Alameda 3363
Estación Central-Santiago
Tel. +56 2 7180765
www.usach.cl

Universidad de Santiago de Chile

Departamento de Economía

Working Paper Series

(Anti-) Coordination in Networks

Autor: J. Kovarik, F. Mengel y J. G. Romero

WP 2011 - Nº 07

(Anti-) Coordination in Networks*

Jaromir Kovarik† University of the Basque Country

Friederike Mengel ‡ Maastricht University

José Gabriel Romero§ University of Santiago de Chile

April 14, 2011

Abstract

We study (anti-) coordination problems in networks in a laboratory experiment. Participants interact with their neighbours in a fixed network to play a bilateral (anti-) coordination game. Our main treatment variable is the topology of the underlying network. We find that while coordination failure is very persistent in some networks, it does not occur in others. In addition, even though there is a large multiplicity of Nash equilibria theoretically, we observe a very sharp selection empirically. The network structure induces regularities on selection which cannot be understood by relying on conventional equilibrium refinements. As a second treatment variation we let agents decide endogenously on the amount of information they would like to have and find that local (endogenous) information is equally effective in ensuring successful coordination as full information. Key to understanding our results is the fact that more connected players have more 'stable best responses'. We introduce this concept formally and provide empirical evidence for it using our data set. JEL-Classification: C72, C90, C91, D85. Keywords: *Game Theory, Networks, Coordination Problems, Learning, Experiments.*

*We wish to thank Jordi Brandts, Sanjeev Goyal, Dan Levin, Pedro Rey-Biel, Marco van der Leij, Ingrid Rohde, Jakub Steiner, Martin Strobel, Aljaz Ule and Alexander Vostroknutov as well as seminar participants at Birkbeck College, Cambridge, Leicester, Maastricht, Microsoft Research, QMU London, Odense and Royal Holloway for valuable comments. Friederike Mengel thanks the European Union (grant PIEF-2009-235973) and Jaromír Kovářík the Basque Government (IT-223-07) and the Spanish Ministry of Science and Innovation (SEJ2006-06309 & ECO2009-09120)for financial support.

†Dpto Fundamentos Análisis Económico I, Universidad del País Vasco/Euskal Herriko Unibersitatea, Av. Lehendakari Aguirre 83, 48015 Bilbao, Spain. jaromir.kovarik@ehu.es

‡Corresponding Author. Department of Economics (AE1), Maastricht University, PO Box 616, 6200 MD Maastricht, The Netherlands; Tel:+31-43-38 83891; Fax:+31-43-38 84878; E-mail: f.mengel@maastrichtuniversity.nl

§Departamento de Economia, USACH, Alameda 3363, Santiago, Chile. gabriel.romero@usach.cl

1 Introduction

Anti-coordination refers to coordination problems where agents have to choose different actions in order to reach a Nash equilibrium.[1] More generally situations of anti-coordination represent interactions where choices are strategic substitutes, i.e. where the incentives of agents to choose a certain strategy decrease as more of their neighbors choose that strategy. Such effects typically appear in contexts such as congestion, pollution, oligopolistic (quantity) competition, immunization, provision of public goods or whenever there are gains from differentiation.[2] Individual differentiation within an organization (developing an expertise not duplicated by others nearby) is one example but also many problems in logistics and operations research can be thought of as Anti-Coordination problems on networks (see e.g. Garey and Johnson, 1979). (Anti-) Coordination in networks also has implications for development economics (Ray, 2007), technological standard-setting (Katz and Shapiro, 1986), financial regulation (Leitner, 2006) and corruption (Cheloukine and King, 2007) among others.

Whenever individuals in these situations interact only with a subset of the population, networks seems the right way to think about the matching structure. A central question is then how the structure of the underlying network affects the coordination problem of the players. Understanding how network structure impacts coordination is important for at least two reasons. It can help designing networks such that desired outcomes are achieved and makes it possible to identify and target the players in a network who have most leverage on coordination process.

Relationships between network structure and behaviour are hard to establish in empirical field studies of existing networks. In such studies, the network structure is fixed and given, thus preventing the investigation of alternatives. A different approach is to conduct controlled laboratory experiments in which the network structure and the information available are deliberately varied. This is the approach we take in this paper.

In our study we focus on one key feature of real life networks, which is network heterogeneity and we ask whether, and if so how exactly, such heterogeneity affects (anti-) coordination. The importance of network heterogeneity has long been recognized in other fields such as e.g. in development economics. Hirschman (1958) for example argued that subsidizing "leading" individuals or industry sectors may be crucial for development if these players have leverage and can trigger a transition to a more favorable equilibrium.[3] Also governments and firms can make use of social networks to increase revenues by targeting certain people in the social network, who have more influence on the outcome of the coordination process (Galeotti and Goyal, 2009).

In our experiment we let participants interact in a 4×4 Anti-Coordination game with their neighbours in a network. The main treatment variation in our study is the topology of the underlying network. Networks are chosen in such a way that a number of characteristics (number of nodes, links, clustering, characteristic path length, average degree etc.) are held constant, while heterogeneity in degree (the variance in the number of connections) changes. Of course there may be potential conflicts between effective coordination and efficiency or redistributive concerns. The design of our study will enable us to identify these.

As a second treatment variation we let agents either decide endogenously on the amount of information they would like to see or provide them with full information about the network. If a certain network structure facilitates coordination only if all agents are fully informed about the entire network structure, then this may be problematic. In real life agents will only have local information

[1] See e.g. Bramoulle (2007) or Vega Redondo (2007).

[2] There are many empirical studies showing such effects. A nice example are Foster and Rosenzweig (1995) who find evidence that farmers tend to experiment less with new technologies if their neighbors experiment more.

[3] See also Ray (2007).

about the network and increasing the amount of information they have may be very costly. In addition it is unclear whether (boundedly rational) agents would actually know how to make use of such additional information. Our second treatment variation hence aims at understanding the amount of information players need in order to coordinate.

We find the following results

1. Network topology has a strong impact on behavior. In particular heterogeneity improves coordination in networks.

2. Even though there is a large multiplicity of Nash equilibria theoretically, a very sharp selection is obtained empirically. In particular: (i) the most connected players can impose their preferred Nash equilibrium almost all the time and (ii) coordination is always such that all connected players choose best responses *bilaterally* (which theoretically is a sufficient but not necessary condition for Nash equilibrium).[4]

3. Local (endogenous) information about the network structure is as effective as full information in ensuring coordination and yields the same selection of Nash equilibria.

Our results show that analyzing the network structure can help us understand when coordination is likely to be successful and when it is likely to fail. Network structure affects convergence to Nash equilibrium, but the regularities on equilibrium selection we observe are the same across the networks we study.

For any given network our study can serve as an example showing how understanding the topology of the network can help with equilibrium selection even in situations where conventional refinements have no bite. Our study is meant to provide insights mainly for small networks. We also provide an explanation for our results based on the concept of "stable best responses" which we introduce formally in Section 3 of the paper. This concept (and some variations of it that we discuss as well) could be used as new equilibrium refinements for network games. It can also be useful to understand to which extent and under which conditions the particular regularities identified here generalize to larger networks.

Key to understanding our results is the fact that highly connected players will tend to have more 'stable best responses', meaning that their best response is less sensitive to the behavior of other agents in the network.[5] If local heterogeneity is high then the neighbours of highly connected agents will tend to have few connections and hence less stable best responses. The differential adaptation which is induced by this difference is what helps coordination. We introduce the concept of 'stable best response' formally and show that there is evidence for such an effect by analyzing the switching behavior in our data. Since coordination propagates along the paths emerging from the players with the most stable best responses, coordination will be such that all connected players choose best responses *bilaterally.* In the treatments with endogenous information highly (little) connected players choose the action corresponding to their preferred Nash equilibrium more (less) often if they are informed about higher order neighbors and hence realize that they have more (less) connections than their neighbours. These effects are strongly significant. Our third result shows on the one hand that very local information is often sufficient to ensure quite successful coordination, but also

[4]We will define this more precisely later. What we mean is that if we take any two players from the network, between which there is a link, then in equilibrium the action choices of these two players will be a Nash equilibrium. We will give an example below of why this is not necessarily true in any Nash equilibrium of the network game.

[5]In the paper we define the concept assuming that agents learn via myopic best responses. We do this because there is evidence in our data (information search) that they do so. The concept is easily extended to cases where agents are forward looking.

that participants are not able to use the additional information in the full information treatments effectively.

These findings have important implications for many economic settings. They can help targeting the right individuals with health policies, such as immunization programs and they can be crucial for the implementation of development policies among others. Our results show that the potential success of such policies may depend crucially on the underlying social structure. Coordination failure may be very persistent in some networks and not occur in others. When a policy maker has the possibility to affect network structure few changes to the network may greatly improve coordination. In addition the regularities on equilibrium selection we identify shed light on potential trade offs between improved coordination, overall efficiency (whether or not the equilibrium played is pareto efficient or not) and potential redistributive concerns (which players get which share of the pie). They also show how the network can be used to induce a certain desired selection of equilibria. This can be very useful knowledge for designing networks in organizations, industrial sectors or even countries.

Previously Cassar (2007) has studied 2×2 coordination games in local, random and small world networks and has found that the tendency of agents to coordinate on efficient outcomes is slightly higher in small world networks than in others. She also found that both higher clustering coefficients and a shorter characteristic path length have a positive impact on agents choosing payoff dominant actions. Since in her study all networks are (ex ante) homogeneous, she cannot study the effect of heterogeneity. Unlike her we are also interested in the effect of the network on coordination per se. In addition we study anti-coordination games rather than coordination games, we control for information and we introduce the new concept of stable best responses to explain our results. Kearns et al (2006) have studied the graph-coloring problem in an experimental setting. In the graph coloring problem the goal for each player is to choose a colour for her node which differs from the colours of all her neighbours. The number of colours made available is the minimum necessary to colour the entire graph without conflict (chromatic number). This problem can be thought of as an Anti-Coordination game whenever there are two colours. They found that networks generated by preferential attachment (typically characterized by higher heterogeneity in connections) make the graph colouring problem *more* difficult than networks based on cyclical structures. Their results are hard to interpret, though, in the light of our questions. The reason is that their networks differ in a number of network characteristics, in the number and structure of "equilibria" as well as in the size of the action set (number of colours available), which makes it impossible to evaluate the effect of heterogeneity. There is also some relation to work My et al. (2001) who study coordination problems with local vs global interaction and find a slight tendency for local interaction to lead to more coordination on risk-dominant outcomes. Corbae and Duffy (2002) study coordination games after endogenous network formation.[6] Cooper et al (1993) study the "battle of sexes game" with pairwise random matching and find that coordination failure is very common, occurring in roughly 60% of the cases.

Theoretical literature on coordination games in networks possibly starts with Blume (1993) and Ellison (1993).[7] Anti-coordination games in networks have been studied theoretically by Bramoullé (2007). He analyzes how many agents will choose a certain strategy in a network as it becomes more advantageous and illustrates that the answer to this question differs across networks.[8] Bramoullé and Kranton (2007) study a public good provision game which has features of Anti-Coordination. This model has been tested by Rosenkrantz and Weitzel (2008). They compare complete networks, star

[6] See also Choi and Lee (2010) who study preplay communication or the survey by Kosfeld (2004).

[7] Other papers include Morris (2000) Alos-Ferrer and Weidenholzer (2008), Jackson and Watts (2002), Goyal and Vega Redondo (2005) and Hojman and Szeidl (2006).

[8] The textbooks by Jackson (2008), Goyal (2008) or Vega Redondo (2007) provide some survey of the literature on games played on graphs.

networks and circle networks. Again, as in Kearns et al (2006), this makes it hard to evaluate the impact of heterogeneity since many network measures differ across these networks. Anti-coordination games in *evolving* networks have been studied theoretically by Bramoullé et al (2004).[9]

The paper is organized as follows. In Section 2 we present the experimental design. In Section 3 we discuss theoretical predictions and in section 4 our main results. Section 6 is dedicated to a discussion of our main results. Some graphs, tables and the experimental instructions can be found in the appendix.

