Economic Networks in the Laboratory: A Survey

MICHAEL KOSFELD*

Institute for Empirical Research in Economics, University of Zurich

Abstract

This paper provides a survey of recent experimental work in economics focusing on social and economic networks. The experiments consider networks of coordination and cooperation, buyer-seller networks, and network formation.

1 Introduction

Economic research on networks has increased tremendously in recent years. Substantial evidence emphasizing the important role of networks on social and economic outcomes exists. Prominent examples have pointed out the function of networks with regard to job search (Holzer, 1987; Montgomery, 1991), trade (Lazerson, 1993; Nishiguchi, 1994), the granting of credit (McMillan and Woodruff, 1999), mutual insurance (Fafchamps and Lund, 2001), and welfare participation (Bertrand, Luttmer and Mullainathan 2000).[1] While theoretical research on economic networks has also received extensive interest (see references in the subsequent sections), no experimental work on networks in economics existed until very recently. The number of network experiments is still very small, but the literature is starting to grow. The aim of this paper is to provide an overview of existing experimental work and to suggest paths for interesting future research in this area.

Laboratory experiments present a useful and powerful technique for analyzing economic questions. The main advantage of experiments lies in the ability to control variables (such as costs, benefits, information, and timing) that could possibly influence individual and aggregate behavior, something which is very hard and sometimes even impossible to achieve in the field. Together with theoretical models, in particular models based on game theory, that provide the language for an exact formulation of hypotheses, controlled experiments are a key element in the transformation of economics to an empirical science. And, as one of the major contributors to economic theory writes in a recent paper, "moving from arm-chair theorizing to controlled laboratory experiments may be as important a step in the development of economics as it once was for the natural

* Institute for Empirical Research in Economics, University of Zurich, Blümlisalpstrasse 10, CH-8006 Zurich, Switzerland. E-mail: kosfeld@iew.unizh.ch I am grateful to Armin Falk, Jacob Goeree, Arno Riedl, Sudipta Sarangi, an anonymous referee and the editor for helpful comments and suggestions.

[1] Anthropologists (e.g., Lévi-Strauss, 1963) and sociologists (e.g., Granovetter, 1974) have emphasized the role of networks and social structure outside of economics.

sciences to move from Aristotelian scholastic speculation to modern empirical science" (Weibull, 2001).

In addition, economic laboratory experiments can also be used to glean concrete public policy advice. Network industries (e.g., water, gas, and electric power) have become the object of major privatization and deregulation movements in many different countries in recent years. Experiments offer a method for testing and evaluating different institutional arrangements before their implementation in the real world. In this case, the laboratory serves as a "wind tunnel" for the detection of unforeseen problems and comparison of different market designs.

Long before experimental economists discovered networks other social scientists had already started investigating the role of networks in various experimental settings. Perhaps the earliest network experiments were the "MIT experiments" by social psychologist Alex Bavelas and his colleagues in the early 50's (Bavelas, 1950; Leavitt, 1951). In these experiments, a group of individuals is assigned a problem to be solved, where typically each individual receives a card showing different symbols. Individuals in a group have only one common symbol and the objective is to discover that common symbol. Individuals can communicate by passing written messages to each other, and communication can only flow along an exogenously imposed network. Bavelas and his colleagues consider four different communication networks: the chain, the circle, the star, and the "Y".[2] They find that groups which communicate via the star or the "Y" solve the given problem most quickly. Moreover, they use the least number of messages and also make the fewest errors. On the other hand, individuals in the circle and the chain report high average satisfaction during the experiment. Only the individual serving at the center in the star reports higher satisfaction. As a consequence of these experiments Bavelas *et al.* emphasize the role of structural centrality for group efficiency. Further studies in later years investigated the effect of centrality on communication and organization networks. See Shaw (1964) for a critical review.

Clearly, sociologists also have an ongoing interest in the role of networks. Cook and Emerson (1978) as well as Cook, Emerson, Gilmore, and Yamagishi (1983) conducted early experiments examining the relationship between power and social structure in exchange networks (see also the related theoretical work of Markovsky, Willer and Patton, 1988). Recent articles by Bienenstock and Bonacich (1993, 1997) incorporate game theoretic concepts in the discussion.

This survey focuses on recent experimental work by economists in this area. It seems convenient to organize the literature into four different categories, each of which will be reviewed in a separate section. Section 2 starts with coordination networks. Section 3 and section 4 contain a discussion of cooperation networks and buyer-seller networks, respectively. Section 5 surveys experimental work on network formation. Obviously, this categorization is not the only one possible. Some of the papers that are presented in one section contain elements that are discussed in some of the other sections. Yet, I think that the categorization is helpful in organizing the existing literature.[3] In each section I first

[2] The "Y" is a five-person network that connects the following pairs of individuals each via a direct link: (12), (23), (34), (35).

[3] Clearly, the discussion in each category touches also aspects that are not directly related to networks. For example, experiments on coordination and cooperation networks contain elements that have been analyzed in other settings, as well. However, since the experimental literature on coordination and cooperation is fairly

give a short overview of the theoretical work and then describe the experiments that were conducted in the particular area.[4]

2 Coordination networks

2.1 Theory

Seminal work by Kandori, Mailath, and Rob (1993) and Young (1993) triggered intense interest in the question of equilibrium selection in coordination games. Important research examines the impact of different network structures on equilibrium selection and, in those cases where players can choose their network partners themselves, whether players will form networks that lead to play of the efficient Nash equilibrium in the coordination game.

Ellison (1993) and Morris (2000) analyze the role of local interaction networks in the spread of particular strategies in 2×2 coordination games, showing, for example, how play converges to the risk-dominant equilibrium if players are located on a circle and interact with their two nearest neighbors. Similarly, Blume (1993) and Kosfeld (2002) prove convergence to the risk-dominant equilibrium in a population of players located on a d-dimensional lattice.

In contrast, Ely (2002) and Bhaskar and Vega-Redondo (2002) show, that once players are allowed to choose their partners themselves, the situation looks very different. They introduce a number of locations where players can meet and play the coordination game with each other. Thus, at any time, players choose both a location and a strategy in the game. Under these conditions, the authors show that risk dominance looses its selection force and that the population is most likely to coordinate on the efficient equilibrium. The reason is intuitive. Since players can freely choose their interaction partners, they are able to find partners that play the efficient equilibrium strategy and can simultaneously avoid players that play the inefficient strategy. Mailath, Samuelson, and Shaked (2001) also emphasize that the latter condition, that is, the ability to avoid bad matches, is crucial.

Since migration is cost free in the models described above, an interesting and important question is which equilibrium will be selected when migration or the formation of links is costly. Goyal and Vega-Redondo (2000) present a first theoretical approach in this direction. They find that the cost of a link to another player plays a decisive – and somewhat counterintuitive – role in the selection of equilibrium. Players will only coordinate on the efficient Nash equilibrium if costs are sufficiently high. In contrast, however, if costs are low the risk-dominant equilibrium will prevail. Droste, Gilles, and Johnson (2000) consider a similar stochastic learning model. In their set-up, players are located on a circle where they can form links to other players and links to more distant players are assumed to be more costly than those to neighbors. Their work shows that while the coexistence of both equilibrium strategies is possible in the medium run, the risk-dominant equilibrium will prevail in the long run.

---

large and has been reviewed elsewhere (e.g., Ochs, 1995; Ledyard, 1995) I do not address these more general issues in this survey but focus on network-related aspects only.

[4] I have included all experimental network papers that I know of. I would, of course, be happy to learn of other papers that were omitted in this survey.

