CO-EVOLUTION OF SOCIAL NETWORKS AND BEHAVIOR IN SOCIAL DILEMMAS

Manuscript committee:	Prof. dr. K. van den Bos	(Utrecht University)
	Prof. dr. K. S. Cook	(Stanford University)
	Prof. dr. A. M. Riedl	(Maastricht University)
	Prof. dr. B. G. M. Völker	(Utrecht University)

Funding for this research was provided by Utrecht University through the High Potentials 2004 subsidy for the research program "Dynamics of Cooperation, Networks, and Institutions."

ISBN 978-90-393-5070-6
Cover design: Pieterjan Corten
Cover photo: Rense Corten
Printed by Ponsen & Looijen bv, Wageningen

Co-evolution of Social Networks and Behavior in Social Dilemmas

Theoretical and Empirical Perspectives

Co-evolutie van Sociale Netwerken en Gedrag in Sociale Dilemma's:
Theoretische en empirische perspectieven
(met een samenvatting in het Nederlands)

Proefschrift

ter verkrijging van de graad van doctor aan de Universiteit Utrecht
op gezag van de rector magnificus, prof. dr. J. C. Stoof,
ingevolge het besluit van het college voor promoties
in het openbaar te verdedigen op
woensdag 24 juni 2009 des ochtends te 10.30 uur

door

Rense Corten

geboren op 25 december 1979
te Groningen

Promotoren: Prof. dr. W. Raub
Prof. dr. S. Rosenkranz

Co-promotor: Dr. ir. V. W. Buskens

Contents

List of Figures

List of Tables

CHAPTER 1

Introduction

Our social environment influences much of what we do. To be more precise, individual behavior is often influenced by the *social network* surrounding the individual. For instance, when we form an opinion about a political issue, we are likely to be influenced by the opinions of our friends, family, and colleagues and likewise, we influence them.

Much sociological research has been devoted to showing how various forms of social influence shape individual action (Marsden and Friedkin, 1993). However, social networks are not always rigid structures imposed on us. Often, we have considerable control over our own social relations. Returning to our example, we may be influenced by our friends when forming political opinions, but we are also, to a large extent, free to choose our own friends. Moreover, it is likely that our decisions in choosing friends are in part related to those same opinions. Thus, social networks and the behavior by individuals within those networks *develop interdependently*, or in other words, *co-evolve*. What social network structures should we expect to emerge, and how will behavior be distributed in those networks? In a nutshell, this is the general type of problem this dissertation is concerned with. The example of political opinions is one in which social networks and individual characteristics co-evolve and the same holds for many other types of opinions and behavior. This dissertation focuses on the co-evolution of networks and behavior of a particular kind, namely, behavior in *social dilemmas*.

1.1. Social dilemmas and social networks

Broadly speaking, a *social dilemma* is a social situation in which individually rational behavior can lead to suboptimal results at the collective level. We encounter many social dilemmas in daily life. For example, when two researchers are working on a joint project, each might be tempted to let the other person do the majority of the work while profiting equally. However, if both follow this reasoning, the project will never get done, and there will be no profit at all. Similarly, if there is a rumor that a bank might go bankrupt, it is perfectly rational for every individual client to go to the bank and try to withdraw her savings. However, if all clients do this, the bank *will* indeed go bankrupt, and most clients will lose their savings. Another social dilemma arises when a group of people want to participate in an event, say, a protest demonstration. For each participant, it is only worth going to the demonstration if others are going as well. By attending, one runs the risk of being the only participant, in which case there will be no demonstration and one will have wasted one's time. Given this risk, it might be wise to stay at home, in which case the demonstration will indeed not occur.

The situations described above have in common that if the individuals try to obtain the most favorable outcome for themselves and behave rationally, the result might be that collectively everyone is worse off than they could have been. In other words, we find a conflict between *individual* rationality and *collective* rationality (Rapoport, 1974). Using the terminology of game theory,[1] we can define a social dilemma more formally as a situation (game) that has at least one Nash equilibrium that is Pareto-suboptimal.[2]

While all of the examples above can be classified as social dilemmas in this sense, there are also some differences between them. Roughly speaking, the first two examples can be described as *cooperation* problems, and the third example can be described as a *coordination* problem. While other types of social dilemmas exist, we only focus on coordination and cooperation problems in this dissertation.

The crucial characteristic of cooperation problems is that although the actors involved can benefit from cooperation, they have an incentive to take advantage of each other, which leads to suboptimal outcomes at the collective level. A game-theoretic model for such situations is the famous *Prisoner's Dilemma* (see Figure 1.1a). In this game, each player has two options: cooperation or defection. The players' payoffs associated with each combination of actions are represented as numbers in a matrix. The actual numbers in Figure 1.1a serve only as an illustration. The relation between the payoffs is what matters. In this game, both players are tempted to play "defect" because this will lead to a higher payoff regardless of what the other player does. If both players defect, they will both earn only 1, which is suboptimal because both could have earned 3 if they had cooperated. However, game theory predicts that goal-directed players will defect because mutual defection is the only Nash equilibrium. Clearly, this equilibrium is suboptimal.[3]

The Prisoner's Dilemma has become the archetypical social dilemma in the literature, and has motivated a vast amount of research. It is typically associated with the problem of *social order*, which has to do with the questions of why people cooperate even when they

[1] We assume here that the reader is somewhat familiar with basic game-theoretic terminology. Good introductions to game theory are Rasmusen (2001) and Binmore (2007), among many others.

[2] The precise definition of the term "social dilemma" is somewhat controversial. As we do not want to argue about definitional issues, it suffices to point out that the definition we use here is less restrictive than the definition used by Dawes (1980), but compatible with the progressively broader definitions by Kollock (1998) and Van de Rijt and Macy 2008. The definition used here is also compatible with the definition by Raub and Voss (1986) of problematic social situations.

[3] The description of a certain outcome as "optimal" does not necessarily mean that this outcome is also desirable from a societal point of view. Some outcomes of social dilemmas that are optimal for the actors involved may actually be quite undesirable for society as a whole; typical examples are collusion between firms or cooperation in criminal gangs.

	Cooperate	Defect
Cooperate	3,3	0,5
Defect	5,0	1,1

(a) A Prisoner's Dilemma

	LEFT	RIGHT
LEFT	4,4	0,2
RIGHT	2,0	3,3

(b) A coordination game

Figure 1.1: Two social dilemma games

have incentives to exploit each other and why society does not collapse into a "war of every man against every man" (Hobbes, [1651] 1988). The Prisoner's Dilemma has been used to model a wide range of social phenomena, including the production of public goods (e.g., Heckathorn, 1996), social exchange (Hardin, 1995), and the emergence of social norms (e.g., Ullmann-Margalit, 1977; Voss, 2001).

In coordination problems, the dilemma is of a different nature. Figure 1.1b shows a *coordination game* (the labels "LEFT" and "RIGHT" in this game are arbitrarily chosen and have no further substantive meaning). In contrast to the Prisoner's Dilemma, the coordination game no longer demands that to maximize her payoffs the player should always perform the same action regardless of what the other does. Rather, each player prefers to perform *the same action as the other player*, that is, to coordinate. Players do not have incentives to exploit one another, but there are incentives to try to work together. Thus, game theory predicts that either both players will play LEFT, or both players will play RIGHT. Once the players have established one of these equilibria, they have no incentive to deviate as long as the other player does not deviate. In this sense, we can consider equilibria in coordination problems as *conventions* (Lewis, 1969).

In so-called *pure coordination games*, actors have no preference for one convention over the other, but this is not the case in Figure 1.1b. If both players play LEFT, they both earn more than if they both play RIGHT. Therefore, the equilibrium (LEFT, LEFT) is the *efficient* equilibrium, also called the *payoff-dominant* equilibrium (Harsanyi and Selten, 1988). At first sight, it may seem obvious that if the players simply play LEFT, the social dilemma is solved. This game, however, also involves an element of risk. If the row player plays LEFT and the column player plays RIGHT, the outcome is suboptimal for both. However, the burden of the suboptimal outcome is not distributed equally among the players: the column player still earns a payoff of 2, whereas the row player earns nothing. Given that the row player does not know in advance what the column player will do, playing LEFT is risky. In fact, if the row player assumes that it is equally likely that column player will play LEFT or RIGHT, the expected payoff of playing LEFT is lower than the expected payoff of playing RIGHT. The same reasoning holds for the column player. This equilibrium (RIGHT, RIGHT) can be classified as the *risk-dominant* equilibrium (Harsanyi and Selten, 1988) because it has less risk.

Although the situation in which both players play LEFT is the efficient equilibrium, it is also the equilibrium with greater risk, and it is therefore not trivial that players will play this equilibrium. This feature makes this game especially interesting for analysis as a social dilemma (Kollock, 1998). Another reason to classify this game as a social dilemma is that the mixed equilibrium (in which the players play each action with some probability) is also inefficient (Harsanyi, 1977).

Referring to one of the examples above, every potential participant of a demonstration would rather join a successful demonstration than stay at home. However, if everyone else stays home, he would rather stay home too. In the case of a coordination failure (some come to the demonstration and some stay home making the demonstration a failure), the outcome is worse for those who did come to the failed demonstration than for those who stayed home. Many forms of collective action share this feature (Hardin, 1995).

Although coordination problems have received less attention in the literature on social dilemmas than cooperation problems, applications in real life are abundant. The analysis relates to many types of conventions, such as etiquette (Elias, 1969), standards of speech, or technological standards (e.g., choice of computer operating systems or GSM frequencies). Generally, the coordination game can be used as a game-theoretic model of *social conformism* whenever actors have strategic reasons to align their behavior. Moreover, it can be argued that many social dilemmas that are commonly viewed as Prisoner's Dilemmas could be more fruitfully analyzed as coordination games (Hardin, 1995; Kollock, 1998). One could say that the value of the coordination game as an explanatory model has been underappreciated in comparison to the enormous amount of attention that the Prisoner's Dilemma has received (Kollock, 1998).

Nevertheless, both dilemmas have been studied extensively in political science, psychology, economics, and sociology (as well as in the life sciences, particularly in the case of the Prisoner's Dilemma). The main question in these studies is, under what conditions will actors behave in such a way that they obtain the socially efficient outcome? Social networks can play an important role in answering this question.

1.1.1. Cooperation and social networks

There are different types of answers to the question of why people cooperate in situations such as in the Prisoner's Dilemma. The first type of answer looks for the solution at the individual level and challenges the assumption that people only care about their *own* payoffs. Proponents of this approach argue that people cooperate because they are motivated by fairness considerations (Rabin, 1993) or inequity aversion (Bolton and Ockenfels, 2000; Fehr and Schmidt, 1999; Kolm and Ythier, 2006).

Another approach does not abandon the assumption that people are selfish, but instead looks for *social* causes of cooperation: social conditions that provide individuals with incentives to cooperate in social dilemmas, even if these individuals only care about their own payoffs. One particular source of such incentives is that cooperative relations often do not occur in isolation, but are *embedded* in a social context (Granovetter, 1985). Such "embeddedness" may take several forms. As argued by Axelrod (1984) and others (Taylor, 1976, 1987), cooperation may emerge if actors interact repeatedly. This type of embeddedness is referred to as *dyadic embeddedness* (Buskens and Raub, 2002). The prospect of a long-term relationship with the same partner may persuade actors to cooperate on the condition that others cooperate as well.

A second type of embeddedness exists in social networks and can be referred to as *network embeddeness* (Buskens and Raub, 2002; Granovetter, 1985). This occurs when interactions are part of a larger network of relations. The presence of third parties further increases the interdependence between interaction partners as compared to dyadic embeddedness because information about what happened in one interaction may spread via the network and influence other interactions. An intuitive and broadly shared view among social scientists is that social cohesion facilitates the emergence of cooperation, trust, and social norms (Coleman, 1990; Homans, 1951; Voss, 2001), a view supported by much qualitative (Ellickson, 1991; Greif, 1989, 1994; Macaulay, 1963; Uzzi, 1996, 1997) and some quantitative (e.g., Burt and Knez, 1996; Buskens, 2002) evidence.

One class of mechanisms through which this information impacts cooperation in social dilemmas is captured under the heading of *reputation effects*. Actors embedded in networks may be more reluctant to defect because word regarding their behavior will spread and lead to retaliation or social sanctions by third parties. In a game-theoretic analysis Raub and Weesie (1990) show that such reputation effects indeed make conditional cooperation by selfish and rational actors more likely. According to this argument, this type of reputation effect is labeled as *control* because cooperation is promoted by actors' concerns about the outcomes of *future* interactions. Another reputation-based mechanism that can facilitate cooperation through networks is *learning*. Actors may be persuaded to cooperate with a given interaction partner because they have received information that cooperation with this partner is profitable (Buskens and Raub, 2002).

1.1.2. Coordination and social networks

While the focus of research on the Prisoner's Dilemma is to recognize under which circumstances people will cooperate, the focus of research on coordination problems is somewhat different. The characteristic feature of a coordination game is that it has *several* Nash equilibria, and a major challenge for game theory is to predict which of these equilibria will be chosen. Theorists have searched for additional mechanisms that can lead to more precise predictions because the standard prediction that rational actors will play a Nash equilibrium is not specific enough. The concepts of *payoff dominance* and *risk dominance* by Harsanyi and Selten (1988) provide some additional guidance, but in many coordination problems payoff-dominant and risk-dominant equilibria do not overlap, and an equilibrium selection problem remains. Analogous to the study of cooperation problems, researchers have turned to *repeated interaction* for a solution (Kandori et al., 1993; Young, 1993). In these models, actors play sequences of coordination games and reach a convention in a stochastic adaptive process. A major result from this research is that risk-dominant conventions are more likely to emerge in the long run.

These earlier models relied on global interaction structures in which every actor interacts with every other actor in the population. Later models (Berninghaus and Schwalbe, 1996; Blume, 1993; Ellison, 1993; Kosfeld, 1999; Young, 1998) introduced social structure by assuming *local interaction* such that actors interact only with some portion of the population. While this can be considered a first step towards introducing social networks into the analysis, these models generally assume only very simple interaction structures such as lattices.

Empirical evidence for network effects in coordination problems comes from experimental studies. Keser et al. (1998) found that it is more likely that subjects coordinate on the risk-dominant equilibrium when they interact in a circle structure than when they interact in small three-person groups. Berninghaus et al. (2002) further investigated local interaction and found that different regular local interaction structures had effects on the likelihood that risk-dominant or payoff-dominant equilibria are chosen. Cassar (2007) compared the effects of different *irregular* network structures on coordination, and found that coordination on the payoff-dominant equilibrium is more likely in "small-world" network structures than in random network structures, or structures with overlapping neighborhoods.

Meanwhile, the sociological literature on social networks has traditionally focused on the effects of social networks on the diffusion of behavior and opinions. This literature includes theoretical studies of *threshold models* (Abrahamson and Rosenkopf, 1997; Centola and Macy, 2007; Granovetter, 1978; Valente, 1996; Watts, 2002), as well as empirical studies on diffusion (e.g., Rogers, 1995) and interpersonal influence in social networks (see Marsden and Friedkin, 1993). Although these studies neglect the strategic interdependence of actors in the game-theoretic sense, they are conceptually very close to coordination problems, as discussed above; this literature generally assumes that actors face incentives to align their behaviors. Generally, it has been found that the structures of social networks can have important consequences for the extent to which behaviors, opinions, and collective action spread in a population.

1.2. Dynamic networks, co-evolution, and research questions

As the above has shown, a wide range of theoretical and empirical research from various disciplines suggests that social networks are relevant in determining the outcomes of cooperation and coordination problems. These findings add to the more general notion that social networks have important effects on many types of social phenomena, including (but not limited to) social inequality (Coleman, 1988; Flap, 2004; Lin, 2001), labor market outcomes, (Granovetter, 1973, 1974), the diffusion of innovations (Coleman et al., 1957), and the spread of diseases (Kretzschmar and Wallinga, 2007; Morris et al., 1995). Naturally, the next question is, where do these network structures come from?

Often, an implicit assumption in theories of network effects is that social networks are *fixed structures*, or *exogenously* imposed on the actors. While in some cases it may be reasonable to assume that people have little or no control over their social environment (e.g., kin relations), many social relations actually result from people's *choices*. We can typically choose, at least to some degree, who our friends are, with whom we share information, and with whom we want to work. These choices are likely to be constrained by the larger social context (e.g., the availability of meeting opportunities; see Fisher, 1982; Mollenhorst et al., 2008; Verbrugge, 1977). Nevertheless, his implies that social network structures develop as the result of individual decisions.

Moreover, the notion that networks have important consequences for behavior suggests that people do not only have the *opportunity* to change their social relations, but also have *incentives* to do so. Given that networks potentially produce benefits for their actors, it seems reasonable that actors will consciously form relationships to optimize their benefits from the networks. Sometimes dense networks will be beneficial, for example, to solve trust- or cooperation problems (Buskens, 2002; Coleman, 1990; Raub and Weesie, 1990). In other settings, open structures are more beneficial, for example, in competitive settings where access and control of information is of crucial importance (Burt, 1992; Granovetter, 1973; for the comparison also see Burt, 2005).

A similar message emerges from the social capital literature, which argues that social inequality can be explained in part by differences in resources that people derive from their personal networks (Coleman, 1988; Flap, 2002, 2004; Lin, 2001). It follows that goal-directed actors would strategically invest in relations that are beneficial and end relations that are not.

However, from the premise that network structures are the results of actors' conscious decisions, it does not follow that socially beneficial network structures will spontaneously emerge (e.g., Dogan et al., 2009; Jackson and Wolinsky, 1996). Although actors may be able to choose their *own* relations, the larger network structure is the result of the combined choices of all actors. Relational choices of one actor may have consequences for other actors. For instance, by breaking just one relation, an actor may interrupt many indirect connections between other pairs of actors, thereby changing the flow of information in the network. Thus, although network structures may be the consequences of individual decisions, they are often *unintended* consequences of individual action (cf. Coleman, 1987; Merton, 1936; Schelling, 1978).

Much literature has emerged in the past decade on *network dynamics* as a response to questions about the origins of social networks. Although the interest in social networks originated within sociology, problems of network dynamics have attracted a great deal of attention from a number of other disciplines, including economics, mathematics, physics, and biology. Problems of network dynamics have also found their way into popular science literature (e.g., Buchanan, 2002). As a result, a number of models have been formulated on the dynamics of "small world" networks (Watts and Strogatz, 1998), scale-free networks (Albert and Barabási, 1999), communication networks (Bala and Goyal, 2000; Buskens and Rijt, 2008), and other topics (see Goyal, 2007, Vega-Redondo, 2007, and Jackson, 2008 for good reviews of this literature).

Many of these models study the formation of network relations *per se*, that is, reasons for network change lie solely in the network structure itself. However, it is likely that the choice of network relations also depends on actual behavior in relevant interactions, namely, on the *content* of relations. After all, one of the reasons to study social networks in the first place is that networks affect *behavior in social dilemmas*. Facing cooperation problems, actors may want to avoid defectors, while in coordination settings, actors may want to avoid those who behave differently and prefer relations with those who behave similarly (cf. McPherson et al., 2001). Thus, on the one hand, networks influence the way people behave in their interactions. On the other hand, individual behavior in interactions also affects the network such that people "themselves constitute each others' changing environment" (Snijders, 2001, p. 363). In other words, one can say that networks and behavior *co-evolve*.

This implies a two-sided causality social scientists have only begun to investigate in recent years. A priori, it is unclear if and how the recognition that networks and behavior may co-evolve changes our expectations on outcomes in social dilemma situations. For instance, most arguments on the positive effects of social networks on cooperation problems crucially rely on the assumption that information is disseminated through the network. However, is the transfer of information still reliable enough for reputation mechanisms to function if actors at the same time change the network? Can we expect that the network structures that are thought to foster cooperation (i.e., dense, close-knit networks) will emerge as a consequence of individual strategic action? Similarly, does the ability to change relations facilitate or hinder the emergence of conventions in coordination problems?

The four studies in this book all revolve around such issues Thus, we can formulate the central questions of this dissertation as follows:

- How do social network structure and behavior co-evolve in different types of social dilemmas?

- Under what circumstances is it more or less likely that networks and behavior evolve into optimal or suboptimal interaction structures?

In summary, we have argued that in order to explain behavior in social dilemmas, we have to study the social networks in which social dilemmas are embedded, and in doing so, account for the evolution of the networks themselves. At the same time, studying social dilemmas in dynamic networks can help us to better understand social networks in general.

Although the field of network dynamics as a research topic is relatively new, it has roots in various disciplines, most prominently in sociology and economics. In the next section, we argue that a fruitful analysis combines insights from both disciplines.

1.3. Social networks and social dilemmas between sociology and economics

Duesenberry (1960) famously stated that "Economics is all about how people make choices; sociology is all about how they don't have any choices to make" (p. 233). Traditionally, economics has taken individual action as its main focus, whereas sociology has emphasized how individual action is shaped by social context, as exemplified by Durkheim's ([1897] 2002) classic study of the social antecedents of suicide.

Over the years, attempts have been made to reconcile these apparently incompatible approaches. Some sociologists have adopted modes of theorizing from economics in the form of methodological individualism, rational choice theory, and game theory (Boudon, 1981; Coleman, 1990; Elster, 1989; Hedström, 2005; Schelling, 1978). These approaches have roots in the early sociology of Parsons (1937) and Weber ([1921] 1976), and are closely tied to the study of the problem of social order (Hobbes, [1651] 1988; Parsons, 1937), which is considered to be one of the central problems of sociology. Along with these concepts, some sociologists have also tried (with lesser success) to import the accompanying *tools* of formal modeling from economics into sociology (Coleman, 1964; Fararo, 1973; Merton, 1968).

On the other hand, attempts have been made to introduce sociological arguments about social structure into economic theory (interestingly, it seems that most of these attempts are made by sociologists). One of the most influential arguments was put forth in the seminal paper by Granovetter (1985) who argued that whereas economic theory suffered from being "undersocialized," sociological theory was "oversocialized" (also see, e.g., Coleman, 1984). Economics, on the one hand, is criticized for modeling human action as overly atomistic and anonymous without the influence of social context. Sociology, on the other hand, too often takes the influence of the social context for granted, to the extent that it is assumed that social norms and values are completely internalized. Ironically enough, this also leads to an overly atomized account of social action.

As a solution, Granovetter proposed to analyze social and economic phenomena as resulting from *individual* (rational) action, which is nevertheless *embedded* in social structures, or social networks.:

> A fruitful analysis of human action requires us to avoid the atomization implicit in the theoretical extremes of under- and oversocialized conceptions. Actors do not behave or decide as atoms outside a social context, nor do they adhere slavishly to a script written for them by the particular intersection of social categories that they happen to occupy. Their attempts at purposive action are instead embedded in concrete, ongoing systems of social relations. (Granovetter, 1985, p. 487)

The collection of studies in this dissertation continues and extends Granovetter's general research program. It is a continuation in the sense that we study goal-directed action as embedded in systems of social relations (i.e., social networks). It extends Granovetter's ideas by applying a similar logic to the system of social relations *itself*. If social relations are crucial for understanding human action, and if social relations are—at least to some extent—the result of purposive individual action, explanations should also account for the emergence of structures of social relations.

Thus, the approach advocated here aims to contribute to the bridging of the "theory-gap" in social network research, which was identified by Granovetter (1979) who states that empirical sociological research on social networks lacks systematic theoretical foundations (see also Flap, 2002). Models that explain *both* causes *and* consequences of social networks as the results of goal-oriented strategic action might provide such theoretical foundations. Other theoretical strategies that also rely on goal-oriented individual action have been proposed (e.g., Snijders, 2001), but these approaches typically neglect the strategic interdependence involved in network formation processes.

Granovetter argues that theories about social networks could serve to bridge the gap between macrolevel and microlevel theories. Following Coleman (1984, 1987, 1990) and Wippler and Lindenberg (1987), we argue that theory on social networks also needs to account for the micro-macro link. Establishing this micro-macro link (or "transformation rules," Wippler and Lindenberg, 1987) requires modeling the collective effects of interdependent individual decisions. For this task, the theoretical tools for studying strategic decision making developed in economics are particularly useful. The theoretical strategy applied in this dissertation can be considered as a synthesis of the typical approaches of sociology and economics. On the one hand, it acknowledges the sociological emphasis on social structure, while on the other hand, the standard economics assumption of goal-driven, *strategic* action is maintained. Not surprisingly, this approach adds further complexity to theoretical models. As a result, the studies in this book do not offer more than piecemeal solutions to the many problems that arise when studying the co-evolution of strategic behavior and social networks.

1.4. Approach

1.4.1. Theoretical approach

In line with the arguments sketched above, the overall theoretical approach relies on the assumption that both behavior in social dilemmas and network formation are driven by goal-directed individual action. In the theoretical models employed in Chapters 2–5, actors choose the actions in social dilemmas and choose the social relations that are beneficial to them. With this approach, we follow the economic literature on networks in which network formation is analyzed as a *strategic decision problem* in the sense that benefits from forming social relations depend not only on an actor's decisions, but also on the decisions of others. *Game theory* was developed to analyze strategic decision problems, and recently, specific game-theoretical tools have been developed to analyze networks (see Jackson, 2008, Ch. 11 for a good overview).

Within this general framework, there are many specific modeling choices to make. What are the action alternatives of the actors? How are costs and benefits of their choices determined? What do they know about the actions of other players? How rational are the actors when making these choices? The chapters in this dissertation apply various approaches to study co-evolution of behavior in social dilemmas and social networks. Nevertheless, the chapters have a number of assumptions in common, which are briefly outlined here.

In every chapter situations are studied in which actors are involved in social dilemmas, modeled as games. Actors can choose *how* to play in these games and *with whom* to play. Their choice of with whom they will play results in a social network. That is, the links in the network consist of interactions in social dilemmas. For example, consider collaborations between researchers in projects. Researchers can choose with whom to work and how much effort to spend on each collaborative project. These choices result in a network of collaborations. In principle, we can also think of situations in which interactions in social dilemmas and relations in social networks exist independently. A neighborhood may face a social dilemma in keeping the neighborhood clean, while there also exists a network of friendship relations between the neighbors. Although this friendship network may be relevant for solving the social dilemma, and vice versa, interactions in this social dilemma and relations in the social network exist independently. Such situations are not analyzed here (see Takács et al., 2008, as an example of a model in this direction).

In choosing how to act in social dilemmas and with whom to play, actors must weigh costs and benefits. In the different models in this dissertation, benefits are determined by the outcomes of interaction in social dilemmas. At the same time, social relations are assumed to be *costly*. The rationale for this assumption is that maintaining social relations takes time and effort, such that people want to maintain only those relations that are *worthwhile* so that the expected net benefits from these relations are positive. The exact specification of these costs and benefits differs between the chapters.