2 Design

The experiment was conducted between May and December 2009 at the BEE-Lab at Maastricht University using the software Z-tree (Fischbacher, 2007). 224 students participated in one of six treatments N-1, N-2, N-3, R-1, R-2 or R-3, i.e. we used a between subject design. In all treatments participants played the symmetric two player game depicted in Table 1 with their neighbors in the network.[10]

	A	B	C	D
A	20, 20	**40, 70**	10, 60	20, 30
B	**70, 40**	10, 10	30, 30	10, 30
C	60, 10	30, 30	10, 10	**30, 40**
D	30, 20	30, 10	**40, 30**	20, 20

Table 1: The Game.

Each player had to choose the *same* action against all her neighbors and payoffs in each round are given by the average payoff obtained in all the games against the neighbors.[11] This is the interesting case to study in networks, since if participants were allowed to choose different actions for their different neighbors, there would essentially be no network effects for coordination. The game was repeatedly played for 20 rounds.

This (Anti-) Coordination game is well suited to study our questions, since it involves a conflict between efficiency and risk-dominance. Equilibria (A,B) and (B,A) are efficient while there is a sense in which strategies C and D are less risky. (C maximizes payoffs if neighbours choose randomly according to a uniform distribution and D is the maxmin choice). Furthermore within each subset $\{A, B\} \times \{A, B\}$ and $\{C, D\} \times \{C, D\}$ each player has an equilibrium which she strictly prefers allowing us to look at issues of distribution. Hence the game allows us to address all aspects of Anti-Coordination on networks: whether or not coordination is effective, efficiency concerns and distributive concerns.

We chose such a game rather than e.g. a simpler 2x2 game for several reasons. First a 2x2 game does not allow us to disentangle the effect of the network structure on efficiency and distribution. Such tensions between equilibria are characteristic of many real life situations, though. Two examples are the following. First consider individual differentiation within an organization. We have to invest in becoming experts in either data-base managing or using text processing software and we want to be experts at something our neighbors are not. We can invest in learning the software by either of two big companies, whose applications are mutually incompatible. As a second example think of a

[9]Their model has been tested by Berninghaus, Ehrhart and Ott (2008).

[10]The game was presented to all participants as if they were row players (see the Instructions in the Appendix).

[11]We chose to pay average payoffs rather than total payoffs to prevent too high inequality in payments among our participants.

joint business project, where we have to decide simultaneously on one of two business strategies - one of which is efficient but also more risky as well as on our level of investment/effort (high/low). An equilibrium is achieved (i) if one of us invests a high and the other a low amount and (ii) if we coordinate on the same business strategy (risk-dominant vs efficient). Such situations can be represented by a game such as ours.

The second reason for choosing a larger 4x4 game is that our results on coordination will be more robust if we obtain them in a richer setting, i.e. in the presence of such conflicts. Finally there were quite some constraints in choosing the exact payoffs which ensure (i) that there is a trade off between efficiency and risk dominance, (ii) that the structure of Nash equilibria across our networks is (roughly) the same. (See Claim 1) and (iii) that standard refinements do not select between equilibria (See Claims 2 and 3).

The three treatments N-1, N-2 and N-3 differed only in the topology of the underlying network. An equilibrium in a network game (in our experiment of 8 players) is obtained whenever all players choose an action that is a best response to whatever their neighbours choose. All our networks are such that many pure strategy equilibria exist. Still, it is clear that coordinating a network of 8 players on any one of the many possible equilibria is difficult, especially since players have conflicting interests. We will describe the Nash equilibria in detail in Section 3. Both the game and the networks were chosen in such a way that (i) there are many pure strategy Nash equilibria and hence coordination is possible but not obvious, (ii) players have conflicting preferences about these Nash equilibria, (iii) Nash equilibria are "qualitatively similar" across all networks (see Section 3) and (iv) networks are comparable in terms of their characteristics except for the degree of heterogeneity. The three networks are depicted in Figures 1-3.

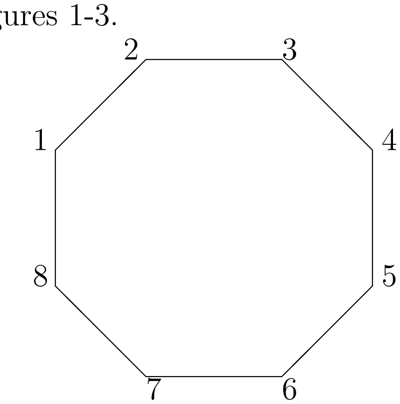

Figure 1: Network 1 (Treatments N-1 and R-1)

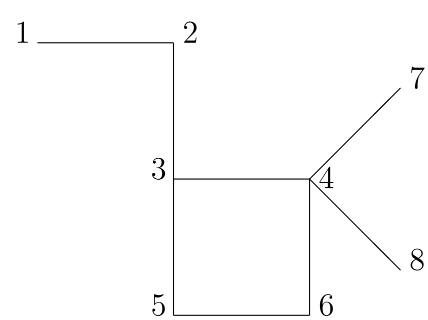

Figure 2: Network 2 (Treatments N-2 and R-2)

The networks were chosen in such a way that starting from the homogeneous network, the circle, heterogeneity in degree is varied, while many other standard network characteristics are constant

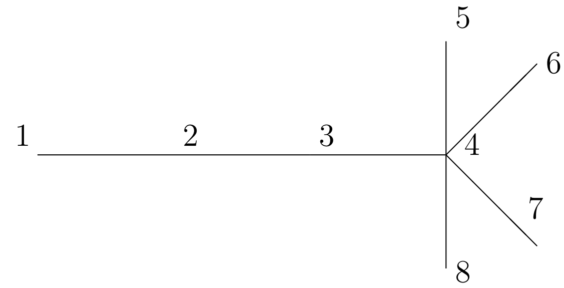

Figure 3: Network 3 (Treatments N-3 and R-3)

across at least two of the three networks (see below). This is why networks 2 and 3 look "non-standard". If we did compare e.g. the circle and the star network, then it would be much harder to trace back the results to particular network characteristic, since those two networks differ across virtually all network characteristics. The labeling of players with numbers (1,2,3,4,5,6,7,8) is arbitrary except for the fact that we have labeled the most connected player in networks 2 and 3 "player 4". Since in the circle all players are the same obviously player labeling does not mean anything in the circle network. Note also that all networks are bi-partite.

Table 2 summarizes some characteristics of these networks. k denotes degree and $\sigma^2(k)$ the variance in degree. The characteristic path length of a network is the average length of the shortest path between any two agents and the betweenness of an agent/node measures the share of shortest paths that pass through her. The average betweenness is the average of the betweenness measures across all nodes. Betweenness is often used to measure the centrality of a node in the network. We consider two measures of heterogeneity in degree. One obvious candidate is certainly the variance in degree $\sigma^2(k) = \sum_{i=1}^{8}(k_i - \overline{k})^2$. This measure of heterogeneity, though, neglects the network structure to the extent that it does not account for correlation between neighbour's degrees. For example it may be possible that variance is very high in one part of a network but very low in another part of the network. Hence as a second measure we use the minimal variance across all (non-trivial) partitions of the network. Consider partitions G of N s.t. (i) each $g \in G$ is connected (i.e. there is a path between any two players in g) and (ii) $|g| \geq 3, \forall g \in G$.[12] Then we can define this second measure as $\sigma^2_{\min} = \min_G \left(\min_{g \in G} \sum_{i \in g}(k_i - \overline{k_i}(g))^2\right)$, where $\overline{k_i}(g) = \sum_{i \in g} \frac{k_i}{|g|}$.

	Network 1	Network 2	Network 3
number of players	8	8	8
number of links	8	8	7
average degree $\overline{k}$	2	2	1.75
diameter	4	4	4
$\sigma^2(k)$	**0**	**8**	**13.5**
$\sigma^2_{\min}$	**0**	**2**	**0.68**
degree assortativity	1	-0.06	-0.14
charact. path length	2.14	2.21	2.21
clustering coeff.	0	0	0
average betweenness	0.21	0.20	0.20
variance betweenness	0	0.06	0.10
existence of cycle (size)	yes (8)	yes (4)	no

Table 2: Network Characteristics.

[12] If a subset of players is not connected, then the variance within this subset is meaningless. The second condition ensures that there are at least three players in a subset, since variance for two players is trivial.

In table 2 we have listed network characteristics referred to in textbooks on networks as the principal characteristics according to which networks can be classified networks.[13] As can be seen from Table 2, the networks are the same or "very similar" in terms of most network characteristics, with the exception of our heterogeneity measures, the variance in degree $\sigma^2(k)$ and the minimal variance across non-trivial partitions $\sigma^2_{\min}$. The networks are also quite different in terms of their degree assortativity (Pearson correlation coefficient of degrees) which is 1 in the circle and lowest (-0.14) in network 3. Of course "very similar" is an imprecise term here, since it is not clear whether e.g. an average betweenness of 0.21 is "similar" to 0.20. Note, though, that no other network characteristics ranks the networks in the same way as $\sigma^2_{\min}$ does. More specifically, network 3 is much more heterogeneous than network 2 which is in turn more heterogeneous than Network 1 according to $\sigma^2(k)$ but in terms of $\sigma^2_{\min}$ Network 2 is more heterogeneous than Network 3 (than 1).

To be able to interpret our results more easily we wanted to control for the amount of information participants use. Hence in treatments N-1 to N-3 we did not simply provide them with all the information about the network, their neighbors actions and payoffs but we asked them at the end of each round which information they would like to obtain. They could choose to have information about the identity, actions and/or payoffs of their first-order, second-order, third-order and/or fourth-order neighbors. Each piece of information had a small cost of 1 ECU (experimental currency unit). Requesting information about the identity of neighbors means learning the label (1,..8) of the player in question and seeing the links between oneself and all the neighbors whose labels have been requested. Requesting such information about the network had a cost of 10 ECU, since this information is permanent and was also permanently displayed to the participants once they had requested it. If a participant which has only first and second order neighbors requested information e.g. about third-order neighbors he was informed that she has no higher order neighbors. In this case she did not have to pay the cost. In treatments R-1, R-2 and R-3 there was no information request stage and all the information was displayed all the time to all participants. Other than that R-1 coincides with N-1, R-2 with N-2 and R-3 with N-3.

	Network 1	Network 2	Network 3
Endogenous Information (N)	800 (5)	1120 (7)	800 (5)
Full Information (R)	480 (3)	480 (3)	480 (3)
Total	1280 (8)	1600 (10)	1280 (8)

Table 3: Treatments and Number of Observations (Independent Observations)

Controlling for the amount of information participants have about the network has two advantages. First it helps us to understand and interpret our data much better. And second it addresses our last research question about the amount of information needed to coordinate.

Finally, the fact that we also wanted to investigate information search was one motivation to use networks of 8 players. In smaller, more stylized, networks most participants have only first-order neighbours and hence it is not possible to investigate information search in a meaningful way. The larger the network, though, the harder it is to achieve coordination for the participants and the noisier our data will become.

At the end of the experiment earnings were converted into Euros according to the exchange rate 1 Euro=75 Ecu. The experiment lasted between 60-90 minutes and participants earned between 7,70 and 16,90 Euros with an average of 11,40 Euro.

[13] See the books by Goyal (2007), Jackson (2008) or Vega-Redondo (2007).

3 Theoretical Predictions

In this section we would like to discuss the predictions of a number of standard learning models in some more detail. Our fundamental view is that participants in our experiment will *learn* to play these games in the laboratory rather than choosing repeated game strategies at the beginning of the experiment. Support for this view is given by the fact that play in the first 10 periods of the game in *any* network is far from play of *any* Nash equilibrium, be it of the repeated game or of the one-shot game. We are then interested in (i) whether play converges to one of the Nash equilibria of the underlying network game and (ii) if so to which Nash equilibrium it does.