2.2 Experiments

Extensive literature exists on experiments on coordination games (see, for example, Ochs, 1995 for a review). Keser, Ehrhart, and Berninghaus (1998) present the first experiment that considers the role of networks in coordination games along the lines discussed above. In this experiment, the authors study the impact of local interaction as analyzed by Ellison (1993) and Morris (2000). They implement two different treatments. In the local-interaction treatment, groups of eight players each interact around a circle. Players that are located next to each other play a coordination game for 20 periods. Keser *et al.* find that play converges to the risk-dominant Nash equilibrium in this treatment, as theory predicts. In contrast, once players interact in isolated groups of three, which corresponds to the neighborhood size in the local-interaction treatment, they find that play converges to the efficient Nash equilibrium. This finding is compatible with a similar experimental result of Van Huyck, Battalio, and Beil (1990) on coordination in small groups. In their second paper, Berninghaus, Ehrhart, and Keser (2002) put their result in a more general framework. In particular, they (i) modify the payoff function in the network-coordination-game, reducing the riskiness of the efficient Nash equilibrium, and (ii) vary the neighborhood structure that determines the way in which players interact locally. With regard to (i), they find that if the efficient Nash equilibrium becomes less risky, populations that interact locally on a circle also converge to efficient play in most cases. Varying the structure of the neighborhood, however, generates a converse effect. In order to address point (ii), Berninghaus *et al.* compare two treatments, where each player interacts locally with his four nearest neighbors in both treatments. In one treatment players are located on a circle, while they are located on a two-dimensional lattice in the second treatment.[5] Thus, the size of a player's neighborhood remains constant across treatments and only the neighborhood structure differs. The authors find that play is more likely to converge to the risk-dominant equilibrium when players interact on the lattice than on the circle. This result is particularly interesting since subjects had exactly the same instructions in both treatments, that is, they were not informed about the precise neighborhood structure of the population. One possible explanation the authors offer is the following: observations of individual play in the lattice treatment show it to be more changing than in the circle treatment. In consequence, risk dominance as an individual motive has more power in the lattice treatment than in the circle treatment.

Boun My, Willinger, and Ziegelmeyer (2001) also analyze the effect of interaction networks on equilibrium selection. Using a setting similar to Keser *et al.* (1998), they compare the repeated play of a 2×2 coordination game under global and under local interaction with varied degrees of risk dominance of the inefficient equilibrium. In contrast to Keser *et al.*, however, they keep the size of the total population fixed rather than the size of a player's neighborhood: groups contain eight players in each treatment. While the authors find that the degree of risk dominance in local interaction has the expected effect on the likelihood of selection of the inefficient equilibrium (that is, players coordinate more often on the inefficient equilibrium if it becomes more risk-dominant), no such effect is found in global interaction. Moreover, the interaction structure itself does not seem to play a significant role in the convergence of play. In particular, contrary to the studies described above, Boun My *et al.* do not find that players who interact locally on the circle coordinate more frequently on the risk-dominant equilibrium.

---

[5] In fact, players interact on a torus, which like a circle has no boundaries.

Corbae and Duffy (2002) also test the local-interaction hypothesis in a recent experiment. They look at groups of four players who operate in a network structure of either global interaction (the complete network), local interaction (the circle), or "marriage interaction". In the latter case, the population is split into two isolated pairs of players, who interact with each other. Independent of the network, subjects play ten periods of a coordination game in each group, where the efficient Nash equilibrium is also risk-dominant. With the exception of one group, play converges to the efficient equilibrium. After these ten periods, subjects face a new game that differs from the previous game only with respect to the off-equilibrium payoff of the efficient Nash equilibrium strategy, which renders the inefficient Nash equilibrium risk-dominant. This game is played for ten periods. The hypothesis is that if no player is forced to play the inefficient equilibrium strategy, players will keep coordinating on the efficient equilibrium irrespective of the underlying interaction structure. The observations in the experiment (three groups for the global interaction network and two groups for both the local and the marriage interaction network) confirm this hypothesis. Yet, if a single player in the group is forced to play the inefficient strategy, the hypothesis is that convergence to the inefficient equilibrium will be observed in the local- and in the marriage-interaction treatment but less so in the global-interaction treatment. Again, the data confirms this hypothesis.

Corbae and Duffy also address the question of which network will form if players can choose it themselves. They again consider groups of four subjects who first interact on an exogenously imposed network structure for an initial phase of five periods. The network structure is either a global-, a local-, or a marriage-interaction network. Subjects play the same coordination game as in the second part of the exogenous-network treatment described above. After this initial phase, subjects can freely decide with which of the other three subjects they want to interact in the subsequent periods. If two subjects agree mutually to interact, a link is formed between them.[6] For the next five periods, all subjects that are directly linked play the coordination game; this procedure is repeated four times. So far, the analysis in the paper is still preliminary. Basically, no stable network structure seems to emerge. Perhaps the main finding is that groups starting with a network of marriage interaction also show a high tendency to form the same network in later periods. Clearly, further investigation of the data but also more experimental evidence is needed.

Cassar (2002) compares convergence to equilibrium across three different network structures: a local interaction network, a random network, and a "small-world" network.[7] Small-world networks are obtained by starting from a circle and rewiring each link with probability p. A random network is a small-world network where each link is rewired with probability 1. Intuitively, small-world networks for very small p possess the nice feature of having a substantial degree of clustering, that is, a large overlap of neighborhoods, and yet only short paths connecting any two individuals in the network. In contrast, random networks have short connecting paths and low clustering, whereas the circle has high clustering but also long connecting paths. Generally, high clustering implies that interaction in the network resembles interaction in a closed group, while short connecting paths suggest that contagious behavior can spread more easily.

In the experiment, Cassar considers groups of size 18, each group being connected according to one of these three network structures. In every group, connected subjects play

[6] This assumption is based on the notion of pairwise stability defined by Jackson and Wolinsky (1996). See the section on network formation in this paper.

[7] Watts and Strogatz (1998) introduce small world networks, see also Watts (1999).

approximately 80 periods of a coordination game with conflicting risk dominance and efficiency of equilibria. The results show that subjects almost always converge to the efficient Nash equilibrium in the small-world network, while convergence is less likely (although still above 60 percent) in the other networks. Moreover, convergence to the efficient equilibrium is fastest in the small-world network, as well. These findings are consistent with the network implications suggested above.

3 Cooperation networks

3.1 Theory

Eshel, Samuelson and Shaked (1998) show that cooperation in the prisoners' dilemma game can survive if players in a population interact locally with each other and adaptation is driven by imitation of successful behavior. The crucial effect of local interaction in this model is that it allows cooperative players to cluster together. The idea is the following: since the positive externalities from cooperation are restricted locally, the interaction network reduces the possibility for other (more distant) players to exploit cooperation. In consequence, cooperators surrounded by other cooperators can earn higher payoffs than defectors primarily surrounded by other defectors. Thus, cooperation can survive if combined with imitation. Related work include Nowak and May (1992) and Kirchkamp (2000).

The interaction network of the players in these studies is exogenously given. A recent theoretical paper considering endogenous cooperation networks is Vega-Redondo (2002), which studies the formation of networks among players bilaterally involved in infinitely repeated prisoners' dilemma games. In addition to specifying which pairs of players in the population play the game, a network also determines how strategic information diffuses among the players and how players find cooperation opportunities. Assuming that payoffs in the prisoners' dilemma game fluctuate over time, Vega-Redondo analyzes the notion of pairwise-stable cooperation networks, where, intuitively, two players are directly connected with each other only if both players have an incentive to use the connection for cooperation in the prisoners' dilemma game. The main results are that players can only sustain a dense social network, that is, a network with sufficiently many individual connections, if payoff volatility is not too high. Moreover, higher payoff volatility increases the cohesiveness of the network, that is, the average distance between two players in the network declines as payoffs fluctuate more strongly.

