the individual and the group levels. To test whether it is really the risk threshold that matters, we include a dummy variable indicating whether the proportion of neighbors who played LEFT in the previous period exceeds the threshold (NEIGHTHRES), in addition to the proportion of the subject's neighbors playing LEFT (NEIGHLEFT) in the previous period. According to Hypothesis 1.1, NEIGHTHRES should have a significant positive effect, but there should be no additional significant effect of the proportion of neighbors playing LEFT (NEIGHLEFT). As control variables, we include a dummy variable for high risk (HIGH RISK), the subject's own behavior in the previous period (EGOLEFT), and the proportion of the whole group playing LEFT (GROUPELEFT). The latter variable is only included in the model for global information, as subjects were not informed about the behavior of the group beyond their own neighbors. Furthermore, we include time (PERIOD), the number of ties a subject has (NUMTIES), and whether the decision was made in the first part or the second part of the experiment (PART). Because this set of variables differs between information regimes, we estimate separate models for each regime (Table 4).

To test Hypothesis 1.1, we compare the effects of NEIGHTHRES and NEIGHLEFT. Under global information, there is a significant positive effect of NEIGHTHRES, as predicted: subjects are more likely to play LEFT if the proportion of neighbors who played LEFT exceeds the risk threshold. However, in contradiction with Hypothesis 1.1, there is an additional effect of neighbors' behavior (NEIGHLEFT). Under local information, there is only an effect of NEIGHLEFT, and no significant effect of the specific threshold. Thus, although these results again show that subjects do strongly react to their neighbors' behavior, their behavior does not conform exactly to the threshold effect as predicted by Hypothesis 1.1. There is also a significant effect of subjects' own behavior in the previous period, indicating some degree of behavioral inertia, especially under the local information condition. Under global information, the average behavior of the group in the previous period has a significant positive effect. Thus, controlling for the behavior of neighbors in the previous period, subjects tend to go along with the rest of the group. A possible explanation for this finding is that subjects are to some degree forward-looking; that is, they adapt to the behavior by non-neighbors in anticipation

of becoming neighbors themselves. Under the global information condition, there is also a weakly significant effect of PART, indicating that subjects were more likely to play LEFT when they were playing the second set of 15 periods, after the reshuffling of groups. Nevertheless, additional analyses (not reported here), in which we estimate the same models but use only observations from the first 15 periods (before rematching) do not show substantive differences.

These findings suggest ways in which our actor model could be improved. One way would be to endow actors with forward-looking capabilities. Deriving precise implications from such an adapted model is beyond the scope of this paper (i.e., it would require another simulation exercise as in Section 2). To speculate, however, we expect that if it is true that efficiency is higher than expected because RIGHT-playing subjects anticipate ties with LEFT-playing non-neighbors, this effect is smaller under local information (because subjects cannot observe non-neighbors). This is consistent with the results in Table 2 showing that efficiency is lower under local information.

4.3. Individual behavior II: linking decisions

Fig. 10 shows the proportion of ties created and dissolved by information regime. The results distinguish between pairs playing similar or dissimilar behavior. Ties between actors playing the same action (similar pairs) are created more frequently than ties between actors playing different actions (dissimilar pairs). Under local information, this relationship is much weaker because actors have to guess who will act similarly.⁴ Ties between dissimilar individuals are more likely to be dissolved than ties between similar individuals. These results are all significant. Thus, subjects tend to sever ties with neighbors who behave differently in either information regime, and create ties more often with subjects who display similar behavior as far as they can observe. Hypotheses 2 and 3 are therefore confirmed.

Fig. 10 does not show how subjects make linking decisions under local information. In the simulation model, we assumed that actors use information on their neighbors' behavior and "project" this onto potential neighbors. To investigate whether subjects use this heuristic, we run a random intercept logistic regression analysis of the decision to create at least one new tie (either doing by making or accepting a proposal). The unit of analysis is the subject-period; the dependent variable is coded 1 if the subject created at least one new tie during the period, and 0 otherwise. Because we are interested in the creation of new ties, we only include cases in which the maximum profitable number of ties (5 or 7) was not yet reached. Moreover, to account for the interdependence of decisions within subjects over periods and between subjects within groups, we again add random intercepts at the subject level (14 periods per subject) and the group level (eight subjects per group).