Arguments about *information* are central to many theories concerning network effects on behavior in social dilemmas. For example, the diffusion of information via social networks is often assumed to be the driving force behind network effects on cooperation (e.g., Raub and Weesie, 1990). Consequently, assumptions on information also play an important role in the studies of this dissertation. The implications of different assumptions on how the diffusion of information depends on the network are studied in Chapters 3 and 5.

The use of the terms "purposive action" and "strategic" above does not imply that actors are necessarily assumed to be perfectly rational, calculating the costs and benefits of every potential situation and fully anticipating the consequences of their actions. Given the complexity of the processes that study, it does not seem likely that human actors involved in these processes would be able to behave perfectly rational. For instance, perfect rationality would imply that actors in a large population would consider the potential costs and benefits of interactions with *all* other actors in the population, anticipate the reactions of these actors, and anticipate the consequences of those reactions for the other actors, the network structure, and so forth. Instead, we model actors as *boundedly* rational (Rubinstein, 1998) and account for actors having limited cognitive capabilities. The different chapters use different approaches to model these capabilities. Note that this does *not* imply that we relax the assumption that actors are goal-directed in the sense that they try to maximize their *own* utility. One might instead, for example, assume that people are to some extent altruistic and are not only motivated by their own payoffs, but also by the payoffs of others (e.g., Fehr and Schmidt, 1999). This would drastically change the nature of our main explanatory problem, namely behavior in social *dilemmas*. Instead, we assume that actors at least *try* to maximize their own payoffs, but may make less than perfect rational choices in doing so.

Game theory and related analytical methods allow for the systematic derivation of propositions on behavior in strategic situations. However, in many cases game-theoretical models predict a wide range of situations as possible outcomes of these processes (e.g., Jackson and Watts, 2002; Jackson and Wolinsky, 1996). For empirical applications, it is preferable to have more informative predictions on which situations are more likely to occur. One way to address this issue is *simulation*. In the social sciences, computer simulation is increasingly used as a theoretical tool to study complex systems (Cederman, 2005; Epstein, 2007; Gilbert and Troitzsch, 2005; Halpin, 1999).[4] The studies presented here make extensive use of computer simulations to study how the emergence of certain outcomes depends on properties of co-evolution processes and the initial conditions of these processes. Results of the simulations play a crucial role in the derivation of testable hypotheses in Chapters 3 and 4.

1.4.2. Empirical approach

The recent flourishing of interest in social network dynamics has resulted in the publication of a large number of theoretical models. Many of these models are designed to explain some "stylized fact" about real social networks, such as the small-world phenomenon (Watts and Strogatz, 1998) or the emergence of "stars" (Bala and Goyal, 2000). However, much less frequently are these models rigorously *tested* empirically. Thus, while Granovetter (1979) identified a lack of theoretical foundations (the "theory-gap") in the very empirically oriented sociological network research at the time, the emergence of a theoretical literature on network dynamics has not yet led to a tight connection between theoretical models and empirical approaches. This is not only a problem for empirical research (as argued by Granovetter, 1979) but also for the development of theory. It is hard to judge the appropriateness of model assumptions because many models go untested, thus making it difficult to judge theoretical progress in the field.

Therefore, in this dissertation we develop theoretical models, and also attempt to test predictions from these models empirically. Two types of research designs are used: *laboratory experiments* and *survey research*. While the traditionally preferred method of social psychology, experimental methods have also emerged as an important paradigm in economics (for good overviews of experimental economics, see Camerer, 2003, Kagel and Roth, 1995, and Plott and Smith, 2008). Experimental studies are less common in sociology (Diekmann, 2008), but do have a long history (e.g., Burgess and Bushell, 1969). In general, the advantage of experimental methods is that the researcher can purposively manipulate conditions of theoretical interest while keeping other conditions as constant as possible. This allows for direct identification of causal mechanisms that are hard to isolate in non-experimental studies in which many other factors may confound the mechanisms of interest. Two additional advantages of experimental studies are particularly relevant for research on network dynamics.

First, laboratory experiments are well suited to study details of *individual* decision making. Theoretical models on network dynamics typically imply highly complex interdependencies between actors, and at the same time make strong assumptions on actors' rationality. In this context it becomes important to study to what extent actual human behavior in complex social settings deviates from what is assumed in theoretical models, and to what extent this affects the predictions from these models.

[4] Despite the growing popularity of simulation, a common misconception among social scientists is that simulation is used as an alternative to empirical research. It is important to realize that simulation is a *theoretical tool*, used to *derive* hypotheses that should subsequently be subjected to empirical tests.

Second, experimental methods allow for relatively easy testing of hypotheses on the *macro*level, that is, hypotheses on differences between groups or networks. In observational studies, this is often problematic because the complexity and time-consuming nature of collecting data on social relations typically imposes limits on the number of groups that can be observed. Moreover, studying network *dynamics* ideally requires longitudinal network data, which further complicates data collection. As a result, relatively few datasets are available to test theories on network dynamics (with some exceptions, such as the dataset collected by Knecht, 2006 that we analyze in Chapter 4). In laboratory experiments, many small groups can be observed in detail in a relatively short amount of time and group-level conditions can be easily manipulated. In this sense, laboratory experiments provide a useful "early testing ground" for network theories.

Such considerations have led to extensive experimental literature in sociology on *exchange networks* (e.g., Willer, 1999). In experimental economics, a number of experiments have been conducted that study either social dilemmas in networks or network formation processes (for a review, see Kosfeld, 2004), but there are only a few experimental studies on the co-evolution of networks and behavior (Corbae and Duffy, 2008; Ule, 2005). Chapter 3 aims to provide further contributions on this topic.

However, experimental methods also have obvious disadvantages. Laboratory experiments typically involve highly abstract settings, use relatively small groups, and are conducted with subjects recruited from a specific subpopulation (i.e., undergraduate university students). As a result, while experiments are highly useful for testing theories "at close range," they risk lacking external validity. It is not always clear to what extent theoretical models that accurately predict behavior in the laboratory can also explain social phenomena outside the laboratory. While lacking the controlled nature of experiments, survey methods conducted in real-life settings are typically stronger in terms of external validity. In sociology, specialized survey methods have been developed to study social networks (Scott, 2000; Wasserman and Faust, 1994). In Chapter 4, such methods are applied to test hypotheses that are developed in the earlier chapters. By using experimental methods *and* survey methods, we aim to benefit from the strengths of both approaches. On the one hand, we are able to test our theoretical predictions in a very direct way in the laboratory, while on the other hand, we can also verify that these mechanisms play a role in more natural social settings that we are substantively interested in.

1.5. Description of the remaining chapters

The approach outlined above is brought into practice in four studies that make up Chapters 2–5 of this dissertation. The first three of these study coordination problems both theoretically and empirically, while Chapter 5 is a theoretical study of cooperation problems.

Chapter 2 presents a theoretical study of coordination in dynamic networks, i.e., situations in which actors not only have a preference to behave as their interaction partners do, but can also choose their interaction partners. A model is proposed in which actors play coordination games with multiple interaction partners (i.e., in networks), but can also choose the partners with whom they interact. We study two types of outcomes in this model. First, emerging behavior and networks might lead to inefficient outcomes for the reasons sketched above. Second, the network might become segregated because actors who do not behave the same choose to give up their relationships rather than adapt their behavior.

Although possible stable states of the co-evolution process can be characterized analytically, there is typically a large variety of stable states. Therefore, we conduct a simulation study to investigate how emerging network structures and conventions depend on the following; the level of risk in the coordination game, the cost of maintaining relations, and the initial conditions (in terms of the properties of the initial network structure and initial propensities for choices in the game). The results of the simulation provide *predictions* in which network structures and conventions are more likely to emerge under various conditions.

In Chapter 3, the results of a laboratory experiment are described, which tests a number of hypotheses from Chapter 2. In addition, a number of new hypotheses are developed, which focus on the effects of the *availability of information* on emerging behavior and stable networks. The model of Chapter 2 assumes, as is common in the literature, that actors are informed about the behavior of *all* other actors in the population. This assumption is obviously unrealistic, especially for larger populations, and is somewhat at odds with the motivations to study *local* interaction. Therefore, in Chapter 3, we relax this assumption and study a more general model in which actors are informed about the behavior of only a *part* of the network. We compare two variants in particular: the case in which actors are informed about the whole population ("global information," as in Chapter 2), and the case in which actors are only informed about their direct neighbors ("local information").

Following the development of theoretical models on coordination in dynamic networks (Chapter 2) and the experimental tests of these models (Chapter 3), Chapter 4 studies coordination problems again, but this time uses an empirical setting outside the lab. We use longitudinal survey data on networks in school classes to study the development of alcohol use among adolescents. We argue that the choice to use alcohol in this setting resembles a coordination game. Based on this argument, hypotheses derived in the earlier chapters are tested again.

In Chapter 5, the attention is shifted to *cooperation problems*. A common finding in sociological research is that cooperation in social dilemmas is more likely if interactions are embedded in cohesive social networks. It is argued that cooperation is promoted through mechanisms of reputation, social control, and learning. An underlying (and often implicit) assumption in these explanations is that networks are exogenously imposed on the actors and are not subject to change. In this chapter we relax that assumption and study a model in which actors play dyadic Prisoner's Dilemmas while also choosing their interaction partners, so that behavior and social networks co-evolve. Can cohesive networks *and* cooperation co-evolve? Is an exogenously imposed cohesive network a condition for cooperation, or are cohesive networks a *result* of high levels of cooperation? We study these questions using a formal model in which actors are modeled as boundedly rational and base their decisions to cooperate with a given partner on their expectations of the partner's behavior. At the same time, they build or dissolve interactions based on the expected utility of these interactions. To form expectations, actors learn from their own experience as well as from third-party information (i.e., reputation) which spreads through the network. After obtaining some basic analytical results on stable states in this model, we apply computer simulations to study the dynamics of the process in detail. Chapter 6 summarizes the results of the preceding chapters and discusses possibilities for both theoretical and empirical future research.

Finally, a note on the setup of this dissertation. The chapters were written as independent research papers and can be read independently as such. Cross-references between the chapters are provided as references to the corresponding articles or working papers. As a consequence of this approach, terminology and notation may vary slightly between the chapters. Moreover, some overlap, especially between the introductory sections of the chapters and this introduction, could not be avoided.

C H A P T E R 2

Consent or Conflict: Co-evolution of Coordination and Networks*

2.1. Introduction

A long research tradition in sociology and social psychology has shown that social networks play an important mediating role in the diffusion of behaviors and opinions through a society. In many different contexts, people are influenced by those with whom they interact (Erickson, 1988; Marsden and Friedkin, 1993; Merton, 1968). Empirical examples of such processes include peer pressure among adolescents (Davies and Kandel, 1981), diffusion of innovations (Coleman et al., 1957; Valente, 2005), rebellion, and collective action (Gould, 1991, 1993; Opp and Gern, 1993). These findings are relevant for the study of *polarization*, described as the social or ideological separation of a society into two or more groups (Esteban and Scheider, 2008), because the adaptation process might increase agreement within groups, while it deepens disagreement between groups. The extent to which a society will polarize into possibly opposing camps is likely to be influenced by the patterns of social relations through which members of the society influence each other, and through which opinions, behaviors, and ideologies diffuse.

It is important to realize that social networks are not always static but can be altered by actors consciously selecting their relations. At least in part, this selection processes is based on behavioral traits of others; sociological research shows that people tend to choose their friends among those who behave and think as themselves (a process known as "homophily," see Lazarsfeld and Merton, 1954; McPherson et al., 2001; Zeggelink et al., 1996). The combination of network influence and selective network formation processes implies that, first, polarization may occur because behavior or opinions cluster locally within networks and, second, a society may segregate socially because people with different behaviors or opinions tend to avoid each other.

Our study of polarization takes into account that social and ideological or behavioral

*This chapter was written in collaboration with Vincent Buskens and Jeroen Weesie. A slightly modified version appeared in the *Journal of Peace Research* (Buskens, Corten, and Weesie, 2008). As compared to that version, some notation was adapted to improve consistency with the other chapters in the dissertation.

alienation between groups develop *interdependently*. In other words, the degree to which polarization on a behavioral trait occurs depends on patterns of social ties, but this social structure itself is also influenced by behavioral choices. We aim at a theoretical understanding of the interplay between polarization of behavior and social structure. We develop a model in which actors are involved in interactions with others with whom they have social relations, while this social network itself is subject to change by the actors. This model predicts how the likelihood of polarization of behavioral outcomes depends on the social structure of a society at the time actors have to decide on a certain form of collective action or have to develop an opinion on some issue that becomes salient. In addition to problems in which persistent disagreement constitutes a clear potential for conflict, the model captures other types of processes from which conventions emerge, such as life style choices of pupils in school classes. In such networks, persistent disagreement is not problematic per se.

2.1.1. Polarization, conflict, and coordination

Group mobilization and group formation have previously been modeled in different ways, for example, as social influence processes (Axelrod, 1997),[1] as multi-person Prisoner's Dilemmas (Takács, 2001), or as collective action problems (Gould, 1999) identification with a group can also be considered a multi-person *coordination problem* in the sense that belonging to a group and making the same choice as others, is more important than what choice is actually made (Hardin, 1995). In group identification, one prefers to join a group if others do the same, because benefits can be expected from group membership itself. There is little sense in speaking English if everyone else speaks French; similarly, it may not be beneficial to identify as Serbian if everyone else identifies as Yugoslavian. However, if enough people start to call themselves Serbs, it becomes attractive to join this group.

It can be argued that not only identification with a group, but also group *action* is mainly a matter of coordination. Usually, group mobilization and other forms of collective action are seen as free-rider problems. According to the "logic of collective action" (Olson, 1971), individual group members should not be expected to contribute to collective efforts unless they have individual (selective) incentives that compensate their efforts. Hence, collective action should not occur in most cases because every individual has reasons to free ride on the other group members. This prediction seems at odds with the real-world observation that group action *does* occur in many instances, from voting to mass demonstrations and collective violence. This led some scholars to argue that coordination rather than cooperation is the basic strategic interaction that underlies group action. According to Hardin (1995), the power resulting from mass action can diminish the costs of joining to a level that is sufficiently low to reduce the free-rider problem to a problem of coordination (see Heckathorn, 1996; Macy,

[1] See also Stokman and Zeggelink (1996) for a model in which opinions and networks are adapted simultaneously

1991; Marwell and Oliver, 1993). Others (Chwe, 2001; Gould, 1995; Klandermans, 1988; Lohmann, 1994) have also pointed at the importance of coordination in collective phenomena such as rebellion, uprisings, and union participation. Therefore, by modeling collective action as a coordination problem we abstract from free-riding problems and focus on settings in which actors have an incentive to join if enough others do so as well.

Modeling group identification and group mobilization as a coordination problem is not to say that actors are indifferent between behavioral alternatives as long as they coordinate with others. The coordination game that is the backbone of our model accommodates for ranking of behavioral alternatives, while coordination with others still has priority. Consider the simple coordination game as displayed in Figure 2.1, with $c = 8$. The actors have two choices, L(EFT) or R(IGHT). The game has two Nash equilibria, (R,R) and (L,L), in each of which actors do not wish to deviate as long as the other actor does not deviate. However, (L,L) yields higher payoffs for both actors and is termed *efficient* or *payoff dominant*. The other equilibrium (R,R) is attractive in the sense that it is less risky: if an actor assumes that the other actor chooses R and L with equal probabilities, the expected payoff of choosing R is higher than of choosing L. Therefore, (R,R) is *risk dominant* (Harsanyi and Selten, 1988). For our applications, choosing L may represent joining an uprising in order to accomplish a regime change, while choosing R is to stick to the status quo. Joining the uprising is risky if you are not sure that others will also do so.

	LEFT	RIGHT
LEFT	d,d	b,c
RIGHT	c,b	a,a

	LEFT	RIGHT
LEFT	20, 20	0, 10
RIGHT	10, 0	14, 14

Figure 2.1: The two-person version of the multi-person coordination game with payoffs as in the simulation $b < c < a < d$; $(a-b) > (d-c)$; $c = 4$ or $c = 8$.

A consequence of choosing a *multi-person version* of the coordination game as described above implies that we cannot only provide predictions on the likelihood and extent of polarization, but also on the extent to which actors coordinate on the efficient equilibrium. However, the model does not predict how the likelihood of the emergence of violent conflict depends on the polarization that might arise in a population. Rather, the theory *assumes* that polarization into separated but internally coordinated groups provides *potential* for violent group conflict, while the model specifies the conditions under which a polarized situation is more or less likely to occur.

As a measure of polarization, we use the two-group version of the index for qualitative variation IQV (Mueller and Schuessler, 1961, pp. 177-179; see also Agresti and Agresti, 1978) which is defined as $4p(\text{L})(1-p(\text{L}))$, in which $p(\text{L})$ is the proportion of actors in the population choosing L. The measure implies that polarization is 0 if $p(\text{L})$ = 0 or $p(\text{L})$ = 1, while it equals its maximal value 1 for $p(\text{L}) = \frac{1}{2}$. This measure is the standardized version of a much older version of a diversity measure that dates back to Gini (see Lieberson, 1969, for an overview). We focus on polarization as defined above because with only two groups we cannot distinguish it from, for instance, fractionalization (see Montalvo and Reynal-Querol, 2005; Reynal-Querol, 2002). The maximum of our polarization measure is reached if both groups are of equal size, which corresponds with the maximum for conflict potential according to Esteban and Ray (1999) for the general polarization measure (see also Esteban and Scheider, 2008).

2.1.2. Coordination and social networks

Earlier theoretical studies of coordination in large groups consider models in which actors interact with all other actors in the population, without assuming any social structure (e.g., Kandori et al., 1993; Young, 1993). However, this assumption seems highly unrealistic for many applications. Actors can often only observe the behavior of those with whom they interact directly; they only observe their own personal network. This is especially true for cases in which no public information is available about the distribution of behavior in the larger population, and so actors really have to rely on their close surroundings for information. They may even use the behavior of the members of their personal network as an approximation of the behavior in the larger population. Opp and Gern (1993) and Lohmann (1994), for example, discuss the importance of personal networks for the "Monday Demonstrations" in Leipzig, 1989, which is a typical case in which public information was highly restricted. Gould (1991, 1993, 1995), Scott (1990), Hardin (1995), and Chwe (2001) similarly emphasize the role of networks in instances of (violent) collective action.

Theoretical models dealing with local interaction have been formulated by (among others) Ellison (1993), Young (1998), and Berninghaus et al. (2002), although they do not consider possible heterogeneity of the network structure. Buskens and Snijders (2008) explicitly deal with coordination in *heterogeneous* social network structures, where heterogeneity refers to actors having different positions in the network, that is, they do not necessarily have the same number of relations. Considering static networks only, Buskens and Snijders (2008) show that denser and less segmented networks reach consensus more easily than less dense and more segmented networks. In addition, segmentation leads to fewer actors choosing the risky option, while density encourages more actors to choose the risky option. Thus, segmentation leads to less and density to more efficiency in the emerging behavior.

However, networks are created by people and change over time (Flap, 2004). Therefore, we consider a model in which actors are organized in a *dynamic* social network and obtain "high" payoffs for relations with actors who behave the same and obtain "low" payoffs for relations with actors who behave differently. Assuming that maintaining social relations is costly, relations have to be chosen cautiously, which implies that relations with actors who behave differently may be terminated.

Recently, the study of dynamic networks with strategic decision making of actors in the network has developed quickly (for an overview see Dutta and Jackson, 2003; Newman et al., 2006). Specific models for coordination games played on dynamic networks are studied by Skyrms and Pemantle (2000), Jackson and Watts (2002), Goyal and Vega-Redondo (2005), and Berninghaus and Vogt (2006). These models focus on which networks are stable and on how behavior is distributed in stable networks. Typically, many configurations are possible. Therefore, we take a different approach and focus on how preconditions determine the emergence of a stable network and the related distribution of behavior. Hence we study the extent to which polarization in segregated groups emerges and inefficient behavior persists as a result of the initial network and the initial distribution of behavior. More specifically, we aim to answer the following research question: *how do polarization and efficiency of the emerging distribution of behavior depend on the initial distribution of behavior, the initial network, the tie costs, and the payoffs in the coordination games?*

In the next section, we present the dynamic model and analytic results on stable states. Subsequently, we describe simulations and derive predictions on the effects of initial conditions on polarization and efficiency. In the final section, we conclude and interpret the results in terms of the likelihood that conflicts may emerge.

2.2. The model

Actors are organized in a network of n actors with the $n \times n$ adjacency matrix $N = (n_{ij})$, i.e., $n_{ij} = 1$ if two actors are connected and $n_{ij} = 0$ otherwise. We assume that relationships are undirected, so $n_{ij} = n_{ji}$. Relations have both benefits and costs. Actors have to choose between two types of behavior (or opinions, attitudes) and their benefits or payoffs depend on their own behavior and the behaviors of the actors with whom they have relations. They cannot differentiate their behavior depending on specific relations. In every existing relation, the payoff is related to the actor's own behavior and the behavior of the other person in correspondence with the coordination game depicted in Figure 2.1. From actors with whom actor i does not have a relation, i obtains a payoff 0. In our multi-person coordination game, this implies that the payoff of an actor i choosing R equals $\sum_{j(R)} n_{ij}a + \sum_{j(L)} n_{ij}c$, in which $j(z)$ runs over other actors who choose z, $z = \text{R}, \text{L}$. If actor i would choose L, the payoff would be $\sum_{j(R)} n_{ij}b + \sum_{j(L)} n_{ij}d$.

We study two variants. Variant one assumes that $(a-b) < (d-c)$, so that choosing L not only can lead to the efficient payoff d, but choosing L is also less risky because the expected payoff from playing L is relatively high if you do not know what others will do.[2] The second variant assumes $(a-b) > (d-c)$ so that choosing L is the risky choice because the expected payoff from playing L is relatively high if you do not know what others will do. The actual payoffs that we use in the simulation are indicated in Figure 2.1 as well. Buskens and Snijders (2008) demonstrate that

$$\text{risk} = \frac{a-b}{a-b-c+d} \tag{2.1}$$

in combination with the behaviors of the actors with whom a focal actor is connected, determines whether or not this focal actor wants to change behavior. In addition, Buskens and Snijders (2008) show that network effects are small for static networks if risk values are far from 0.5, but that there are major differences in network effects depending on whether risk is just above or just below 0.5. Therefore, in the simulation risk varies between 0.467 and 0.538.

In addition, we assume that ties are costly. We abstract from separate costs for the creation or the deletion of ties. Thus, the costs of existing ties have to be paid in every period of interaction, but ties can be deleted without any cost. We assume increasing marginal tie costs, and so the more ties one has, the more "effort" is required for another tie. Increasing marginal tie costs can also be interpreted as diminishing marginal returns of relationships; an equivalent model assumes constant tie costs and benefits that decrease in an actor's number of interactions. The total costs for an actor of having t ties are

$$k(t) = \alpha t + \beta t^2 \tag{2.2}$$

in which $\alpha > 0$ and $\beta \geq 0$. If $\beta = 0$, the tie costs are linear in the number of ties, and there are no increasing marginal costs of having more ties. Otherwise, there is a maximum number of network ties one can maintain given the payoffs one can obtain related to relationships. In one period of interactions the total payoff of an actor equals $\sum_{i \neq j} n_{ij} p_{ij} - k(n_{i+})$, where p_{ij} is the payoff i receives as a result of his own and j's behavior and $n_{i+} = \sum_{i \neq j} n_{ij}$ is the number of ties of actor i.[3] Finally, we define how the network is changed. Actors can add and sever ties. Because of the undirected nature of the ties, ties can only be created with the consent of both partners, but can be removed unilaterally. In other words, we assume a two-sided link formation process (Jackson and Wolinsky, 1996). We also assume that actors have full information on the behavior of all actors in the network.

[2] In terms of the two-person game, this would be the condition that (L,L) is not only the payoff-dominant equilibrium, but also the risk-dominant equilibrium (in the sense of Harsanyi and Selten, 1988).

[3] Although our cost function is provided in a very explicit form, the only crucial aspect is the number of ties an actor wants to have depending on whether he obtains mostly a or d in his relations. The possible number of ties is varied in the simulation over the complete relevant range.

2.3. Stable states

In line with related papers (Berninghaus and Vogt, 2006; Goyal and Vega-Redondo, 2005; Jackson and Watts, 2002), we search for stable networks first, i.e., networks in which

- No actor wants to change behavior,
- No actor wants to sever a tie, and
- No pair of actors wants to add a tie given the behaviors in the network.

This corresponds to *pairwise stability* (Jackson and Wolinsky, 1996) in networks in which only ties can be changed. The condition on behavior is added to guarantee stability in terms of the behavior chosen by the actors. In order to characterize stable networks, two definitions are helpful.

Definition 2.1. A (sub)network is t-full if and only if none of the actors have more than t ties and either (a) the addition of a tie causes at least one actor to have more than t ties, or (b) no ties can be added to the (sub)network.

Definition 2.2. For $z = \text{R}, \text{L}$, $\bar{t}_z$ is the maximum number of ties an actor wants to have if he chooses z, and all actors with whom he has a relation choose z as well.

These definitions are useful because the number of relations an actor wants to maintain is determined by the tie costs and the payoffs an actor can obtain. More specifically, a new relation with someone with whom an actor i can earn p_{ij} is only started if the number of ties actor i has will be less than or equal to $(p_{ij} - \alpha + \beta)/\beta$ after adding this new relation. Otherwise the marginal costs of adding this new tie are larger than its benefits. This implies that

$$\bar{t}_R = \left\lfloor \frac{a - \alpha + \beta}{2\beta} \right\rfloor \tag{2.3}$$

$$\bar{t}_L = \left\lfloor \frac{d - \alpha + \beta}{2\beta} \right\rfloor \tag{2.4}$$

Theorem 2.1. If tie costs are equal to $k(t) = \alpha t + \beta t^2$ ($\alpha > b$, and $\beta > 0$), networks are stable if and only if one of these conditions holds:

1. All actors choose the same behavior Z, where $Z = \text{R}, \text{L}$ and the network is $\bar{t}_Z$-full.
2. The network consists of two subgroups of actors choosing R and L and these subgroups are $\bar{t}_R$-full and $\bar{t}_L$-full, respectively, and there are no ties between the two subgroups.