3.2 Experiments

Although the experimental literature on cooperation in prisoners' dilemma games is extensive, only few recent papers consider the role of networks in cooperation.[8] These are Kirchkamp and Nagel (2001), Cassar (2002), and Riedl and Ule (2002).

Kirchkamp and Nagel (2001) are interested in the prediction of Eshel *et al.* (1998) showing that cooperation can be sustained by local interaction and imitation. They consider a similar experimental design to that of Keser *et al.* (1998), which consists of two treatments. In the first treatment, 18 subjects interact around a circle, each subject playing

[8] See Ledyard (1995) for a general review of cooperation experiments.

the prisoners' dilemma game with his four nearest neighbors, that is, two neighbors to the left and two neighbors to the right. In the second treatment, subjects interact in isolated groups of size five. Subjects play 80 periods in each treatment and observe the strategies and payoffs of each of their interaction partners in each period. They must use the same strategy against all partners.

In contrast to the theoretical prediction, which says that there should be more cooperation in the local interaction treatment since subjects can learn from neighbors, Kirchkamp and Nagel find that cooperation rates are higher if subjects interact in isolated groups than if they interact locally on a circle. While initial cooperation rates are close to 30 percent in both treatments, they decline to below five percent in the local interaction treatment. In contrast, cooperation rates remain at about the same level in the group treatment. This result clearly contradicts the theoretical prediction.

Cassar (2002) reports a similar decline in cooperation in the local interaction network, the small-world network, and the random network and finds no major differences between these networks.

A possible explanation for the instability of cooperation in these networks is that subjects do not learn as models of imitation assume. Indeed, Kirchkamp and Nagel find that learning is not driven by imitation of neighbors' successful strategies in either treatment, but mainly by positive reinforcement of one's own successful strategies. Hence, the main mechanism which makes cooperation survive on the circle but not in isolated groups does not seem to be at work in the laboratory. This alone, however, does not explain why subjects cooperate more in the group treatment, since the success, and hence the reinforcement of one's own cooperation strategy should be the same (namely, zero) in both treatments. Yet, the surprising and interesting finding is that the cooperation strategy increases one's own payoff in groups but not on the circle. Subjects seem to interact more in a reciprocal manner, that is, they cooperate more if others cooperate, when they interact in isolated groups than if they interact in locally overlapping groups of the same size. So far it remains unclear what precisely generates this result, whether some sort of strategic reasoning plays a role, as suggested by Kirchkamp and Nagel, or if other more psychological factors are involved.

The experiment of Riedl and Ule (2002) considers a different question: endogenous network formation when players play a repeated prisoners' dilemma game, such as in the model of Vega-Redondo (2002). Riedl and Ule use groups of six. The network structure is exogenously fixed to the complete network in their control treatment, that is, each subject plays the prisoners' dilemma game against every other group member. Subjects decide whether to form links with other group members in the other treatments. A link is established if both parties agree to form one (pairwise stability, cf. Jackson and Wolinsky, 1996). Links are cost free and each pair that is linked plays the prisoners' dilemma game. If some party rejects a link both parties earn an outside-option payoff. Subjects have to choose the same strategy against all partners. Every treatment consists of 60 periods of play.

The analysis presented in Riedl and Ule (2002) is still preliminary. However, some interesting results can already be summarized. First, cooperation rates are significantly higher in the endogenous-network treatments compared to the exogenous-network treatment. While initial cooperation rates equal or exceed fifty percent in all treatments, cooperation declines if the network is fixed, whereas cooperation remains stable almost until the end if subjects can choose their partners themselves. Cooperation in the endogenous-network treatments only decreases in the final five periods. Second, cooperation rates are highest if the outside-option payoff lies between the Nash-equilibrium and the cooperation payoff in the prisoners' dilemma game and if subjects can observe the strategies of all other subjects in the group.[9] Assuming that some of the subjects are reciprocators, who cooperate if others cooperate as well, this suggests that the value of the outside option might serve as a signaling device for cooperative play. Third, cooperators are more likely to propose links to subjects who cooperated in the previous period than to previous-period defectors. This holds even if the outside-option payoff is lower than the Nash-equilibrium payoff, that is, when exclusion of defectors is costly. This result calls an interesting finding of Ehrhart and Keser (1999) to mind, who study a public goods experiment where subjects have the possibility of changing groups. They observe that the more cooperative subjects are on the run from the less cooperative subjects who follow the former around.

The results of Riedl and Ule clearly show that the fact whether the interaction network is exogenous or endogenous plays an important role in the degree of cooperation in prisoners' dilemma like situations. In particular, the possibility of excluding players who defect, even if this is costly, turns out to be a powerful instrument with striking effects on behavior. It is illuminating to compare this finding to a recent experiment by Brown, Falk, and Fehr (2002), who do not study network formation explicitly, but allow subjects to form bilateral relations in a two-person incomplete contract setting. They also find that the threat to terminate a relation represents a powerful discipline device, which induces partners to cooperate and consequently enhances efficiency.

4 Buyer-seller networks

4.1 Theory

Buyer-seller networks represent an area where both theoretical and experimental research in economics has been recently initiated.

Kranton and Minehart (2001) examine the individual motives of buyers and sellers in forming particular networks, as documented in the Japanese electronics industry (Nishigushi 1994) and the Italian garment industry (Lazerson 1993), for example. In particular, they ask what might drive buyers and sellers to establish links to multiple trading partners, and whether these networks can be expected to be efficient.

To answer these questions, they consider a number of buyers and sellers, where each seller has an indivisible object for sale and buyers have i.i.d. random valuations for the object. A buyer can purchase from a seller if and only if the two are linked. The buyers establish the links and face costs for these links.[10] Transactions and prices are determined by an English, that is, ascending-bid, auction. Buyers drop out of the bidding as the price exceeds their valuation. This process continues until demand equals supply. Kranton and Minehart show that competition generates an efficient allocation of goods in the network. Furthermore, prices reflect the link pattern in the sense that a buyer's profit equals the

[9] Hauk and Nagel (2001) find a similar increase in cooperation rates in a related prisoners' dilemma experiment with the possibility of partner choice. In this experiment, however, no networks are considered.
[10] Jackson (2001) generalizes the model in such a way that links are costly to buyers and to sellers.

marginal social value of his participation in the network. From this it follows that efficient network structures are always an equilibrium outcome.

The model of Kranton and Minehart (2001) emphasizes two reasons why buyer-seller networks may emerge, one economic, the other strategic. First, networks may allow buyers and sellers to pool uncertainty in demand, which buyers' random valuations cause in the present model. Second, a trader's multiple links can enhance his competitive position.

Network consequences on competition are also addressed in the network model of Corominas-Bosch (1999). However, in contrast to the model of Kranton and Minehart (2001), a bargaining process rather than an English auction determines prices in this model, and a buyer's valuation of the seller's good is certain. Again, a link between two trading partners is necessary for possible transaction. Consequently, if an individual has several links, he has several possible partners he can trade with. Thus, the network structure directly determines the individual players' bargaining power.

Bargaining follows a variation of the (infinite-horizon) Rubinstein alternating-offer protocol. Each seller calls out a price in the first period. Buyers then simultaneously choose to accept at most one of the prices offered by a seller to whom they are linked, with ties broken randomly. If a buyer and a seller trade, their links are removed from the network. The situation reverses in the next period and buyers call out prices, which the sellers connected to them then accept or reject. This process repeats itself until all remaining buyers and sellers are no longer linked. Future periods are discounted according to a common discount factor.