We include as independent variables the number of ties the subject already has (NUMTIES), the subject's behavior in the previous period (EGOLEFT), the average behavior of the subject's neighbors in the previous period (NEIGHLEFT), and the proportion of neighbors acting the same as the focal subject (NEIGHSIM). We hypothesize that the effect of NEIGHSIM is positive: the more a subject experiences that his local environment acts like him- or herself. Lastly, we add the period (PERIOD), risk level (HIGH RISK), and whether the decision was made in the first or second part of the experiment (PART).

Table 5 shows the results. First, the significant negative effect of NUMTIES reflects the increase in marginal tie costs as implemented

⁴ The small difference still noticeable under local information might be due to subjects avoiding other subjects with whom they just dissolved a tie: these subjects are more likely to play dissimilar behavior.

Fig. 10. Proportions of ties changed.

Table 5
Random intercept logistic regression analysis on creating at least one new tie with local information.

	Coeff.	S.e.	<i>p</i>
EGOLEFT	0.01	0.17	0.95
NEIGHLEFT	0.05	0.25	0.84
NEIGHSIM	−2.15	0.30	0.00
NUMTIES	−0.39	0.05	0.00
HIGH RISK	−1.02	0.27	0.00
PERIOD	−0.08	0.02	0.00
PART	0.48	0.22	0.03
Constant	4.87	0.67	0.00
Var. ind. level	0.01	0.07	
Var. group level	0.24	0.09	
Number of obs.	1509		
Log likelihood	−869.29		

in the cost function: the more ties a subject already has, the smaller the probability that she will form another one. Moreover, the likelihood of creating a new tie decreases with PERIOD and RISK. Contrary to expectations, the effect of NEIGHSIM is *negative*. Subjects are *less* likely to create new ties when they are more similar to their neighbors, and they thus are not using the heuristic assumed in the simulation model. Rather, subjects seem to assume that unknown potential neighbors are playing the alternative action. It might be that this unexpected effect is the result of subjects remembering interactions in previous periods, for which the theoretical model does not allow. Fig. 10 shows that subjects tend to sever ties with neighbors playing the alternative behavior, and it might be that after such a deletion they conclude that these previous neighbors (now invisible under local information) will persist in playing the different behavior. This logic would result in a negative effect of NEIGHSIM. Additional analyses show that the negative effect of NEIGHSIM becomes stronger in later periods of the experiment. This does indeed suggest that subjects make use of their knowledge of the history of play, albeit in a manner different from the manner proposed before.

5. Conclusions and discussion

We studied coordination in dynamic networks, focusing on efficiency and heterogeneity of emergent behavior, and on the influence of information availability. We specified a game-theoretic model and used simulation methods to generate specific predictions about the effects of initial conditions and limited information on the efficiency and heterogeneity of emergent behavior. We tested micro- and macro-level hypotheses in a laboratory exper-

iment. Three macro-level hypotheses were confirmed. First, the behavior of subjects in the first period determines to a large extent their behavior in the final period. Second, efficiency is lower if risk is higher. Third, efficiency is lower if the network is initially denser. For the remaining three macro-level hypotheses on the effects of initial network structure and information availability, the results were always in the expected direction but not significant.

At the micro level, we found that, by and large, individual behavior appears to resemble behavior as assumed in the model, at least when information on the behavior of non-neighbors is available. Subjects adapt their behavior to that of their neighbors in the previous period (Hypothesis 1.1). In their choices of network relations, subjects have a clear preference for subjects who play the same action as themselves (Hypothesis 1.3) and exclude those who play the alternative action (Hypothesis 1.2). Moreover, many experimental groups managed to converge on stable constellations that were theoretically predicted involving both efficient conventions and situations in which no single convention was reached.

In some respects, the behavior of subjects also deviates from the behavior assumed in the model: subjects more easily opt for the payoff-dominant action than would be expected from myopic best-reply behavior (Hypothesis 1.1). Moreover, if subjects are informed on the behavior of the whole group rather than only their neighbors, they are also influenced by the behavior of those who are not their neighbors, which suggests some anticipation of future interactions with these other subjects. For local information, it was theorized that subjects use the average behavior of their neighbors as a predictor for the behavior of potential neighbors (Hypothesis 1.4). However, subjects seem to use the behavior of their neighbors as a predictor for what their neighbors are *not* doing. Furthermore, we showed that this effect becomes stronger in later periods of the experiment, which suggests that subjects use the history of interactions in their decisions.