We assume that participants learn about successful actions and use the information available in order to do so. In the R-treatments (full information) there may actually be more information available than participants use in their learning rules. In the N-treatments (endogenous information) participants can choose which information they want to consider. We are aware that in the N-treatments one could also analyze the Nash equilibria of a larger (incomplete information) game which includes information search in the strategy set. Doing this would not affect the equilibrium action distribution as described in Claim 1. We do not perform such analysis here since it is orthogonal to our research questions. In Section 4.5 we present empirical results on information search.[14]

Before we mentioned that theoretically it is not required for an equilibrium of the network game that two connected players choose best response *bilaterally* or in other words that all links are in Nash equilibrium. This means that if we take any two players from the network, between which there is a link, then in equilibrium the action choices of these two players need *not* be mutual best responses. This may be the case even though the network is in equilibrium. The reason is that in Nash equilibrium each player has to choose a best response to the *distribution* of her neighbours choices. Hence for example a player with two neighbors one of which choosing C and one D will choose D as a best response in which case *bilaterally* she is not choosing a best response to the neighbor choosing D. Hence this link will not be in *bilateral* Nash equilibrium, while the network overall may well be in equilibrium.

We next describe the strict Nash equilibria of the one-shot network game.[15]

Claim 1 (Strict Nash equilibria) *Irrespective of the network: (i) there are four strict Nash equilibria-profiles where all links are in bilateral Nash equilibrium, and (ii) there are no additional strict Nash equilibria where all players choose the "efficient" actions, A or B. In Network 1 there are eight more strict Nash equilibria and in Networks 2 and 3 six more strict Nash equilibria where some players choose C and some D.*

Proof. Appendix. ■

The claim shows that all of our networks have many strict Nash equilibria and that there are no systematic differences in the structure of Nash equilibria (in the sense of Claim 1). We have focused on strict Nash equilibria in Claim 1 because those are equivalent to the asymptotically stable points under the standard Replicator Dynamic. The standard replicator dynamic in turn approximates many learning models, such as e.g. reinforcement learning or some versions of fictitious play.[16]

Claim 2 (Asymptotically Stable Nash equilibria) *Irrespective of the network all the strict Nash equilibria mentioned in Claim 1 are asymptotically stable to the evolutionary Replicator Dynamics. There are no other asymptotically stable points.*

[14] This analysis also seems to confirm the view that participants most likely simply select the information they need given the learning rule they employ (rather than e.g. being strategic about it or learning which information is useful).

[15] All strict Nash equilibria are explicitly described in the Appendix.

[16] See e.g. Boergers and Sarin (1996) or Hopkins (2002).

Proof. This is a standard result which can be found in e.g in Chapter 5 of the textbook by Weibull (1995). ■

A selection criterion often employed in network games is stochastic stability. To employ this criterion let us assume that our agents are myopic best response learners, i.e. they best respond in each period to whatever their neighbours did in the previous period. Such a model of learning is largely consistent with the information search we observe, even though few agents also look for information about second or higher order neighbors which suggests that not all agents act entirely myopically (see Section 4.5).[17] Myopic best responses, though, is the standard model employed in such contexts.[18]

Define the state at time $t = 1, ..20$ to be the vector of actions of all players at time t: $a^t = (a_1^t, ..a_8^t)$. The myopic best response process then defines a Markov chain, whose absorbing states are given by the Nash equilibria of the (one-shot) network game. Stochastic Stability then, loosely speaking, asks how likely it is - given that players make errors with small probability ϵ - to move from one absorbing state to the basin of attraction of another absorbing state. We follow Jackson and Watts (2002) and Goyal and Vega Redondo (2005) in assuming that errors are uniform. The absorbing states which are most easily reached from other states and most difficult to exit are called stochastically stable. We can show the following Claim.

Claim 3 (Stochastic Stability) *In all networks the absorbing states induced by any of the Nash equilibria identified in Claim 1 and Table A-1 are stochastically stable under myopic best responses.*

The proof of this claim as well as a more technical exposition of the concept of stochastic stability can be found in the Appendix. Stochastic stability does not select between our candidate equilibria.

We hope to have convinced the reader that standard learning models do not select between the Nash equilibria of our network games.[19] Obviously it is never possible to be exhaustive, but we believe we have covered the standard models used in these contexts. This section shows also that there is a large multiplicity of strict Nash equilibria, which means that theory makes only weak predictions about the outcome of the network games. The actual behavior we observe in the experiment, though, eliminates this multiplicity (whenever play converges). As we will see below, out of the 10-12 theoretically possible (one-shot) Nash equilibria we only observe only one or two empirically.

4 Results

We will now proceed to present our experimental results and we will give an explanation for them in Section 5.

4.1 Coordination

Let us start with our results on coordination. Since it is easily possible in a network of eight players that one or two players make mistakes, we consider three measures for coordination to Nash equilibrium in order to account for this possibility. Our measures are the percentage of cases (across the last 5 rounds) in which the entire network is in Nash equilibrium and the percentage of cases

[17] In a companion paper, Kovarik, Mengel and Romero (2010), we give many more details about learning.

[18] See e.g. Bramoulle et al (2004), Jackson and Watts (2002) or Goyal and Vega Redondo (2005).

[19] Bramoulle (2007) shows that stochastic stability assuming logit (rather than uniform) errors may select in $2x2$ Anti-Coordination games. We come back to this result when we discuss our empirical results.

(also across the last 5 rounds) in which the entire network would be in Nash equilibrium if one (two) players changed their action.

Table 4 shows the percentage of successful coordination in the last 5 rounds. The most interesting line in the table is certainly the first row (Nash equilibrium). Nash equilibrium-1 and Nash equilibrium-2 are harder to interpret, but we think they are useful to understand how robust our qualitative results are (in terms of the comparison of the three networks) to mistakes by one or two players. After all it is not that unlikely that one out of eight participants in a network plays a bit randomly. The results are presented in the following tables.[20]

To help interpreting the tables also note that if players did choose actions uniformly at random the probability to coordinate on a Nash equilibrium would be 0.00018 for network 1 and 0.00015 for networks 2 and 3. For the measures Nash equilibrium except 1 player these probabilities are 0.0059 and 0.0049 respectively and for Nash equilibrium except 2 players they are 0.16 and 0.14 respectively. It should also be noted that once a network has reached a Nash equilibrium after round 15, they stay at this same Nash equilibrium for all remaining periods more than 90% of the time.

	Network 1	Network 2	Network 3
Nash equilibrium	0.12 $(0.12, 0.13)$	0.44 $(0.43, 0.46)$	0.20 $(0.24, 0.13)$
Nash equilibrium except 1 player	0.22 $(0.28, 0.13)$	0.70 $(0.71, 0.67)$	0.45 $(0.40, 0.53)$
Nash equilibrium except 2 players	0.37 $(0.36, 0.40)$	0.89 $(0.89, 0.87)$	0.72 $(0.72, 0.73)$

Table 4: Percentage of successful coordination in last 5 rounds. In brackets below are separate averages for the treatments with endogenous (N) and full info (R). All differences are pairwise highly significant across networks using a two-sided Mann-Whitney test with each matching group as unit of independent observation ($p < 0.0001$).

The table illustrates that network structure dramatically impacts Coordination. In fact the rate of successful coordination is twice as high in Network 3 than in Network 1 and almost three times higher in Network 2 compared to Network 1. In order to test for differences in the distribution of rates of succesful coordination across types of networks and/or treatments we use a two-sided rank sum (Mann-Whitney) test with each network (matching group) as a unit of independent observation. All differences are pairwise highly significant across networks (Mann-Whitney, $p < 0.0001$). The difference between "Nash equilibrium except 1 player" in N-1 and R-1 is marginally significant ($p = 0.0864$). Other differences between N-1 and R-1 or differences between N-2 and R-2 as well as between N-3 and R-3 are not statistically significant. (Mann-Whitney, $p > 0.1987$).

Result 1 *Coordination to Nash equilibrium is significantly better in Network 2 compared to Network 3 and significantly better in Network 3 compared to Network 1.*

One conclusion from this data is that our more "local" measure of the variance $\sigma^2_{\min}$ describes the relevant heterogeneity better than the simple variance $\sigma^2(k)$. We will present a more detailed explanation of Result 1 in Section 5.1. Note that network 2 can be obtained from Network 3 by rewiring the link $5 \leftrightarrow 4$ to $5 \leftrightarrow 3$ and adding the link $5 \leftrightarrow 6$. An implication of Result 1 is hence

[20]Note also that deviations of a single player from Nash equilibrium need not always represent errors. In one network (in N-2) we observed for example that starting from the Nash equilibrium where player 4 chooses D (and all links are in Nash equilibrium), player 4 sometimes started to switch to B for some rounds, possibly to induce a transition to the efficient Nash equilibrium. No other player reacted though for three rounds and player 4 switched back to D.

that a designer facing network 3 may even want to add a link to improve coordination. (Of course our result suggests that adding the link $5 \leftrightarrow 6$ will not be necessary but rewiring $5 \leftrightarrow 4$ to $5 \leftrightarrow 3$ will be sufficient, since this alone increases local heterogeneity $\sigma^2_{\min}$.)

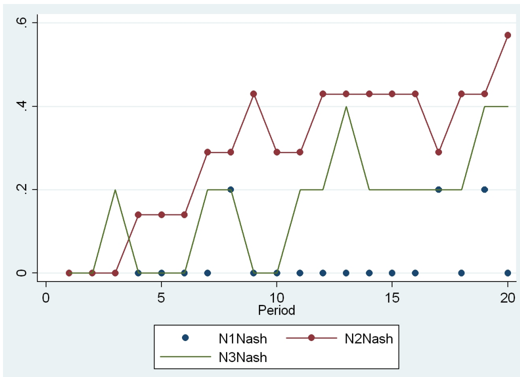

(a) Nash Equilibrium

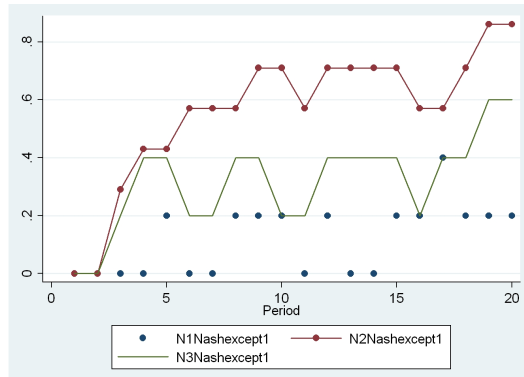

(b) Nash equilibrium except 1 player

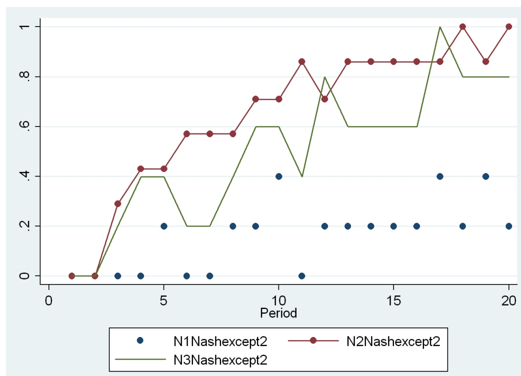

(c) Nash equilibrium except 2 players

Figure 4: Coordination to Nash equilibrium over time.

Figures 4a-4c show the percentage of coordination over time. The figures illustrate clearly the difference in coordination rates across the three networks. They also show that the coordination failure in Network 1 seems persistent, i.e. not simply a matter of the speed of convergence.

Note also that - if players did choose uniform randomly either of two actions - the probability for a given link (rather than a network) to be coordinated on a Nash equilibrium would be 50%. A binomial test rejects the hypothesis that 50% of links are coordinated in the last 5 rounds for N-2 and N-3 ($p < 0.000001$) but not for N-1 ($p = 0.137367$). Hence coordination in the homogeneous network is not much better than if all players did choose randomly among two actions. (As we will see in the following subsection 90% of choices in Network 1 are only either C or D).