Corominas-Bosch shows that, depending on the given network structure, the subgame-perfect equilibrium of the bargaining game has the following properties. If the network is "competitive", the short side of the market receives all surplus. For example, if there is only one buyer who is linked to two sellers, competition between the sellers will reduce the price to the extent that the buyer extracts all surplus. This is reversed for a single seller who is linked to two buyers. If, however, the network is "even", that is, the number of buyers and sellers linked to each other is the same, traders split the surplus evenly. In the case of a single seller linked to a single buyer, this result corresponds exactly to the Rubinstein bargaining model. Corominas-Bosch then shows that any network can be decomposed into a union of subnetworks that are either competitive or even, plus some extra links. Furthermore, subgame-perfect equilibrium outcomes of the bargaining game are such that all the surplus is given to the short side of the market in every competitive subnetwork, while the surplus is divided evenly in every even subnetwork.

4.2 Experiments

The model of Corominas-Bosch (1999) was recently tested by Charness, Corominas-Bosch, and Frechette (2001), who are particularly interested in the predictive power of the subgame-perfect equilibrium outcome in competitive versus even buyer-seller networks. They extend the infinite-horizon Corominas-Bosch model to a finite-horizon game. In the initial phase of their experiment, subjects interact on one of two different network structures: a three-player network, where a single buyer is connected to two sellers, or a four-player network, where two buyers and two sellers are each connected to each other. The three-player network is competitive whereas the four-player network is even. Subjects are engaged in alternating bargaining over five to six rounds in both networks. Sellers start offering a proposal to divide 2500 with any of their linked buyers, who then simultaneously decide to accept at most one of the proposals. If a transaction is made, the corresponding pair of a buyer and a seller is removed from the network. The game proceeds to the second round if any links remain. In the second round buyers offer a division of 2400 to all of their linked sellers, who choose to accept or not to accept any of these offers. If needed, a third round is implemented, etc.. A coin is flipped to determine whether bargaining ends after five or six rounds. All unmatched subjects receive 200. This bargaining protocol is repeated for four periods, with subjects' roles reversed after each period.

The second phase of the experiment begins after the fourth period, where the two separate networks are merged to obtain a single seven-player network. This is done by adding a link either from the short (that is, buyer) side or from the long (that is, seller) side of the competitive network to the corresponding other side of the even network, respectively. If the two networks are connected via the short side of the competitive network, theory predicts that equilibrium-outcomes remain the same for each player. If, however, the two networks are connected via the long side of the competitive network, the new seven-player network becomes competitive. Consequently, equilibrium-outcome predictions change for all subjects from the former even subnetwork. The second phase of the experiment consists of six periods of bargaining, each containing again up to six alternating rounds.

Charness *et al.* find that while exact equilibrium-point predictions fail, the qualitative behavior in the experiment is consistent with theoretical predictions. For example, seller's payoffs are lower in the competitive three-player network than in the even four-player network. Hence, buyers seem to make successful use of their bargaining power. Furthermore, the way in which the two subnetworks are connected with each other has a significant effect on the bargaining outcomes realized in subsequent periods. Theory predicts that payoffs between buyers and sellers in the four-player subnetwork should diverge if the large network becomes competitive, while they should be roughly the same if the subnetwork remains even. The data clearly confirms this prediction.

The fact that any precise equilibrium prediction fails may be due to insufficient time for the subjects to learn. However, as Charness *et al.* show, allocations in the past strongly affect the willingness to offer or accept certain shares. Subjects learn from each other and develop a social norm for appropriate bargaining outcomes. This suggests that learning may not necessarily lead in the direction the theory predicts. Finally, the fact that bargaining outcomes tend to be less extreme than predicted by standard economic theory is consistent with extensive experimental research on ultimatum and dictator games (see Roth, 1995; Camerer and Thaler, 1995). Observations in these games can be explained by the assumption that considerations for fairness and equity drive players' behavior (e.g., Fehr and Schmidt, 1999; Bolton and Ockenfels, 2000). It seems obvious that fairness considerations may also play a role in the present experiment of a buyer-seller network. However, the precise interaction between network structure, bargaining power, and fairness considerations still needs to be studied (with this regard see also the related experimental work by Roth, Prasnikar, Okuno-Fujiwara and Zamir, 1991; and Fischbacher, Fong and Fehr 2003).

Market design experiments for network industries, like, e.g., water, electric power, and passenger rail service, represent a different literature of experiments with buyer-seller networks. Recent work in this area is Murphy, Dinar, Howitt, Rassenti, and Smith (2000), Rassenti, Smith and Wilson (2003) and Cox, Offerman, Olson, and Schram (2002). These papers use laboratory experiments as a "wind tunnel" for the actual design of market institutions in various network industries. In particular, they show how laboratory experiments can not only be used to answer scientific economic questions but also to produce relevant insights for important real-life problems, like the practical design of market institutions.[11]

Cox *et al.* (2002), for example, present a study that the Dutch Ministry of Transport explicitly commissioned. In their experiment, the authors compare two possibilities for the privatization of passenger rail service in the Netherlands. The first possibility ("competition for the rails") allows different operators to bid in an auction for the monopoly right to provide passenger rail service in a particular region for a limited time. Service provided by the winning operator has to meet at least a minimum schedule the government defines. Operators bid passenger-ticket prices to be charged and the operator willing to charge the lowest price receives the monopoly right. With regard to the second possibility ("competition on the rails"), the government first distinguishes a number of individual route/time slots (e.g., Rotterdam-Amsterdam at 7:10 am) and then allocates the rights to these slots using a combinatorial auction. In this case, there is no minimum schedule. Cox *et al.* compare the two competition frameworks in a laboratory experiment, where subjects act in the roles of different operators who interact in any of the auctions described above. The network in the laboratory, though obviously much simpler, maps the basic structure of the Dutch railway network. The authors derive several results from their experiment; most of them are directly related to particular questions raised by the Dutch government. For example, they find that rail transportation prices are significantly lower under competition for the rails than under competition on the rails. Next, the number of transported passengers is higher under competition for the rails and economic efficiency, measured by government income plus consumer and producer surplus, is larger as well. If any of these criteria are the objectives of the government, the results of the study clearly favor competition for the rails. In fact, the authors write that their research "led to a recommendation (of the Ministry) to introduce competition in a 'for the rails' design instead of an 'on the rails' design" and a corresponding law "that is currently under consideration by Parliament" (Cox *et al.* 2002, p732).

Murphy *et al.* (2000) report laboratory experiments designed to test different institutional arrangements for so-called "smart", that is, computer-assisted water markets. They use a sealed-bid uniform price double auction for the allocation of water and transportation capacity rights among different buyers, sellers, and transporters. Water markets in their study are based on the California water network. Three measures are considered: efficiency, distribution of surplus, and price volatility. The authors find that although prices do not follow the (exogenously determined) changes in water supply very well, smart water markets are able to generate highly efficient outcomes (about 90 percent efficiency on average). In addition, competitive cotenancy, that is, shared property rights for transportation capacity, generates a transfer of surplus from sellers to buyers in the experiment.

Finally, Rassenti *et al.* (2003) compare different market institutions for electricity networks. They introduce demand-side bidding in an experimental market for electricity,

[11] Due to the nature of the problem, most of these experiments are very complex and involve distinctive features that depend on the particular national situation. In addition, the focus differs clearly from that of most other papers reviewed in this article, which take a more theoretical standpoint. I therefore give only a short introduction into this literature and leave details, in particular with regard to the experimental design, to another, more thorough, discussion.

comparing a situation with supply-side market power to one with no supply-side market power. Their main finding is that demand-side bidding completely neutralizes the exercise of market power and eliminates price spikes. The study thus shows that the empowering of wholesale buyers in electricity markets offers an effective decentralized approach to control supply-side market power and price volatility.