Proof. As tie costs are always larger than b, a tie between actors with different behavior is never sustainable because one of the actors wants to sever the tie or change behavior. It is easily checked that the definition of $\bar{t}_Z$-full implies that no pair of actors want to add a tie and no actor wants to sever a tie. Because actors are not connected with actors who behave differently, they do not want to change behavior. All other networks are unstable, because in (sub)networks that are neither $\bar{t}_L$-full nor $\bar{t}_R$-full, some actors want to remove or add ties. $\square$

The theorem is a reformulation of the corresponding theorem in Berninghaus and Vogt (2006) for costs that are non-linear in the number of ties and for the conditions we want to consider. We extend the results of Berninghaus and Vogt (2006) not only by allowing for non-linear costs but also by studying the likelihood of the emergence of different structures depending on initial conditions; see the simulation section below. Jackson and Watts (2002) study coordination and endogenous formation of networks. They analyze which networks are stochastically stable in a specific dynamic context, showing analytically that homogeneous networks emerge in which all the actors coordinate on one behavior (cf. Young, 1998). Which behavior is chosen depends on the tie costs. In their analysis, conflict situations are excluded as possible long-term outcomes because they are less stable than non-conflict situations. In contrast with their study, we do not include “trembles,” but analyze how the likelihood of emerging structures in a deterministic dynamic environment depends on initial conditions. In a deterministic environment, networks with polarized behavior can be stable. This enables addressing the likelihood of conflict. Goyal and Vega-Redondo (2005) analyze a similar model, but assume one-sided link-formation cost. More importantly, they characterize only stable states without analyzing the likelihood that certain stable networks emerge.

The condition in the theorem that (sub)networks should be $\bar{t}_R$-full or $\bar{t}_L$-full suggests that the payoff and cost structure does not allow much variation in network structure. This is true in the sense that the density (proportion of ties present in the network) of the emerging network hardly varies with the cost function, the payoffs, and the distribution of behavior in the emerging network. Some variation is still possible. Let us consider as an example a nine-actor network in which everybody wants to have two ties. If behavior is homogeneous, both three closed triads and one circle of nine actors are stable networks. Another type of variation is related to the possibility that one actor may still have fewer ties than he wants to have but all other actors in his group have their maximum number of ties (see Jackson and Watts, 2002, for more details on variations in network structures). The simulation results presented below indeed confirm that the density of the emerging network is almost perfectly determined by the emerging distribution of behavior, the payoffs in the game, and the tie costs. Other network measures turn out to be hard to predict from the initial network structure. Therefore, the analyses below focus on predicting how the emerging distribution of behavior depends on initial conditions.

2.4. Simulation design

To analyze the effects of the model parameters on the emergence of stable networks by means of computer simulation, we systematically vary the initial conditions of the dynamic process, and relate the outcomes of the process to these conditions. The conditions include the initial network, the initial distribution of behavior, the payoff structure of the coordination game, the tie costs, and the adaptation rules in the dynamic process. Network size ranges from 2 to 50 actors. For networks of 2 to 8 actors, we include all 13,597 possible non-isomorphic networks. For networks of size 9 to 50, the number of possible networks becomes extremely large, and we take a sample stratified on the size (number of actors) and density of the network. In other words, we draw a set of random networks while density and size are fixed such that for each size there are about the same number of networks. Also for each density within each size a similar number of networks are drawn. Extreme densities for which the number of non-isomorphic networks is small are excluded. This results in a set of 95,729 networks. The probability for each actor to initially choose L equals 0.25, 0.50, or 0.75. This results in a wide range of distributions of actual initial behavior.[4] For reasons mention above, we vary only payoff c (see Figure 2.1) such that risk takes the values 0.467 ($c = 4$) and 0.538 ($c = 8$), fixing the other payoffs at $b = 0$, $a = 14$, and $d = 20$. With regard to the tie costs, we vary both the linear component and the quadratic component such that all $0 \leq \bar{t}_R \leq \bar{t}_L \leq n-1$ are possible. Linear cost α is chosen as an integer number from 1 through 16 (excluding c and 14 to avoid equalities with payoffs) each with an equal probability. For $\alpha < 14$, β is chosen such that all values of $\bar{t}_R$ ($0 < \bar{t}_R \leq n-1$) are equally likely by choosing $\beta = (14-\alpha)/(1+r(n-1))$, where $r \sim U[0,2]$. If $\alpha > 14$, we have $\bar{t}_R = 0$. In these situations, β is chosen such that all values of $\bar{t}_L$ ($0 < \bar{t}_L \leq n-1$) are equally likely by choosing $\beta = (20-\alpha)/(1+r(n-1))$, and again $r \sim U[0,2]$.

The dynamic model assumes discrete time. All actors simultaneously choose behavior in each period. We distinguish three methods for how actors change behavior and ties between periods. All methods assume actors to be myopic, i.e., optimizing under the restriction that the behavior and network of the previous period persists. It is also assumed that all actors know the behavior of all other actors.[5] The three methods are different in the relative adjustment rates of behavior and ties. Since these adjustment rates can be expected to affect outcomes (Skyrms and Pemantle, 2000), we compare three approaches:

[4] Using this procedure we obtain satisfactory amounts of variation in the independent variables that we want to use as independent variables to explain polarization and efficiency. Selectivity on the independent variables should in principle not effect the estimations of the regression models we use below. Nevertheless, we did some checks to ensure that the sampling procedure does not affect the substantive conclusions and this does not seem to be the case.

[5] A natural extension of the model is to relax this assumption and assume that actors only know behavior of actors they are connected to, or that are within a given distance in the network. First analyses of such extensions (to be reported elsewhere) however suggest that although assuming limited information in itself has interesting effects, the conclusions as reached in the current study are not undermined.

1. "Actor-based:" actors decide themselves which type of change they prefer. After every period, a random actor is allowed to change either behavior or *one* network tie, whatever is most beneficial. We assume the actor to choose the alternative with the highest payoff to him given the network and behavior of other actors in the preceding period. If multiple tie changes yield the same maximal benefit, one is selected at random.

2. "Alternating:" after every period, the actors in a random dyad decide whether they want to add or remove a tie between them. In addition, one random actor decides whether or not he wants to change behavior. This mechanism is similar to the one applied by Jackson and Watts (2002).

3. "Fast network:" now the network is allowed to change relatively fast compared to behavior. In comparison to the "alternating" version, not one but $[n/2]$ dyads are considered to change their tie after every period, and, next, one random actor is allowed to change behavior.

In any of the three methods, actors are informed about previous changes before they have to decide. If actors are indifferent between a change and the existing situation, we assume that actors do not change. We assume that actors do not change ties if both actors are indifferent between having and not having a tie. One additional assumption is needed here to handle actors who are not connected to anyone at all. These actors cannot adapt their behavior to any connected actor; moreover, in some cases, they are not able to connect to anyone else, regardless of the behavior they choose, because the other actors already have the optimal number of ties. In such cases, the unconnected actor changes behavior only if that might create opportunities to become connected in the next period. Otherwise, the actor does not change behavior. Clearly, other rules for changing behavior and relations can be conceived. We have chosen some of the most straightforward options on who might change what and when. However, further investigation how the dynamics depend on these options is only called for if outcomes differ dramatically between them. As we will see below, this is not the case.[6]

For each of the $13,597 + 95,729 = 109,326$ networks and each of the three versions of the dynamics, we varied the initial conditions in the following way. The simulation was done four times with different values of the quadratic cost component for each version of the dynamic process and for 13,597 networks of size up to 8. We distinguished only two levels of the quadratic cost component for the $95,729$ larger networks. One random choice was made to select the other conditions. Then, each of the initial conditions was repeated four times. These repetitions with the same set of initial conditions enable us to distinguish between stochastic variation in the outcomes that is related to variations in initial conditions and

[6] Additional simulations with simultaneous updating of behavior (not reported here) also lead to similar results.

randomness caused by the dynamic process. At each repetition, we let the process continue until convergence. This leads to $13,597 \times 48 + 95,729 \times 24 = 2,950,152$ simulation runs.

To analyze the effects of the initial network and to evaluate the emerging networks, we need to define and compute some key network characteristics:

- *Size* represents the number of actors.
- *Density* is the number of existing ties divided by $\text{size} \times (\text{size} - 1)/2$, the possible number of ties in the network (Wasserman and Faust, 1994, p. 101).
- *Degree* is the number of ties of an actor divided by $(\text{size} - 1)$. This "relative" definition differs slightly from the more common "absolute" definition (e.g., Wasserman and Faust, 1994, p. 100), but facilitates comparison across network sizes.
- *Centralization* is measured through the standard deviation of the degrees as defined above. This measure is derived directly from the variance in degrees as proposed by Snijders (1981); see also Wasserman and Faust (1994, p. 101). Other centralization measures (Freeman, 1979) lead to similar results in the analyses that we present below.
- *Distance*: a dyadic measure, the minimum path length between two actors (Wasserman and Faust, 1994, p. 110).
- *Segmentation* is the proportion of dyads at distance larger or equal to 3 of all dyads at distance larger or equal to 2 (Baerveldt and Snijders, 1994). In accordance with Baerveldt and Snijders, we count disconnected dyads as having a distance larger than 3. In the complete network, there is clearly no segmentation, which implies that this measure is equal to 0.
- *Segregation* measures the extent to which ties are limited to actors with the same behavior, rather than between actors with different behavior. It is defined as the expected number of between-group ties minus the observed number of between-group ties, divided by the expected number of between-group ties, assuming random matching (Freeman, 1978). This is the only measure that combines behavior and network structure. Because there are no ties between actors with the same behavior in the emerging networks, this measure always equals 1 for the emerging network.
- *Number of components* is the number of connected subgraphs (including isolates) that are not connected to the rest of the network. If the number of ties actors can have is low due to the high tie costs, groups of actors that behave in the same way might still fall apart into different components.

Table 2.1: Summary statistics of initial conditions in the simulation ($N = 1,544,100$)

Variable	Mean	SD	Min	Max
Size	24.506	13.757	2	50
Proportion actors playing L	0.500	0.225	0.022	0.974
$\bar{t}_R$	12.279	10.487	1	49
$\bar{t}_L$	18.282	13.407	1	49
Density	0.500	0.244	0	1
Centralization	0.238	0.130	0	0.873
Segregation	0.000	0.166	−3	1
Segmentation	0.164	0.267	0	1
Number of components	1.541	2.582	1	42

Size, density, and centralization are included because these represent the most basic network measures: the number of actors, the number of ties in the network, and the variation of ties between actors. The measures for segregation, segmentation, and components are included because they are particularly relevant for group formation processes in general and polarization in particular. While polarization indicates the extent to which actors are divided in subgroups due to their behavior, segmentation and components indicate the extent to which actors are divided in subgroups due to the network structure. Segregation is a measure for the extent to which behavioral and structural groupings coincide. The other measures are described only to define the network measures that we need in the analyses below. We also record the distribution of behavior in the initial and the emerging networks, the number of tie changes until convergence, and the number of behavior changes until convergence.

Some parameter values among the initial conditions are less interesting to analyze and are difficult to compare with the large set of cases. These are the following subsets:

- Cases for which $a > 14$. If $a > 14$, actors who choose R cannot maintain any ties. As a result, in over 80% of the cases all actors choose L. Actors will only stick to R if no one chooses L or all L-choosing actors do not want to add ties.

- Cases that start in a situation in which everyone chooses R or everyone chooses L. Behavior does not change in these cases and only the network ties are adapted until the network is $\bar{t}_R$-full or $\bar{t}_L$-full. Excluding these cases, $1,544,100$ cases remain for the analyses. Summary statistics for the initial conditions are shown in Table 2.1.

2.5. Simulation results

In this section, we explain properties of the stable networks in terms of polarization and efficiency by the model parameters, using statistical regression analysis. Table 2.2 presents

Table 2.2: Summary statistics of stable states ($N = 1,544,100$)

Variable	Mean	SD	Min	Max
Behavior changes per actor	0.250	0.161	0	4.897
Tie changes per dyad	0.492	0.191	0	2.333
Proportion actors playing L	0.518	0.423	0	1
Polarization of behavior	0.271	0.387	0	1
Density	0.591	0.295	0.02	1
Segmentation	0.379	0.430	0	1
Centralization	0.158	0.185	0	0.873
Number of components	1.664	1.712	1	25

summary statistics for the stable states. One result is that the relative number of behavioral changes (number of changes per actor) is considerably lower than the relative number of network changes (number of changes per dyad). This can be understood by recognizing that changing one's behavior has much more impact than changing one tie: in case of a behavioral change, the payoffs resulting from all interactions are affected, while changing a tie affects only one relation. Therefore, an actor mostly does not change behavior more than once. Usually actors need to adapt multiple relationships before they cannot improve their situations anymore.

The type of network dynamics does not have a large influence on the stable networks. In both "alternating" and "fast-network" dynamics, there are somewhat more changes in behavior and ties than in "actor-based" dynamics. Surprisingly, the number of tie changes is only marginally larger in "fast network" dynamics as in "alternating" dynamics although there are much more opportunities to change ties in the first one. Because of the limited differences between the types of dynamics, we only provide joint summary statistics in Table 2.2.

Polarization is our first important dependent variable. Recall that polarization is defined as $4p(\text{L})(1-p(\text{L}))$, in which $p(\text{L})$ is the proportion of actors choosing L. In more than 60% of the stable states behavior is homogeneous and, thus, without polarization. *Efficiency*, as expressed by the proportion of actors choosing the L-behavior, is a little above 50%, which is only slightly higher than the average efficiency in initial networks (see Table 2.1). The standard deviation of efficiency in stable states is rather large, reflecting the large proportion of homogeneous stable states. These percentages are not well-interpretable, because they depend to a large extent on our choice of initial conditions. Therefore, we focus below on how polarization and efficiency change with initial conditions.

2.5.1. *Predicting stable states I: Polarization*

To examine how polarization depends on the parameters of the simulation, we use linear regression models, with polarization as the dependent variable and the initial conditions (the network structure, the initial distribution of behavior, the tie costs, and the dynamics version) as independent variables. It is important to realize that standard errors reflect uncertainties due to the randomness of the dynamic process and misspecification of the model, not "sampling." In addition, with more than a million cases, even very small effects are significant. Therefore, we restrict ourselves to comparisons of relative sizes of standard errors between variables to provide information on the relative stability of the effects and we do not report significance levels. Standard errors are adapted for clustering within initial conditions (Rogers, 1993).

Polarization in stable states has an extremely skewed distribution, with over 60% of the stable networks showing homogeneous behavior where the value of polarization is zero. Our statistical model has to take this unusual distribution into account. We decide to model polarization with two separate analyses. First, we estimate a model predicting whether stable states are heterogeneous or not, using logistic regression. Then, we apply a linear regression model to predict the *extent* of polarization in the cases with polarization. The tie costs are included in the analysis as the difference in the number of ties an actor can have while choosing L as compared to the number of ties he can have while choosing R divided by size. The effect of tie costs is small and only the quadratic cost component turns out to be relevant for predicting behavior in the analysis. This component can be adequately summarized by $(\bar{t}_R - \bar{t}_L)/\text{size}$. We add dummies for the different types of dynamics.

Because the outcomes of the analyses strongly depend on initial polarization, we ran separate analyses for different categories of initial polarization. These analyses suggested adding interactions of initial polarization with density and segregation to our model. To facilitate the interpretation of the main effects, the density is centered at the mean before taking the interaction. Initial polarization and segregation already have a mean equal to 0.

In Table 2.3, we report the results separately for high and low risk, although the models do not differ substantially between risk values, to facilitate comparison with the results on efficiency reported below. The initial distribution of behavior, operationalized as the polarization at the start of the process, has a large positive effect on polarization in stable networks. There is a slightly higher probability of persisting polarization for the "alternating" version of the dynamics compared to the "actor-based," which is the reference category. The probability of persisting polarization is highest for the "fast-network" version, which is understandable since many network changes result in a network falling apart in groups with different behaviors before actors have an opportunity to adapt behavior. Substantively, this leads to the plausible interpretation that in societies in which relations are more volatile while people tend to stick more strongly to their behaviors or opinions, it is more likely that persistent disagreement arises that may induce conflicts. The difference between how many ties one wants to have choosing L compared to choosing R has a negative effect on the probability of any polarization as well as the extent of polarization. The probability of persistent polarization and

Table 2.3: Logistic and linear regression of behavioral polarization on simulation parameters with standard errors adapted for clustering within initial condition

	risk $= 0.467, c = 4$				risk $= 0.538, c = 8$			
	Probability of any polarization		Extent of polarization		Probability of any polarization		Extent of polarization	
Variable in initial condition	Coeff.	SE	Coeff.	SE	Coeff.	SE	Coeff.	SE
Polarization	5.739	(0.025)	0.820	(0.005)	6.500	(0.028)	0.926	(0.004)
Version "alternating"	0.094	(0.008)	0.023	(0.001)	0.282	(0.008)	0.014	(0.001)
Version "fast network"	0.601	(0.009)	0.023	(0.001)	0.801	(0.009)	0.015	(0.001)
$(\bar{t}_R - \bar{t}_L)/\text{size}$	−0.383	(0.020)	−0.013	(0.002)	−0.419	(0.020)	−0.030	(0.002)
Size	−0.016	(0.000)	−0.005	(0.000)	−0.018	(0.000)	−0.004	(0.000)
Density	−4.166	(0.051)	−0.003	(0.007)	−4.500	(0.056)	−0.083	(0.007)
Centralization	−1.150	(0.041)	0.182	(0.005)	−1.154	(0.041)	0.184	(0.005)
Segregation	0.301	(0.028)	−0.083	(0.004)	0.258	(0.031)	−0.074	(0.005)
Segmentation	0.636	(0.024)	0.047	(0.003)	0.680	(0.025)	0.037	(0.003)
Number of components	0.134	(0.003)	0.004	(0.000)	0.147	(0.003)	0.005	(0.000)
Polarization ×density	−0.402	(0.113)	−0.516	(0.014)	0.292	(0.122)	−0.389	(0.014)
Polarization ×centralization	3.468	(0.091)	0.315	(0.012)	3.447	(0.098)	0.275	(0.012)
Constant	−3.096	(0.033)	0.037	(0.005)	−3.666	(0.034)	−0.019	(0.004)
Number of observations	767,946		277,578		776,154		282,468	
(Pseudo) R^2	0.283		0.394		0.305		0.425	
Log-likelihood	-360,165				-353,668			

the extent of polarization decreases for larger networks.

Considering the network effects, density has a large negative effect on the probability of any polarization as well as on the extent of polarization (given positive polarization). Segregation has a small positive effect on the probability of any emerging polarization for average initial polarization and a small negative effect on the extent of emerging polarization. Segmentation increases the likelihood of persistent polarization as well as the extent of polarization. Centralization promotes the probability of homogeneity although given that there is some polarization, polarization will be larger if initial centralization is higher.

Due to the interaction of polarization and density, the effect of density is stronger for higher initial polarization. This implies that density really helps to solve the polarization problem, and if some disagreement persists, at most a small minority will stick to the "deviant" behavior, and will be excluded by the majority. Although the main effect of segregation is small the interaction effect with initial polarization is substantial. It shows that segregation enhances polarization if the initial network is rather polarized. Thus, if differences in behavior are aligned with initial group boundaries, the situation likely becomes only worse due to network dynamics.

2.5.2. Predicting stable states II: Efficiency

Efficiency has a somewhat peculiar distribution. In more than 60 percent of the stable states, behavior is homogeneous (efficiency is 0 or 1), while the remaining cases are more or less evenly distributed between the two extremes, such that the total distribution is U-shaped. An appropriate way to analyze such a dependent variable is logistic regression for grouped data, which predicts the number of successes (i.e., the number of actors choosing L), given network size. As the size of some network effects depend strongly on risk, we conduct the analyses separately for low risk and high risk (Table 2.4).

The coefficients refer to the log-odds of the predicted proportion. Clearly, the initial proportion L determines to a large extent the emerging distribution of behavior. If the group starts with a majority of actors choosing L, there is a large probability that an even larger majority or the whole group will ultimately choose L. The other effects have to be interpreted in terms of the extent to which they affect this baseline dynamic. The "alternating" and "fast network" dynamics have higher rates of efficiency as compared to "actor-based" dynamics for low risk, while this is the other way around for high risk. This can be interpreted in the following way. Changes in behavior mostly have the largest impact. Thus, if actors behave differently from most actors they interact with, these actors start changing behavior and thereafter they optimize their ties. As we know from models with static networks, adaptation of behavior leads to attraction to the risk-dominant equilibrium. In the "alternating" and "fast network" versions of the dynamics, changing ties often has to be done before behavioral changes. This

Table 2.4: Logistic regression for grouped data on the proportion of actors choosing L with standard errors adapted for clustering within initial condition

	risk $= 0.467, c = 4$		risk $= 0.538, c = 8$	
Variable in initial condition	Coeff.	SE	Coeff.	SE
Proportion L	16.728	(0.033)	16.778	(0.033)
Version "alternating"	0.129	(0.006)	−0.129	(0.006)
Version "fast network"	0.107	(0.006)	−0.023	(0.006)
$(\bar{t}_R - \bar{t}_L)/\text{size}$	0.189	(0.014)	0.357	(0.013)
Size	−0.004	(0.000)	−0.006	(0.000)
Density	1.479	(0.019)	−1.414	(0.018)
Centralization	−0.413	(0.032)	0.453	(0.030)
Segregation	−0.281	(0.009)	0.134	(0.009)
Segmentation	−0.058	(0.013)	−0.182	(0.012)
Number of components	−0.004	(0.000)	0.004	(0.000)
Prop. L × density	31.080	(0.120)	30.845	(0.122)
Prop. L × centralization	−23.000	(0.143)	−22.538	(0.144)
Constant	−8.144	(0.026)	−8.196	(0.025)
Number of observations	767,946		776,154	
Log-likelihood	-5,304,285		-5,373,063	

apparently decreases the attraction to the risk-dominant equilibrium. Except for the situation in which no ties were possible between actors choosing R, which are excluded from the analysis, tie costs have only small effects.

The network effects vary not only greatly with risk, but some also with the initial distribution of behavior. Therefore, we consider again interaction effects between the initial distribution of behavior and other variables. Especially the density and centralization effects depend considerably on initial behavior. In Table 2.4, interactions of the centered variables were included. In both models, strong and positive interaction effects exist between density and the initial distribution of behavior, such that the effect of density is negative for low initial proportions and positive for high initial proportions. For centralization and initial behavior, the interaction effect points in the other direction. The effects of density and centralization at their means are in different direction and they also switch if one compares high and low risk. At the means, density promotes efficiency for low risk, but hampers efficiency for high risk. In contrast, centralization hampers efficiency under low risk, but promotes efficiency under high risk. It is important to realize that the effects at the mean are relatively small compared to the interaction effects. Segregation has a negative effect on efficiency under low risk and a positive effect under high risk. This result indicates that segregation is to some extent able to stabilize the more risky equilibrium. Segmentation has a negative effect for low and high risk. Network size has only a small negative effect on efficiency.

To sum up, the strongest effect of the initial network structure on efficiency is that density galvanizes initial behavioral tendencies. If we start at a rather high level of efficiency, higher density leads to the emergence of even more efficient networks. However, if initial behavior is inefficient, the situation is likely to worsen with dense networks. Other network effects have a modest size. In the low-risk situation, segregation prevents the efficient behavior for diffusion through the whole network and therefore has a negative effect on total efficiency. However, segregation helps to maintain efficient behavior in parts of the network in the high-risk situation. On average, density decreases efficiency in the high-risk situation, which is surprising because density increased efficiency in high-risk situations for static networks (Buskens and Snijders, 2008). Centralization, in contrast to density, favors minority behavior in high-risk situations. An interpretation for this would be that if this minority has a central position in the network, they can induce the majority to change their behavior.

2.6. Conclusions and discussion

We develop a model explaining opinion formation or mass mobilization processes as, for example, occurred during the 2004 elections in the Ukraine and the "Monday Demonstrations" in Leipzig 1989–1991. The model formalizes the co-evolution of coordination and networks to study under what conditions it is more or less likely that the emergence of stable states leads to inefficient situations or situations with considerable conflict potential. We assume that actors are organized in a specific network in which coordination problems emerge. Initially, all actors behave in a certain way (or have a certain opinion or attitude toward a given issue). Depending on their own behavior and the behavior of the actors with whom they have relations, actors change their behavior and their relationships. After the network has developed into a stable situation we consider behavioral polarization and efficiency in the emerged stable situation.

It turns out that the initial network structure might affect the emerging distribution of behavior. The most important result with respect to polarization is that dense networks lead to more homogeneous behavior, while more segmented and segregated networks have the opposite effect. The latter effect becomes especially important if the initial behavior is already polarized. If polarized societies are indeed more prone to end up in conflicts, conflicts are less likely in dense cohesive societies, while conflicts are more likely in segregated and segmented societies, especially if the initial attitudes in sensitive issues are correlated with initial groups in social networks. The effect of centralization is multifold: centralization of the initial network decreases the probability that the network polarizes, but if there is polarization, centralization slightly promotes the extent of polarization.

The most salient finding on efficiency is that network density amplifies the effect of the initial distribution of behavior. The higher the density, the larger the effect of initial inefficient behavior on the inefficiency of the emerging network. In addition, if the initial behavior is equally distributed, density still increases the likelihood that the emerging behavior will be inefficient if the efficient behavior is risky. A similar effect is found for centralization, although smaller and in the opposite direction: centralization has a positive effect on efficiency if initial efficiency is low and a negative effect if initial efficiency is high. In addition, in larger, more segmented, and more segregated networks, behavior tends to become less efficient. These results are consistent with the fact that dictatorial states often survive for quite some time without having to cope with major mass demonstrations (cf. Hardin, 1995). As soon as a status quo with no major opposition is reached, it is difficult to turn this situation around. The centralization result shows that the best opportunity to escape from such an inefficient situation should come from central people who can mobilize others to start a revolt.

Although this study provides an extension to existing models on network formation in coordination problems and provides more insight into the relation between initial conditions and the emerging networks, there are still a number of limitations. First, our model assumes extreme opinion formation problems in which you can basically have only two opinions and if you do not agree a relationship would be very unattractive. In some of the examples mentioned such as the ones where you have to choose between standing up against the regime or remaining quiet, these assumptions are clearly more realistic than in less extreme situations. Second, while coordination problems represent the evolution of norms or opinions many social interactions might lead to conflicts of a different character. Such situations can also be related to, say, trust, cooperation, or distribution problems (Heckathorn, 1996). Then, it becomes likely that actors differentiate behavior between their partners. People might trust some people and distrust others. Therefore, the study of the evolution of networks and the possible emergence of related social problems can be extended to other types of interactions in settings where behavior co-evolves with networks. In addition, the theory developed has to be tested in experimental and real-world settings to corroborate the hypotheses developed in this chapter.