4.2 Efficiency

Table 5 lists action choices across the last 10 periods of the experiment for the different treatments. The risk-dominant actions C and D are chosen in about 90% of all rounds in all treatments. The differing percentages of C and D choices in the homogeneous network 1 reflect the extent of coordination failure illustrated in the previous subsection. A Mann-Whitney test rejects the hypothesis that the distribution of efficient vs non efficient choices are equal in N-3 and N-1 (N-2) ($p = 0.0330$, $p = 0.0043$) only if we use each individual as independent unit of observation. If we use networks as independent unit of observation there are no significant treatment differences. There are no significant differences between N-1 and N-2 (Mann-Whitney, $p = 0.5964$) even if we use individuals as independent unit of observation.

	Network 1	Network 2	Network 3
A	0.04 $(0.04, 0.05)$	0.03 $(0.04, 0.00)$	0.08 $(0.09, 0.06)$
B	0.08 $(0.06, 0.13)$	0.05 $(0.05, 0.03)$	0.07 $(0.08, 0.06)$
C	0.36 $(0.37, 0.33)$	0.45 $(0.45, 0.46)$	0.36 $(0.35, 0.37)$
D	0.51 $(0.53, 0.49)$	0.47 $(0.46, 0.51)$	0.49 $(0.48, 0.51)$

Table 5: Distribution of Choices last 10 Periods. In brackets below are separate averages for the treatments with endogenous and full info.

Result 2 *In all treatments agents (try to) coordinate on risk-dominant outcomes.*

We also observe that *if* coordination occurs it is always on a Nash equilibrium where all players choose risk-dominant actions. Hence efficient choices are clearly "off equilibrium" phenomena in our experiment. Also note that the fact that risk dominant rather than efficient equilibria are observed depends clearly on the payoff matrix chosen and is hence to be enjoyed with care. Still we can make even much stronger statements regarding selection as the next subsections illustrate.

4.3 Bilateral Coordination

As mentioned above theoretically not all links need to be in Nash equilibrium in order to have a Nash equilibrium of the network game. (See the discussion before Claim 1). Empirically, though, we observe strong regularities. In fact among those profiles which were a Nash equilibrium in the last 5 rounds 100% were Nash equilibria in which all links are coordinated in all networks N-1, N-2, N-3, R-2 and R-3. For the measure Nash equilibrium-1 the percentage is 100% (100%,80%,100%, 100%) for N-1 (N-2,N-3, R-2 and R-3) and for Nash equilibrium-2 it is 100% (100%, 84%,100%,100%). In R-1 we have only very few instances of convergence. In this treatment 0% of Nash equilibrium (Nash equilibrium-1) were such that all links are coordinated and 67% of Nash equilibrium-2 were such that all links are coordinated.

Result 3 *In all treatments N-1, N-2, N-3, R-2 and R-3, if there is convergence to Nash equilibrium, then participants coordinate always on a Nash equilibrium where all links are coordinated.*

This result suggests that coordination to such equilibria is easier. We will give an intuitive explanation for this fact in Section 5. Interestingly Bramoulle (2007) has shown in a related context that with logit errors stochastic stability will select for equilibria where all links are (bilaterally) coordinated. However, his results are not directly applicable to our games. As we have shown above with uniform errors stochastic stability does not select in our context. Since in each network there are exactly four equilibria where all links are coordinated and since efficient equilibria do not occur, we have narrowed down the multiplicity of equilibrium outcomes that arises theoretically to only two Nash equilibria arising empirically. In fact, as the following subsection shows, we can make even more precise predictions in heterogeneous networks.

4.4 Distribution

In this subsection we would like to know whether more connected agents can impose their "preferred" Nash equilibrium more often than others or more generally how the topology of the network affects potential redistributive concerns. With 'preferred' Nash equilibrium here we mean equilibria where the agent in question chooses one of the hawkish actions B or D. Since successful Coordination always involved only C and D choices, essentially this boils down to asking whether an agent chose D in Nash equilibrium. To analyze this question we had a closer look at which Nash equilibrium players coordinated on in the heterogeneous networks. Conditional on being coordinated on a Nash equilibrium at all (or being coordinated except for one or two players) we then have the following share of coordination on player 4's preferred Nash equilibrium.

	Network 2	Network 3
Nash equilibrium	0.87 $(0.82, 1.00)$	0.82 $(0.71, 1.00)$
Nash equilibrium except 1 player	0.85 $(0.78, 1.00)$	0.80 $(0.80, 0.78)$
Nash equilibrium except 2 players	0.83 $(0.78, 0.93)$	0.87 $(0.90, 0.82)$

Table 6: Nash equilibrium preferred by player with highest degree. In brackets below are separate averages for the treatments with endogenous and full info.

Pairwise treatment differences (between any two treatments) are not significant according to a two-sided Mann-Whitney test using matching groups as independent unit of observation.[21] This result shows that the most connected players can impose their preferred Nash equilibrium most of the time and in the case of full information even all the time. Together with the previous results this allows us to make sharp predictions about the outcome of the game in spite of the large multiplicity of Nash equilibria. This is most striking in the case of full information (R-2) where one particular Nash equilibrium out of 12 possible equilibria is always observed. In N-2 and N-3 two Nash equilibria are observed. In both risk dominant actions are chosen and all links are coordinated and in more than 80% of the time the Nash equilibrium which player 4 prefers is observed.

Result 4 *In the heterogeneous networks players coordinate most of the time on a Nash equilibrium preferred by the most connected player.*

We will provide an explanation of this fact in Section 5. Note that this result is obtained in a context where (i) networks are small and (ii) where there is one player having most connections. In larger networks it may be necessary to consider the characteristics of sets of players rather than individual players. This selection of Nash equilibria we observe is also reflected in the payoffs.

4.5 Payoffs

The following table summarizes the average payoffs of participants in Network 2 conditional on their position in the network 1,..8.

In network 2, players 2, 4 and 5 make the highest payoffs. The reason that players 2 and 5 make higher payoffs is of course that their preferred Nash equilibrium is played more often. In addition -

[21] If we use individuals rather than networks as unit of independent observation - hence increasing the likelihood to find an effect - we find a marginally significant difference between N-2 and R-2 (Mann-Whitney, $p = 0.0725$) for Nash equilibrium.

	Network 2	Network 3
Player 1	659	684
Treatments (N,R)	(647,690)	(662,720)
Player 2	756	734
Treatments (N,R)	(734, 810)	(694, 810)
Player 3	674	667
Treatments (N,R)	(673, 677)	(650, 696)
Player 4	718	706
Treatments (N,R)	(694, 774)	(694, 725)
Player 5	733	631
Treatments (N,R)	(637, 696)	(626, 640)
Player 6	654	631
Treatments (N,R)	(637, 696)	(626, 640)
Player 7	646	648
Treatments (N,R)	(622, 703)	(658, 633)
Player 8	602	636
Treatments (N,R)	(585, 643)	(626, 653)

Table 7: Average Payoffs per Player position and Network (treatment).

unlike player 4 - they do not have the burden to "ensure" coordination, as we will see in Section 5. In N-3 and R-3 it is Players 2 and 4 that make the highest profits.

4.6 Information

To conclude this section let us have a look at what information agents asked for in treatments N-1 to N-3. Since there were no significant differences in information search across those treatments we report aggregate data in Figure 5.[22] Figure 5(a) shows the percentage of participants per round that requested info about their first- (second-, third-, fourth-) order neighbours conditional on having a first- (second-, third-, fourth-) order neighbour. Figure 5(b) shows the percentage of participants per round that request information about their first- (second-, third-, fourth-) order neighbours 'actions again conditional on having a first- (second-, third-, fourth-) order neighbour.

The vast majority of agents ($> 90\%$) request information about who their first-order neighbours are and most also want to know what their first-order neighbours chose in at least half of the rounds. Out of those that do not check their neighbours actions still quite some check their own payoff and less than 10% of agents do not request any information at all. Around 45% of agents request information about their second-order neighbours in addition. There is a slightly decreasing trend in the frequency with which participants check their neighbor's action choices, which is consistent with the fact that there is much less variation in choices over the last ten period compared to the first ten periods and hence less can be learned from this information. The information search we observe is consistent with most of our participants being myopic best response learners.[23]

Overall we can see that participants in the N-treatments rely mostly on local information about their first or second order neighbours. The fact that very much the same behavior is observed in the R-treatments as in the N-treatments provides strong evidence that participants are not able to use higher order information effectively. On the other hand we see that very local information (about

[22]Separate graphs per treatment are available upon request.

[23]In Kovarik, Mengel and Romero (2010) we analyze learning rules in more detail.

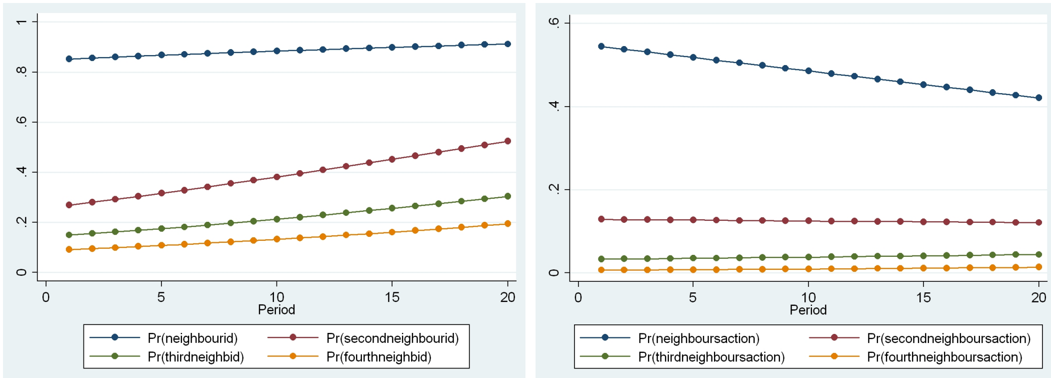

(a) Information about Neighbours

(b) Information about Action Choices

Figure 5: Information requested over time. Note that (by design) the highest line is first-order, then 2nd order, third-order and fourth-order neighbors.

mostly first and sometimes second order neighbours) is sufficient to guarantee successful coordination.

Result 5 *Local (endogenous) information is as effective as full information in ensuring successful coordination.*

The fact that we are able to observe which information participants used, enables us to say something more about the underlying reasons for our results. This will be the objective of the next section.

5 Explanation and Discussion

In this section we will present an intuitive explanation for why we observe much better coordination in network 2 compared to Network 3 and Network 1 and why we observe such a strong equilibrium selection.

5.1 The Notion of Stable Best Responses

We start by defining the notion of stable best responses. We will assume that agents are myopic, because this is most consistent with the information search we observe. The definition could be extended though to allow for forward looking agents as well.

Remember that we defined the state at time $t = 1, ..20$ to be the vector of actions of all players at time t: $a^t = (a_1^t, ..a_8^t)$. Let us define the distance $d(a, a')$ between any two states a and a' to be the number of differing elements between them. Denote by $BR_i(a)$ the myopic best response of player i at state a and define $x_i(a) = \min_{\{a': BR_i(a) \neq BR_i(a')\}} d(a, a')$. In words $x_i(a)$ is the minimal number of changes needed starting from state a such that player i's best response changes. If there is no a' such that $BR_i(a) \neq BR_i(a')$ then we set $x_i(a) = \infty$. (This happens for example if a player has a dominant strategy and is not relevant for our game.)

Definition We say that BR_i is more stable than BR_j if $\forall a : x_i(a) \geq x_j(a)$ with strict inequality for at least one a.