5 Network formation

5.1 Theory

Next to implications on coordination, cooperation, and competition, one of the most important questions is how networks emerge. In recent years, several theoretical approaches have been proposed to address this question, using techniques from cooperative and non-cooperative game theory.[12]

The paper of Myerson (1977) is, probably, one of the first important contributions to this literature. Myerson analyzes a cooperative game that is enriched by a network structure describing the possibilities for communication or cooperation among different players. Individuals can act as a coalition if and only if they are connected through links in the network. While this idea constitutes an important step forward, it leaves several issues unsolved. Because the value function is still defined on coalitions and not on the network directly, the theory does not distinguish between different networks that connect the same players but differ in the way these players are connected. In consequence, many interesting details of the network formation process, for example costs and benefits of particular links, can not be analyzed in the model.

Jackson and Wolinsky (1996) follow a different approach considering value functions that are defined on networks directly. The main issue they address is the conflict between efficiency, that is, value maximization, and stability. With regard to the latter, they analyze the notion of pairwise stability, which assumes that links are formed if and only if both players connected by a link actually agree to form the link. On the other hand, links are severed if either of the two individuals decides to do so.

A particular network model they consider is the so-called connections model, where individuals receive benefits from being connected to other individuals and bear costs for maintaining direct links. Jackson and Wolinsky show that the set of efficient networks reduces to only three different types of networks: the complete network if costs are low, the star network if costs are intermediate, and the empty network if costs are high. While the complete and the empty network are pairwise stable if they are efficient, the star network may fail to be pairwise stable. As Jackson and Wolinsky show, this conflict between efficiency on the one hand and pairwise stability on the other, is not a unique feature of the connections model but extends to more general network settings, as well.

The work of Myerson (1977) and Jackson and Wolinsky (1996) has attracted much interest in economic models of network formation. Subsequent models include Aumann and Myerson (1988), Dutta and Mutuswami (1997), Dutta, van den Nouweland, and Tijs (1998), Slikker and van den Nouweland (2000) and Johnson and Gilles (2000). All these models study the formation of networks in a static setting. Watts (2001) departs from this

[12] See Jackson (2003) for a detailed survey of the theoretical literature on network formation.

assumption and analyzes the connections model in a dynamic setting, where individuals meet over time and decide to form or sever links between each other. Similarly, Jackson and Watts (2001) consider the evolution of more general network models.

Somewhat parallel to the literature above with clear origins in cooperative game theory, Bala and Goyal (2000a,b) develop models of network formation that use tools from non-cooperative game theory. Rather than considering pairwise stability, Bala and Goyal assume that individuals can form and sever links unilaterally, that is, in particular no mutual consent is needed to form a link between two individuals. Clearly, this assumption changes the incentives of the players; hence the analysis in Bala and Goyal (2000a,b) differs substantially from the analysis in the models mentioned above. A central implication of unilateral link formation is that with regard to stability it leads to the concept of Nash equilibrium. Papers that follow a similar methodological approach include Goyal and Moraga-Gonzalez (2001), Goyal and Joshi (2002), Haller and Sarangi (2001) and Sarangi, Kannan, and Ray (2003).

The main idea of the network model in Bala and Goyal (2000a) is similar to the connections model of Jackson and Wolinsky (1996): players earn benefits from being connected to other players and bear costs for maintaining direct links. Benefits result from valuable, non-rival information that flows through the network. Bala and Goyal distinguish between two different scenarios of information flow. In the first scenario (the 1-way flow model), information only flows to the player who maintains the link. In the second scenario (the 2-way flow model), information flows both ways. Independent of the information flow, Bala and Goyal assume that players simultaneously decide with whom to form a direct link, a link being costly to the individual who forms it.

Assuming that information flows through the network with no decay, Bala and Goyal prove that Nash equilibria of the network-formation game are the following.[13] In the 1-way flow model a Nash equilibrium is either the empty network, where no player maintains any connection to any other player, or minimally connected, that is, it has a unique component that splits if one link is severed. Analogously, a Nash equilibrium in the 2-way flow model is either the empty network or minimally 2-way connected, that is, it has a unique component, no cycle, and no two individuals both maintain a link with each other.[14] Intuitively, a network is Nash in both models if (i) either none or all players are connected and (ii) no redundant links are maintained.

Depending on the number of players, the number of Nash (equilibrium) networks can be quite large. Therefore, a reasonable refinement to be considered is the notion of strict Nash equilibrium, where each player plays his unique best response to the strategy profile of the other players. As it turns out, the set of strict Nash networks is much more restrictive. Bala and Goyal show in the 1-way flow model that the only strict Nash networks are the empty network and the circle (or, as Bala and Goyal call it, the wheel). Only the empty network and the center-sponsored star are strict Nash networks in the 2-way flow model. The center-sponsored star is the network, where one individual (the center) maintains a direct link to every other individual, and no other individual maintains any link.

[13] Bala and Goyal prove also results for the general case, where they allow for decay. However, results are more clear-cut if no decay exists.

[14] A cycle is a n-player network with $n \geq 3$ that connects the following pairs of players each via a direct link: (12), (23), ..., (n-1n), (n1).

Both the circle and the center-sponsored star are shown to be efficient networks, where efficiency is defined in terms of maximizing the sum of players' payoffs. Thus, contrary to the model of Jackson and Wolinsky (1996), no conflict between efficiency and stability exists in the model of Bala and Goyal (2000a). This suggests that the circle and the center-sponsored star may serve as a powerful prediction for network formation.

5.2 Experiments

Recent economic experiments that consider network formation are Deck and Johnson (2002), Callander and Plott (2003), and Falk and Kosfeld (2003). Vanin (2002) presents a pilot study for the model of Jackson and Wolinsky (1996). Three groups of four subjects each collectively bargain about what network to form in his experiment. Links between two subjects are established based on pairwise stability, that is, bilateral agreement is required. Each group forms a network in three different scenarios: the connections model with and without side payments and the co-author model.[15] Vanin finds that groups tend to form pairwise unstable but efficient networks both in the co-author model and, two out of three times, also in the connections model with side payments. In the connections model without side payments, groups form networks that are inefficient but equalize payoffs among the subjects.

The experimental study of Deck and Johnson (2002) is inspired by the network-formation model of Johnson and Gilles (2000), which introduces a spatial cost topology in the connections model of Jackson and Wolinsky (1996). In this model, players are located on a line and the cost for a direct connection between two players monotonically increases with the distance between the two players. Deck and Johnson examine a comparison of three different institutions for network formation in this model, and consequently implement one treatment for each institution. In the first institution, called "Split", players select those direct links for which they are willing to pay exactly half of the cost. If both players agree to pay half of the cost of their connecting link, the link is formed. The second institution, "Primary", allows players to bid between zero and the total cost of a link for all their direct links. Finally, players can bid for all possible links in the institution "Secondary". In particular, they can bid for a link between any two of the other players. Under the latter institutions, a link is formed if and only if the total sum of players' bids for that link exceeds the cost of the link. Deck and Johnson use the notion of Nash equilibrium as a stability concept for a network.

In each treatment of the experiment subjects interact consecutively in three different payoff environments, two of them involving groups of five players while players interact in pairs in the third environment. In the following, I focus on the five player environments only. A special feature of the experimental design of Deck and Johnson is that the situation is described to the subjects as a decision task faced by managers of a train station. In particular, a subject in the experiment is in the role of a station manager, who has to decide (that is, to bid) on the connections between different stations. Just like in the connections model, the minimum number of stops a fictitious passenger would have to make on his trips from the manager's station to the other stations determines the value of a network to a station manager. The parameterization in the first payoff environment is such that the unique efficient network requires each player to be directly connected to his nearest and to

[15] The co-author model assumes that an additional link generates a negative externality on individuals already connected. Jackson and Wolinsky (1996) show that pairwise stable networks are generally over-connected and hence inefficient in this model.

his second nearest neighbor(s). This network can be supported by a Nash equilibrium under all three institutions. Parameters in the second payoff environment are chosen such that the chain is the unique efficient network, where each player is only connected to his nearest neighbor. The chain can only be supported by a Nash equilibrium under the Primary and the Secondary institution, but not under the Split institution.