Before we move to broader conclusions based on the results, let us briefly discuss some limitations of the experimental design and analyses. First, the number of groups was relatively low given the number of experimental conditions. This is the result of a practical trade-off between group size and the number of groups. We feel that for network experiments to capture the notion that individuals cannot observe or influence the network as a whole (as is mostly the case in real life), groups should be relatively large. Our experiment is one of the few in which groups larger than six are used. The price, due to practical and financial reasons, is fewer groups. The lack of significant results at the group level (even though results tend to be in the expected direction) might be due to this low number of observations.

Second, the choice of payoffs led subjects to choose the payoff-dominant action relatively easily, resulting in little variation in macro-level outcomes, especially under the low-risk condition. The results might have been stronger if the differences between the payoff-dominant equilibrium and the risk-dominant equilibrium had been more pronounced.

The many interdependencies, especially with linking behavior, posed significant methodological challenges that we dealt with only in part. Multinomial logistic regression models for tie decisions could have been used, which resemble other more sophisticated methods for longitudinal network data (i.e., Snijders, 2001; Snijders et al., 2007). Considering the already rather extensive theoretical and empirical analyses, we chose not to introduce these further complexities.

Our results indicate that people are able to coordinate on efficient behavior if the interaction structure is not exogenously determined, but rather co-evolves with behavioral choices. We also found that the initial network structure matters: if the network is initially denser, emerging conventions are more likely to be inefficient. We did not find convincing evidence that the emergence of conventions is very dependent on information availability: also if subjects possessed only local information, they reached high efficiency levels.

A further conclusion from our results is that the simple model in which actors play the best reply against their neighbors' behavior in the previous period is too simple to adequately capture real individual behavior in situations represented by the model. We saw signs of both forward-looking and backward-looking behavior. This seems to lead to a more frequent emergence of efficient behavior than theoretically predicted. The discrepancy between the model and the actual behavior of our subjects might be a reason for deviations between predictions and observations at the macro-level as well, although it is not yet clear how micro-level differences affect macro-level outcomes. Therefore, extending the existing models to incorporate more complex decision-making processes and derive new macro-level implications is a desirable direction for further research.

However, one might also argue that such effects are typical for the relatively small group sizes used in our experiment. In some real-life applications, where groups tend to be much larger, both remembering previous interactions and anticipating behavior in the population as a whole would be much harder. From this perspective, our model might be judged as fairly appropriate for modeling coordination in real networks. This is because behavior on the individual level in the model approximates the empirical behavior of subjects in the experiment fairly well. It remains to be seen (in experiments with more observations on the group level or in field applications) whether the predictions of the model also hold on the macrolevel. Moreover, a remaining challenge is to understand choices in network formation when limited information is available.

Appendix A. Supplementary data

Supplementary data associated with this article can be found, in the online version, at doi:10.1016/j.socnet.2009.04.002.