The concept of stable best responses simply captures the intuitive idea that an agent with a more stable best response is less likely to react to changes in the behaviour of others. Now heterogeneity in connections typically induces heterogeneity in the stability of best responses. An agent i with more connections (a higher degree k_i) will typically have more stable best responses, since changes in any given neighbour's action have a smaller effect on payoffs of such highly connected agents. (In our game for example players with only one neighbour have $x_j(a) = 1, \forall a$). Hence the concept is connected in a very natural way to networks. More precisely we can state the following

Claim 4 *(i) If agent i has a more stable best response than agent j, then $k_i > k_j$. (ii) In Network 2 players 3 and 4 have the most stable best responses and in Network 3 player 4 has the most stable best response.*

Two remarks are at order. First note that one could strengthen the concept by requiring a to be an absorbing state. We choose to work with the weaker concept here, but this could be a straightforward extension. Secondly note that the concept of stable best responses is very different than many concepts referring to the stability of equilibria (such as asymptotic stability, evolutionary stability, stochastic stability etc..). While the latter concepts ask whether small perturbations can lead away from equilibrium, stable best responses compares different players and asks how variable best response correspondences are to changes in other player's choices.

Note also that having a "stable best response" *per se* does not say much about how many players one is able to influence in the network. For example it may be that I need many players to change their action in order for me to change mine, but my neighbors need even more players to change their action. In this case I will have relatively little influence. Hence it is important to note that the concept makes most sense when viewed *relatively*, i.e. when asking the question which of two players has the more stable best response. Still the concept is probably best suited to explain the behavior of myopic agents, i.e. agents that in each period best respond to the behavior of their neighbors in the previous period and do not entertain the possibility that they are able to influence the best response of their neighbor at $t+1$ by changing their action at time t. For such forward looking agents one may define the concept of "Influence" as follows.

Denote by $M_i(a)$ the set of states that can be reached from a via only one change in player i's action. Then we can define $y_i(a)$ as follows: $y_i(a) = \max_{a' \in M_i(a)} \{j | BR_j(a') \neq BR_j(a)\}$. Hence $y_i(a)$ measures how many agents player i can "influence" by changing his action. This is a measure of influence for one-period forward looking agents. If agents are two or more periods forward looking the concept could be extended in a straightforward manner.

Definition We say that player i has more influence than player j in the network if $\forall a : y_i(a) \geq y_j(a)$ with strict inequality for at least one a.

A player's influence at a given state depends on her degree (how many neighbors she has) and on the stability of her neighbors best responses (how easily their best responses are affected). Given this the following claim should come as no surprise.

Claim 5 *In both Network 2 and Network 3 player 4 has the most influence.*

Proof. Appendix. ■

The intuition is quite simple. In particular in network 3 where player 4 has four neighbors whose best response changes every time player 4 switches it is clear that no other player can have higher influence. This intuition makes clear as well why it is interesting to look also at "higher order" influence. The influence player 4 has in Network 3 over players 5,6,7,8 does not translate in any influence on the rest of the network, since the change in best response induced for those players will not propagate through the network. Note also, that while in our setting it is player 4 who has both most influence and the most stable best response in general the player with the most stable best response can be different from the player with the highest influence.

5.2 Coordination

Our conjecture is that coordination is better in Network 2 than in the other networks 1 and 3 because there is high enough variance in the stability of best responses in this network. The most connected player (player 4) will be crucial for coordination, since she has a more stable best response. If the variance in degree is high, then in addition the neighbors of the highly connected agents will have a smaller degree and hence are more likely to adapt their action to the highly connected agents choice. This differential reaction to each others choices is what avoids or breaks coordination cycles. This also means that player 4 will have high influence on her neighbours. If local variance $\sigma_{\min}$ is high then the influence of the most connected player is able to spread effectively through the network. If variance is low in some part of the network (as it is the case in Network 3), then the higher-order influence of the most connected player will be limited and coordination may fail.

To see whether this explanation has empirical support we will study two empirical measures. First we will study the number of times each player switches their action in the experiment (Switching Rate) which is a measure of the stability of best responses. Second we study the empirical frequency with which switches by a player trigger switches by any given of his neighbours. The latter is a measure of Influence.

Switching Rates

Table 8 shows the switching rates across our networks. Switching rates are higher in Network 1 compared to Network 2 and 3 reflecting worse convergence. Note also that - in spite of there being very little successful coordination in Network 1 - switching rates decline substantially in the second half of the experiment even in Network 1. This indicates that play does stabilize in Network 1, even though, as we have seen in Section 4, there is almost 90 percent coordination failure in this network. It can also be seen that player 4 in both networks 2 and 3 switches less often than all other players on average. Note that the fact that a player has a more stable best response does not imply that she switches *much* less often than others. If convergence is perfect than the neighbors of player 4 should switch only once more than player 4 in expectation.

Switching Rates	Network 1	Network 2	Network 3
overall	0.47	0.41	0.42
periods 11-20	0.34	0.29	0.31
player 4 (11-20)	–	0.24	0.28
other players (11-20)	–	0.30	0.32

Table 8: Stability of best responses: Switching rates across networks.

Influence

Table 10 shows a measure of "Influence". For every period t in which player i switched her action, we computed the frequency with which any given first-order neighbour switches her action in $t+1$. We interpret this number as a measure of the "average influence" player i has over any one of his neighbors. We then also measure how influence propagates, i.e. the frequency of the joint event that a first-order neighbour switches at $t+1$ and one of her neighbours (i.e. a second-order neighbour of i) at $t+2$ conditional on i switching at t etc... Of course players may switch their action for any other reason not connected to player i's switching. Such random switches should become less frequent over time, though. This is why we present evidence from the last ten periods 11-20. Note also that the table does not capture delayed reactions to any players switches.

Influence	$N-2$ 1st	2nd	3rd	4th	$N-3$ 1st	2nd	3rd	4th
Player 1	0.17	0.04	0.00	0.00	0.16	0.02	0.00	0.00
Player 2	0.36	0.06	0.04	-	0.17	0.01	0.00	-
Player 3	0.29	0.06	-	-	0.12	0.03	-	-
Player 4	**0.45**	**0.14**	**0.15**	-	**0.43**	**0.08**	**0.04**	-
Player 5	0.17	0.00	0.00	-	0.10	0.03	0.00	0.00
Player 6	0.33	0.06	0.02	0.00	0.00	0.00	0.00	0.00
Player 7	0.07	0.00	0.00	0.00	0.12	0.04	0.00	0.00
Player 8	0.02	0.00	0.00	0.00	0.17	0.04	0.00	0.00

Table 9: Influence per neighbor: Empirical frequency with which any given first-order neighbor of player i switches her action at $t+1$ after i has switched at t. Empirical frequency of the joint event that any given 2nd-order neighbor j switches at $t+2$ and that the 1st order neighbor linking i and j switches at $t+1$ conditional on i switching at t etc.. The first four columns are N-2. Columns 5-8 are N-3. Periods 11-20.

The table shows that - consistently with the theory - player 4 seems to have a higher influence on any given neighbor than any other player the network. This is true for both networks 2 and 3. In both networks a switch of action by player 4 induces a switch of any given of her neighbours in 40-45 percent of the cases. In both networks players who are only connected to player 4 exert very little influence.

A fundamental difference between the two networks appears, though in the influence players 2 and 3 have on their neighbours. This influence is much higher in Network 2 compared to Network 3. The reason is that in the subnetwork of players $\{1,2,3\}$ players 2 and 3 have the same number of connections, while in Network 2 there is clear hierarchy. (Player 3 having more connections than player 2 who has more than player 1). But this lower influence of player 3 over his neighbour 2 and the lower influence of player 2 over player 1 in Network 3 immediately translates into less *higher-order influence* by player 4. This may make it harder for the network to coordinate.

If this explanation is correct, then coordination should be much worse in the part of the network 3 which has a smaller variance. In the following table we list the percentage of time in the last 5 rounds in which the "homogeneous" subset $\{1,2,3\}$ and the "heterogeneous subset" $\{4,5,6,7,8\}$ in N-3 are in Nash equilibrium. Note that since the first set is smaller we should expect a higher percentage of coordination there. Moreover if we want to use the measure Nash equilibrium-1, the analysis becomes trivial in a line of three players. Hence we also report coordination in the subset $\{1,2,3,4\}$.

N-3	Nash equilibrium	Nash equilibrium except 1 player	Nash equilibrium except 2 players
$\{1,2,3\}$	0.24	–	–
$\{1,2,3,4\}$	0.24	0.44	
$\{4,5,6,7,8\}$	0.40	0.76	0.92

Table 10: Coordination in subgraphs N-3

Table 11 illustrates the origin of the coordination failure in network 3. While in the subnetwork $\{4,5,6,7,8\}$ the rates of successful coordination are as high (or even higher) as in Network 2, coordination is much worse in the subnetwork $\{1,2,3\}$, even though it consists of 3 players only. The table also illustrates that the subnetwork $\{1, 2, 3\}$ is almost entirely responsible for the coordination failure in network 3 and fully responsible for the difference between Network 3 and Network 2. This gives support to our conjecture.

Our explanation can also explain why in the Nash equilibria which are empirically observed almost always all links are in Nash equilibrium. Successful coordination is characterized by some players (such as players 4 in Network 2 and Network 3) emerging as leaders and others adjusting to their behavior subsequently along the paths leading away from these agents.

Summary *Successful coordination is driven by the most connected player displaying more stable best responses, while her less connected neighbours adapt quickly to her choices. The most connected player emerges as a "leader" in coordination. Only if local variance is high can the influence of the most connected player propagate effectively through the network. This implies (i) that coordination is best if the minimal variance across connected subnetworks is high and (ii) that Nash equilibrium will be such that all links are coordinated*

5.3 Information Search and "Hawk" Choices

A second observation we made is that the most connected agents are almost always able to impose their preferred Nash equilibrium (See results 4 and 5). In this subsection we will evaluate two candidate explanations. The first is that action D is simply perceived as the better response for a more connected player, since it might be consistent with player 4 facing more or less risk than other players. The second is that the most connected players realize their importance in the coordination process and hence impose their preferred Nash equilibrium.

It is the second explanation which is supported by our data. The minimization of risk hypothesis seems not to work well. Let's have a look at the difference (in percentage points) between action choices among players 4 and other players in the early periods 1-10.

"Player4-others"	**N-2**	**N-3**	**R-2**	**R-3**
A	-0.02	-0.05	0.00	-0.04
B	**0.13**	**0.30**	**0.16**	**0.06**
C	-0.13	-0.09	-0.20	-0.13
D	0.02	-0.16	0.04	0.11

Table 11: Difference in Choice between Player 4 and others in early periods (1-10)

In early periods players 4 do not choose D much more often than others, but they choose B much more often than other players. Clearly choosing B is not consistent with an hypothesis that players 4 face more risk, since C obviously dominates B in terms of risk. One may sustain hence that players 4 actually face less risk than other players and hence are more willing to choose risky actions. If this were the case, then it should also be true (to some extent) for player 3 in networks 2 and R-2 at least in early periods, since they have more neighbours than other players (except for player 4). We compare hence action choices of player 3 in N-2 and R-2 to those of *other neighbours of player 4*.

It can be seen that player 3, having more neighbours, does choose C more often in early periods than other neighbours of player 4 which have a lower degree. This is inconsistent with the explanation that highly connected players choose aggressive actions more often because they face less risk.

On the other hand, if the second explanation is right, then it seems that highly connected agents should play D with higher probability if they do have information about second (or higher) - order neighbors and hence realize that they have a higher degree than others. Without such information C seems to be the more likely choice (especially in early periods) since it is risk-dominant and hence

"Player3-other Neighb."	**N-2**	**R-2**
A	-0.04	0.00
B	-0.04	-0.03
C	**0.23**	**0.11**
D	-0.15	-0.08

Table 12: Difference in Choice between Player 3 and other neighbours of player 4 in periods 1-10

one may expect coordination to an equilibrium where the most connected player chooses C. We conducted a Spearman test to find out the correlation between the "hawkish" B- and D-choices and information search (about second neighbors). The results are presented in the following table

Player	**N-2**	**N-3**
4	$\rho = \mathbf{0.2470}^{***}$	$\rho = \mathbf{0.1701}^{*}$
3	$\rho = 0.0428$	$\rho = -0.0153$
5	$-$	$\rho = -0.2200^{**}$
6	$\rho = -0.4145^{***}$	$\rho = 0.0116$
7	$\rho = -0.0749^{**}$	$\rho = -0.4203^{***}$
8	$\rho = -0.2256^{***}$	$\rho = -0.1670^{*}$

Table 13: Correlation betw. "Hawk Choices" and Information Search by network position.