Subjects play 15 periods of the first environment, followed by 10 periods of the second environment. The results reported in the paper are based on the last two-thirds of all periods in each environment. Deck and Johnson find that the Primary institution performs best in the first environment, yielding an average level of efficiency of 89 percent while Split and Secondary, on average, achieve an efficiency of 83 percent and 81 percent, respectively. The authors explain their result as follows: while subjects form too many of the long links, that is, connecting stations that are far away from each other in all treatments, subjects are more successful in forming the necessary short links in the Primary treatment. All three institutions show a similarly poor performance in the second environment. No group achieves a positive surplus under any institution. Note that the chain is the unique efficient network in this environment and indirect links provide a significant percentage of the network's value. Since subjects do not succeed in coordinating on this network, the authors conclude that network formation driven by individual decision-making is a difficult process that is likely to produce inefficient outcomes. However, the individual strategy of forming a chain in the analyzed environment is a high-risk strategy that is likely to generate a negative payoff to each player if players do not successfully coordinate. It remains to be shown whether the individual decision-making itself or rather the riskiness of the efficient equilibrium-strategy leads to the coordination failure. In fact, the studies of Callander and Plott (2003) and Falk and Kosfeld (2003) show that if this riskiness is reduced, individuals are well able to coordinate on an efficient network. These studies are discussed in the remainder of the paper.

Callander and Plott (2003) report on – as they call it – an "exploratory" experiment in which they study the evolution of information networks under various treatment conditions. They consider groups of six players in each treatment. The first treatment involves between 10 and 20 periods of a 1-way flow network-formation game à la Bala and Goyal (2000a), using a random stopping rule to determine the end of the experiment. In this experiment, subjects sit together in a room. In each period, after each subject has recorded his direct connections payoffs in the resulting network are calculated. Precisely, different physical signs are placed in front of each subject; these signs represent individual links and are used to determine subjects' profits. Costs and benefits in this treatment are such that the unique strict Nash network prediction is the circle, which is also efficient. In the second treatment that Callander and Plott study, individual decisions are made via the computer. Moreover, whereas the subjects' decision making is simultaneous in the first treatment, decisions are made continuously over two minute rounds in the second treatment, where choices can be adjusted repeatedly in real time. Subjects are continuously updated about the choices of the other group members and each choice adjustment is costly. Each session lasts between 15 and 20 rounds, using again a random stopping rule to determine the end of a session. Different games are played in this treatment. Each session starts with a 1-way flow network-formation game with a slight variation in benefits stemming from the game played in the first treatment but the same theoretical prediction (that is, the circle). In case subjects form the same network for three consecutive periods, the game is changed in such a way that from then on direct links to nearest neighbors are twice as costly as links to other subjects.[16] This modification has the effect that only particular circles, where no nearest neighbors are directly connected, are efficient. If subjects form again the same network for three consecutive periods under the new setting, the game is changed yet another time. In particular, one subject is selected as the recipient of free direct links, meaning that links to and from him do not bear any cost. This renders the circle inefficient and predicts a star network as the unique efficient (and strict Nash) network. Subjects are not informed about the potential change of the games at the beginning of the experiment.

The main findings of Callander and Plott are as follows: first, networks often converge to Nash equilibrium in both treatments. If convergence is observed, subjects form the efficient and strict Nash network, which is the circle, in all cases except for one group.[17] Second, decision making with real-time choice adjustment facilitates the coordination on a Nash network. While two out of five groups eventually converge to Nash in the first treatment, six out of seven groups eventually form a Nash network in the second treatment. At the individual level, Callander and Plott reject the hypothesis that subjects adjust their strategies based on a best-response rule with an additional error term. Instead, the authors conclude that many subjects seem to exhibit so-called "simple strategic behavior", meaning that subjects form exactly one direct link that is part of a focal (e.g., a clockwise) circle network.

The experiment of Callander and Plott (2003) is inspired by the network-formation model of Bala and Goyal (2000a). The authors, however, only consider the 1-way flow model and in this model only those treatments where the circle is the unique efficient strict Nash network. Falk and Kosfeld (2003) present a general analysis of the Bala-Goyal model, where both 1-way and 2-way flow networks are studied and several treatment conditions are implemented yielding different theoretical predictions.

Falk and Kosfeld (2003) consider groups of four in their experiment. Subjects decide anonymously and independently in every period with which other group members they want to form a direct link. After decisions are made, subjects learn the realized network and earn a payoff that depends on the number of links they form and the number of group members to whom they are connected in that period. Groups stay together for five periods, after which they are randomly recomposed. Overall, subjects play 15 periods of the network formation game and participate in three different groups.

The experiment consists of five treatments, three 1-way flow and two 2-way flow treatments. The benefit from an additional (direct or indirect) connection is equal to 10 in every treatment. The cost of a direct link is either 5, 15, or 25 in the 1-way flow model, while it amounts to 5 or 15 in the 2-way flow model. The circle is a strict Nash network in all 1-way flow treatments. It is the unique strict Nash network when costs of a direct link equal 5, while the empty network is also a strict Nash network when costs equal 15 or 25. The circle is the unique efficient network in all 1-way flow treatments. In the 2-way flow model, the center-sponsored star is the unique strict Nash network if costs equal 5 and the empty network is the unique strict Nash network if costs equal 15. In both treatments, efficient networks are those that are minimally 2-way connected. This includes the center-sponsored star but not the empty network.

[16] This is the opposite of what is assumed in the model of Johnson and Gilles (2000).

[17] Only one group participated in the final parameter setting of the second treatment. No convergence was observed in this case.

Falk and Kosfeld (2003) derive several results. Their main finding is the following: while more than fifty percent of the networks formed by the subjects are strict Nash networks (that is, either the circle or the empty network) in the 1-way flow model, no strict Nash network (that is, neither the center-sponsored star nor the empty network) ever arises in the 2-way flow model. Hence, there is a significant difference between the two scenarios of information flow. Whereas the notion of strict Nash equilibrium serves as a good prediction in the 1-way flow scenario, it has no predictive power at all in the 2-way flow scenario. This difference is no artifact of the refinement of strict Nash equilibrium, but remains true if we look at all Nash networks instead. Although subjects form Nash networks in the 2-way flow model, they do so significantly less than in the 1-way flow model.

What explanation can be offered for this finding? One reason might be that the 2-way flow model is more complex and that it might take more time for the subjects to coordinate on equilibrium. However, simulations of Bala and Goyal (2000a) suggest that convergence times are similar in both models if subjects have the chance to revise their strategy in every period, as is the case in the experiment. Moreover, following the complexity argument, one should not see any difference in behavior once an equilibrium is reached. However, Falk and Kosfeld do find this difference in the data. Given that subjects form a Nash network in period t, the probability of forming a Nash network in period $t+1$ is significantly smaller in the 2-way flow model than in the 1-way flow model. Thus, even if subjects succeed in coordinating on a Nash network, equilibria are less stable in the 2-way flow model.