CHAPTER 3

Co-evolution of Conventions and Networks: An Experimental Study*

3.1. Introduction

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

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

3.1.1. Coordination, conventions, and networks

Often, coordination problems are resolved by conventions, i.e., behavioral patterns that are mutually expected and self-reinforcing (Lewis, 1969). Everyday conventions include traffic rules (driving left or right), technological standards (GSM frequencies), spelling standards, and rules for appropriate behavior, such as dress codes and table manners (Coleman, 1990; Elias, 1969; Ullmann-Margalit, 1977). These conventions share the feature that, once estab-

*This chapter was written in collaboration with Vincent Buskens. A slightly modified version is forthcoming in *Social Networks*.

lished, none of the actors involved has an incentive to deviate from the convention, provided that others do not deviate.

In some coordination problems, there is no reason to prefer one convention over another. The driving problem is a prominent example, but the situation probably also holds for many etiquette rules (e.g., does anyone really prefer “black tie” over “blue tie”?). In other cases, possible conventions are ranked according to their utility for all actors involved. For instance, we may choose between everyone being self-sufficient or everyone specializing in one type of labor, where the latter is more efficient (provided that others do the same). A further classification might be made according to the consequences of coordination failure. While some actions have higher or lower payoffs if they are also chosen by others, actions may have different consequences when they are not chosen by others. When two people fail to coordinate in trying to lift a heavy object, the one who does not lift is better off than the one who unsuccessfully lifts. Still, both would have preferred to lift the object together. Similar risks apply to many collective action problems (cf. Hardin, 1995). Situations where it is problematic to reach socially and individually optimal conventions are central to this chapter. The coordination game in Figure 3.1 represents such a situation for two actors.

	LEFT	RIGHT
LEFT	d,d	b,c
RIGHT	c,b	a,a

	LEFT	RIGHT
LEFT	20, 20	0, 10
RIGHT	10, 0	14, 14

Figure 3.1: Two-person version of the multi-person coordination game, in general form and with payoffs as in the simulation. $b < c < a < d$; $(a-b) > (d-c)$; $c = 4$ or $c = 8$.

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

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

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

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

A number of studies address theoretical perspectives on this question. Jackson and Watts (2002) propose a game-theoretic model in which actors play coordination games in an endogenous network and derive conditions under which constellations of networks and behavior are stochastically stable depending on the cost of maintaining ties. Their main finding is that, whereas various network structures are possible, behavioral choices in the coordination game become homogeneous. Berninghaus and Vogt (2006) analyze a similar (though deterministic) model, and find that networks can emerge consisting of multiple unconnected groups, while different conventions are maintained within various groups. (See also Goyal and Vega-Redondo, 2005; Skyrms and Pemantle, 2000, for related models.) These studies provide general characterizations of networks that might emerge, but many different constellations are still usually at least theoretically possible. To examine which stable structures are

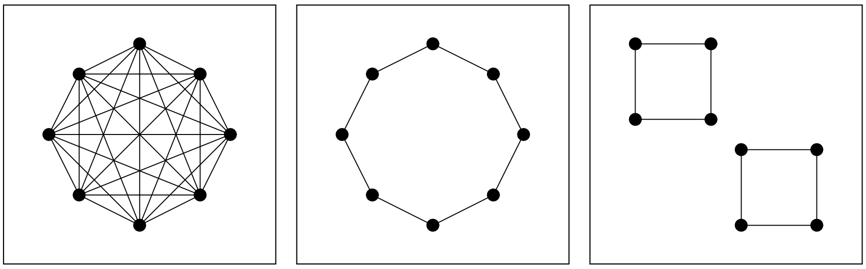

(a) Full **(b)** Circle **(c)** Two squares

Figure 3.2: Initial networks used in the experiment

more likely depending on various starting conditions, Buskens et al. (2008) apply computer simulations. They find that the *density* of the initial network has a strong influence on the way behavior develops: the higher the density, the stronger the influence of the initial behavioral distribution on the behavior's emergent distribution. Moreover, they find that if the initial network is more *segmented*, the higher the likelihood that two groups with different behavior will emerge. More generally, it is found that in a majority of the cases, a single convention is reached.

Our first aim is to empirically study which outcomes of the co-evolution process are more or less likely given initial conditions, such as initial network structures. For this purpose, our experimental setup includes three different initial eight-person networks: the full network, the circle network, and the "two-squares" network (see Figure 3.2). By choosing these three network structures, we vary both network density and segmentation.

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

3.1.2. An experimental approach

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

As an alternative, we suggest laboratory experiments. Experiments have a number of well-known advantages that make them the preferred research method for behavioral approaches: the experimenter is in considerable control of incentive structures, information availability, and other ingredients of game-theoretic models that are hard to measure in real-life situations (Camerer, 2003; Crawford, 1997). Moreover, behavior can be unambiguously observed in the laboratory. Accurate information on relations and behaviors can be recorded at every time-point, and one can relatively easily observe multiple networks and then examine the effects of various conditions at the network level.

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

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

We aim to take maximal advantage of experimental methods' benefits by explicitly making the experimental design and theoretical model as similar as possible. We use computer simulations to generate predictions tailor-made to our experimental conditions. In this way, we hope to minimize the misfit between the model and experimental conditions, and thereby obtain strong tests of our hypotheses.

While experimental research on network dynamics is not abundant, it has seen a considerable increase in the past years (see Kosfeld, 2004, for a survey). However, experiments on coordination in dynamic networks specifically are rather scarce. Corbae and Duffy (2008) conducted one other experiment on coordination games in dynamic 4-person networks. They find that, in the presence of shocks, only networks consisting of pairs are stable. By comparison, we use 8-person networks and focus on the efficiency of emerging behavior, the influence of initial network structures, and information availability.

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

3.2. Model and simulation

3.2.1. The model

First, we define the underlying game: a coordination game played in a network. Actors interact if there exists a tie between them. The actors with whom an actor interacts are called *neighbors*. Actors play a repeated multi-person coordination game as shown in Figure 3.1 with their neighbors, choosing only one of two actions against all neighbors. In each period, they receive payoffs from all interactions with their neighbors.

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

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

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

Both in the simulation and the experiment, the dynamic process described above runs in three phases per period:

1. Each actor initiates at most one change in ties. Either a new tie is proposed, an existing tie is severed, or nothing is changed. Bilateral proposals immediately result in ties, and removals are immediately implemented.

2. Actors choose to accept or reject incoming (unilateral) proposals from phase 1.

3. Actors choose an action in the multi-person coordination game in the network that results from phases 1 and 2. They receive their payoffs for this period.

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

Local Information: Actors only observe the behavior of their neighbors;

Global Information: Actors observe the behavior of all actors.

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

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

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

In its basic setup, our model is very similar to that analyzed by Jackson and Watts (2002), but there are some important differences. The first difference is that Jackson and Watts include random noise in their model, which allows them to derive general analytical results on stochastically stable networks. In contrast, we employ simulation methods to derive hypotheses on the effects of specific starting conditions on pairwise stable emerging network-behavior constellations. The inclusion of noise has some consequences for predictions on the emergence of conventions: while Jackson and Watts find that, in the limit, all actors will play the same strategy (one convention), Buskens et al. (2008) find that, without noise, constellations in which two different conventions persist in unconnected groups of actors are pairwise stable. In a sense, one may interpret these different findings as long-run predictions versus short-run predictions: while Jackson and Watts find that eventually, given enough external shocks, differences in behavior will vanish, Buskens et al. (2008) find that in the shorter run, where external shocks play a much smaller role, differences may persist. For many real social situations, as well as for for our experimental setting, these short-run predictions seem to be more informative. A second difference is that Jackson and Watts (2002) use sequential updating in their model, understandably for reasons of analytical tractability. In our model, actors update their behavior simultaneously, which corresponds with our experimental setup. Lastly, Jackson and Watts (2002) assume global information throughout, while we study local information in addition.

3.2.2. Analytic results

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

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

Definition 3.1. A (sub)network is t-full if and only if none of the actors have more than t ties and either (a) the addition of a tie causes at least one actor to have more than t ties, or (b) no ties can be added to the (sub)network.

Definition 3.2. $\bar{t}_z$ is the maximum number of ties an actor wants to have if she chooses action Z, where $Z \in \{\text{L(EFT)}, \text{R(IGHT)}\}$.

Theorem 3.1. Assume that tie costs are equal to $k(t) = \alpha t + \beta t^2$, where $\alpha > b$ and $\beta > 0$. Under global information, networks are stable if and only if they satisfy one of the following conditions:

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

Theorem 3.2. Assume that tie costs are equal to $k(t) = \alpha t + \beta t^2$, where $\alpha > b$ and $\beta > 0$. Under local information, networks are stable if and only if one of the conditions specified in Theorem 3.1 are satisfied, and if condition (2) is satisfied, there is at most one action $Z \in \{L, R\}$ for which at most one actor plays action Z and has less than $\bar{t}_z$ ties.

Theorem 3.1 states that, in any stable state, the network consists of one or more such subnetworks, *within* which all actors play the same action. The number of ties within each component is the maximum that the actors can afford given the payoffs to their actions and the tie costs. Within these boundaries, however, many different constellations are possible, such as only one component in which a single convention is played, several components playing the same convention, or several components playing different conventions (cf. Jackson and Watts, 2002). For the proof of Theorem 3.1, we refer to Buskens et al. (2008).

Under local information, our decision rule implies that actors in a homogeneous component guess that actors in other components also behave the same as themselves. The reason for this is that actors only observe their own behavior in the component. Therefore, we need the extra condition in Theorem 3.2. The argumentation for this extra condition is rather obvious given that, when this condition is not fulfilled, at least two actors will want to establish a new tie. After the actors discover that they play different actions, two things might happen. Either the tie is severed again, or one of the actors changes behavior and becomes part of the other subnetwork. In the last situation, the two actors continue to create and sever the tie.

For our experimental conditions, we characterize the general results on stable states more precisely. The experiment is conducted with groups of eight players, who play one of the two games shown in Figure 3.3. The tie costs are defined as $k(t_i) = 6t_i + t_i^2$, such that, for both games, $t_L = 7$ and $t_R = 5$. That is, given this cost function and the payoffs in the games, actors playing LEFT can profitably maintain, at most, seven ties (with other actors also playing

LEFT), while actors playing RIGHT can profitably maintain, at most, five ties (with other actors also playing RIGHT).

According to Theorem 3.1, this means that in our eight-actor setup the following types of constellations are stable under global information:

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

Under local information, both homogeneous constellations are still stable, but most of the heterogeneous constellations do not fulfill Theorem 3.2. Consider the case where there are two fully connected components: one component of three players playing LEFT, and one of five players playing RIGHT. Because behavior within the subnetworks is homogeneous, these actors conclude that the actors in the other subnetwork also play the same action, in which case it would be profitable to form ties. Thus, ties are formed. Subsequently, depending on the specific payoffs in the game, LEFT-players might switch to RIGHT. Alternatively, the tie is severed again in the next period. As a result, the only heterogeneous stable constellations under local information consist of six RIGHT-players in a full component and a dyad of LEFT-players.

3.2.3. Simulation

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

We study the following conditions:

- Two payoff sets for the coordination game as shown in Figure 3.3) that only vary in the risk involved in playing LEFT. Because the risk is lower in Game I than in Game II, we label them “low risk” and “high risk,” respectively; in both games, playing RIGHT is risk-dominant.

- Two information regimes as described above, in which actors can either observe the actions of all other actors or their neighbors' actions only.
- Three initial networks with eight actors: the full network, the circle network, and the two-squares network (see Figure 3.2).
- The propensity to play LEFT in the first period, using $\frac{2}{8}$, $\frac{3}{8}$, $\frac{4}{8}$, $\frac{5}{8}$, and $\frac{6}{8}$.

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

	LEFT	RIGHT
LEFT	20,20	0,10
RIGHT	10,0	14,14

(a) Game I (low risk)

	LEFT	RIGHT
LEFT	20,20	0,14
RIGHT	14,0	16,16

(b) Game II (high risk)

Figure 3.3: The coordination games as used in the experiment

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

For a more detailed examination of our results, we examined average efficiency levels and *heterogeneity* in stable states per information regime and initial network, separately for the two risk levels. To quantify the extent of heterogeneity, let π_L indicate the proportion of actors playing LEFT. Heterogeneity is defined as $h = 4\pi_L(1-\pi_L)$. The measure varies between 0 and 1, and equals 1 when π_L is exactly $\frac{1}{2}$ (indicating maximal heterogeneity), and equals 0 when π_L is either 0 or 1 (indicating homogeneity). The results are shown in Figures 3.4 and 3.5.

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

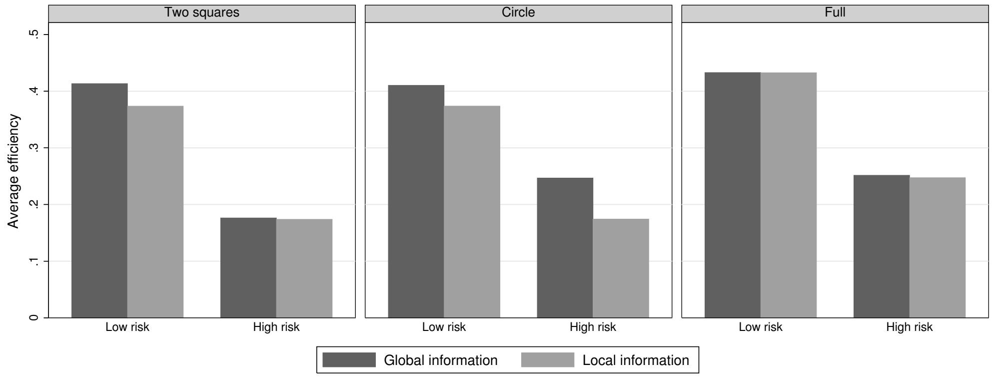

Figure 3.4: Average efficiency by risk level, initial network, and information regime (simulation results)

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

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

To estimate the effects of the independent variables on the proportion of actors playing LEFT, we treat each case (i.e., a group of eight actors) as a number of successes (actors playing LEFT), and apply logistic regression to predict the likelihood of success at the actor level. Because this approach inflates the number of observations, standard errors are adjusted

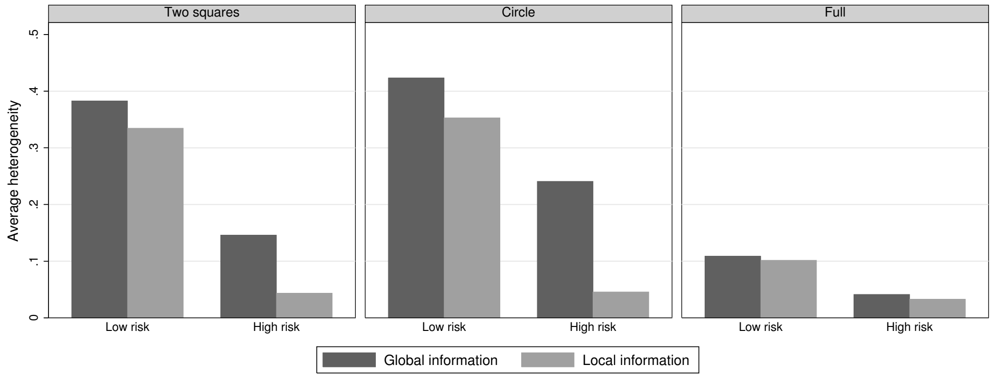

Figure 3.5: Average heterogeneity by risk level, initial network, and information regime (simulation results)

Table 3.1: Logistic regression for grouped data on efficiency (simulation results), per information regime and risk level

	Low risk				High risk			
	Information regime							
	Local		Global		Local		Global	
	Coeff.	SE	Coeff.	SE	Coeff.	SE	Coeff.	SE
PLEFT	11.30	(0.16)	10.65	(0.14)	17.64	(0.56)	12.74	(0.27)
FULL	-0.36	(0.08)	-0.71	(0.08)	-2.61	(0.37)	-3.39	(0.32)
FULL×PLEFT	17.34	(1.03)	18.49	(0.95)	22.72	(2.04)	22.64	(1.73)
Constant	-6.79	(0.10)	-6.04	(0.08)	-13.46	(0.41)	-9.44	(0.19)
Number of groups	6000		6000		6000		6000	
Log pseudolikelihood	-13461.51		-14032.85		-7097.06		-9985.47	
McFadden's pseudo R^2	0.58		0.57		0.70		0.61	

accordingly.[1] We conduct the analysis for the two risk levels and the two information regimes separately (Table 3.1).

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

[1] To estimate this model, we use the "blogit" procedure in Stata 9 (StataCorp, 2005).

3.2.4. Overview of microlevel and macrolevel hypotheses

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

The first microlevel hypothesis describes how actors decide on their actions when playing the coordination game. The model assumes that actors play according to a best-reply logic: they adapt their behavior to what their current neighbors played during the previous period. As stated in Section 3.2.1, actors only play the payoff-dominant action LEFT if the proportion of their neighbors also playing LEFT is at least as large as the risk threshold. These assumptions translate directly into the following hypothesis:

Hypothesis 3.1. Actors play LEFT if and only if the proportion of their neighbors who played LEFT in the previous period exceeds the risk threshold.

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

Hypothesis 3.2. Actors sever ties with neighbors who played an action different from their own in the previous period.

Hypothesis 3.2 applies to both global and local information, because when deleting a tie, actors are already aware of a specific actor's previous behavior. When creating *new* ties, the situation is different. Actors can only observe the behavior of potential neighbors under global information, leading to the following hypothesis:

Hypothesis 3.3. Under global information, actors create new ties with other actors who played the same action as their own action in the previous period.

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

Hypothesis 3.4. Under local information, the higher the proportion of an actor's neighbors who play the same action as this actor, the higher the likelihood that the actor proposes a new tie.

The macrolevel hypotheses are based on the results of the simulation. The first macrolevel hypothesis concerns the effect of risk on efficiency, and follows from Figure 3.4 and Table 3.1.

Hypothesis 3.5. The higher the risk involved in playing an efficient action, the lower the efficiency in stable states.

A second set of hypotheses about the effects of the initial network and information regime on efficiency follows from Figure 3.4 and Table 3.1.

Hypothesis 3.6. The higher the initial efficiency, the higher the efficiency in stable states.

Hypothesis 3.7. The higher the initial network's density, the lower the efficiency in stable states.

Hypothesis 3.8. The higher the initial network's density, the stronger the effect of the initial behavioral distribution on the emerging distribution of behavior.

Lastly, we derive a set of hypotheses that are concerned with effects on heterogeneity following from the results on heterogeneity in Figure 3.5.

Hypothesis 3.9. Higher initial network density leads to lower heterogeneity in stable states.

Hypothesis 3.10. In initially sparse networks, more information leads to higher heterogeneity in stable states.

3.3. Experimental design

We test these hypotheses in a computer-aided experiment, designed to reflect both the assumptions of the theoretical model and its implementation in the simulation model as closely as possible. Subjects played the repeated coordination games used in the simulation, choosing one strategy with all neighbors, while also having the opportunity to choose with whom they interacted. The experimental conditions included the three initial networks used in the simulation (see Figure 3.2), the two information regimes, and the two risk levels. The games used in the experiment are shown in Figure 3.3. Groups of eight subjects played one of the games for 15 subsequent periods. In each period, they faced the following decisions (in this order):

1. Decide whether to change one relation: that is, propose one new link to another subject or unilaterally sever one existing link;

2. Decide whether to accept incoming proposals from other subjects;
3. Choose their behavior (LEFT or RIGHT).

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

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

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

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

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

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

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

3.4. Results

We first present some general macrolevel results, and then proceed to test the macrolevel hypotheses (Section 3.4.1). In Section 3.4.2, we test the microlevel hypotheses on individual decisions in the coordination game. In Section 3.4.3, we test the microlevel hypotheses on linking decisions.

3.4.1. Macrolevel results

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

Figure 3.6 presents the average proportion of people choosing LEFT per initial network per period for each information regime. From this figure it is immediately clear that subjects had a tendency to choose LEFT. A majority of subjects chose LEFT in all but one condition from the first period. This is largely consistent with earlier experimental research on coordination games (cf. Straub, 1995). In the final period, 74% of all subjects played LEFT. Of all groups, fifteen converged on playing LEFT and only four converged on playing RIGHT. Eight of the fifteen LEFT-playing groups also reached a stable state in tie choices; that is, they had established a full network after fifteen periods. None of the RIGHT-playing groups managed to reach a stable state in tie choices, although their behavior was stable. One group reached a stable state with heterogeneous behavior under global information: after fifteen periods, the network consisted of a full component of five subjects playing LEFT, and a full component of three subjects playing RIGHT, with no links between the two components. A second group converged on this constellation under local information, although this was not stable given model assumptions.

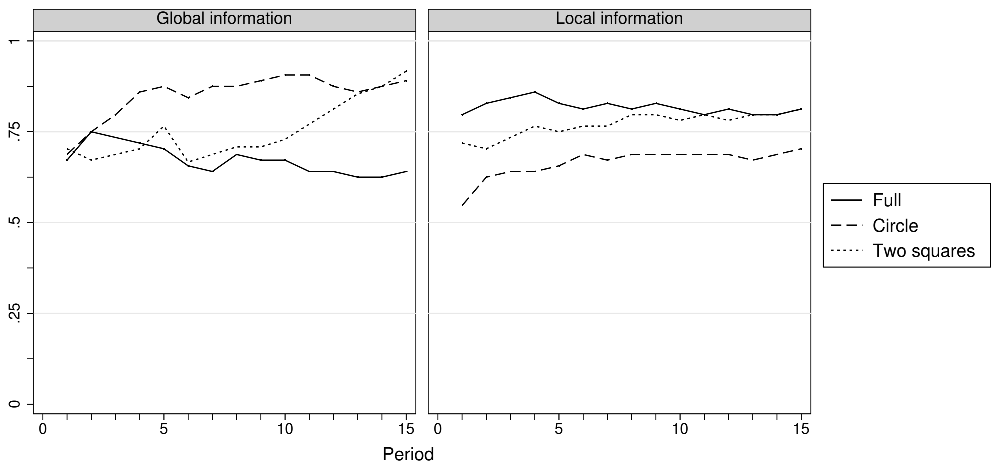

Figure 3.6: Average proportion LEFT-choices in six conditions

To investigate the macrolevel outcomes of the co-evolution process, we take all observations from the 15th period. Table 3.2 shows the *efficiency* and *heterogeneity* for each experimental condition, and also shows totals for the three initial networks and two information regimes. Figures 3.7 and 3.8 present a graphical impression of these results in a similar fashion as that for the simulation (see Figures 3.4 and 3.5).

To test Hypothesis 3.5, we compare the average efficiency between risk levels. Clearly, efficiency is lower in the high-risk condition (.92 vs. .67), which is significant using a Mann-Whitney test (p = .006) and supports Hypothesis 3.5.

In accordance with Hypothesis 3.7, efficiency is lowest for the full initial network (rightmost column of Table 3.2). This difference, however, is not significant. Moreover, closer inspection reveals that the difference only occurs under global information; under local information, the difference is either zero or reversed.

To test Hypotheses 3.6, 3.7, and 3.8 on efficiency, we run a regression analysis similar to that in the simulated data (see Table 3.1): logistic regression for grouped data on the number of LEFT-choices in period 15 as the dependent variable. The unit of observation in this analysis is the individual decision; standard errors are adjusted to account for the fact that individuals are clustered in 46 groups. To improve statistical power, we pool observations from different conditions and analyze them simultaneously, using control variables for the various conditions. As independent variables, we use the distribution of behavior in period 1 (PLEFT), a dummy variable indicating whether the initial network was the full network (FULL), and the interaction between the two. The main effect of the network dummy refers to the situation where efficiency in period 1 is .5, because this variable is centered at .5. Moreover, we include a dummy for local information (LOCAL) and a dummy for high-risk

Table 3.2: Efficiency and heterogeneity per experimental condition ($N = 46$)

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

Eff. = efficiency, Het. = heterogeneity

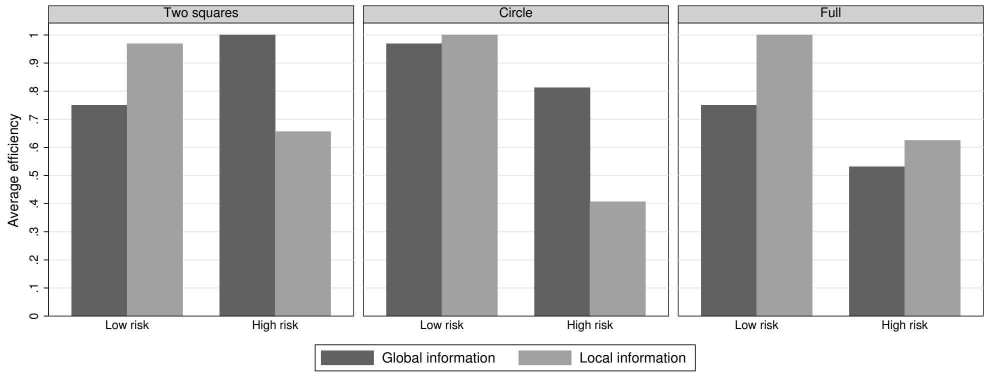

Figure 3.7: Average efficiency in the last period by risk level, initial network, and information regime (experimental results)

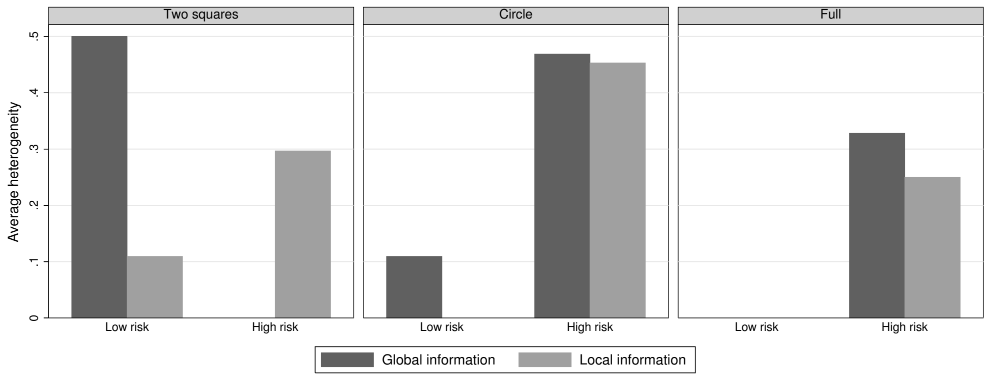

Figure 3.8: Average heterogeneity in the last period by risk level, initial network, and information regime (experimental results)

level (HIGH RISK). Table 3.3 presents the results of the analysis. There is a strong positive effect of initial behavior on behavior in the last period, which confirms Hypothesis 3.6. Also, there is a smaller but significant negative effect of starting in the full network, which supports Hypothesis 3.7.

There is a positive but insignificant interaction effect between the initial proportion playing LEFT and the full network. Hypothesis 3.8 can therefore not be confirmed.