Table 14 clearly illustrates that those player 4 who are informed about second-order neighbours and hence about the fact that they have a relatively higher degree than others choose D significantly more often than those who are not. On the other hand player 4's immediate neighbours (with the exception of Player 3) choose D significantly less often if they are informed about second-order neighbours and hence about their relatively lower degree compared to player 4. For player 3 we do not find a significant effect which is probably due to the fact that player 3 herself has a relatively high degree in Network 2. This result can clearly not be explained by differences in risk.

This result also means that agents are very well aware of the leading role highly connected agents play in the coordination process. Loosely speaking, the more stable best responses they have the longer they are able to hold out in a 'Hawk fight' along any given link. The highly connected players seem to use this knowledge to impose their preferred Nash equilibrium while their neighbours are willing to "give in" more easily if they are aware of the special role of player 4.[24]

Summary *If the most connected players are aware of their importance for the coordination process they tend to choose "hawkish" actions B and D significantly more often. They are able to impose their preferred Nash equilibria because of their leading role in the coordination process. Differences in risk cannot explain this phenomenon.*

5.4 Information Search and Payoffs

Let us see whether those agents that did not check any information (in most of the rounds, i.e. more than 10 times) received lower or higher average payoffs than those that did in the N-treatments. This can be an indicator of whether it was worthwhile to look for information. Non-surprisingly we

[24]We also checked whether the amount of coordination itself in the last 10 periods is correlated with whether agents have information about their second order neighbours and find no significant effects, which is consistent with the explanation given above. This also means that the correlations identified above are not simply a by-product of better coordination in these cases.

find that in N-2 this is not the case (two sided Mann-Whitney, $p = 0.5378$), since in this network convergence is very good. If we focus uniquely on the first five rounds of play (before convergence occurred) there is a negative effect of not searching information, though, which is significant at the 5% level. Our evidence from N-1 and N-3 illustrates that if convergence is not that good, then information acquisition is worthwhile. In N-1, where convergence to Nash equilibrium is very bad, we find that those agents which check information very rarely make substantially smaller profits (with an average of 27) than those which check information regularly, which average 34.7. (Mann-Whitney, $p = 0.0018$). In N-3 we find a similar effect. Not searching information yields significantly lower payoffs especially across the first ten rounds (Mann-Whitney, $p = 0.0006$). Overall hence we find that information search was worthwhile for participants but that local information proved enough to ensure successful coordination.

5.5 Questionnaire Data

In this last subsection let's have a look at our questionnaire data, where we elicited risk attitudes. Those data can serve as a check of whether our treatments were balanced in terms of risk attitudes and whether we interpret the riskiness of actions in the same way participants do. We asked participants three questions regarding their risk attitudes in a post-experimental questionnaire:

1. We toss a coin in the air. Choose one of the following options: (i) Receive 1000€, independently of the outcome of the toss coin. (ii) Receive 2000 € if head turns up or 0€ if tail turns up.
2. Choose between the two following options the one you prefer: (i) Play a lottery ticket, in which you win 45€ with probability 80%, or 0€ with probability 20%. (ii) Receive 30€.
3. We toss a fair coin. Do you accept the following deal: If head turns up you get 150€, while if tail turns up you loose 100€. Yes/No

Not a single one of our participants chose the coin toss in question 1, 48.5% of our participants preferred the lottery in question 2 and 28.5% accepted the coin toss in question 3. We then categorize our participants into three categories: the most risk averse participants, which chose the safe outcome each time ($r = 0$), those that accepted one of the three lotteries ($r = 1$) and those that accepted two lotteries ($r = 2$). The distribution across treatments of these types was as follows

	N-1	N-2	N-3	R-1	R-2	R-3
$r = 0$	0.45	0.53	0.45	0.46	0.41	0.33
$r = 1$	0.38	0.31	0.38	0.30	0.38	0.33
$r = 2$	0.17	0.16	0.17	0.24	0.22	0.33

We find no significant differences in the pairwise comparison between treatments in the proportion of "least risk averse" ($r = 2$) participants (Mann-Whitney, $p > 0.2816$) or "most risk averse" ($r = 0$) participants (Mann-Whitney, $p > 0.1217$) with the exception of the comparison between N-2 and R-1 for $r = 0$ (Mann-Whitney, $p = 0.0462$).

Next we check whether participants classified as "most risk averse" ($r = 0$) choose the efficient (but more risky) actions A and B more or less often. The results can be found in the following table.

N-1	N-2	N-3	R-1	R-2	R-3	**overall**
-0.0170	0.0004	-0.0841^{**}	-0.1401^{**}	-0.0461	-0.0905^{**}	$\mathbf{-0.0515}^{***}$

Clearly the results show that risk-averse participants choose the risky actions significantly less often overall. If we separate data per treatment, the effect is only significant in N-3, R-3 and R-1, which is where efficient actions where played more often. We also check whether those players 4 who are classified as "least risk averse" choose the risk-dominant action C more or less often and again we find strongly significant results.

Player 4	N-2	N-3	R-2	R-3
$r = 2$ vs C	–	-0.2004^{**}	-0.3350^{***}	-0.1825
$r = 0$ vs C	0.1212	–	0.3050^{***}	0.1265

Since actions C and D are chosen with much higher probability in all treatments the effect appears much stronger here compared to the table above. Players 4 classified as "least risk averse" choose action C much less often. This illustrates that C is perceived as less risky which is consistent with our analysis in subsection 5.2. (In treatment N-2 all players 4 where classified as $r = 1$ or $r = 0$). The opposite result is also true. Players 4 classified as very risk-averse ($r = 0$) choose C much more often than others. (Again in N-3 no player 4 was in this category).

Overall the results from the questionnaire show that we can be confident in the quality of our data as well as in the fact that our definition of risky actions coincided with that of the participants.

6 Conclusions

This paper provides experimental evidence on how network architecture affects play in coordination games. Subjects are distributed on a fixed social network and matched with all their neighbors to play a 4×4 game, which exhibits features of both Coordination and (Anti-)Coordination.

We observe that some network heterogeneity in terms of connectivity can help to overcome coordination failure. The underlying intuition is that more connected individuals display more stable best responses. We introduce this concept formally and use evidence from our data to show that behavior of more connected individuals is indeed more stable empirically and that more connected agents have more influence on others.

Despite the multiplicity of equilibria in the network game, we observe a sharp selection empirically. People generally coordinate on risk dominant actions and on the equilibria, in which the most connected player imposes her preferred action. Furthermore observed equilibria are almost always such that all links are in bilateral Nash equilibrium which is not needed theoretically. Local (endogenous) and global information lead to the same equilibrium selection. Global information does not improve the rate of succesful coordination.

These findings have implications for modeling social structures and for choosing which agents to target in situations where coordination is desirable, such as health and development programs, coordination of behavior in corporations, or market coordination on standards and norms.

A large literature in Economics, Management Studies and Political Sciences has acknowledged that leadership is crucial to overcome coordination problems and has tried to understand how such leadership can emerge (Brandts and Cooper, 2006; Weber et al., 2001; Calvert, 1992). Our results show that leadership can emerge in a natural way from the learning process of individuals who differ only in their position in the social network (and not e.g. in other characteristics). Some networks allow such natural leadership to be successful, while in others even highly connected individuals may fail to be effective leaders.

Future research could explore whether and if so how endogenous network formation interplays with coordination. Also the concept of stable best responses seems crucially important for coordination and certainly worthwhile exploring theoretically. Network heterogeneity is probably one of the most natural mechanisms (though certainly not the only one) which induces differential stability in best responses.

References

[1] Alos-Ferrer, C. and S. Weidenholzer (2008), Contagion and Efficiency, Journal of Economic Theory, forthcoming.

[2] Berninghaus, S., K. Ehrhart and M.Ott (2008), Myopically Forward Looking Agents in a Network Formation Game, working paper SFB 504.

[3] Blume, L. (1993), The Statistical Mechanics of Strategic Interaction, Games and Economic Behavior 5(3), 387-424.

[4] Bramoullé, Y. (2007), Anti-Coordination in Networks, Games and Economic Behavior 58, 30-49.

[5] Bramoullé, Y., S. Goyal, D. Lopez-Pintado and F. Vega-Redondo (2004), Network Formation and Anti-Coordination Games, International Journal of Game Theory 33, 1-19.

[6] Bramoullé, Y. and R. Kranton (2007), Public Goods in Networks , Journal of Economic Theory 135, 478-494.

[7] Brandts, J. and D.J. Cooper (2006), A change would do you good ... An Experimental study on now to overcome coordination failure in organizations, American Economic Review 96(3), 669-693.

[8] Calvert, R.L. (1992), Leadership and its basis in problems of social coordination, International Political Science Review 13(1), 7-24.

[9] Cassar, A., (2007), Coordination and cooperation in local, random and small world networks: Experimental Evidence, Games and Economic Behavior 58(2), 209-231.

[10] Cheloukine, S. and J. King (2007), Corruption networks as a sphere of investment activities in modern Russia, Communist and post-communist studies 40(1), 107-122.

[11] Choi, S. and J. Lee (2010), Communication, Coordination and Networks, mimeo UCL.

[12] Cooper, R., D.V. DeJong, R. Forsythe and T.W. Ross (1993), Forward induction in the battle-of-the-sexes games. American Economic Review 83, pp. 1303–1316.

[13] Corbae, D. and J. Duffy (2008), Experiments with Network Formation, Games and Economic Behavior 64, 81-120.

[14] Ellison, G. (1993), Learning, Local Interaction and Coordination, Econometrica 61(5), 1047-1071.

[15] Fischbacher, U. (2007), z-Tree: Zurich toolbox for ready-made economic experiments. Experimental Economics 10 (2): 171-178.

[16] Foster, A. and M. Rosenzweig (1995), Learning by Doing and Learning from Others: Human Capital and Technical Change in Agriculture, Journal of Political Economy 103(6), 1176-1209.

[17] Galeotti, A. and S. Goyal (2009), Influencing the Influencers: A theory of strategic diffusion, RAND Journal of Economics 40(3), 509-532.

[18] Garey, M.R. and D.S. Johnson (1979), Computers and Intractability, W.H. Freeman, New York.

[19] Goyal, S. and F. Vega-Redondo (2005), Network formation and social coordination, Games and Economic Behavior 50, 178-207.

[20] Goyal, S. (2007), Connections: An Introduction to the Economics of Networks, Princeton: Princeton University Press.

[21] Hirschman, A. (1958), The Strategy of Economic Development, New Haven, CT: Yale University Press.

[22] Hojman, D. and A. Szeidl (2006), Endogenous Networks, Social Games and Evolution, Games and Economic Behavior 55 (1), 112-130.

[23] Jackson, M. O. and A. Watts (2002), On the formation of interaction networks in social coordination games, Games and Economic Behavior 41, 265-291.

[24] Jackson, M.O. (2008), Social and Economic Networks, Princeton University Press.

[25] Katz, M.L. and C. Shapiro (1986), Technology Adoption in the Presence of Network Externalities, Journal of Political Economy 94(4), 822-842.

[26] Kearns, M., S. Suri and N. Montfort (2006), An Experimental Study of the Coloring Problem on Human Subject Networks, Science 313, 824-827.

[27] Kosfeld, M. (2004), Economic Networks in the Laboratory: A survey, Review of Network Economics 3: 20-41.

[28] Kovarik, J., F. Mengel and J.G. Romero (2010), Learning in Networks, mimeo Maastricht University.

[29] Leitner, Y. (2006), Financial Networks: Contagion, Commitment and Private Sector Bailouts, The Journal of Finance 60:6, 2925-2953.

[30] Morris, S. (2000), Contagion, Review of Economic Studies 67, 57-78.

[31] My, K.B., M. Willinger and A. Ziegelmeyer (2001), Global versus Local Interaction in Coordination Games: An Experimental Investigation, BETA-working paper.