Another explanation might be that fairness considerations drive the subjects' behavior in the experiment. Note that a crucial difference between the 1-way and the 2-flow model is that every subject has to form a link in the first model if he wants to receive any benefits from the network. In contrast, a subject can earn large benefits in the 2-way flow model even if he does not form a link himself. As a consequence, payoffs are very unequal in equilibrium. For example, the individual in the center of center-sponsored star, who maintains all links, earns a payoff of 25, while individuals on the periphery earn a payoff of 40. If subjects dislike unequal payoffs, they may be unwilling to form such networks.

In Falk and Kosfeld (2003) this argument is tested by a probit regression measuring the impact of payoff inequity on inertia. The latter is defined as any instance where a subject plays the same strategy as in the previous period. Payoff inequity is measured by the sum of absolute payoff differences between own and the other players' payoff. Controlling for best replies, it is found that payoff inequity has a significant negative impact on inertia. The higher the payoff inequity between subjects in a given network is, the less willing subjects are to maintain that network, even if (according to monetary rewards) the maintenance of the network is an individual best reply. As Nash networks in the 2-way flow model generate higher payoff inequity than in the 1-way flow model, this implies that subjects form fewer Nash networks in the 2-way compared to the 1-way flow model, which the experimental data confirms.

One conclusion from this result is that networks have to be fairness compatible in order to be stable. Otherwise, the notion of Nash equilibrium together with standard money-maximizing preferences generates incorrect predictions. A second implication is that mechanisms that help overcome the conflict between fairness and stability (e.g., compensation of the central player or rotation within the network) may play an important role in the formation and maintenance of social networks. Subjects did not have such possibilities in the experiment. In reality, however, there are many of these possibilities, and, as evidence from sociology, psychology, and anthropology suggests, these mechanisms are widely used. Clearly, these mechanisms also provide an interesting area for further experimental research.

6 Conclusion

The experiments discussed in this paper belong to a recent wave in experimental economics focusing on social and economic networks. Present work emphasizes individual incentives for network formation, as well as the impact of networks on equilibrium selection, competition, and cooperation. Given the interest in the topic and the existing theoretical literature in this area, it seems clear that more experimental studies are on the way.

So far, the research has shown that networks and the formation of networks by individual decision-making represent an interesting and important research topic that generally contains a very high degree of complexity. In this situation, laboratory experiments offer a useful method to reduce some of this complexity, because important elements concerning the institutional environment can be controlled explicitly and possible effects of these elements can be studied using a clear ceteris-paribus approach.

While it is, probably, too early to claim general agreement on individual aspects, the results of the experiments so far have shown that network configurations have important effects on economic outcomes, such as the convergence towards equilibria, the support of Pareto superior states, or the distribution of surplus among economic agents. Future experiments will, hopefully, build on the present work generating a more detailed understanding of these effects and the underlying mechanisms that generate them.

7 References

Aumann, R. and R. Myerson (1988) "Endogenous Formation of Links Between Players and of Coalitions: an Application of the Shapley Value," in Roth, A.E., ed., *The Shapley Value,* 175-191. Cambridge University Press.

Bala, V. and S. Goyal (2000a) "A Noncooperative Model of Network Formation," *Econometrica*, 68: 1181-1229.

Bala, V. and S. Goyal (2000b) "A Strategic Analysis of Network Reliability," *Review of Economic Design*, 5: 205-228.

Bavelas, A. (1950) "Communication Patterns in Task-Oriented Groups," *Journal of the Acoustical Society of America*, 22: 725-730.

Berninghaus, S.K., K.-M. Ehrhart, and C. Keser (2002) "Conventions and Local Interaction Structures: Experimental Evidence," *Games and Economic Behavior*, 39: 177-205.

Bertrand, M., E.F.P. Luttmer, and S. Mullainathan (2000) "Network Effects and Welfare Cultures," *Quarterly Journal of Economics*, 115: 1019-1055.

Bhaskar, V. and Vega-Redondo, F. (2002) "Migration and the Evolution of Conventions," *Journal of Economic Behavior and Organization*, forthcoming.

Bienenstock, E.J. and P. Bonacich (1993) "Game-Theory Models for Exchange Networks: Experimental Results," *Sociological Perspectives*, 36: 117-135.

Bienenstock, E.J. and P. Bonacich (1997) "Network Exchange as a Cooperative Game," *Rationality and Society*, 9: 37-65.

Blume, L.E. (1993) "The Statistical Mechanics of Strategic Interaction," *Games and Economic Behavior*, 5: 387-424.

Bolton, G.E. and A. Ockenfels (2000) "ERC – A Theory of Equity, Reciprocity and Competition," *American Economic Review*, 90: 166-193.

Boun My, K., Willinger, M., and Ziegelmayer, A. (2001) "Global versus Local Interaction in Coordination Games: An Experimental Investigation," mimeo, Université Louis Pasteur.

Brown, M., A. Falk, and E. Fehr (2002) "Relational Contracts and the Nature of Market Interactions," *Econometrica*, forthcoming.

Callander, S. and C.R. Plott (2003) "Networks: An Experimental Study," mimeo, Northwestern University.

Camerer, C. and R. Thaler (1995) "Ultimatums, Dictators, and Manners," *Journal of Economic Perspectives*, 9: 209-219.

Cassar, A. (2002) "Coordination and Cooperation in Local, Random and Small World Networks: Experimental Evidence," in ed. D.K. Levine, W. Zame, L. Ausubel, P.A. Chiappori, B. Ellickson, A. Rubinstein and L. Samuelson in *Proceedings of the 2002 North American Summer Meetings of the Econometric Society: Game Theory*, http://lev0201.dklevine.com/proceedings/game-theory.htm.

Charness G., M. Corominas-Bosch, and G.R. Frechette (2001) "Bargaining and Network Structure: An Experiment," mimeo, University of California, Santa Barbara.

Cook, K.S. and R.M. Emerson (1978) "Power, Equity, and Commitment in Exchange Networks," *American Sociological Review*, 43: 721-739.

Cook, K.S., R.M. Emerson, M.R. Gillmore, and T. Yamagishi (1983) "The Distribution of Power in Exchange Networks: Theory and Experimental Results," *American Journal of Sociology*, 89: 275-305.

Corbae, D. and J. Duffy (2002) "Experiments with Network Economies," mimeo, University of Pittsburgh.

Corominas-Bosch, M. (1999) "On Two-sided Network Markets," PhD Thesis, Universitat Pompeu Fabra.

Cox, J.C., Offerman, T., Olson, M.A., and Schram, A.J.H.C. (2002) “Competition for versus on the Rails: A Laboratory Experiment,” *International Economic Review*, 43: 709-736.

Deck, C. and C. Johnson (2002) “Link Bidding in a Laboratory Experiment,” *Review of Economic Design*, forthcoming.

Droste, E., R.P. Gilles, and C. Johnson (2000) “Evolution of Conventions in Endogenous Social Networks,” mimeo, Virginia Tech.

Dutta, B. and S. Mutuswami (1997) “Stable Networks,” *Journal of Economic Theory*, 76: 322-344.

Dutta, B., A. van den Nouweland, and S. Tijs (1998) “Link Formation in Cooperative Structures,” *International Journal of Game Theory*, 27: 245-256.

Ehrhart, K.-M. and C. Keser (1999) “Mobility and Cooperation: On the Run,” mimeo, CIRANO.

Ellison, G. (1993) “Learning, Local Interaction, and Coordination,” *Econometrica*, 61: 1047-1071.

Ely, J.C. (2002) “Local Conventions,” *Advances in Theoretical Economics*, 2.

Eshel, I., L. Samuelson, and A. Shaked (1998) “Altruists, Egoists, and Hooligans in a Local Interaction Model,” *American Economic Review*, 88: 157-179.