Unlike in Table 3.2, there is no significant effect of HIGH RISK after controlling for the behavior in the first period. This suggests that differences between risk conditions in the final period are caused by differences in subjects' decisions in the first period, when they are not yet reacting to other subjects, and not by differences in the co-evolution process.

Table 3.2 shows average heterogeneity in the last period by experimental condition. Hy-

Table 3.3: Logistic regression for grouped data of the proportion playing LEFT (experimental results)

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

** $p < 0.05$

pothesis 3.9 predicts that a higher density of the initial network leads to lower heterogeneity. Heterogeneity is indeed lower in the full network condition than in the other two conditions combined. This difference, however, is not significant using a Mann-Whitney test. Thus, Hypothesis 3.9 cannot be supported.

Hypothesis 3.10 predicted that more information leads to greater heterogeneity, especially in low-density networks. Heterogeneity is lower under local information in the circle network (.23 vs. .29), but not in the two-squares network (.20 vs. .17). Moreover, these differences are not significant. Also, if we compare heterogeneity between local and global information over all networks, there is no significant difference.[2]

3.4.2. Individual behavior I: Decisions in the coordination game

We analyze individual behavior to assess the extent to which the model reflects actual decision-making by subjects. Hypothesis 3.1 states that subjects should play LEFT only if the proportion of their neighbors who played LEFT in the previous period exceeds the risk threshold. First, we plot average efficiency against the distribution of neighbors' behavior in the previous period for both risk levels separately. Under the assumption that people exclusively play a best reply against what their neighbors did in the previous period, subjects are expected to play RIGHT under low risk as long as less than 58% of their neighbors play LEFT. Under high risk, this percentage is 73%.

Figure 3.9 shows a somewhat more complicated picture. Clearly, subjects switch to LEFT at lower proportions than the .58 and .73 thresholds. Already, the proportion of LEFT-choices increases considerably in both cases at levels above.35. This indicates that subjects tend to

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

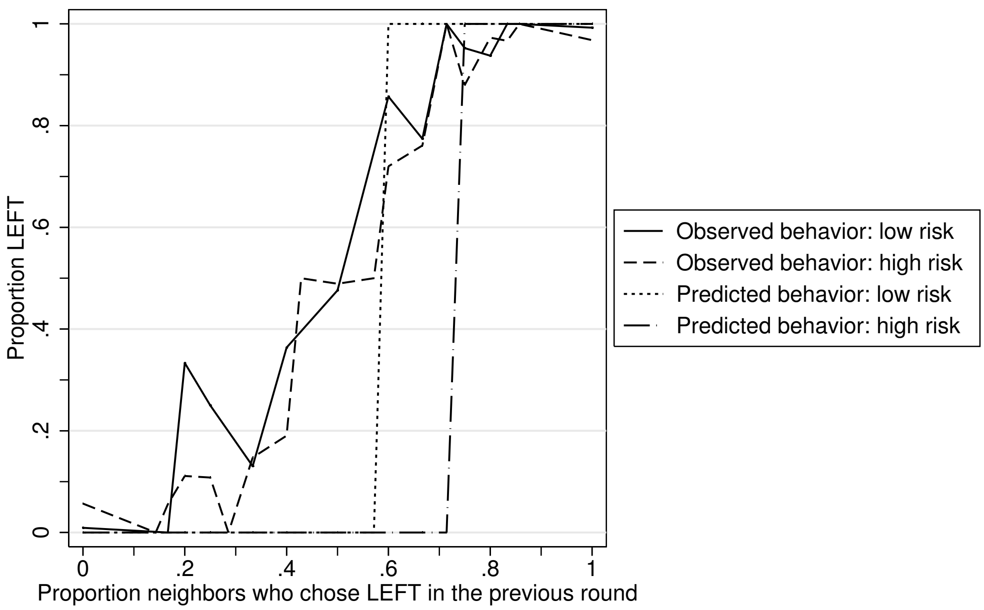

Figure 3.9: Proportion of subjects playing LEFT by proportion of their neighbors playing LEFT in the previous period

take on more risk than the simple myopic best-reply heuristic of the model assumes. However, subjects' behavior *is* strongly associated with their neighbors' behavior in the previous period. For the majority of subjects, the threshold for switching lies between .4 and .6. Although the payoff functions are different in both risk conditions, subjects' reactions to their neighbors are very similar.

Figure 3.9 shows only the bivariate relation between neighbors' behavior and the subjects' own behavior. To isolate the effect of neighbors' behavior from those of other effects, we conduct a logistic regression analysis with behavior as the binary dependent variable and subjects' decisions in every period as the unit of observation. Because observations within and between subjects are not independent, we use a model with a random intercept at the individual and the group levels. To test whether it is really the risk threshold that matters, we include a dummy variable indicating whether the proportion of neighbors who played LEFT in the previous period exceeds the threshold (NEIGTHRES), in addition to the proportion of the subject's neighbors playing LEFT (NEIGHLEFT) in the previous period. According to Hypothesis 3.1, NEIGTHRES should have a significant positive effect, but there should be no additional significant effect of the proportion of neighbors playing LEFT (NEIGHLEFT). As control variables, we include a dummy variable for high risk (HIGH RISK), the subject's own behavior in the previous period (EGOLEFT), and the proportion of the whole group playing

LEFT (GROUPLEFT). The latter variable is only included in the model for global information, as subjects were not informed about the behavior of the group beyond their own neighbors. Furthermore, we include time (PERIOD), the number of ties a subject has (NUMTIES), and whether the decision was made in the first part or the second part of the experiment (PART). Because this set of variables differs between information regimes, we estimate separate models for each regime (Table 3.4).

To test Hypothesis 3.1, we compare the effects of NEIGTHRES and NEIGHLEFT. Under global information, there is a significant positive effect of NEIGTHRES, as predicted: subjects are more likely to play LEFT if the proportion of neighbors who played LEFT exceeds the risk threshold. However, in contradiction with Hypothesis 3.1, there is an additional effect of neighbors' behavior (NEIGHLEFT). Under local information, there is only an effect of NEIGHLEFT, and no significant effect of the specific threshold. Thus, although these results again show that subjects do strongly react to their neighbors' behavior, their behavior does not conform exactly to the threshold effect as predicted by Hypothesis 3.1. There is also a significant effect of subjects' own behavior in the previous period, indicating some degree of behavioral inertia, especially under the local information condition. Under global information, the average behavior of the group in the previous period has a significant positive effect. Thus, controlling for the behavior of neighbors in the previous period, subjects tend to go along with the rest of the group. A possible explanation for this finding is that subjects are to some degree forward-looking; that is, they adapt to the behavior by non-neighbors in anticipation of becoming neighbors themselves. Under the global information condition, there is also a weakly significant effect of PART, indicating that subjects were more likely to play LEFT when they were playing the second set of 15 periods, after the reshuffling of groups. Nevertheless, additional analyses (not reported here), in which we estimate the same models but use only observations from the first 15 periods (before rematching) do not show substantive differences.

These findings suggest ways in which our actor model could be improved. One way would be to endow actors with forward-looking capabilities. Deriving precise implications from such an adapted model is beyond the scope of this chapter (i.e., it would require another simulation exercise as in Section 3.2). To speculate, however, we expect that if it is true that efficiency is higher than expected because RIGHT-playing subjects anticipate ties with LEFT-playing non-neighbors, this effect is smaller under local information (because subjects cannot observe non-neighbors). This is consistent with the results in Table 3.2 showing that efficiency is lower under local information.

Table 3.4: Logistic random intercept regression on playing LEFT, per information condition

	Information regime					
	Global			Local		
	Coeff.	SE	p	Coeff.	SE	p
NEIGTHRES	1.01^{**}	(0.45)	0.03	-0.06	(0.44)	0.89
NEIGHLEFT	3.78^{***}	(0.67)	0.00	5.98^{***}	(0.77)	0.00
EGOLEFT	1.52^{***}	(0.29)	0.00	2.88^{***}	(0.26)	0.00
GROUPLEFT	2.11^{**}	(0.77)	0.01			
HIGH RISK	-0.16	(0.37)	0.67	-1.10^{***}	(0.38)	0.00
PERIOD	0.01	(0.03)	0.69	0.06^{*}	(0.03)	0.07
NUMTIES	0.06	(0.08)	0.49	0.07	(0.08)	0.37
PART	0.76^{*}	(0.39)	0.05	0.25	(0.30)	0.40
Constant	-4.07^{***}	(0.55)	0.00	-4.05^{***}	(0.57)	0.00
Var. ind. level	1.21	(0.53)		0.85	(0.45)	
Var. group level	0.00	(0.00)		0.00	(0.00)	
Number of obs.	2520			22677		
Log likelihood	-361.25			-324.22		

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

3.4.3. Individual behavior II: Linking decisions

Figure 3.10 shows the proportion of ties created and dissolved by information regime. The results distinguish between pairs playing similar or dissimilar behavior. Ties between actors playing the same action (similar pairs) are created more frequently than ties between actors playing different actions (dissimilar pairs). Under local information, this relationship is much weaker because actors have to guess who will act similarly .[3] Ties between dissimilar individuals are more likely to be dissolved than ties between similar individuals. These results are all significant. Thus, subjects tend to sever ties with neighbors who behave differently in either information regime, and create ties more often with subjects who display similar behavior as far as they can observe. Hypotheses 3.2 and 3.3 are therefore confirmed.

Observations during the experiment illustrate this finding. At numerous occasions, it would happen that a subject in a group that had already reached a convention changed his or her behavior — possibly accidentally by pushing a wrong button on the screen, or for other reasons. In such cases, the neighbors of this subject would usually immediately start to dissolve the ties with this subject, such that after a few rounds she would we isolated. Often, the "deviant" subject then reconsidered, changed back to the conventional behavior and was readmitted into the group.

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

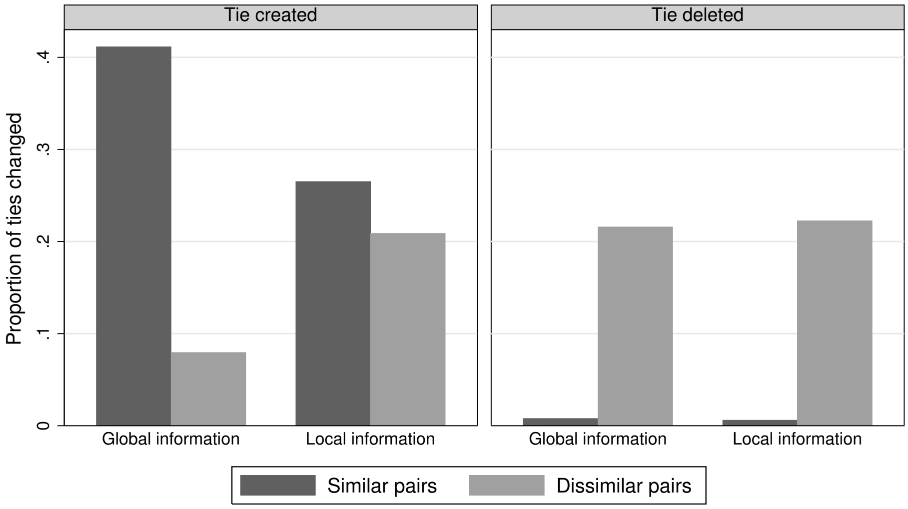

Figure 3.10: Proportions of ties changed

Figure 3.10 does not show how subjects make linking decisions under local information. In the simulation model, we assumed that actors use information on their neighbors' behavior and "project" this onto potential neighbors. To investigate whether subjects use this heuristic, we run a random intercept logistic regression analysis of the decision to create at least one new tie (either doing by making or accepting a proposal). The unit of analysis is the subject-period; the dependent variable is coded 1 if the subject created at least one new tie during the period, and 0 otherwise. Because we are interested in the creation of new ties, we only include cases in which the maximum profitable number of ties (5 or 7) was not yet reached. Moreover, to account for the interdependence of decisions within subjects over periods and between subjects within groups, we again add random intercepts at the subject level (14 periods per subject) and the group level (8 subjects per group).

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

Table 3.5 shows the results. First, the significant negative effect of NUMTIES reflects the increase in marginal tie costs as implemented in the cost function: the more ties a subject already has, the smaller the probability that she will form another one. Moreover, the likelihood of creating a new tie decreases with PERIOD and RISK. Contrary to expectations, the effect of NEIGHSIM is *negative*. Subjects are *less* likely to create new ties when they are more similar to their neighbors, and they thus are not using the heuristic assumed in the simulation model. Rather, subjects seem to assume that unknown potential neighbors are playing the alternative action.

Before we accept this conclusion,we test an alternative explanation in Model 2. It might be that this unexpected effect is the result of subjects remembering interactions in previous periods, for which the theoretical model does not allow. Figure 3.10 shows that subjects tend to sever ties with neighbors playing the alternative behavior, and it might be that after such a deletion they conclude that these previous neighbors (now invisible under local information) will persist in playing the different behavior. This logic would result in a negative effect of NEIGHSIM. To test this alternative explanation we introduce a new variable into the model, measuring the *number of times that a subject lost (deleted) a tie with a neighbor playing different behavior* (NUMDISSIM), and interact it with the existing variable NEIGHSIM. If the negative effect of NEIGHSIM does indeed depend on the history of play as sketched above, this interaction effect should be negative, and the main effect of NEIGHSIM should become positive. Because NUMDISSIM will correlate with time, we also include an interaction effect between PERIOD and NEIGHSIM to control for this. The interaction term with PERIOD is centered at PERIOD = 2. NUMDISSIM is not centered. This means that the main effect of NEIGHSIM now refers to the situation in which PERIOD = 2 and NUMDISSIM = 0; this is the effect of similarity to neighbors in period 2 before any ties with neighbors playing the different behavior were deleted.

The results for Model 2 in Table 3.5 show that this alternative explanation does not hold. The interaction term with NUMDISSIM is indeed negative, but not significant. The main effect of NEIGHSIM is now no longer significant either, but is still negative. The interaction effect with PERIOD however is significantly negative, indicating that the negative effect of NEIGHSIM becomes stronger in later periods of the experiment. This would indeed suggest that subjects make use of their knowledge of the history of play so far, albeit in a way different from the one proposed before.

3.5. Conclusions and discussion

We studied coordination in dynamic networks, focusing on efficiency and heterogeneity of emergent behavior, and on the influence of information availability. We specified a game-theoretic model and used simulation methods to generate specific predictions about the effects of initial conditions and limited information on the efficiency and heterogeneity of emergent behavior. We tested micro- and macrolevel hypotheses in a laboratory experiment. Three macrolevel hypotheses were confirmed. First, the behavior of subjects in the first period de-

Table 3.5: Random intercept logistic regression analysis on creating at least one new tie with local information

	Model 1			Model 2		
	Coeff.	SE	p	Coeff.	SE	p
EGOLEFT	0.01	(0.17)	0.95	0.02	(0.18)	0.91
NEIGHLEFT	0.05	(0.25)	0.84	0.09	(0.26)	0.72
NEIGHSIM	-2.15^{***}	(0.30)	0.00	-0.33	(0.43)	0.45
NUMTIES	-0.39^{***}	(0.05)	0.00	-0.50^{***}	(0.07)	0.00
PERIOD	-0.08^{***}	(0.02)	0.00	0.21^{***}	(0.07)	0.00
RISK	-1.02^{***}	(0.27)	0.00	-0.94^{***}	(0.26)	0.00
PART	0.48^{**}	(0.22)	0.03	0.33	(0.23)	0.15
NUMDISSIM				-0.05	(0.40)	0.91
NUMDISSIM×NEIGHSIM				-0.24^{***}	(0.06)	0.00
PERIOD×NEIGHSIM				-0.27^{***}	(0.08)	0.00
Constant	4.87^{***}	(0.67)	0.00	3.38^{***}	(0.77)	0.00
Var. ind. level	0.01	(0.07)		0.17	(0.11)	
Var. group level	0.24	(0.09)		0.15	(0.08)	
Number of obs.	1509			1509		
Log likelihood	-869.29			-850.30		

** $p < 0.05$; *** $p < 0.01$

termines to a large extent their behavior in the final period. Second, efficiency is lower if risk is higher. Third, efficiency is lower if the network is initially denser. For the remaining three macrolevel hypotheses on the effects of initial network structure and information availability, the results were always in the expected direction but not significant.

At the micro level, we found that, by and large, individual behavior appears to resemble behavior as assumed in the model, at least when information on the behavior of non-neighbors is available. Subjects adapt their behavior to that of their neighbors in the previous period (Hypothesis 3.1). In their choices of network relations, subjects have a clear preference for subjects who play the same action as themselves (Hypothesis 3.3) and exclude those who play the alternative action (Hypothesis 3.2). Moreover, many experimental groups managed to converge on stable constellations that were theoretically predicted involving both efficient conventions and situations in which no single convention was reached.

In some respects, the behavior of subjects also deviates from the behavior assumed in the model: subjects more easily opt for the payoff-dominant action than would be expected from myopic best-reply behavior (Hypothesis 3.1). Moreover, if subjects are informed on the behavior of the whole group rather than only their neighbors, they are also influenced by the behavior of those who are not their neighbors, which suggests some anticipation of future interactions with these other subjects. For local information, it was theorized that subjects use the average behavior of their neighbors as a predictor for the behavior of potential neighbors

(Hypothesis 3.4). However, subjects seem to use the behavior of their neighbors as a predictor for what their neighbors are *not* doing. Furthermore, we showed that this effect becomes stronger in later periods of the experiment, which suggests that subjects use the history of interactions in their decisions.

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

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

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

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

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

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

CHAPTER 4

Alcohol Use among Adolescents as a Coordination Problem in a Dynamic Network*

4.1. Introduction

Adolescence is a life stage in which many forms of problematic behavior reach their peak (Steinberg and Morris, 2001), among which delinquency and substance abuse. Even though there is little evidence that problematic behavior in adulthood *originates* from behavior during adolescence (Moffitt, 1993), these types of behavior may have considerable impact on adolescents themselves and on society in general. Especially substance abuse, on which we focus in this chapter, is associated with problems in other areas, including delinquency, mental health problems, and problems with educational attainment (Newcomb and Bentler, 1989). Social influence by peer groups has often been named as one of the important factors that can trigger various types of problematic behavior, including alcohol and drug abuse, (e.g., Hawkins et al., 1996; Moffitt, 1993; Newcomb and Bentler, 1989). Consequently, relationships among adolescents have been the focus of a considerable body of literature (Giordano, 2003).

Issues of peer influence and selection have been of major concern within this context. On the one hand, it has been found that adolescents are sensitive to the influence of peers (Bot et al., 2005; Graham et al., 1991; Swadi, 1999). On the other hand, it has also been recognized that not only are adolescents influenced by their social environment, but also choose peers as friends who are similar to themselves, leading to network homophily (Lazarsfeld and Merton, 1954; McPherson et al., 2001). Disentangling these two simultaneous processes poses an ongoing theoretical and methodological challenge (Bauman and Ennett, 1996; Kirke, 2004).

Small and medium-sized groups of adolescents have since long been a popular setting for sociologists to study social networks, and the emergence of cultures and norms. Schools, in particular, are an attractive setting to study these topics because they constitute relatively well-delineated social contexts, in which complex processes can be observed relatively easily

*This chapter was written in collaboration with Andrea Knecht.

(some prominent examples include Bearman et al., 2004; Coleman, 1961; Epstein and Karweit, 1983). In this sense, a school constitutes a kind of "social microcosm," or, as Coleman (1961, p. 9) put it, a "society within society."

Research on diffusion dynamics has shown that the overall structure of a network has important consequences for emerging patterns of behavior (e.g., Centola and Macy, 2007; Granovetter, 1973; Watts, 2002). Yet, while studies on adolescent behavior often emphasize the importance of social networks, most research focuses on *individual-level* explanatory factors such as attributes of personal networks (e.g., Graham et al., 1991) or individual network positions (e.g., Ennett and Bauman, 1993). A possible reason for this divergence is that theorizing on the effects of the macrolevel network structure requires the specification of the mechanisms that connect the macrolevel network structure with microlevel individual behavior, and conversely, microlevel behavior with macrolevel collective outcomes. Given the interdependencies involved in interpersonal influence processes, specifying these macro-to-micro and micro-to-macro mechanisms is not trivial (cf. Coleman, 1990), and the inclusion of co-evolving network complicates matters further.

Theoretical tools that were particularly suited to deal with interdependencies of individual action and the transitions between different levels of aggregation, are game theory and agent-based modelling. In this chapter, we follow an approach based on these tools to study effects of the macrolevel network structure on alcohol use. We aim to explain differences in average alcohol use between *groups* of adolescents (in this case, school classes) by means of the social network structure in the group at the start of the influence-selection process. We formulate the following research question:

To what extent can properties of the initial overall network structure explain differences in average alcohol use between school classes?

In answering this question, we use a theoretical model based on *strategic* interaction, in which we model social influence as a coordination game. In this model, we explicitly account for endogenous evolution of the network. Before we explain how our approach complements existing approaches to co-evolution of networks and behavior, we outline our theoretical model.

4.1.1. Coordination, influence, and alcohol use

In developing our model, we assume that when deciding whether to use alcohol, adolescents have incentives to choose the same behavior as the peers they interact with, for reasons that we outline below. This implies an interdependence between adolescents' decisions resembling the strategic structure of a *coordination game*. We assume that the outcome of this "game" determines the utility that a student derives from each of his friendship relations. The outcome of the game, in turn, depends on the behavioral choices of both students involved in a friendship relation. The game is shown in Figure 4.1, in general form and with numerical payoffs (only the relative values of the payoffs matter). In a friendship *network*, students play this game with multiple friends simultaneously. For ease of exposition, however, we first discuss the two-player setup and afterwards generalize this to a network setting.

Each player in the game has two options: either to drink or not to drink alcohol. This game is a coordination game because it has two Nash equilibria in which both players choose the same action. Thus, the players prefer to play the same action as their interaction partner, reflecting the basic idea that adolescents face a *pressure for conformity* in their behavior.

There are several reasons why such a peer influence might exist. First, there may be some intrinsic reason why an activity brings more utility if it is coordinated with others. A trivial example is games: playing ball is more fun if you can coordinate with someone to play with you. Similarly, drinking alcohol is most likely a social behavior, in the sense that the utility of use is higher if it is shared with someone else. A second incentive for coordination is *imitation*. During adolescence, people go through important changes, and consequently face many uncertainties. As a result, adolescents may look at their peers as a reference to help them determine which behavior is appropriate (Marsden and Friedkin, 1993). Third, there may be norms among groups of adolescents that promote conformity in general, also in the area of substance use (Sherif and Sherif, 1964). These three distinct pressures all lead to incentives for adolescents to coordinate their behavior. In the coordination game in Figure 4.1, this is represented by the fact that the payoffs for both players are higher when they choose the same action than when they choose different actions (i.e., $a > b$ and $d > c$).

The preference to coordinate does not imply that students are necessarily indifferent between using alcohol or not. In the structure of the game, we assume that students prefer abstinence to using alcohol, given that they coordinate their behavior. In Figure 4.1, this is reflected by the fact that $d > a$. This reflects the assumption that the disadvantages of using alcohol in terms of the financial costs, long-term health risks, and possible sanctions by parents and teachers are higher than the short-term gains.

Another feature of this game is that the "punishment" for failure of coordination differs between the actions. In our numerical example, if Player 1 chooses to drink and Player 2 chooses not to drink, the payoff of Player 1 is 8 while the payoff of Player 2 is 0. In other words, in a situation where one drinks and the other does not, this is worse for the one who does not drink. This assumption signifies that the use of alcohol has a number of effects on social behavior that have a negative impact on especially the social environment of the user, rather than on the user herself. For instance, drinking can lead to inflation of the ego, increased risk-taking, or downright aggression (Steele and Josephs, 1990). In situations where some use alcohol and others don't, such behaviors can have negative consequences especially for the non-users. In this sense, we can say that abstinence involves a higher risk on a lower payoff, such that the equilibrium in which both players drink can be classified in game- theoretic terms as a *risk-dominant* equilibrium (Harsanyi and Selten, 1988). In Figure 4.1, this feature is established by the assumption that $(a-b) > (d-c)$.

	Drink	Not drink
Drink	a,a	c,b
Not Drink	b,c	d,d

	Drink	Not drink
Drink	$14,14$	$8,0$
Not Drink	$0,8$	$20,20$

Figure 4.1: Alcohol use as a coordination game, in general form and with numerical payoffs.
$b < c < a < d$; $(a-b) > (d-c)$.

The game in Figure 4.1 provides a simple model for two actors in a friendship relation, deciding whether or not to drink. In reality, such choices take place in friendship *networks*, in which students maintain relations with several friends. We can extend our two-player model to a network model by assuming that every player plays with *several* interaction partners simultaneously. Each player can choose one action against all interaction partners (i.e., either consume alcohol or not), and receive the payoff as in Figure 4.1 from every interaction. Thus, students receive utility from every friendship relation separately, but must also adjust their behavior to several friends simultaneously.

A crucial assumption of this model is that actors cannot differentiate their behavior between different interaction partners: a student cannot choose to drink with one friend and not with the other. The rationale for this assumption is the idea that by choosing to drink in some situations, students make a general decision to "be a drinker," and thereby influence *all* their relations. Obviously, this is a simplification. In reality, it is very well conceivable that students behave differently with one friend than with the other. However, such differentiation is likely to be more costly in terms of effort than the situation in which one can simply use one mode of behavior. By considering alcohol use as an individual attribute rather than as relational choice, we also conform to the standard in the literature (e.g., Bot et al., 2005; Graham et al., 1991; Kirke, 2004), in which substance use is typically analyzed as an individual characteristic and not as a relational attribute.

In order to include selection, we assume that actors can choose with whom to interact. We assume that the coordination game is played repeatedly, and that in every iteration of the game, actors can choose both their behavior in the game, and with whom to play. When updating their behavior or their network links, actors choose the optimal response to what their interaction partners did in the previous interaction. Figure 4.2 illustrates this process. Actor i plays the coordination game with her two "neighbors" j and k. Actors i and j play one type of behavior, while k and l play the other type of behavior. At the same time, i and l have the opportunity to form a new link. We assume that maintaining social relations costs time and effort, such that every actor has to pay some cost for every link she maintains. The link between i and l is only formed if the expected benefits will outweigh the costs for both i and l.