[32] Ray, D. (2007), Development Economics, New Palgrave Dictionary of Economics eds L.Blume and S.Durlauf.

[33] Rosenkrantz, S. and U. Weitzel (2008), Network Structure and Strategic Investments: An Experimental Analysis, working paper Utrecht University.

[34] Vega-Redondo, F. (2007), Complex Social Networks, Cambridge University Press.

[35] Weber, R., C. Camerer, Y. Rottenstreich, and M. Knez (2001), The illusion of leadership: misattribution of cause in coordination games, Organizational Science 12(5), 582-598.

Influence	1st	2nd	3rd	4th	1st (N-3)	2nd (N-3)	3rd (N-3)	4th (N-3)
Player 1	0.30	0.18	0.08	0.05	0.21	0.04	0.00	0.00
Player 2	0.51	0.17	0.11	-	0.22	0.04	0.00	-
Player 3	0.49	0.30	-	-	0.28	0.06	-	-
Player 4	**0.55**	**0.27**	**0.19**	-	**0.45**	**0.08**	**0.02**	-
Player 5	0.27	0.15	0.07	-	0.31	0.08	0.00	0.00
Player 6	0.30	0.22	0.09	0.04	0.22	0.04	0.00	0.00
Player 7	0.26	0.16	0.04	0.02	0.20	0.06	0.00	0.00
Player 8	0.28	0.16	0.07	0.03	0.19	0.08	0.01	0.00

Table 14: Influence per neighbor: (All periods). Average number of first-order neighbors of player i that switch actions at $t+1$ after i has switched at t as a percentage of the number of first-order neighbors of i. Average number of 2nd order neighbors that switch at $t+2$ after 1st order neighbors switch at $t+1$ and i at t etc..

A Appendix

A.1 Additional Regression Tables

A.2 Remarks and Proofs for Claims 1,3 and 4

Proof of Claim 1

Proof. The following table shows the Nash equilibria referred to in Claim 1 in the format $(a_1, a_2, ...a_8)$.

Nash equilibrium	Network 1	Network 2	Network 3
	(A,B,A,B,A,B,A,B)	**(A,B,A,B,B,A,A,A)**	**(A,B,A,B,A,A,A,A)**
	(B,A,B,A,B,A,B,A)	**(B,A,B,A,A,B,B,B)**	**(B,A,B,A,B,B,B,B)**
	(C,D,C,D,C,D,C,D)	**(C,D,C,D,D,C,C,C)**	**(C,D,C,D,C,C,C,C)**
	(D,C,D,C,D,C,D,C)	**(D,C,D,C,C,D,D,D)**	**(D,C,D,C,D,D,D,D)**
	(D,D,C,D,C,D,D,C)	(C,D,D,D,C,D,C,C)	(D,C,D,D,C,C,C,C)
	(D,C,D,C,D,D,C,D)	(D,C,D,D,C,D,C,C)	(C,D,D,C,D,D,D,D)
	(C,D,C,D,D,C,D,D)	(D,C,D,D,D,C,C,C)	(A,B,C,A,B,B,B,B)
	(D,C,D,D,C,D,D,C)	(C,D,D,C,C,D,D,D)	(D,C,D,B,A,A,A,A)
	(C,D,D,C,D,D,C,D)	(A,B,C,D,D,C,C,C)	(C,D,C,A,B,B,B,B)
	(D,D,C,D,D,C,D,C)	(C,D,D,D,C,D,C,C)	(B,A,B,C,D,D,D,D)
	(D,C,D,D,C,D,C,D)		
	(C,D,D,C,D,C,D,D)		

Table A-1: Strict Nash equilibria. The format is $(a_1, ..., a_8)$ where $a_i, i = 1, ..8$ is the action of player i. Nash equilibria marked in bold have the property that all links are bilaterally in Nash equilibrium.

(i) The four Nash equilibria where all links are bilaterally in Nash equilibrium are the first four profiles in each column of the table.

(ii) Other than those there are no Nash equilibria where all agents choose A or B. The reason is that this would involve either two neighboring agents choosing A or two neighboring agents choosing B. Now for any pair of agents across all our networks, the agent with less connections among the two always has at most two connections with the exception of one pair, namely players 3 and 4 in Network 2. But for all pairs where one agent has one or two connections there can be no additional equilibria since the best response to one neighbor choosing A and one neighbor B is given by $BR(A, B) = C$

(and obviously $BR(A) = B$ and $BR(B) = A$. Consider now agents 3 and 4 in network 2. We have $BR(A, B, B) = C$ and $BR(A, A, B) = B, C$. Hence the only such candidate equilibrium is the one where agent 3 chooses B and agent 4 chooses B as well. Then player 3's neighbors 2 and 5 would need to choose A. But note that in this case player 6 (who is connected only to player 5 (choosing A) and to player 4 (choosing B) would best respond with C. Hence no additional equilibria exist where all agents choose A or B.

(iii) We now show that set of strict Nash equilibria presented in the table above is exhaustive. **Network 1**: Remember that in Network 1 each player has two neighbours. In Network 1 there are no additional strict Nash equilibria where *any* agent chooses A or B. To see this it suffices to realize that this would involve an agent choosing B with neighbors choosing A and C. (A is not a strict best response to (B,D) and B is not a strict best response to either (A,B) or (A,D)). This yields the string (..,-,A,B,C,A,-,..), since C is the only a strict best response to (B,A). But A is not a best response to any of the following: (C,A), (C,B), (C,C) or (C,D).

Network 2: Players 7 and 8 will always choose a best response to whatever player 4 chooses. Let us assume player 4 chooses A (and hence players 7 and 8 choose B). Then both players 3 and 6 have to choose B for A to be a strict best response to (-,-,B,B). The only candidate profile is (-,-,B,A,-,B,B,B). Clearly player 5 will choose A. B is strict best response for player 3 to (-,A,A) whenever player 2 chooses A or C. But BR(D,B)=$\{A, C\}$. Hence there is no additional strict Nash equilibrium.

Assume next that player 4 chooses B (and hence players 7 and 8 choose A). This is a strict best response to (-,-,A,A) whenever either players 3 and 6 choose (A,A) or else whenever they choose (B,C), (C,D) or (C,C). This yields the following candidate equilibria (-,-,B,B,D,C,A,A), (-,-,C,B,D,B,A,A), (-,-,C,B,D,D,A,A), (-,-,D,B,D,C,A,A), and (-,-,C,B,D,C,A,A), where we have entered player 5's strict best response in each case already. Let us check for whether player 6 is choosing a strict best response. None of the candidate equilibria survives this check. Hence there is no additional strict equilibrium where player 4 chooses B.

Let us next assume that player 4 chooses C (and that players 7 and 8 choose D). C is a best response to (-,-,D,D) whenever players 3 and 6 choose (D,D), (A,A), (A,D), (A,C) or (B,D). After checking player 5's and 6's best responses the following candidates survive: (-,-,A,C,D,D,D,D) and (-,-,D,C,C,D,D,D). But since A is never a strict best response to (-,C,D) for player 3 the only additional equilibrium is (C,D,D,C,C,D,D,D).

Finally assume that player 4 chooses D (and hence players 7 and 8 choose C). This is a strict best response to (-,-,C,C) whenever players 3 and 6 choose (C,C) (C,D) (D,D), (B,C), (B,B) or (B,D) yielding candidates (-,-,C,D,C,D,C,C), (-,-,D,D,D,C,C,C), (-,-,D,D,C,D,C,C), (-,-,B,D,D,C,C,C), (-,-,C,D,-,B,C,C), (-,-,B,D,-,B,C,C),(-,-,B,D,C,D,C,C) and (-,-,D,D,-,B,C,C), where we have filled in player 5's strict response wherever possible. Checking whether player 6 is choosing a strict best response returns candidates (-,-,C,D,C,D,C,C), (-,-,D,D,D,C,C,C), (-,-,D,D,C,D,C,C), (-,-,C,D,D,C,C,C) and (-,-,B,D,C,D,C,C). Verifying player 3's best response in each case yields (D,C,D,D,C,D,C,C), (C,D,D,D,C,D,C,C), (D,C,D,D,D,C,C,C) and (A,B,C,D,D,C,C,C) as additional equilibria.

Network 3: Players 5,6,7 and 8 will always choose a best response to whatever player 4 chooses. Let us hence focus on the first 4 players. Assume player 4 chooses A (then players 5-8 will choose B). Player 3 can best respond to this obviously with B. If player 3 best responds with A, then player 4 has no strict best response. Assume player 3 best responds with C. Then player 4's best response is indeed A. We have the string (-,-,C,A,B,B,B,B). Now C is a best response to either (A,B) or (A,D). Hence we have two equilibria where player 4 chooses A: (A,B,C,A,B,B,B,B) and (C,D,C,A,B,B,B,B). Assume that player 4 chooses B (and hence players 5-8 choose A). Player 4 has no incentives to deviate irrespective of what player 3 chooses. Consider then the four candidate strings (-,-,A,B,A,A,A,A),(-,-,B,B,A,A,A,A),(-,-,C,B,A,A,A,A) and (-,-,D,B,A,A,A,A). Now A is a strict best response only to (B,B). B is never a strict best response to (B,-). C is a strict best response to (B,A) and hence we get (B,A,C,B,A,A,A,A). But A is not a best response for player 2 to (B,C). D is strict best response only to (B,C) and hence we get (D,C,D,B,A,A,A,A). Assume next that player 4 chooses C (and hence players 5-8 D). Consider candidate strings (-,-,A,C,D,D,D,D), (-,-,B,C,D,D,D,D),(-,-,C,C,D,D,D,D) and (-,-,D,C,D,D,D,D). A is never a best response to (C,-). B is a best response to (A,C) yielding (B,A,B,C,D,D,D,D). C is never a best response to (C,-) and D is a best response to (B,C) or (D,C) yielding (A,B,D,C,D,D,D,D) and (C,D,D,C,D,D,D,D). The former is not an equilibrium since B is not a best response for player 2 to (A,D). Finally assume that player 4 chooses D (and hence players 5-8 C). Player 3 then has to choose B, C or D. Consider candidate strings (-,-,B,D,C,C,C,C),(-,-,C,D,C,C,C,C) and (-,-,D,D,C,C,C,C). Now B is never a best response to (D,-). C is a strict best response to (A,D) yielding ((B,A,C,D,C,C,C,C), which is not an equilibrium since A is not a best response for player 2 to (B,C). D is a strict best response to (D,C) yielding (D,C,D,D,C,C,C,C).

■

Stochastic Stability

Before we discuss stochastic stability first note that the myopic best response dynamics defines a Markov chain whose absorbing states are exactly the one-shot Nash equilibria of the network game. Stochastic Stability is a refinement of the notion of absorbing state, which is often used in the literature on games played on networks. (See e.g. Jackson and Watts, 2002; Goyal and Vega Redondo, 2005; Bramoulle et al, 2004 etc..). The underlying idea is to check how robust absorbing states are towards small mistakes by a player.

Trees Maybe the underlying intuition is best described by relying on the graph-theoretic techniques developed by Freidlin and Wentzell (1984). They can be summarized as follows. For any state a an $a-$tree is a directed network on the set of absorbing states, whose root is a and such that there is a unique directed path joining any other a' to a. For each arrow $a' \to a''$ in any given $a-$tree the "cost" of the arrow is defined as the minimum number of simultaneous trembles (mistakes) necessary to reach a'' from a'. The cost of the tree is obtained by adding up the costs of all its arrows and the stochastic potential of a state a is defined as the minimum cost across all $a-$trees.

Result (Young 1993) State a^* is stochastically stable if it has minimal stochastic potential.

The intuition behind Young's result is simple. In the long run the process will spend most of the time in one of its absorbing states. The stochastic potential of any state a is a measure of how easy it is to jump from the basin of attraction of other absorbing states to the basin of attraction of state a by perturbing the process through trembles.