Fafchamps, M. and S. Lund (2001) “Risk-Sharing Networks in Rural Philippines,” mimeo, Oxford University.

Falk, A. and M. Kosfeld (2003) “It’s all about Connections: Evidence on Network Formation,” mimeo, University of Zurich.

Fehr, E. and K.M. Schmidt (1999) “A Theory of Fairness, Competition and Cooperation,” *Quarterly Journal of Economics*, 114: 817-868.

Fischbacher, U., C.M. Fong, and E. Fehr (2003) “Fairness and the Power of Competition,” mimeo, University of Zurich.

Goyal, S. and S. Joshi (2002) “Networks of Collaboration in Oligopoly,” mimeo, Queen Mary, University of London.

Goyal, S. and J.L. Moraga-González (2001) “R&D Networks,” *RAND Journal of Economics*, 32: 686-707.

Goyal, S. and F. Vega-Redondo (2000) “Learning, Network Formation and Coordination,” mimeo, University of Alicante.

Granovetter, M. (1974) *Getting a Job*, University of Chicago Press, Chicago.

Haller, H. and S. Sarangi (2001) “Nash Networks with Heterogeneous Agents,” mimeo, Virginia Polytechnic Institute.

Hauk, E. and R. Nagel (2001) “Choice of Partners in Multiple Two-Person Prisoner’s Dilemma Games,” *Journal of Conflict Resolution*, 45: 770-793.

Holzer, H. (1987) “Hiring Procedures in the Firm: Their Economic Determinants and Outcomes,” NBER Working Paper No. 2185.

Jackson, M.O. (2001) “The Stability and Efficiency of Economic and Social Networks,” in eds. Sertel, M. and S. Koray, *Advances in Economic Design*, forthcoming Springer, Heidelberg.

Jackson, M.O. (2003) “A Survey of Models of Network Formation: Stability and Efficiency,” mimeo, California Institute of Technology.

Jackson, M.O. and A. Watts (2001) “The Evolution of Social and Economic Networks,” *Journal of Economic Theory*, forthcoming.

Jackson, M.O. and A. Wolinsky (1996) “A Strategic Model of Social and Economic Networks,” *Journal of Economic Theory*, 71: 44-74.

Johnson, C. and R.P. Gilles (2000) “Spatial Social Networks,” *Review of Economic Design*, 5: 273-300.

Kandori, M., G.J., Mailath, and R. Rob (1993) “Learning, Mutation and Long Run Equilibria in Games,” *Econometrica*, 61: 29-56.

Keser, C., K.-M. Ehrhart, and S.K. Berninghaus (1998) “Coordination and Local Interaction: Experimental Evidence,” *Economics Letters*, 58: 269-275.

Kirchkamp, O. (2000) “Spatial Evolution of Automata in the Prisoners’ Dilemma,” *Journal of Economic Behaviour and Organization*, 43: 239-262.

Kirchkamp, O. and R. Nagel (2001) “Local and Group Interaction in Prisoners’ Dilemmas,” mimeo, University of Mannheim.

Kosfeld, M. (2002) “Stochastic Strategy Adjustment in Coordination Games,” *Economic Theory*, 20: 321-339.

Kranton, R. and D. Minehart (2001) “A Theory of Buyer-Seller Networks,” *American Economic Review*, 91: 485-508.

Lazerson, M. (1993) “Factory or Putting-Out? Knitting Networks in Modena,” in Grabher G., ed., *The Embedded Firm: On the Socioeconomics of Industrial Networks*, 203-226. New York: Routledge.

Leavitt, H.J. (1951) “Some Effects of Certain Communication Patterns on Group Performance,” *Journal of Abnormal and Social Psychology*, 46: 38-50.

Ledyard, J.O. (1995) "Public Goods: A Survey of Experimental Research," in Kagel, J.H. and A.E. Roth, eds., *Handbook of Experimental Economics*, 111-194. Princeton University Press.

Lévi-Strauss, C. (1963) *Structural Anthropology*. New York: Basic Books.

Mailath, G., L. Samuelson, and A. Shaked (2001) "Endogenous Interactions," in eds. Nicita, A. and U. Pagano, *The Evolution of Economic Diversity*, 300-324. New York: Routledge.

Markovsky, B., D. Willer, and T. Patton (1988) "Power Relations in Exchange Networks," *American Sociological Review*, 53: 220-236.

McMillan, J. and C. Woodruff (1999) "Interfirm Relationships and Informal Credit in Vietnam," *Quarterly Journal of Economics*, 114: 1285-1320.

Montgomery, J. (1991) "Social Networks and Labor Market Outcomes: Towards an Economic Analysis," *American Economic Review*, 81: 1408-1418.

Morris, S. (2000) "Contagion," *Review of Economic Studies*, 67: 57-79.

Murphy, J.J., Dinar, A., Howitt, R.E., Rassenti, S.J., and Smith, V.L. (2000) "The Design of 'Smart' Water Market Institutions Using Laboratory Experiments," *Environmental and Resource Economics*, 17: 375-394.

Myerson, R. (1977) "Graphs and Cooperation in Games," *Mathematics of Operations Research*, 2: 225-229.

Nishiguchi, T. (1994) *Strategic Industrial Sourcing*. New York: Oxford University Press.

Nowak, M.A. and R.M. May (1992) "Evolutionary Games and Spatial Chaos," *Nature*, 359: 826-829.

Ochs, J. (1995) "Coordination Problems," in Kagel, J.H. and A.E. Roth, eds., *Handbook of Experimental Economics*, 195-251. Princeton University Press.

Rassenti, S.J., Smith, V.L., and Wilson, B.J. (2003) "Controlling Market Power and Price Spikes in Electricity Networks: Demand-Side Bidding," *PNAS*, 100: 2998-3003.

Riedl, A. and A. Ule (2002) "Exclusion and Cooperation in Social Network Experiments," mimeo, University of Amsterdam.

Roth, A.E. (1995) "Bargaining Experiments," in eds. Kagel, J.E. and A.E. Roth, *Handbook of Experimental Economics*. Princeton University Press.

Roth, A.E., V. Prasnikar, M. Okuno-Fujiwara and S. Zamir (1991) "Bargaining and Market Behavior in Jerusalem, Ljubljana, Pittsburgh, and Tokyo: An Experimental Study," *American Economic Review*, 81: 1068-1095.

Sarangi, S., R. Kannan, and L. Ray (2003) "The Structure of Information Networks," mimeo, Louisiana State University.

Slikker M. and A. van den Nouweland (2000) "Network Formation Models with Costs for Establishing Links," *Review of Economic Design*, 5: 333-362.

Shaw, M.E. (1964) "Communication Networks," in Berkowitz, L., ed., *Advances in Experimental Social Psychology*, 111-147. New York: Academic Press.

Van Huyck, J.B., R.C. Battalio, and R.O. Beil (1990) "Tacit Coordination Games, Strategic Uncertainty, and Coordination Failure," *American Economic Review*, 80: 234-249.

Vanin, P. (2002) "Network Formation in the Lab: A Pilot Experiment," mimeo, Universitat Pompeu Fabra.

Vega-Redondo, F. (2002) "Building Up Social Capital in a Changing World," mimeo, University of Alicante.

Watts, A. (2001) "A Dynamic Model of Network Formation," *Games and Economic Behavior*, 34: 331-341.

Watts, D.C. (1999) *Small Worlds*. Princeton University Press.

Watts, D.C. and S.H. Strogatz (1998) "Collective Dynamics of 'Small-World' Networks," *Nature*, 393: 440-442.

Weibull, J.W. (2001) "Testing Game Theory," mimeo, Stockholm School of Economics.

Young, H.P. (1993) "The Evolution of Conventions," *Econometrica*, 61: 57-84.