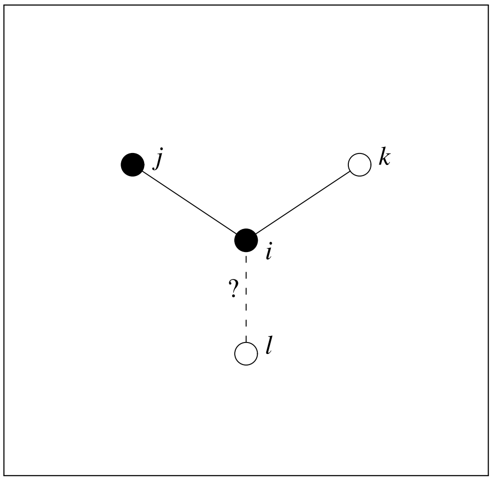

Figure 4.2: A coordination game in a dynamic network

Thus, we have sketched the outlines of a simple game-theoretic model for selection and influence in a dynamic network. We use a slightly more complex version of this basic model to derive specific hypotheses to be tested in the context of alcohol use. We first discuss how such a model could contribute to understanding influence and selection processes in comparison to earlier approaches.

4.1.2. Approaches to the study of selection and influence

Disentangling the simultaneous effects of influence and selection has been the focus of considerable research effort in the past decade. The basic problem was already recognized by Cohen (1977) and Kandel (1978), who noted that the effect of influence is likely to be overestimated if selection effects are ignored (the reverse also holds). A major breakthrough in the study of influence and selection was achieved by the introduction of actor-oriented, simulation-driven statistical models for longitudinal network data, implemented in the SIENA software (Snijders, 2001). In brief, this method works as follows: Given a set of subsequent observations of a social network, these panel data are considered to constitute "snapshots" of an underlying dynamic process. It is assumed that this process is driven by actors trying to optimize an *objective function* (roughly, a random utility function), both with regard to their own network position and with regard to their own behavior. The aim of the method is to estimate the components of this objective function based on the observed "snapshots" of this process. This is achieved by simulating the underlying process, and optimizing the fit between the simulated process and the observations based on maximum likelihood criteria. The effects of influence and selection can be identified as separate coefficients in the estimated objective function. Detailed expositions of the SIENA method can be found in Snijders (2006) and Snijders et al. (2007). Examples of applications of the method studying co-evolution of behavior and networks can be found in Steglich et al. (2006), Light and Dishion (2007), and Knecht (2008), among others.

Using the data we also analyze in this chapter, Knecht (2008) applied SIENA to study the co-evolution of friendship networks and alcohol use to disentangle influence and selection effects. The results showed a clear selection effect (adolescents select friends based on similarity in drinking behavior), but provided only weak evidence for social influence.

In this chapter, we take a somewhat different approach. We first use the coordination model to derive predictions about the relation between the initial state of a group (in this application, a school class) in terms of network structure and behavior and aggregate behavior at a later stage. We compute aggregate statistics on the groups in our data to create a dataset of *groups*, each observed at two different time-points. We then test statistically whether the groups in the data developed in the way that was predicted by the model.

To explain how this approach compares to SIENA, we highlight the most important differences and similarities. First and foremost, our method tests predictions at the *macrolevel* (that is, at the level of a whole network), while SIENA tests predictions at the individual level (namely, hypotheses about components of the utility function). In a sense, we can say that SIENA tests hypotheses on how the process works at the microlevel, while we test hypotheses on the *outcomes* at the macrolevel, assuming a specific theory on how the process works at the microlevel. We explicitly use initial states of the co-evolution process to predict outcomes, while SIENA merely conditions its simulations on the initial states. In this sense, the approaches are complementary. In the treatment of the data, the main difference is that we use only aggregate measures of network structure and behavior to test our predictions, while SIENA needs individual-level information.

A second difference between our approach and the SIENA approach is that rather than estimating properties of the network dynamics from the data, we assume a very *explicit* model of network formation and co-evolution. That is, we specify in detail the precise strategic nature of the interaction by means of the coordination game. In this respect, our model is more detailed in the specification of actors' incentives than SIENA. Similarly, we do not estimate the rate at which network changes can occur (the *rate function* in SIENA terminology), but instead make specific assumptions on this rate (in Buskens et al., 2008, we show that outcomes are robust under different assumptions on the speed of network dynamics).

Besides differences, there are also a number of similarities between our analysis and a typical SIENA analysis. First, and most importantly, both methods assume the possibility of an underlying co-evolution process in which both individual characteristics and the network change. Second, both methods rely on simulation to handle the complexity implied by a co-evolution process. Third, both methods assume that changes take place in "microsteps": only one link, or one actor's behavior, changes at a time.

By testing macrolevel hypotheses, we avoid two disadvantages of SIENA. The first problem concerns the theoretical interpretation of SIENA results; the second (related) problem concerns data requirements. Consider a co-evolution process based on coordination in a dynamic network. Suppose that, at some point, the network has reached a stable state in terms of behavior: no actor can improve her utility by changing her behavior, but some *can* improve their utility by changing ties. We observe the network for a few more "snapshots," in which very little change in behavior is observed (because it is already in, or close to, a stable state), but some tie changes are observed. SIENA results on these observations would probably indicate that there is no influence going on, but only selection. The theoretical implication of this would be that only selection plays a role in the considerations of the actors, even though stability in behavior is implied by a coordination game. In other words, social influence can only be identified as change, although theoretically, social influence could also be expressed by *stability*. The second problem is related to the first. Estimating effects of selection and influence in a SIENA model requires that enough changes in both the network and behavior are observed in the data. If an observed network is close to stability on one of the dimensions, this may lead to estimation problems in SIENA, forcing the analyst to drop these observations from the data (cf. Knecht, 2008, who can analyze only 78 out of 120 school classes).

The alternative approach we use here suffers less from these issues. Because our approach relies on examining *outcomes* of an underlying co-evolution process, it circumvents the problem of the interpretation of results that we identified above. Even if, in a given network, only relational changes are observed, our model still provides predictions that follow a model that assumes both selection and influence. Thus, also networks that are observed in a stable state can be compared with the predictions, because the theoretical model *predicts* what stable states should look like. As a result, also networks that are relatively stable contribute to the test of the hypotheses, while they would lead to estimation problems in SIENA. In the following, we apply our model to predict properties of stable states from initial conditions in terms of behavior and network structure. Results on this relation between initial conditions and outcomes can be used to identify selection- and influence effects. As we will show in the analysis, we find evidence that the emerging distribution of alcohol use is influenced by the *initial* network structure. We argue that this is a strong indication for the existence of influence: if the process would be driven only by selection, the network should only adapt to the distribution of behavior, and not vice versa. In principle, this logic could also be applied to identify selection effects: in that case, the model would predict an effect of the initial distribution of behavior on the emerging network structure. In this chapter, however, we restrict the analysis to emerging behavior as the dependent variable because we are mainly interested in explaining differences in alcohol use.

Because our model predicts properties of stable states, our method does not require that enough changes in ties and behavior are observed in each network to make estimation possible. As a result, we are able to use a larger share of the available data to test our hypotheses. However, our approach *does* require that enough variation in initial conditions and outcomes is present in the data. In this study we meet this demand by using data on a large number of groups. Another disadvantage of our approach is that we cannot directly test hypotheses on individual decision processes, as is possible with SIENA.

Although our theoretical model takes into account that behavior and the network co-evolve, it provides the most informative predictions on the emerging distribution of *behavior*, and less in terms of the emerging *network*. The analyses in Buskens et al. (2008) show that the emerging network is almost perfectly determined by the emerging distribution of behavior: actors maintain only links with other actors with the same behavior. These predictions can be tested on the individual level, and therefore an empirical analysis of network formation based on our model would not provide addition insight as compared to a SIENA analysis. Moreover, from a substantive point of view, we are mainly interested in explaining differences in alcohol use, and less in explaining network structures.

4.2. Predictions

The theoretical model outlined above describes a dynamic process in which the network and behavior co-evolve. It can be shown analytically that this process may converge to a large variety of stable states (Berninghaus and Vogt, 2006; Buskens et al., 2008; Jackson and Watts, 2002). Thus, by itself, this model does not yet provide precise predictions on which stable states will occur. To obtain more informative predictions on which stable states are more likely to occur than others, one can use computer simulations. Buskens et al. (2008) ran extensive computer simulations of the same model as we use here. The simulations resulted in a large dataset of initial conditions and resulting outcomes. This dataset was subsequently analyzed using conventional regression analysis methods, yielding predictions on how outcomes in terms of aggregate behavior depend on initial conditions, in terms of the initial distribution of behavior and the initial network structure.

We use these results to derive specific hypotheses on development of alcohol use among adolescents in school classes. To connect the simulation model with our empirical setting, we need to make a number of assumptions. We assume that alcohol use is for adolescents essentially a coordination game, as explained above: adolescents have incentives to display the same behavior as those they interact with; jointly using alcohol has (*ceteris paribus*) a lower utility than abstinence; and the risk involved in unilaterally using alcohol is lower than the risk of not unilateral abstinence. The network is the friendship network between adolescents in a school class. The underlying assumption is that the group of classmates constitutes a salient interaction context for adolescents. Because they spend a considerable share of their time at school among peers, we expect that they adjust their behavior to interaction partners from this group. This implies that we also assume that adolescents are not influenced by relations they might have *outside* their class. Admittedly, this is probably an unrealistic simplification of the situation. However, there is evidence that the students have most of their friendships and also their most important friendships in school classes (Knecht and Friemel, 2008).

We observe school classes at four different points in time. Applying our model, we aim to predict the behavior in the last observed period from the first observed period. Accordingly, all hypotheses are formulated in terms of effects of properties of a group (school class) at t_1 on properties of the group at t_4. A key group-level property at t_1 is the initial distribution of the *propensities* to choose alcohol use or abstinence. This propensity determines the likelihood that a student will use alcohol in the first "round of the game." The distribution of propensities determines only how the process starts; in subsequent time points, actions are exclusively the result of interaction in the coordination game.

As to effects of the initial network structure, we focus on the effects of network *density* and network *centralization*. In the simulation analyses by Buskens et al. (2008), these measures proved to have the largest effects on emerging behavior. *Density* refers to the extent to proportion of ties resent in the network, given the number of members of the network. *Centralization* is the extent to which ties are concentrated with relatively few individuals, rather than distributed uniformly among the network members (Snijders, 1981).

The first hypothesis serves as a "baseline" hypothesis, and relates the initial propensity to use alcohol to the resulting behavior at the end of the process:

Hypothesis 4.1. The higher the average propensity to use alcohol in a class at t_1, the higher the proportion using alcohol at t_4.

The next two hypotheses concern effects of the *initial density* of the network. Buskens et al. (2008) report that a higher initial density leads to a higher proportion of actors choosing the risk-dominant action, when starting from a situation in which the initial propensity is 50%. The intuition is that because the risk-dominant action is a stronger "attractor" (Young, 1998), more interaction at the start of the process (i.e., a higher density) leads more easily to convergence to the risk-dominant equilibrium. However, when the initial propensity is skewed, a higher initial density favors the action towards which the process already tended from the start.

This implies a hypothesis on a *main effect* of initial network density, and a hypothesis on an *interaction effect* between initial density and the initial propensity to use alcohol.

Hypothesis 4.2. The higher the density of the network in a class with an equal distribution of initial propensity to use alcohol at t_1, the higher the proportion of students using alcohol at

t_4.

Hypothesis 4.3. The higher the density of the network in a class at t_1, the stronger the effect of the proportion of students using alcohol at t_1 on alcohol use at t_4.

We expect similar effects of *centralization* of the initial network as for density, but in the opposite direction:

Hypothesis 4.4. The higher the centralization of the network in a class class with an equal distribution of initial propensity to use alcohol at t_1, the lower the proportion of students using alcohol at t_4.

Hypothesis 4.5. The higher the centralization of the network in a class at t_1, the weaker the effect of the propensity to use alcohol at t_1 on alcohol use at t_4.

These two hypotheses signify that centralization of the network helps to counter the forces of the initial distribution of behavior and risk dominance. If the network is initially more centralized, there are actors in the network having relatively many interactions. Those actors are more influential, and if those actors happen to choose the risk-*dominated* behavior (i.e., *not* drinking), they are more likely to pull the rest of the network in this direction. Thereby, it is easier to "escape" from the "pull" of the risk-dominant behavior if the network is more centralized initially.

4.3. Data

4.3.1. Data collection

The data for this study were collected in a longitudinal survey project on 14 Dutch secondary schools, conducted in 2003 and 2004 (Knecht, 2006). From each school, all first-year classes were selected (between 5 and 14 classes per school, with an average of 9), and in each of these classes, all students were surveyed at regular intervals using written questionnaires. The first measurement took place shortly after the students entered the secondary school from primary education. The students were then surveyed again after three months, for a third time after another three months, and for a fourth and last time after another three months, resulting in a total of four waves. In total, 120 classes participated in all four waves. The survey included questions on personal characteristics of students, on various types of behavior (including alcohol use) and opinions, and various network measures.

The questionnaires were administered at school, with the students from each class together in a classroom. A researcher or research assistant was present at each session. Because the survey sessions were held during normal school hours, it could happen that not all students of one class were present. Moreover, some students may have joined a class between the moments of observation. Because of this, the number of students per class in the data may slightly differ between the different waves.

4.3.2. Variables and measures

Individual level measures

Personal networks Social relations in classes were measured using various *name generators*. In each wave, students were asked to name their best friends in class, the classmates with whom they spent leisure time, and those with whom they discussed personal matters. For each of these questions, they were allowed to name up to twelve classmates, using a list of codes for all classmates provided with the questionnaire. The maximum of twelve nominations was used only very rarely (up to 2% of the observations for any of the network measures), which indicates that this maximum was not a limitation on the measurement of nominations.

On the basis of these individual level measures, we construct networks at the class level. To verify that the results do not depend too much on the specific construction method, we use two different methods and report results using both methods.

For the first method, we combine measurements on three different name generators to identify interactions. We use nominations of "best friends," spending leisure time together and discussing personal matters. We require that, *taken together*, nominations on the three variables are reciprocated. Thus, we assume that two students interact if they nominate each other, each on at least one of the three variables. This method takes into account that interpretations of friendly relations may differ between students: while student i might consider student j as one of her best friends, j might nominate i merely as someone with whom she discusses personal matters. As we are only interested in the extent to which students interact, we think that such mutual nominations, even though the interpretations of the relation slightly differ, can be interpreted as mutual interactions.

For the second method, we use only nominations from one name generator of best friends (as is most common in the network literature), and assume that if one student nominates another, the two interact. That is, we do not require that nominations are reciprocated, and we interpret every directed tie as a symmetric relation. Bilateral nominations are treated the same as unilateral nominations.

Both methods result, at the aggregate level, in a *non-directed* network in which all ties are bilateral. This is required to adequately test the predictions form our model, which explicitly assumes that adolescents have incentives to coordinate their behavior if they *interact*. By definition, interaction is non-directed, which implies a non-directed network.

Alcohol use The use of alcohol by the students was measured in different ways in the different waves of data collection. In the first wave, students were asked how often they used alcohol in the preceding three months. Answers could be given on a five-point scale: "never," "once," "2 to 4 times," "5 to 10 times," and "more than 10 times." In waves 2, 3, and 4, students were asked how often they had used alcohol in the preceding three months *with friends*, with the same answer categories as in the first wave. Thus, the measurement differs between the first wave and the other three waves in that the question of the first wave does not ask specifically about drinking with friends, but rather about drinking in general.

For this reason, Knecht (2008), who analyzes the same dataset, can use only the last three waves. Here we take the difference to be an advantage. In the context of our model, the measure in waves 2 to 4 represents alcohol use as far as it happens in a context of interaction with friends. This fits well in our theoretical framework, in which we assumed that alcohol use is a choice in a game of social interaction.[1] The measure used in wave 1, in contrast, we interpret as an indicator for the *individual propensity* to use alcohol before the influence/selection process that takes place among the students within one class starts. This corresponds with initial alcohol use at t_1 in our theoretical model. We feel that this is appropriate because the data collection in the first wave took place shortly after the start of the first year in secondary school. Thus, the three months mentioned in the question would refer for the largest part to the period just before the students entered secondary school, before they were influenced by the friendship network in their class. For this reason, it is not problematic that the measure of the first wave does not measure alcohol use with friends only.

Interpreting the measures in this way has the advantage that the complete time span between the four waves of data collection can be used in the analysis. The disadvantage is that the measure of wave 1 cannot be directly compared with the measures in the later waves. We can, however, use the measure of wave 1 as a predictor of alcohol use as measured in the later waves in a multivariate analysis (as we explain below).

Network level measures

Aggregate network measures Using the two operationalizations of interaction between students, we can construct a friendship network for each class at each time point. To be able to test our hypotheses, we compute network measures for each of these networks. The hypotheses are concerned with density and centralization. *Density* is defined as the number of existing ties divided by the number of possible ties, given the size of the network (Wasserman and Faust, 1994, p. 101). For *centralization*, we use the measure proposed by Snijders (1981), which is based on the (normalized) degree variance. Besides the measures needed for testing

[1] This assumption does not imply that we believe that there are no individual factors influencing alcohol use. In this study, however, we largely disregard these factors because we are interested in the *social dynamics* of alcohol use.

the hypotheses, we compute a measure of *relative network change* for descriptive purposes. This measure describes the extent to which the networks change between the different waves and is defined as the proportion of dyads in a network that have changed status (i.e., created a tie or deleted a tie) from one time point to the other. The measure is only defined as long as the set of nodes in the network does not change. Consequently, we cannot compute this measure for all networks on all time points. The number of networks for which we can compute the measure is at least 90 on each time point. Moreover, for the networks of "non-reciprocated friendship ties," (method 2) we also report the proportion of nominations in these networks that are actually reciprocated.

Aggregate measures of alcohol use We aggregate the individual measures of alcohol use into aggregate measures of alcohol use per class in two steps. First, we dichotomize the individual level measures between "1" (never) and "2" or higher (once or more). This is done to make the measures more consistent with the theoretical model, which assumes that only two different actions are possible. We choose this specific dichotomization because it is substantively clear: we now distinguish between those who do not drink at all and those who drink sometimes. Moreover, the empirical distribution on this variable is such that the large majority of the students does not drink. Taking this group as a distinct category therefore seems most appropriate. In the second step, we calculate the proportion of students drinking ("1" on the dichotomized variable) per class.

4.4. Methods of analysis

Our analytical strategy is set up as follows. We start the analysis with some descriptive statistics on the development of behavior and the network across the four waves. We then turn to regression analysis to test the hypotheses. In line with the analyses in Buskens et al. (2008) and Corten and Buskens (2008), we conduct an analysis at the macrolevel, using classes as the unit of analysis. The basic aim is to explain the level of alcohol use at the last observed time point, using measures characterizing the initial state per class as predictors. We use (linear) regression analysis, with the proportion of students that uses alcohol per class as the dependent variable. Because the classes were not independently sampled but are nested within schools, we use multilevel random intercept regression (Snijders and Bosker, 1999) with a random intercept at the school level.

Using a linear regression model to analyze a dependent variable that is a proportion is not without problems. We see a number of reasons why, in this case, using a linear model is not problematic. First, the distribution of our dependent variable does not show peaks at the edges of the distribution. In fact, the distribution closely resembles a normal distribution. Second, our models do not predict impossible outcomes, that is, values below 0 or higher than 1. Third, standard regression diagnostics do not indicate severe violations of model assumptions. Fourth, additional analyses (not reported here) using a logistic transformation of the original dependent variable do not lead to qualitatively different results. Buskens et al. (2008) and Corten and Buskens (2008) use logistic regression for grouped data for their analyses on the macrolevel. In those studies, use of logistic models was necessary because of the highly skewed distribution of the dependent variable. Because our dependent variable here is approximately normally distributed, we do not suffer from this problem. We instead prefer to use the somewhat simpler linear model, which allows for better treatment of the multilevel structure of the data. To verify whether the results are robust against different specifications of the network variables, we repeat the regression analyses for the two specifications discussed in Section 4.3.2.

To test hypotheses 4.3 and 4.5, we construct two interaction terms by multiplying the initial proportion drinking with density and centralization of the initial network, respectively. To facilitate the interpretation of respective main effects, we subtract 0.5 from the initial propensity before multiplication. This is necessary to test hypotheses 4.2 and 4.4, because these hypotheses predict effects of the network structure given that the initial propensity is 0.5. To ensure that the main effect of the initial propensity can be meaningfully interpreted, we center the values of initial density and centralization at their respective means before multiplication.

Thus, the interaction between the initial propensity to use alcohol and initial network density is computed as

$$\text{Interaction} = (\text{initial propensity alcohol use} - 0.5) \times (\text{density} - \text{mean(density)})$$

This construction ensures that, in the above case, the main effect of density in the regression equation can be meaningfully interpreted as referring to the situation in which the initial propensity is 0.5, while the main effect of the initial propensity can be interpreted as referring to the situation with average initial density.

4.5. Results

4.5.1. Descriptive results

Table 4.1 provides means and standard deviations on key measures at the individual level: the original five-point measure on drinking behavior, the dichotomized version of this measure, and the number of nominations of best friends, classmates with whom the respondent discusses personal issues, and classmates with whom the respondent spends leisure time. These three network variables are, at the individual level, *directed* measures; they measure the number of “outgoing” ties of a student.

The two measures of alcohol use show a rather consistent pattern: the average alcohol use decreases from wave 1 to wave 2, and then steadily increases. The initial decrease reflects the difference in measurement between wave 1 and 2 (see Section 4.3.2): because the initial measurement in wave 1 does not focus exclusively on alcohol use with friends but also captures drinking in other situations, the figures are somewhat higher. Overall, the averages are fairly low.

The average number of "best friends" nominations shows a slight increase over the first three waves, and then decreases again in the fourth wave. The other two network measures increase consistently over the four waves, but are clearly lower than the number of friends nominations.

In Table 4.2, we summarize the trends in measures on the aggregated (network-)level. Thus, in this table, the unit of analysis is the *class* rather than the individual student as in Table 4.1. The proportion of students drinking per class (as measured by the dichotomized variable) naturally shows the same pattern as the individual level statistics, but has a smaller standard deviation (the aggregate measure might be interpreted as a weighted average of the individual level measure). Both density and centralization are computed on the two different types of networks constructed by the methods described in Section 4.3.2. Averaged over classes, we do not see much of a trend in either of the two density measures. If anything, the figure shows a small increase over the waves 1 to 3, and a decrease in wave 4, which is consistent with the results at the individual level with regard to friends in Table 4.1. Centralization is rather low on both measures, and does not show much of a trend. Either network size is stable over time, and shows very little difference between the two methods. We report also the average change per class compared to the network in the preceding wave. The results on network change indicate that most of the changes occur between waves 1 and 2, and that the friendship network is slightly more dynamic than the network according to the combined measures.

In Table 4.3, we report the pairwise correlations between the dependent and independent variables to be used in the regression analyses. Because the two size measures are nearly identical, we report only results on the size of the combined networks. The results show that there are weak to modest significant correlations between the various network variables and across the behavioral variables. The fact that the measures for density and centralization are correlated between the different construction methods suggests that these construction methods do not lead to very different results. The correlation between density and centralization within each construction method is remarkably high, given that the measure for centralization is controlled for density. Closer inspection of the data shows that these correlations are caused by a small number of classes that have both high density and high centralization. Exclusion of these outliers, however, does not lead to different results of the regression analyses. Network size correlates significantly with all the other network measures, but not with behavior.

Table 4.1: Descriptive statistics at the individual level, per wave

	Wave 1		Wave 2		Wave 3		Wave 4	
Variable	Mean	SD	Mean	SD	Mean	SD	Mean	SD
Drinking (5 point scale)	1.854	(1.205)	1.459	(0.945)	1.636	(1.128)	1.756	(1.230)
Drinking (dichotomous)	0.424	(0.494)	0.242	(0.429)	0.301	(0.459)	0.342	(0.475)
“Friends” nominations	3.591	(2.585)	3.849	(2.743)	3.976	(2.793)	3.772	(2.638)
“Personal” nominations	1.225	(1.571)	1.676	(1.839)	1.882	(1.892)	1.899	(1.963)
“Leisure” nominations	1.466	(1.488)	2.006	(1.816)	2.320	(1.997)	2.444	(2.125)
Number of obs. (listwise)	2,826		2,723		2,768		2,781	

Table 4.2: Descriptive statistics at the class level, per wave (N=120)

	Wave 1		Wave 2		Wave 3		Wave 4	
Variable	Mean	SD	Mean	SD	Mean	SD	Mean	SD
Prop. drinking	0.427	(0.133)	0.245	(0.120)	0.306	(0.136)	0.347	(0.151)
Combined network								
Density	0.092	(0.026)	0.103	(0.032)	0.110	(0.028)	0.107	(0.033)
Centralization	0.112	(0.054)	0.128	(0.056)	0.132	(0.057)	0.124	(0.053)
Network size	25.300	(4.076)	25.233	(4.140)	25.275	(4.122)	25.225	(4.103)
Change	-	-	0.094	(0.026)	0.082	(0.025)	0.079	(0.025)
Friendship network								
Density	0.202	(0.053)	0.226	(0.052)	0.232	(0.050)	0.222	(0.052)
Centralization	0.115	(0.073)	0.152	(0.071)	0.154	(0.064)	0.146	(0.059)
Network size	25.275	(4.079)	25.217	(4.141)	25.183	(4.095)	25.200	(4.095)
Change	-	-	0.167	(0.044)	0.133	(0.036)	0.127	(0.038)
Reciprocity	0.579	(0.084)	0.586	(0.090)	0.571	(0.090)	0.581	(0.086)

Table 4.3: Pairwise correlations between dependent and independent variables (N=120)

	1	2	3	4	5	6
1. Prop. drinking, Wave 1	-					
2. Prop. drinking, Wave 4	0.528**	-				
3. Density (combined)	0.100	0.098	-			
4. Centralization (combined)	0.089	0.233**	0.464**	-		
5. Density (friendship)	0.083	0.057	0.769**	0.522**	-	
6. Centralization (friendship)	0.028	0.103	0.493**	0.658**	0.633**	-
7. Network size	−0.034	−0.120	−0.591**	−0.593**	−0.628**	−0.544**

** $p < 0.05$

4.5.2. Multilevel regression using combined network measures

In this analysis, we use the networks as constructed by our first method, in which several types of name generators are combined. We conduct random intercept regression with average drinking behavior in wave 4 as the dependent variable, and class-level properties in wave 1 as predictors. We estimate three different models. In Model 1, we include only main effects of the initial proportion using alcohol, the initial density, and also control for network size. In Model 2, we add a term for the interaction between the two predictors. In Model 3, we add a main effect and interaction effect of centralization.