It is a well known result that if state a can be reached from a' via one tremble only and if a' is stochastically stable then a must also be stochastically stable. (The reason is that starting from an a'-tree we can simply rewire the arrow leading away from a which yields an a-tree with at most the same stochastic potential as that of the original tree). But then whenever any number of absorbing states is connected via a sequence of one-trembles then either all such states will be stochastically stable or none.

Proof of Claim 3

Proof. In the following we will show that in all our networks all absorbing states are connected via a sequence of one-trembles. Using the result abovem this then shows that all our absorbing states are also stochastically stable. We will keep the exposition short but the reader should easily be able to verify the claims we make.

Network 1 First note that all absorbing states where not all links are in bilateral equilibrium are connected via a sequence of one-trembles where one of two neighbouring D-players switches to

C inducing her C-neighbour to switch to D. Furthermore from any state where all links are in Nash equilibrium (and all players choose C or D) one such state can be reached via one tremble from a C-player to D, after which her neighbours will have incentives to choose D, leading to the new absorbing state. On the other hand from any state where not all links are in bilateral Nash equilibrium either state (C,D,C,D,C,D,C,D) or state (D,C,D,C,D,C,D) can be reached via one tremble from a C player to D as well. From state (A,B,A,B,A,B,A,B,) state (D,C,D,C,D,C,D,C) can be reached via one tremble of an A-player to B after which her neighbors (facing one B and one A neighbour) will choose C as best response. In the next period their neighbors (facing one C and one B neighbour) will choose D etc.. From state (B,A,B,A,...) state (C,D,C,D...) can be reached in the same way. And state (A,B,A,B...) can be reached from (C,D,C,D...) via one tremble by a C-player to A. Her neighbours then have incentives to choose B and their neighbours A etc...

Network 2 Here again (D,C,D,C...) can be reached from (A,B,A,B,..) via one tremble by player 3 to D and (B,A,B,A,...) is reached from (D,C,D,C...) via one tremble by player 3 to C. (C,D,C,D...) is reached from (B,A,B,A...) via one tremble by player 4 to D and (A,B,A,B,...) is reached from (C,D,C,D,...) via one tremble by player 4 to B. Furthermore all states where not all links are in bilateral Nash equilibrium are connected to at least one of these states via one tremble as the reader can easily verify.

Network 3 In this network it is easy to see that a tremble by player 4 connectes all possible states where all links are in Nash equilibrium and all other states are connected to one of those states via one tremble by player 2.

■

Proof of Claim 4 Proof. (i) We prove the first part by contradiction. Assume that $k_i \leq k_j$, but that there is a state a from which more changes are needed to change i's best response rather than j's best response. Assume that in this state player i chooses a_i and player i's neighbors (denoted by $(i1, ..ik_i)$) actions $a_{i1}, .., a_{ik_i}$. Assume furthermore that player j chooses a_j and that player j's neighbors choose $(a_{j1}, .., a_{jk_j})$. Assume that the neighbors are ordered in such a way that the minimal number of changes needed to induce a change in the best response of player j involves a change to player j's first $x_j(a)$ neighbors. Clearly we must have $x_j(a) < k_i$, since if all neighbors of player i change her actions her best response will change (unless the game has a strictly dominant strategy in which case we have reached a contradiction). Now consider the state a' where $a'_i, a'_{i1}, .., a'_{ik_i} = (a_j, a_{j1}, .., a_{jk_i})$ and where $(a'_j, a'_{j1}, .., a'_{jk_j}) = (a_i, a_{i1}, .., a_{ik_i}, ...a'_{jk_j})$. By construction $x_i(a') \leq x_j(a) < x_i(a) \leq x_j(a')$. Hence we have constructed a state from which strictly more changes are needed to change j's best response rather than i's best response contradicting the assumption that i has a more stable best response than j.

(iii) First note that for any player who has at most 2 neighbours, one change in any state suffices to change her best response. For one neighbour this is obvious. For two neighbours note that the best response to a neighborhood AB is C, to a neighborhood AC is B, to AD is C, to BC is D to BD is A or C and to CD is D. Hence in Network 2 we have that $\forall i = 1, 2, 5, 6, 7, 8 : x_i(a) = 1, \forall a$. To complete the proof for Network 2 we have to find two states: one from which player 4's best response can be changed by one mistake, but not player 3's and one where this is the other way round. An example of the first is the state (B, A, B, A, A, B, B, D). Starting from this state player 4's best response changes by having player 7 choose A instead of B, while two changes are needed to affect player 3's best response. An example of the latter is (A, B, A, C, B, A, A, A). In network 3 the proof is immediate since all players have at most 2 neighbours except for player 4 who has 5 neighbours and hence more than one change is needed e.g. in the state (A, B, A, B, A, A, A, A). ■

Proof of Claim 5 Proof. (i) For Network 3 the claim is obvious, since any switch by player 4 changes the best response of her neighbours (5,6,7,8) and hence $y_4(a) \geq 4, \forall a$. But no other player has even 4 neighbours.

(ii) For Network 2 note first that $y_4(a) \geq 2, \forall a$. Hence $y_4(a) > y_j(a)$ for $j = 1, 7, 8$. We also have $y_4(a) \leq y_j(a)$ for $j = 2, 5, 6$. Hence it suffices to find one state where $y_4(a) > y_j(a)$ for $j = 2, 5, 6$, which could be state $a' = (A, B, A, B, B, A, A, A)$. Note that from this state also $y_4(a') > y_3(a')$. It remains to show that the reverse inequality cannot hold for any state. To see this note that $y_3(a) = 3$ implies that a switch by player 3 has to change player 4's and player 5's best response. But one switch by player 3 only can shift player 4's best response only if the link 46 is *not* in bilateral equilibrium. But if this is the case then one switch by player 4 can induce player 6 to switch implying $y_3(a) \leq 3$. ■

A.3 Sample Instructions (Treatments N-1, N-2 and N-3)

Welcome and thanks for participating at this experiment. Please read these instructions carefully. They are identical for all the participants with whom you will interact during this experiment.

If you have any questions please raise your hand. One of the experimenters will come to you and answer your questions. From now on communication with other participants is not allowed. If you do not conform to these rules we are sorry to have to exclude you from the experiment. Please do also switch off your mobile phone at this moment.

For your participation you will receive 2 Euros. During the experiment you can earn more. How much depends on your behavior and the behavior of the other participants. During the experiment we will use ECU (Experimental Currency Units) and at the end we will pay you in Euros according to the exchange rate 1 Euro = 75 ECU. All your decisions will be treated confidentially.

THE EXPERIMENT

In the experiment you are linked up with some other participants in this room, which we will call your neighbours. You will play a game with your neighbours that we will describe below. Your neighbours in turn are of course linked up with you, but (possibly) also with other participants in the room. And their neighbours again are linked up with other participants and so on...

Note that your neighbours are not necessarily the participants who are located to your left and right in the physical layout of the computer laboratory.

During the experiment, you will be able to find out how many neighbours you have as well as their experimental identity, but not who they really are. This also means, of course, that your neighbours will not know your real identity.

The experiment lasts for 20 rounds. In each round you play a game with each of your neighbours. Your payoff in each round is the sum of payoffs obtained in all the games with your neighbours.

Each round consists of three stages, which we will describe in detail below. Here is a summary:

1. In the first stage you choose an action in the game. Note that you have to choose the same action for all your neighbours.
2. In the second stage you can request information about your neighbours, your neighbours' neighbours etc... the actions they chose in the past period and the payoff they obtained in the past period, as well as about your own payoff.
3. In the third stage, the information you requested is displayed on the computer screen.

We will now describe the different stages in more detail.

Stage 1 (Action Choice)

In the first stage you have to choose one action in the game, which is described by the following table, which will be shown to you every time you choose an action.

	A	B	C	D
A	20,20	40,70	10,60	20,30
B	70,40	10,10	30,30	10,30
C	60,10	30,30	10,10	30,40
D	30,20	30,10	40,30	20,20

In the table your actions and payoffs are given in dark grey and your neighbour's actions and payoffs in light grey. The table is read as follows (dark payoffs):

- If you choose A and your neighbour A, you receive 20
- If you choose A and your neighbour B, you receive 40
- If you choose A and your neighbour C, you receive 10
- If you choose A and your neighbour D, you receive 20
- If you choose B and your neighbour A, you receive 70
- If you choose B and your neighbour B, you receive 10
- If you choose B and your neighbour C, you receive 30
- If you choose B and your neighbour D, you receive 10
- If you choose C and your neighbour A, you receive 60
- If you choose C and your neighbour B, you receive 30
- If you choose C and your neighbour C, you receive 10
- If you choose C and your neighbour D, you receive 30
- If you choose D and your neighbour A, you receive 30
- If you choose D and your neighbour B, you receive 30
- If you choose D and your neighbour C, you receive 40
- If you choose D and your neighbour D, you receive 20

Note that your neighbour (light payoffs) is in the same situation as you are. This means that for your neighbour:

- If your neighbour chooses A and you A, your neighbour receives 20
- If your neighbour chooses A and you B, your neighbour receives 40
- If your neighbour chooses A and you C, your neighbour receives 10
- If your neighbour chooses A and you D, your neighbour receives 20
- If your neighbour chooses B and you A, your neighbour receives 70
- If your neighbour chooses B and you B, your neighbour receives 10
- If your neighbour chooses B and you C, your neighbour receives 30
- If your neighbour chooses B and you D, your neighbour receives 10
- If your neighbour chooses C and you A, your neighbour receives 60
- If your neighbour chooses C and you B, your neighbour receives 30
- If your neighbour chooses C and you C, your neighbour receives 10
- If your neighbour chooses C and you D, your neighbour receives 30
- If your neighbour chooses D and you A, your neighbour receives 30
- If your neighbour chooses D and you B, your neighbour receives 30
- If your neighbour chooses D and you C, your neighbour receives 40
- If your neighbour chooses D and you D, your neighbour receives 20

Remember that you have to choose the same action for all your neighbours. Your gross payoffs in each round are given by the sum of payoffs you have obtained in all games against your neighbours divided by the number of neighbours you have.

Stage 2 (Information Request)

In the second stage you can indicate which of the following pieces of information you would like to obtain

- the experimental identity of your neighbours
- the experimental identity of your neighbours' neighbours (2nd order neighbours)
- the experimental identity of your neighbours' neighbours' neighbours (3rd order)
- the experimental identity of your neighbours' neighbour's neighbours' neighbours (4th order neighbours)

Note that who is a neighbour of you does not change during the experiment. Hence once you have asked for this information in some round, it will be displayed in all future rounds. Note also that in order to receive information about your neighbours' neighbours' you first need to request information about your neighbours etc... The cost of requesting each of these pieces of information is 10. You only have to pay this cost once. In addition you can request information about the following items which (in principle) can change in every round.

- the actions chosen by your neighbours
- the actions chosen by your neighbours' neighbours
- the actions chosen by your neighbours' neighbours' neighbors
- the actions chosen by your neighbours' neighbour's neighbours' neighbours
- the payoffs obtained by your neighbours
- the payoffs obtained by your neighbours' neighbours
- the payoffs obtained by your neighbours' neighbours' neighbors
- the payoffs obtained by your neighbours' neighbour's neighbours' neighbours
- your own payoffs

Obviously, in order to receive information about your neighbours (or neighbours' neighbours') actions or payoffs you first need to request information about the experimental identity of your neighbours (neighbours' neighbours) etc...The cost of requesting each of these pieces of this information is 1 and you have to pay it each time you request this information anew. Your net payoffs in a round are your gross payoffs minus the cost of the information you requested.

Stage 3 (Information Display)

The information you have requested in Stage 2 is displayed on the screen for 40 seconds.

Control Questions

Before we start the experiment please answer the following control questions on your screen.

1. Assume you have only one neighbour. She chooses action B and you action D. Which gross payoff will you get in this round?

2. Assume you have three neighbours and they choose action A, B and A. You choose action D. Which gross payoff will you get in this round?

3. True or False: My neighbours change in every round of the game.
4. True or False: My neighbours face the same payoff table as I do.
5. True or False: My neighbours are those sitting in the cubicles to my left and right.