References

- Abrahamson, E., Rosenkopf, L., 1997. Social network effects on the extent of innovation diffusion: a computer simulation. *Organization Science* 8, 289–309.
- Berninghaus, S., Schwalbe, U., 1996. Conventions, local interaction, and automata networks. *Journal of Evolutionary Economics* 6, 231–324.
- Berninghaus, S., Vogt, B., 2006. Network formation in symmetric 2×2 games. *Homo Oeconomicus* 23, 421–466.
- Berninghaus, S.K., Ehrhardt, K.-M., 1998. Time horizon and equilibrium selection in tacit coordination games: experimental results. *Journal of Economic Behavior and Organisation* 37, 231–248.
- Buskens, V., Corten, R., Weesie, J., 2008. Consent or conflict: coevolution of coordination and networks. *Journal of Peace Research* 45, 205–222.
- Buskens, V., Snijders, C., 2008. Effects of network characteristics on reaching the payoff-dominant equilibrium in coordination games: a simulation study. ISCORE paper 232, Utrecht University.
- Camerer, C.F., 2003. *Behavioral Game Theory: Experiments in Strategic Interaction*. Princeton University Press, Princeton, NJ.
- Centola, D., Macy, M.W., 2007. Complex contagions and the weakness of long ties. *American Journal of Sociology* 113, 702–734.
- Coleman, J.S., 1990. *Foundations of Social Theory*. Belknap, Cambridge, MA.
- Cooper, R., 1990. Selection in coordination games: some experimental results. *American Economic Review* 80, 218–233.
- Corbae, D., Duffy, J., 2008. Experiments with network formation. *Games and Economic Behavior* 64, 81–120.
- Crawford, V., 1997. Theory and experiment in the analysis of strategic interaction. In: Kreps, D.M., Wallis, K. (Eds.), *Advances in Economics and Econometrics: Theory and Applications*, vol. I. Cambridge University Press, Cambridge, pp. 206–242.
- Ehrhardt, G., Marsili, M., Vega-Redondo, F., 2006. Phenomenological models of socioeconomic networks. *Physical Review E*, 74.
- Elias, N., 1969. *The Civilizing Process*, vol. I. The History of Manners. Blackwell, Oxford.
- Fischbacher, U., 2007. *z-Tree – Zurich toolbox for readymade economic experiments – experimenter's manual*. Institute for Empirical Research in Economics, University of Zurich.
- Goyal, S., Vega-Redondo, F., 2005. Network formation and social coordination. *Games and Economic Behavior* 50, 178–207.
- Granovetter, M.S., 1978. Threshold models of collective behavior. *American Journal of Sociology* 83, 1420–1443.
- Greiner, B., 2004. The online recruitment system ORSEE 2.0—a guide for the organization of experiments in economics. Working Paper Series in Economics. University of Cologne.
- Hardin, R., 1995. *One for All: The Logic of Group Conflict*. Princeton University Press, Princeton, NJ.
- Hardin, R., 2007. *David Hume: Moral and Political Theorist*. Oxford University Press, Oxford.
- Harsanyi, J.C., Selten, R., 1988. *A General Theory of Equilibrium Selection in Games*. MIT Press, Cambridge, MA.
- Hume, D., [1739–40], 1978. *A Treatise of Human Nature*. Clarendon Press, Oxford.
- Jackson, M.O., Watts, A., 2002. On the formation of interaction networks in social coordination games. *Games and Economic Behavior* 41, 265–291.
- Jackson, M.O., Wolinsky, A., 1996. A strategic model of social and economic networks. *Journal of Economic Theory* 71, 44–74.
- Kandori, M., Mailath, G.J., Rob, R., 1993. Learning, mutation and long run equilibria in games. *Econometrica* 61, 29–56.
- Knecht, A.B., 2008. Friendship selection and friends' influence: dynamics of networks and actor attributes in early adolescence. ICS Dissertation Series. Utrecht University, Utrecht.
- Lewis, D.K., 1969. *Convention*. Harvard University Press, Cambridge, MA.
- Marsden, P.V., Friedkin, N.E., 1993. Network studies of social influence. *Sociological Methods and Research* 22, 127–151.
- McPherson, M., Smith-Lovin, L., Cook, J.M., 2001. Birds of a feather: homophily in social networks. *Annual Review of Sociology* 27, 415–444.
- Skyrms, B., Pemantle, R., 2000. A dynamic model of social network formation. *Proceedings of the National Academy of Sciences* 97, 9340–9346.
- Snijders, T.A.B., 2001. The statistical evaluation of social network dynamics. In: Sobel, M.E., Becker, M.P. (Eds.), *Sociological Methodology*. Blackwell, Boston, MA, pp. 361–395.
- Snijders, T.A.B., Steglich, C., Schweinberger, M., 2007. Modelling the co-evolution of networks and behavior. In: van Montfort, K., Oud, H., Sattora, A. (Eds.), *Longitudinal Models in the Behavioral and Related Sciences*. Lawrence Erlbaum, Mahwah, NJ, pp. 41–71.
- StataCorp, 2005. *Stata Statistical Software: Release 9*. StataCorp LP, College Station, TX.
- Straub, P.S., 1995. Risk dominance and coordination failure in static games. *The Quarterly Review of Economics and Finance* 35, 339–363.
- Ullmann-Margalit, E., 1977. *The Emergence of Norms*. Clarendon Press, Oxford.
- Van Huyck, J.B., Battalio, R.C., Beil, R.O., 1990. Tacit coordination games, strategic uncertainty and coordination failure. *American Economic Review* 80, 234–248.
- Watts, D.J., 1999. Networks, dynamics, and the small-world phenomenon. *American Journal of Sociology* 105, 493–527.
- Young, H.P., 1993. The evolution of conventions. *Econometrica* 61, 57–84.
- Young, H.P., 1998. *Individual Strategy and Social Structure*. Princeton University Press, Princeton, NJ.