The results are displayed in Table 4.4. Model 1 shows a positive and strongly significant effect of drinking behavior in wave 1 as expected (Hypothesis 4.1), but no significant effect of initial density. We also find no significant effect of network size. In Model 2, the additional interaction effect is positive and significant, in accordance with Hypothesis 4.3. The main effect of initial density can in this model be interpreted as the effect of density for cases in which the initial proportion using alcohol is .5. This effect was expected to be positive (Hypothesis 4.2). The coefficient, however, is the opposite direction as expected but not significant. In Model 3, we add both the main effect of centralization and the interaction effect of centralization with drinking behavior in wave 1. Although both effects are in the expected direction (Hypotheses 4.4 and 4.5), they are not significant. Moreover, the likelihood ratio tests also indicates that although Model 2 is a significant improvement over Model 1, Model 3 does not further improve on Model 2. We therefore rely on Model 2, and conclude that only Hypotheses 4.1 and 4.3 are confirmed in this analysis.

4.5.3. Multilevel regression using non-reciprocated friendship ties

To examine to what extent our results depend on the specific network construction method, we repeat the analysis of the previous section using our second construction method, using (unreciprocated) friendship nominations as ties. Apart from this, the two analyses are identical.

Table 4.5 presents the results of this analysis. Overall, the results are consistent with the results in Table 4.4. We again find highly significant effects in the expected direction of the initial propensity and the interaction effect with initial density. Also, we again find no significant effects of initial density (main effect), centralization, or size, while the main effect of initial density is again in the opposite direction as expected.

We also find a number of differences as compared to the previous analysis. First, In Model 7, with effects of centralization included, the interaction term with density remains significant, in contrast with Model 3. A likelihood ratio test, however, indicates that Model 7 is no significant improvement over Model 6. Second, the effect of the interaction term

Table 4.4: Multilevel regression predicting average drinking behavior in wave 4 using combined network measures

		Model 1			Model 2			Model 3		
Variables	Prediction	Coeff.	SE	p	Coeff.	SE	p	Coeff.	SE	p
Prop. drinking, Wave 1	+	0.613***	(0.079)	0.000	0.609***	(0.077)	0.000	0.765**	(0.378)	0.043
Density	+	−0.833	(0.525)	0.112	−0.300	(0.557)	0.591	−0.296	(0.723)	0.683
Network size		−0.004	(0.004)	0.266	−0.003	(0.003)	0.333	−0.004	(0.004)	0.330
Drinking × Density	+				7.917***	(3.155)	0.012	6.575	(4.616)	0.154
Centralization	−							−0.038	(0.766)	0.961
Drinking × Centralization	+							2.134	(5.094)	0.675
Constant		0.257**	(0.127)	0.043	0.192	(0.126)	0.128	0.119	(0.221)	0.588
Variance components										
Var(Constant)		0.004	(0.002)		0.005	(0.002)		0.005	(0.002)	
Var(Residual)		0.012	(0.002)		0.011	(0.002)		0.011	(0.002)	
Log Likelihood		84.424962			87.442315			87.529721		
Likelihood ratio test: χ^2		48.20***			5.99**			0.22		
Number of observations		120			120			120		

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$. Likelihood ratio tests refer to the comparison with the previously estimated model; the likelihood ratio test for Model 5 refers to the comparison with the unrestricted model.

with density in Model 7 is considerably smaller than in Model 2. Third, the main effect of centralization in Model 7 is positive, contrary to expectations, but not significant.

All in all, the results of this analysis are comparable to the results of the analysis with combined network measures. The likelihood ratio tests again indicate that Model 6, which includes the interaction term with density but no effects of centralization, is the preferred model. Most importantly, these results do not lead to different conclusions with regard to the hypotheses. The results suggest that the substantive conclusions are robust against the different specifications of the network variables.

4.5.4. Additional analyses

The analysis in Section 4.5.3 differs from the analysis in Section 4.5.2 in two respects: we changed from the combination of various network measures to friendship nominations only, and from reciprocated nominations to non-reciprocated nominations. For two combinations on these two dimensions, we found no substantive differences in results, but there are two combinations remaining: reciprocated friendship ties, and non-reciprocated combined ties. Results on these two combinations (not reported here) do also show no qualitative difference with the two analysis just discussed.

Our results consistently show an effect of the initial network structure, which leads to conclusions different from those by Knecht (2008) (see the concluding section). A possible explanation of this difference is the different use of data: because alcohol use was measured in a sightly different way in the first wave of data collection, Knecht (2008) could not use this part of the data and instead used the second wave as the first observation. In contrast, we interpreted the measure of alcohol use in the first wave as a pre-existing tendency for drinking, and used the first wave as the first time point. To check the validity of this explanation, we replicated the analyses of Section 4.5.2, using the data of the *second* wave of data collection. Although coefficients are in the same direction and of comparable magnitude, we find no significant results. Therefore, we cannot rule out that the differences in results are driven by a different choice of data, that is, by the difference in the measure of alcohol use or the fact that we look at a longer period.

4.6. Conclusions

In this chapter we aimed to contribute to the understanding of selection and influence processes in the dynamics of alcohol use among adolescents. We did so by adopting a theoretical approach that interprets alcohol use as a coordination game in a dynamic network. Relying on simulation analyses of a game-theoretic model, we formulated hypotheses on the effects of initial conditions in terms of network structure and initial tendencies for alcohol use on re-

Table 4.5: Multilevel regression predicting average drinking behavior in wave 4 using unreciprocated friendship ties

		Model 5			Model 6			Model 7		
Variables	Prediction	Coeff.	SE	*p*	Coeff.	SE	*p*	Coeff.	SE	*p*
Prop. drinking, Wave 1	+	0.610***	(0.080)	0.000	0.593***	(0.077)	0.000	0.591***	(0.076)	0.000
Density	+	−0.352	(0.257)	0.171	−0.094	(0.262)	0.719	−0.137	(0.288)	0.635
Network size		−0.004	(0.004)	0.311	−0.003	(0.003)	0.330	−0.002	(0.004)	0.494
Drinking × Density	+				4.518***	(1.540)	0.003	6.193***	(1.874)	0.001
Centralization	−							0.081	(0.203)	0.689
Drinking × Centralization	+							−1.972	(1.423)	0.166
Constant		0.246*	(0.132)	0.063	0.192	(0.129)	0.136	0.163	(0.129)	0.207
Variance components										
Var(Constant)		0.004	(0.017)		0.004	(0.002)		0.004	(0.002)	
Var(Residual)		0.012	(0.008)		0.011	(0.002)		0.011	(0.002)	
Log Likelihood		84.141062			88.204715			89.298062		
Likelihood ratio test: χ^2		47.63***			8.13**			2.9		
Number of observations		120			120			120		

* $p < 0.1$;** $p < 0.05$; *** $p < 0.01$. Likelihood ratio tests refer to the comparison with the previously estimated model; the likelihood ratio test for Model 1 refers to the comparison with the unrestricted model.

sulting levels of alcohol use per school class. We tested these hypotheses using longitudinal data on alcohol use and social networks in Dutch high schools. Using various specifications of the independent variables, we were able to consistently confirm two hypotheses.

First, we find that the average initial propensity to use alcohol per class has a positive effect on average alcohol per class at a later stage. Second, we find that this effect becomes stronger, the higher the initial density of the social network in a school class. That is, in line with expectations, the density of the network *amplifies* the initial tendency of behavior.

We also predicted that initial network density should have a positive effect on alcohol use for classes that start with a average propensity to use alcohol of 0.5. However, we did not obtain any significant results on this hypothesis, and moreover, the estimated effect was consistently in the opposite direction than expected. Also, we were not able to obtain significant results on the effect of initial network *centralization*, where we predicted that centralization should have a negative main effect on alcohol use and should negatively interact with the initial behavioral tendency. While the direction of the estimated main effect is in some analyses opposite to the prediction, the interaction effect is in the expected direction.

How should our findings be interpreted? Although we were not able to test all hypotheses thoroughly, the fact that we found effects of the initial network structure on resulting behavior is revealing. The implication of this finding is that *influence* must play a role in the co-evolution process of alcohol use and network formation. If only selection would drive the process, then network formation would depend on the distribution of behavior, but not vice versa: the emerging behavior should be independent of the initial network. Instead, we find that there *is* an effect of the initial network. This conclusion contrasts with the findings of Knecht (2008), who found only evidence for selection using the same dataset. We return to this issue below.

The predictions on direction of the main effects of density and centralization depend on the assumption that using alcohol is the risk-dominant action in the coordination game. The facts that both effects were not significant, and that the main effect of density was consistently estimated in the opposite direction, suggest that this assumption might need to be revised. Such a revision might take two directions. First, the finding that the effect of density was consistently negative suggests that abstinence, rather than using alcohol, is the risk-dominant action. On the basis of this assumption, we would expect that a larger number of observations at the class level would yield a significant negative main effect of density. However, in that case we would also expect a significant *positive* effect of centralization, which we did not consistently find in our analyses. A second alternative assumption is that this coordination game is actually risk neutral, in the sense that neither of the two equilibria is risk dominant. In that case, we would indeed expect no main effects of density or centrality.

Such speculations, however, must be made with caution. The predictions on these main effects all refer to the situation in which the initial propensity is 50%. This means that the sizes and directions of these effects rely crucially on the exact definition of this majority, and therefore on the dichotomization of the dependent variable. While we think that our particular specification is well founded, arguments for different specifications are certainly conceivable. Therefore, one should be careful to draw strong conclusions based on the results on these main effects. Note, however, that the predictions on the *interaction effects* depend much less on the exact specification of the dependent variable.

Why do we find evidence for influence effects, while Knecht (2008) found only weak evidence? Part of the explanation could be that our approach indeed solves some of the problems of the SIENA approach applied by Knecht (2008), as we outlined in Section 4.1.2. That is, we are able to use also relatively stable classes for testing the hypotheses, and could therefore use more data. Besides the general methodological approach, however, there are some other differences between the two studies that could potentially account for the different findings. A first difference we already discussed concerns the use of data. Additional analyses showed that we cannot rule out this explanation of the differences in findings. Thus, it could be that the time from wave 2 to wave 4 is too short to observe network influence effects. Another possible source of that differences is that we were able to analyze a larger number of classes.

Second, the theoretical interpretation of influence is somewhat different between the two studies. Knecht (2008) analyzed the *directed* network, assuming that an adolescent is influenced by those peers she nominates as a friend. Thus, the influence is assumed to work in only one direction. In our model, we assume that influence takes place through *interaction*, which is by its nature undirected. While this difference is theoretically important, it should be noted that it is not exclusively a consequence of the different methodological approaches; a model with "two-way influence" would also be possible within the SIENA framework.

Third, partly as a consequence of the differing theoretical conceptualization of influence, our measures of the network are somewhat different. Knecht (2008) uses directed "best friend" nominations, while we use various combinations of network variables, including "best friend" nominations but also measures of spending leisure time together and discussing personal matters.

Given these differences, it is not clear how the differences in findings should be judged. We clearly encounter a discrepancy between findings at the macrolevel, where we do find evidence for influence effects, and the microlevel, where such effects could not be observed. This discrepancy poses a new puzzle that deserves more attention in future research on this topic. Such research should focus on the development of theoretical models that are consistent with empirical research on microlevel behavior and the macrolevel findings as presented in this study (see also Corten and Buskens, 2008). Overall, we conclude that the focus on effects of macrolevel network effects contributes to the explanation of emerging differences between classes and adds interesting new insights to the study of co-evolution processes.

CHAPTER 5

Cooperation and Reputation in Dynamic Networks*

5.1. Introduction

Dyadic cooperation is one of the building blocks of human societies. Whether this cooperation be colleagues working on a project together, friends providing social support, neighbors providing practical support, or firms involved in R&D collaboration, people interact to produce something they could not have produced alone. In many cases, both parties can benefit from mutual cooperation, but both are also tempted to take advantage of the other party for an even larger benefit. In such cases, cooperation takes the form of the well-known two-person Prisoner's Dilemma (see Figure 5.1). The standard game-theoretic prediction is that actors will not cooperate in such a dilemma. Consequently, a very large body of literature has evolved to address the question: under what conditions is cooperation in the Prisoner's Dilemma possible? One major finding is that cooperation is possible (but not guaranteed) when an interaction is repeated.

5.1.1. Cooperation and network effects

A classic sociological argument holds that embeddedness in social networks can promote cooperation in situations involving social dilemmas (e.g., Coleman, 1990; Granovetter, 1985; Homans, 1951). Casual observation suggests that cooperation is more frequent in close-knit groups than in anonymous interactions, and indeed, there are many empirical examples in the literature of the association between network cohesion and cooperation. For instance, Uzzi (1996, 1997) shows that small firms are more likely to cooperate if their interactions are embedded in a dense network.

There are a number mechanisms that have been proposed to produce higher rates of cooperation in dense networks. Buskens and Raub (2002) distinguish between *control* and *learning* effects in social networks. Control, also known as "social control" (e.g., Homans,

*This chapter was written in collaboration with Karen Cook. A shortened version appeared in the *Proceedings of the First International Conference on Reputation: Theory and Technology - ICORE 09* (Corten and Cook, 2009).

1951), refers to the idea that actors are more inclined to cooperate if information about defection can spread through a given network in a way that leads to sanctions by third parties. Thus, actors cooperate because they are concerned with the *future consequences* of defection. Raub and Weesie (1990) develop a formal model that shows that network embeddedness can encourage cooperation in the Prisoner's Dilemma; Buskens (2002) similarly analyzes how network embeddedness can promote cooperation in one-sided Prisoner's Dilemmas or Trust Games. Empirical illustrations of control effects are described by Ellickson (1991), among others.

The second mechanism through which networks affect cooperation is *learning*. Here, it is information on *past* interactions that matters. An actor may be more inclined to cooperate with another if she learns from a third partly that cooperation with this other actor may be more fruitful than defection, because, for example, this other actor appears to be playing according to a tit-for-tat strategy. Other useful information that actors may learn through a network concerns the payoffs of the other player (Buskens, 2003; Kreps and Wilson, 1982). While the focus in the literature is mainly on the positive effects of network density on cooperation, this effect may also be negative. Burt and Knez (1995), for instance, argue that higher network density tends to make actors more certain of their perceptions regarding interaction partners. If this perception happens to be negative, higher network density will only amplify the tendency to defect. Rapoport et al. (1995) similarly report negative effects of third-party information in an experimental study.

Both the control and the learning mechanisms rely on the transfer of information about the behavior of an actor to third parties. In the case of control, this information allows a third party to threaten an actor with sanctions if she does not cooperate. In the learning case, this information allows a third party to decide on the fruitfulness of cooperation. That is, both mechanisms are examples of *reputation* effects. Reputation is one of those concepts that is so widely used that its precise meaning has become obscured. Generally, it refers to an attribute of an actor ascribed to him by other actors (Raub and Weesie, 1990; Wilson, 1985). Raub and Weesie (1990) further distinguish between reputation in the narrow sense and reputation in the broad sense. Reputation in the narrow sense refers to situations in which the behavior of an actor influences his reputation in *this* situation, while reputation in the broad sense refers to situations in which behavior in one interaction influences *other* interactions. We are concerned with reputation in the broad sense.

5.1.2. The case for network dynamics

Summarizing the theoretical ideas described above, we can say that networks are presumed to affect cooperation because they facilitate reputation effects by allowing for the transfer of information. Theories on reputation effects share the assumption that social networks are determined exogenously. Generally, the social network is considered to be a *stable* social context that is imposed on actors and that provides the means for the spread of information through which the mechanisms of control and learning can work. In recent years, this assumption has been challenged in network research. More and more, it is recognized that social networks themselves are also the result of interactions in which actors make conscious decisions about their social relations. The recognition of this fact has led to an explosion of interest in social network dynamics over the past several decades in the field of sociology (e.g., Doreian and Stokman, 1997; Snijders, 2001; Weesie and Flap, 1990), economics (e.g., Dutta and Jackson, 2003; Jackson and Wolinsky, 1996), mathematics, and physics (e.g., Newman et al., 2006; Watts and Strogatz, 1998).

There are at least two reasons to suspect that relaxing the assumption that networks are exogenously fixed has implications for the theoretical predictions about the effects of reputation. First, it is not self-evident that reputation is equally effective if the network through which information is transferred changes as a result of what occurs during actual interactions. For instance, if an actor experiences defection by an interaction partner, she might "spread the word" through her network connections such that the defector can be sanctioned. Alternatively, she might end the interaction with this partner altogether, because she prefers not to interact with partners who defect. In that case, she unintentionally also changes the possibilities for the spread of reputations for other actors in the network. Conversely, network decisions themselves may be affected by reputation effects; actors might be more willing to initiate interactions with potential partners who have a cooperative reputation and may be less willing to interact with actors who are known to defect. It is theoretically unclear whether and how reputation mechanisms function in a dynamic context. The main aim of this chapter is to shed some light on this question.

A second problem is that understanding networks as dynamic raises a possible problem of causal order. Traditionally, it is assumed that the network comes first in the relationship between network density and cooperation; cooperation is seen as the result of the network structure. If we assume that networks may be dynamic, however, we must also acknowledge that the causation could also operate in the other direction. If interactions are more likely between actors who tend to cooperate than between actors who defect (for instance, because mutually cooperative interactions are more rewarding than relationships of mutual defection), then high density would be the *result* of an already-high level of cooperation rather than its cause. Because most empirical studies of network effects on cooperation are cross-sectional, they cannot distinguish between these two possible causal directions.

5.1.3. Learning in networks

In this chapter, we limit the model of reputation effects to learning and neglect control effects. Thus far, combining control and learning effects in one model has proven to be difficult, even for fixed networks (Buskens and Raub, 2002). For dynamic networks, the complexity of these models would likely increase even further. We believe that more theoretical progress can be made at this point by studying the effects of network dynamics on one of the two mechanisms than by immediately trying to combine learning, control, and network dynamics into one model.

Aside from this somewhat pragmatic consideration, there are also substantive reasons to first focus on learning. Extending the cooperation problem with network dynamics increases complexity not only for the modeler but also for the actor we are trying to model. She now finds herself in a situation in which she has repeated interactions with multiple partners and in which an action may have complex externalities due to reputation effects. Models that assume that actors are perfectly rational (as game-theoretic models of control effects typically do), and thereby already make quite strong demands on the cognitive capabilities of the actors. If the network is also dynamic, this means that actors are required to anticipate the consequences of their actions by taking into account that the information structure itself may also change as a result of their actions. This makes it increasingly unlikely, in our opinion, that actors will display the rational forward-looking behavior assumed by game-theoretic models of control. Learning models, in contrast, rely on a backward-looking logic and, as a consequence, must make much less demanding assumptions regarding the rationality of actors (Macy and Flache, 2002). For this reason, we feel that a learning approach would be most useful for studying our problem.

There are different ways of modeling learning in social dilemmas. One broad class of models concerns *reinforcement learning* (e.g., Camerer and Ho, 1999; Erev and Roth, 1998). According to this method, actors learn to choose actions that result in high payoffs based on previous experiences. These experiences may include their own experiences as well as the experiences of others. In the latter case, actors might imitate successful others by adopting their strategies. This model is often used as a social analogue of evolutionary reproduction dynamics (see Cohen et al., 2001 and Ohtsuki et al., 2006 for models incorporating networks). Such models often make the rather strong assumption that actors are informed regarding both the relative success of other actors and the strategies used by these actors. We feel that these assumptions are too strong for most social interactions (cf. Weesie, 1996).

A second class of models comprises *belief-based learning* (Camerer and Ho, 1999). In these models, actors form beliefs about the behavior of their interaction partners and then choose the best response against this behavior. Beliefs can be based on an actor's own past experience, but information can also come from other sources. Given our conception of reputation, belief-based learning provides a natural approach to study the effects of reputation in networks. Reputation in a belief-based learning context enters the decision-making of actors as information learned from third parties.

5.1.4. Related theoretical literature

Studies that model repeated Prisoner's Dilemmas (RPDs) using some sort of partner choice can be roughly distinguished into two categories. On the one hand, there are models in which actors play *dyadic* games, as in our model. This means that actors can discriminate between different interaction partners through their behavior in the game. On the other hand, there are models in which actors choose *one* action for all interaction partners and thus cannot discriminate. This difference changes the strategic situation considerably; if one cannot specify different actions for different partners, it is impossible to sanction or reward specific partners because every choice affects all partners. Models of cooperation in dynamic networks without discrimination are studied in different setups by Eguíluz et al. (2005), Ule (2005), and Biely et al. (2007). Of these, Ule (2005) includes a reputation mechanism for control, yet in this model, reputation does not depend on the network. Biely et al. (2007) assume the development of some type of reputation via learning, but do not explicitly model its underlying mechanism.

Models that study *dyadic* RPDs with partner choice can be traced back to Schuessler (1989), who studies the effects of an "exit-option" in a computational tournament using a method similar to that of Axelrod (1984). Vanberg and Congleton (1992), Stanley et al. (1994), Yamagishi et al. (1994), Weesie (1996), the EdK-Group (2000), and Hauk (2001) all conduct similar analyses using various set-ups.

None of these studies, however, includes a reputation system. In effect, interaction still takes place in independent dyads, and although in principle we can speak of a network of interactions, there are no effects of network *structure*. The only study that we are aware of that does discuss a network-based reputation mechanism is Vega-Redondo (2006), who presents a game-theoretic analysis of the RPD in an endogenous network. Reputation (albeit not by that name) enters the analysis as players punish interaction partners when they learn of defection by a partner via the network. This is indeed reputation in the sense of *control*, which is different from the reputation included in our approach. Moreover, reputation only plays a limited role in network formation. The analysis of this model focuses mainly on the volatility effects of the environment and shows that more volatility in the long-run leads to lower network density as well as to lower average distance in the network.

Finally, Pujol et al. (2005) study dyadic *support games* in a dynamic network setting with reputation effects. Although reputation is network-dependent in their model, its role is limited. Third-party information comes only from direct neighbors and, moreover, is only

	Cooperate	Defect		Cooperate	Defect
Cooperate	R,R	S,T	Cooperate	$3,3$	$0,5$
Defect	T,S	P,P	Defect	$5,0$	$1,1$

Figure 5.1: The Prisoner's Dilemma, in general form and with numerical payoffs, with $T > R > P > S$

used when first-hand experience is not available. A second drawback of their model for our purposes is that free-riding is excluded by assumption; actors are modeled as "benevolent" (Pujol et al., 2005, §2.20). This partly assumes away the very problem we want to address, namely, the emergence of cooperation among egoistic actors.

In the following section, we develop a model for cooperation in dynamic networks with reputation effects based on belief-based learning. We then present some analytical results on stable states for different levels of reputation effects in both fixed and dynamic networks. Because the resulting theorems characterize stable states in rather general terms and say nothing about the generating process, we extend these formal analyses via a computer simulation of our model.

5.2. The model

5.2.1. Formalization of the problem

The stage game

The basic interaction is modeled as an infinitely repeated two-person Prisoner's Dilemma (PD). Thus, in every stage of the process, actors play a game, as illustrated in Figure 5.1. In this game, actors can collectively benefit from mutual cooperation, but they also have an individual incentive to free-ride on the efforts of the other player. Future payoffs are discounted by a discount parameter w.

Games in networks

There is a finite set of actors $N = \{1,2,3,\dots,n\}$. Actors play two-person RPDs, and can be involved in multiple games at the same time. They can also behave differently with different partners. For a population of actors, the collection of dyadic relationships results in a *network* of relations. A network of n actors can be represented by the $n \times n$ adjacency matrix g, where $g(i,j) = 1$ if there is a link from i to j, and $g(i,j) = 0$, otherwise. By assumption, reflexive links are ruled out such that $g(i,i) = 0$ for all i. The network of interactions is undirected by nature, and therefore, $g(i,j) = g(j,i)$ for all i and j. The set of actors with whom actor i has a link is formally denoted as $N_i(g) = \{j \in N | g_{ij} = 1\}$, and these actors will be referred to as the *neighbors* of i.

Information

We make restrictive assumptions regarding the information available to the actors. Actors are informed of the actions chosen by their interaction partners, but they are not informed of the *strategies* according to which their partners are playing the game. Moreover, actors are not aware of the structure of the *network* beyond their own connections. This also implies that actors are not informed about the payoffs their partners receive.

However, while information available to individual actors is thus rather limited, the network structure allows for diffusion of information among actors, which makes the emergence of reputations possible. However, the diffusion of information is assumed to have its limits. It does not flow effortlessly through the network but rather decreases in reliability with each "step" through the network. We model this process explicitly when we discuss how actors use information.

Network dynamics

A key property of our model is that the interaction network is not exogenously imposed, but can rather be changed by the actors. A simple way to build this into the model is to consider a number η of randomly-drawn pairs of actors who can decide whether they want to change their relation, that is, create a new tie between them or delete an existing tie. To have a network tie means to interact, and so it seems natural to allow these ties to be created with mutual consent, though they can be deleted unilaterally.

Another important assumption is that maintaining network ties is costly. The underlying reasoning for this assumption is that the maintenance of social interactions requires some effort, and therefore, actors may want to end relationships that are less profitable. Assigning a fixed cost to every interaction is a convenient way to model this (cf. Jackson and Watts, 2002).

Formally, the total cost k for maintaining z ties for actor i in each round of play is given by the simple linear function

$$k_i(g) = \alpha z_i \tag{5.1}$$

in which z_i denotes the number of ties actor i is involved in (i.e., $z_i = |N_i(g)|$) and $\alpha \geq 0$.

The dynamic process

The various components of the model are combined into a *repeated* game in the following way. Each period of the process consists of three phases:

1. *Network formation phase*: A number of pairs of actors (denoted by η) are randomly chosen to update their relationship (i.e., create a new tie if there is none or remove a tie if there is one) using the information available to them.

2. *Game play phase*: All actors play the PD simultaneously with each of their neighbors, given the network resulting from phase 1.

3. *Information phase*: The actors are informed of the outcomes of phase 2, information spreads through the network by a mechanism we explain below, and actors update their beliefs about each other.

After phase 3, we return to phase 1, and the process repeats. In phases 1 and 2, actors rely exclusively on the information they received in phase 3, from their own experiences, or through reputation. Thus, after new ties are formed in phase 1, there is no additional spread of information, and decisions during the game are based on the same information as the decisions made during the network formation phase.

5.2.2. Individual strategies

The description of the situation above shows that the decision problem for the individual actor is rather complex. How does one decide which strategy to choose in this repeated game? In principal, the number of possible strategies from which to choose is infinite. This is also true for the possible strategies of the opponent, which makes the situation even more complicated. Moreover, our network setting in which information can spread, as well as the choice of relationships may introduce further externalities with regard to this spread of information, thus leading to an extremely complex decision situation with a very large range of action alternatives and possible outcomes. This situation would demand such extreme requirements in terms of information processing and computation that we think it is unlikely that human actors would be capable of acting completely rationally in such a context (cf. Conlisk, 1996).

Instead, we prefer to model actors as *boundedly rational* in the sense that actors make a number of simplifying assumptions about the world around them as well as make use of information in a possibly suboptimal manner (Rubinstein, 1998; Simon, 1956).

More specifically, we assume that actors consider only "$t-1$ matching strategies" by the opponent (Downing, 1975). That is, actors assume that their opponent's action in the game is a response to their own action in the previous period. A simple and famous example of such a $t-1$ matching strategy is the tit-for-tat strategy in which an actor always copies the last move of her opponent, though many other strategies can be expressed in this fashion. Generally, a $t-1$ matching strategy of actor j against i can be described in terms of two probabilities:

$p(C_j \mid C_i)_t$: The probability that j will cooperate at time t after i cooperated at time $t-1$;

$p(C_j \mid D_i)_t$**:** The probability that j will cooperate at time t after i defected at time $t-1$.

We propose that actors try to maximize utility in the repeated game by assuming they interact with an opponent who uses a $t-1$ matching strategy with unknown response probabilities.[1] As the precise strategy of the opponent is not revealed, maximizing utility involves making an assessment of the most likely values of $p(C_j \mid C_i)$ and $p(C_j \mid D_i)$ with regard to the strategy of the opponent. In effect, this means that actors assume they are playing against a probabilistic *automaton* that is driven by two conditional probabilities. This leaves the actor with two tasks: first, to figure out the probabilities according to which the opponent is playing and, subsequently, to determine the optimal response given these probabilities.

To obtain an assessment of the behavior of the opponent in terms of $p(C_j \mid C_i)$ and $p(C_j \mid D_i)$, actors simply use the frequency distribution of the opponent's responses in the game so far. That is,

$$p(C_j \mid a_i)_t = \frac{C^a_{ijt}}{T^a_{ijt}} \tag{5.2}$$

in which T^a_{ijt} is the total number of times that i played action a (C or D) against j until time t, and C^a_{ijt} is the total number of times that j reacted with cooperation to action a by i until time t. We assume that at $t=0$, actors have a prior belief about the behavior of each opponent in terms of a "fictional" T^a_{ij0} and C^a_{ij0}.

The next task is to determine the optimal response when playing against an automaton with two given response probabilities $p(C_j \mid C_i)$ and $p(C_j \mid D_i)$. It can be shown that to maximize utility in such a situation it is sufficient to consider only three different courses of action, which we call *sub-strategies* for convenience. These are:

1. Cooperate in every round, labeled as ALLC;
2. Defect in every round, labeled as ALLD;
3. Alternate between cooperation and defection, labeled as ALT.

For $w \rightarrow 1$, which we assume throughout for simplicity, the expected payoff π_{ij} for actor i interacting with actor j for each of the three courses of action can be written as:

$$\pi(ALLC)_{ij} = p(C_j \mid C_i)R + (1 - p(C_j \mid C_i))S \tag{5.3}$$

$$\pi(ALLD)_{ij} = p(C_j \mid D_i)T + (1 - p(C_j \mid D_i))P \tag{5.4}$$

$$\pi(ALT)_{ij} = \frac{1}{2}\big(p(C_j \mid C_i)T + (1 - p(C_j \mid C_i))P + p(C_j \mid D_i)R + (1 - p(C_j \mid D_i))S\big) \tag{5.5}$$

[1] This approach is closely related to *fictitious play* (Fudenberg and Levine, 1998).

The actor's behavior as described so far is equivalent to the strategy known as the "DOWNING" strategy, which was originally proposed by Downing (1975) and was subsequently the name of one of the competitors in the famous computer tournament conducted by Axelrod (1984) From here on, we use the label "DOWNING" to refer to the individual strategy used in this chapter.

5.2.3. Reputation

According to the original DOWNING strategy, the actor bases her behavior on expectations she forms from her own experience with an opponent, but it is relatively straightforward to include reputation (i.e., third-party information) in the model. Experiences of other actors with whom an actor is directly or indirectly connected can be used to estimate the probabilities of the opponent's expected behavior as discussed above. Since it is likely that information from third parties is treated differently (Burt and Knez, 1995; Granovetter, 1985), the weight an actor assigns to third party information, as opposed to information from her own experience, is a parameter of the model.[2] Another parameter pertaining to reputation may be the extent to which information spreads through the network, for example, the maximum distance that information can travel.

In everyday language, reputation is often conceived as an attribute of a single actor. For example, we say that a certain person "has a good reputation" or has "lost his good reputation." In contrast, in our model, "the" reputation of an actor may consist of a set of reputations that are attributes of *other* actors. The reputation of actor i is the set of beliefs in the minds of other actors about the behavior of i learned from sources other than their own experience. These beliefs do not need to be the same; different actors can have different beliefs about the behavior of i.[3]

More specifically, the reputation of actor j with actor i (that is, the information that i has about the behavior of j in interactions not with i) consists of:

- The total number of times actors other than i played C to j to the knowledge of i;
- The total number of times j responded to C with C in interactions not with i to the knowledge of i;
- The total number of times that actors other than i played D to j to the knowledge of i;

[2] Alternatively, one might also think of this parameter as being part of a *strategy*, i.e., as a conscious choice by the actor to use third-party information or not. While that would be an interesting extension of the current model, we treat the weight of third-party information as an environmental variable here.

[3] The fact that in ordinary language, reputation is usually thought of as an attribute of a single "target" actor suggests that in reality, beliefs about the target actor tend to be—or at least seem to be—so similar that it becomes justifiable to speak of "the" reputation of the target actor. As we will see, this also occurs in our model under certain conditions.

- The total number of times, to the knowledge of i, that j responded to D with C in interactions not with i to the knowledge of i.

The phrase "to the knowledge of i" is meant to indicate that i does not necessarily have information on other interactions in which j is involved; that is, i has this information only if i is directly or indirectly connected to other actors interacting with j. Moreover, it is likely that information received by third parties has a smaller influence on i's decision than her own experience. For instance, i might have less confidence in third-party information, because it is more likely to be distorted the farther it travels through the network. To model such effects, we assume that the weight of the third-party information that i receives about j depends on the *network distance* δ through which this information is transferred. If i and j are connected, then the reputation of j with i consists of the expectations of j's neighbors k about j weighted by the network distance δ_{ik} between i and every k. δ_{ik} is defined as the *shortest path* in the network between i and k (see Wasserman and Faust, 1994, p. 110). The information is weighted in the sense that the information obtained from k is subject to network *decay*, and it is considered less important to i the larger the shortest distance from k to j becomes. These quantities, in combination with those obtained from one's own experience, are used to compute the probabilities that the DOWNING uses to assess an opponent's future behavior. We modify equation (5.2) in the following way: Recall that $N_j(g)$ denotes the set of j's neighbors. Let ω ($0 \leq \omega \leq 1$) denote the extent of network decay of information. Then, the probability $p(C_j \mid a_i)$ that j will cooperate at time t after i played action a is computed as

$$p(C_j \mid a_i)_t = \frac{C^a_{ijt} + \sum\limits_k \omega^{\delta_{ikt}} C^a_{kjt}}{T^a_{ijt} + \sum\limits_k \omega^{\delta_{ikt}} T^a_{kjt}} \tag{5.6}$$

with $k \in N_{jt}(g), k \neq i$. In the numerator of this equation, C^a_{ijt} represents i's own experiences with j, while the term on right-hand side denotes the combined experiences of j's other neighbors; the terms in the denominator are likewise interpreted. The experiences of j's neighbors are weighted by ω^δ. Thus, if $\omega = \frac{1}{2}$, information learned via reputation is discounted with a factor $\frac{1}{2}$ with every step it travels through the network. When $\omega = 0$, reputation does not play any role; when $\omega = 1$, information from any source is given the same weight.

A few additional remarks are needed before we proceed. First, note that the reliability of different kinds of information is modeled by counting events. If information from third parties is based on more observations (e.g., from sources with a long interaction history, a large number of sources, or both), then this information might have more influence on $p(C_j \mid a_i)$ than information from one's own experience (given ω). In this sense, this reputation system is similar to the systems that are implemented by online auction websites such as eBay in which a potential buyer can verify the reputation of a seller in terms of both the number of positive ratings by previous buyers and the total number of transactions in which the seller has been involved (Snijders and Zijdeman, 2004).

Second, it is important to realize that information does not “accumulate” when it travels through the network. That is, when an actor reports to another actor about the behavior of a third actor, she reveals only information from her own experience and does not include the information she has obtained from her own sources. Again, compare this with eBay’s system. There, buyers are supposed to base their feedback only on their own transactions and not to include previous ratings by other buyers of the same seller in their own ratings.

5.2.4. Network decisions

How do actors decide on the creation of a new tie or the deletion of an existing tie? Basically, a tie is only created or maintained if the expected payoff from the interaction exceeds the cost of maintaining the tie. The expected payoff is the payoff the actor would obtain *if* he interacted with the potential partner under consideration. Thus, when deciding on whether to change the status of a tie, an actor i:

1. Computes $p(C_j \mid C_i)$ and $p(C_j \mid D_i)$ as if he were interacting with j using available information;

2. Decides which sub-strategy he would use if he were interacting with j (equations (5.3) to (5.5));

3. Compares the expected payoff from using this sub-strategy (which is normalized to one period) with the cost of maintaining the relation.

If the result of this evaluation is nonnegative for *both* actors, the tie is created or kept; if the result is negative for one or both of the actors, it is either not created or is deleted.

Because we focus this chapter on reputation effects in dynamic networks, we abstract here from strategic considerations concerning the transmission of information That is, we assume that information flows through the network, without modeling the decisions of actors to actually pass information. Lippert and Spagnolo (2005) discuss this issue, but abstract from network formation (see also Buskens and Raub, 2002 and Buskens, 2002, Ch. 7). Similarly, we assume that actors do not purposefully try to obtain strategic information.[4] In some applications, it is likely that actors maintain certain relationships regardless of the payoffs, only because a relationship is a source of important information that can be used in other interactions. However, for now, we consider information transmission to be an unintended consequence of network formation. The situation in which actors form links with the explicit aim of obtaining information is captured in general terms by the “connections model” presented in Bala and Goyal (2000).

[4] We hope to relax this assumption in future developments of this model.

5.2.5. Convergence

We define the dynamic process as converged when two conditions are met. First, there is no pair of actors willing to create a new tie, and there is no single actor willing to remove a tie. This criterion conforms to the notion of *pair-wise stability* of networks, as formulated by Jackson and Wolinsky (1996). Second, the beliefs of actors are stable in the sense that they converge. The convergence of beliefs implies that sub-strategies are also stable. Looking at behavior in the game directly (as is done in, for example, Jackson and Watts, 2002 and Buskens et al., 2008) is not a sufficient criterion for stability of the process under study, because alternation is one of the possible sub-strategies, and moreover, stability in behavior for some period of time does not imply stability of beliefs.

5.3. Analysis of the model

In this section, we derive some analytical results for the model described above. The main focus is on *stable states* of the co-evolution process, that is, converged states of the process. First, however, we briefly discuss the behavior of the DOWNING strategy in the two-person case, as this will be helpful in the subsequent analyses of fixed and dynamic networks.

5.3.1. Dynamics of behavior with two actors

In this section, we analyze the dynamics of play when two players using the DOWNING strategy are interacting, and when neither player can end the interaction. Examining this relatively simple case provides a useful basis from which to move forward to the more complex case involving reputation effects and network dynamics. The main result for the two-player case can be summarized as follows:

Theorem 5.1. If two actors are both using the DOWNING strategy in a two-person RPD, the only stable sub-strategy-combinations are (ALLC,ALLC) and (ALLD,ALLD).

We omit the complete proof here, instead providing a sketch of the argument. Generally, the actor playing the DOWNING strategy cooperates if she finds that the opponent is sufficiently *reactive* in the sense that the opponent is sufficiently more likely to cooperate after cooperation than after defection. This implies that $p(C_j \mid C_i)$ must be sufficiently larger than $p(C_j \mid D_i)$ for i to play ALLC. If, however, $p(C_j \mid C_i)$ and $p(C_j \mid D_i)$ are equal, the actor playing DOWNING concludes that the opponent is not reactive at all. That is, the behavior of the opponent does not depend upon what the actor does. In that case, it is more lucrative to play ALLD, because each period can then be considered effectively independent. Thus, in a sense, one could say that the actor playing DOWNING defects when she can, but cooperates when she has to.

More precisely, from equations (5.4) and (5.5), we can derive that an actor i playing DOWNING will switch from ALLD to ALT if $\pi(\text{ALT})_{ij} > \pi(\text{ALLD})_{ij}$, that is, if

$$p(C_j \mid C_i) > p(C_j \mid D_i)\frac{2T-2P-R+S}{T-P} + \frac{P-S}{T-P} \tag{5.7}$$

Similarly, actor i switches from ALT to ALLC if

$$p(C_j \mid C_i) > p(C_j \mid D_i)\frac{R-S}{2R-2S-T+P} + \frac{P-S}{2R-2S-T+P} \tag{5.8}$$

This implies that the conditions under which DOWNING plays ALLC are rather limited. Regarding the numerical example in Figure 5.1, this can be illustrated by plotting $p(C_j \mid C_i)$ as a function of $p(C_j \mid D_i)$ according to equations (5.7) and (5.8) (see Figure 5.2).

From Figure 5.2, it can also be inferred that in the absence of reputation effects, the only combinations of sub-strategies that are stable are (ALLC,ALLC) and (ALLD,ALLD). Table 5.1 summarizes the sequences of play for the different combinations of sub-strategies. Suppose that two actors are both using ALLC, which implies that they are both in the region above the solid line in Figure 5.2. In that case, $p(C_j|C_i)$ will converge to 1 for both as t goes to infinity, while $p(C_j|D_i)$ will not change, because D is never played. This implies that both will remain in the ALLC region and continue playing ALLC. Therefore, this combination is stable. If both are playing ALLD, both $p(C_j|D_i)$ and $p(C_i|D_j)$ will converge to zero. In Figure 5.2, this means that the beliefs of both actors would move leftwards parallel to the horizontal axis, resulting in the continued play of ALLD or a change to ALT by one or both actors as soon as the dotted line is crossed. However, from Table 5.1, it can be verified that no combination of sub-strategies that involves ALT can be stable; in all cases, the resulting beliefs converge to a combination that does not support ALT. Therefore, (ALLC,ALLC) and (ALLD,ALLD) are the only two stable combinations.

5.3.2. *Stable states in fixed networks*

We now extend the analysis to a "network" setting in which actors play with multiple partners, but we limit ourselves to *fixed* networks ($\eta = 0$). The crucial difference between the network setting and the two-player setting analyzed in the previous section is that actors can share information, and reputations can thus emerge. In the model, the extent to which reputation plays a role is determined by the parameter ω. We first look at two extreme cases, namely, $\omega = 0$ (no diffusion of reputations) and $\omega = 1$ (perfect diffusion of reputations). Let σ_{ij}

Table 5.1: Sequences of play for different combinations of sub-strategies of DOWNING and convergence of actor beliefs for $t \to \infty$

Actor	Sub-strategy	Sequence of play				Convergence of beliefs	
$i:$	ALLC	C	C	C	...	$p(C_j\|C_i) \to 1$	-
$j:$	ALLC	C	C	C	...	$p(C_i\|C_j) \to 1$	-
$i:$	ALLD	D	D	D	...	-	$p(C_j\|D_i) \to 0$
$j:$	ALLD	D	D	D	...	-	$p(C_i\|D_j) \to 0$
$i:$	ALLC	C	C	C	...	$p(C_j\|C_i) \to 0$	-
$j:$	ALLD	D	D	D	...	-	$p(C_i\|D_j) \to 1$
$i:$	ALLD	D	D	D	...	-	$p(C_j\|D_i) \to \frac{1}{2}$
$j:$	ALT	C	D	C	...	$p(C_i\|C_j) \to 0$	$p(C_i\|D_j) \to 0$
$i:$	ALLC	C	C	C	...	$p(C_j\|C_i) \to \frac{1}{2}$	-
$j:$	ALT	C	D	C	...	$p(C_i\|C_j) \to 1$	$p(C_i\|D_j) \to 1$
$i:$	ALT	C	D	C	...	$p(C_j\|C_i) \to 0$	$p(C_j\|D_i) \to 1$
$j:$	ALT	C	D	C	...	$p(C_i\|C_j) \to 0$	$p(C_i\|D_j) \to 1$
$i:$	ALT	C	D	C	...	$p(C_j\|C_i) \to 1$	$p(C_j\|D_i) \to 0$
$j:$	ALT	D	C	D	...	$p(C_i\|C_j) \to 1$	$p(C_i\|D_j) \to 0$

denote the sub-strategy used by i against a neighbor j; σ_{ij} may be ALLC, ALLD, or ALT.

Theorem 5.2.

(i) If $\omega = 0$ and $\eta = 0$, then $\forall j \in N, \forall i \in N_j(g), \sigma_{ij} = \sigma_{ji} \wedge (\sigma_{ij} = \text{ALLC} \vee \sigma_{ij} = \text{ALLD})$ in stable networks.

(ii) If $\omega = 1$ and $\eta = 0$, then $\forall j \in g, \forall i,k \in N_j(g), \sigma_{ij} = \sigma_{kj}$ in stable networks.

Proof. Case (i) of Theorem 5.2 is simply a reiteration of Theorem 5.1. In the absence of reputation effects ($\omega = 0$), the network setting is equivalent to the two-person setting. If $\omega = 1$ (case (ii)), then equation (5.6) reduces to

$$p(C_j \mid a_i) = \frac{C^a_{ij} + \sum\limits_k C^a_{kj}}{T^a_{ij} + \sum\limits_k T^a_{kj}} \tag{5.9}$$

(omitting the subscript t; $k \in N_j(g)$, $k \neq i$) such that $p(C_j \mid a_i) = p(C_j \mid a_k)$ for all i, j, k who are directly or indirectly connected. This implies that if two actors are (directly or indirectly) connected, they are acting on the same information. Because the choice of a sub-strategy depends exclusively on the conditional probabilities of cooperation, it follows that all i and k will choose the same sub-strategy against j. □

Theorem 5.2 states that in the absence of reputation effects, interaction in fixed networks does not differ from interaction in isolated dyads, and behavior converges to mutual ALLC

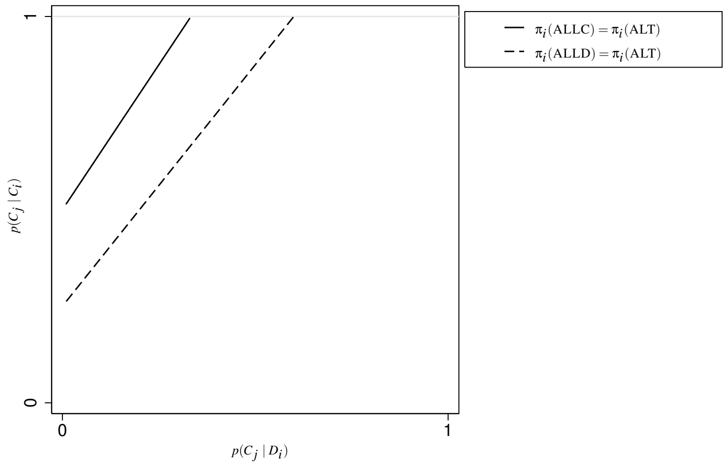

Figure 5.2: Conditions for switching between sub-strategies, with $T = 5, R = 3,$ $P = 1, S = 0$

or mutual ALLD. In the presence of reputation effects, however, behavior within dyads is no longer necessarily symmetric, because observations of the behavior of a partner can be "compensated" by observations of his behavior in interactions with other actors. In the extreme case of perfect information transmission (case (ii)), all partners of a given actor will choose the same strategy against this actor, but this does not imply that all actors choose the same behavior in interactions with *all* their partners. Examples can be constructed in which a group of actors cooperate with each other, but defect against a single actor, who, in turn, cooperates with all of them. In this case, the defections of these actors are "offset" by the observation that they cooperate in all other relations.

5.3.3. Stable states in dynamic networks

Finally, we turn to the situation in which games are played in *dynamic* networks such that actors can choose not only their behavior in the games, but also with whom they play. The important difference as compared to the fixed network case is that certain types of interactions can be discontinued if the expected reward from the interaction is less than the cost of its maintenance. We distinguish between two different scenarios with regard to the cost of network formation: the case in which $T > R > \alpha > P > S$ and the case in which $T > R > P > \alpha > S$, where α is the "maintenance cost" of a tie. In the first case, only interactions in which cooperation takes place are attractive. In the second case, relationships with mutual defection are also attractive, but actors will prefer isolation over exploitation ($\alpha > S$).

Let g^* denote a *component* of network g, that is, a sub-network of g consisting of a maximal subset of nodes and links such that all nodes are directly or indirectly connected.

Theorem 5.3.

(i) If $\omega = 0$ and $R > \alpha > P$, then $\forall i \in N, \forall j \in N_i(g), \sigma_{ij} = \text{ALLC}$ in stable networks. Any network configuration is possible.

(ii) If $\omega = 0$ and $P > \alpha > S$, then the network is complete and $\forall i, j \in N, \sigma_{ij} = \sigma_{ji} \wedge (\sigma_{ij} = \text{ALLC} \vee \sigma_{ij} = \text{ALLD})$ in stable networks.

(iii) If $\omega = 1$ and $R > \alpha > P$, then the network consists of one or more complete components g^*, while actors may also be isolated. $\forall g^* \subset g, (\forall i \in N(g^*), \forall j, k \in N_i(g^*), \sigma_{ji} = \sigma_{ki}) \wedge \neg(\forall i \in N(g^*), \forall j \in N_i(g^*), \sigma_{ij} = \text{ALLD})$.

(iv) If $\omega = 1$ and $P > \alpha > S$, all links are present in the network, and $\forall i \in N, \forall j, k \in N_i(g), \sigma_{ji} = \sigma_{ki}$.

Proof. Case (i) is a reiteration of case (i) of Theorem 5.2 with the addition that interactions in which both actors use ALLD are no longer stable, as $\alpha > P$ (note that the expected payoff per round in that case converges to P). Case (ii) differs from case (i) in that interactions in which both actors use ALLD are also stable, as the expected payoff in these interactions converges to P and $P > \alpha$. Moreover, since P is the minimal expected payoff, all links are created. Case (iii) partly relies on the same argumentation as case (ii) of Theorem 5.2. All actors who are directly or indirectly connected will base their actions regarding a given actor i on the same information. Therefore, *within* a component g^*, all links must be present, because if it is profitable for any actor i to connect to j, it must be profitable for all actors who are directly or indirectly connected to i to connect to j. For the same reason, all neighbors of j use the same sub-strategy in interactions with j. It is not possible that $\sigma_{ij} = \sigma_{ji} = \text{ALLD}$ for all $i, j \in g^*$. In that case, $p(C_j|D_i)$ would tend toward zero for all $i, j \in g^*$, and the expected payoff of all interactions would tend toward P, which is lower than α. However, it is possible to construct examples such that some of the actors in a component play ALLD and others play ALLC. Case (iv) differs from case (iii) in that all links must be created such that any stable network must be complete, because $\alpha < P$ and P is the lowest possible payoff. This implies that the situation in which $\sigma_{ij} = \sigma_{ji} = \text{ALLD}$ for all $i, j \in g$ is also stable. □

In brief, Theorem 5.3 states that if the cost of tie maintenance is low enough (cases (ii) and (iv)), the complete network will form. If the cost of tie maintenance is high, either any network configuration is possible (case (i)), or the network will consist of fully-connected components (case (iii)). Perfect transfer of information (cases (iii) and (iv)) has the effect that components must be fully connected. Next, let us compare the situation without reputation ($\omega = 0$) with the situation with full reputation ($\omega = 1$). When costs are low (cases (ii) and (iv)), the complete network will form, and there are no clear effects of reputation. When costs are high, however, (cases (i) and (iii)), we see that the presence of reputation effects clearly has an effect on the possible stable distribution of behavior. Without reputation (case (i)), only mutually-cooperative ties are stable. *With* reputation (case (iii)), the constellations in which some actors defect are also stable, as long as it is not the case that *all* actors in a component play ALLD against all their neighbors. Thus, we see that when there is a high cost of relations, reputation opens the door for defection. However, this does not necessarily mean that reputation always leads to lower overall cooperation. While case (i) states that only cooperative interactions are stable, the emerging network of interactions may be very sparse or even empty, yielding a very low level of cooperation. The question whether processes with and without reputation effects lead to different levels of overall cooperation will be addressed in the next section through the use of computer simulations.

Finally, we compare our formal results regarding the effect of reputation on cooperation between fixed and dynamic networks. In fixed networks, we did not observe a clear effect of reputation on the level of cooperation, although we did find an effect on the combinations of sub-strategies that can be stable. In dynamic networks, we found similar effects on possible combinations of sub-strategies, but we also found that in dynamic networks, reputation increases the potential for defection.

The theorems characterize stable networks for the extreme cases in which $\omega = 0$ or $\omega = 1$, but they do not characterize such networks for intermediate values of ω. In many cases, the characterizations are rather general and allow for many combinations of network configurations and behaviors to be stable. Moreover, the theorems do not provide any insight into which of the possible stable states are more or less likely to emerge given some present state of the process. That is, the analyses say nothing about the *dynamic process* which brings about stable states. In the following sections, we conduct computer simulations to address these issues.

5.4. Set-up of the simulation

In this section, we describe the setup of the simulation (or "computational experiment") that we use to explore the properties of our model with regard to the questions formulated in Section 5.1. In the simulation, we vary the parameters of the model in a systematic way to assess their effects on the outcomes of the process.

The simulation input consists of a set of *parameters* and a set of *initial conditions*. The parameters determine how the process works and are not changed in the process. These include the payoffs of the game, the strength of the reputation effect (ω), the speed of network adaptation (η), and the cost of maintaining network ties (α).

The initial conditions constitute the state in which the process starts. They include the initial network structure and the initial distribution of the beliefs of the actors. From a theoretical point of view, these parameters are somewhat less interesting, because they change as the process proceeds and can therefore not be used to make general claims about the typical outcomes of the process. However, their role is not merely instrumental. It is not uncommon for complex dynamic systems, such as the one under study here, to have a wide range of stable states, even for one set of exogenous parameter values. In such cases, studying which equilibria are more likely given certain initial conditions increases the empirical applicability of the model (cf. Buskens et al., 2008; Jackson and Watts, 2002).

Before we turn to the actual values of exogenous and endogenous parameters, however, we discuss the outcome variables in the simulation.

5.4.1. Dependent variables

There are basically two types of outcomes that are of interest for our inquiry: those related to cooperation and those related to the emerging network. To express the amount of cooperation in the process, we define two different measures.

The proportion of cooperation in interactions This is the proportion of cooperative actions (actors choosing C) among all *existing* interactions, that is, among connected actors. This measure is conditional upon the emerging network. If the network is very sparse, the level of cooperation *within* this network can still be high, even though very few actors are actually cooperating. For an empty network, this measure is not defined.

The total proportion of cooperation This is the proportion of cooperative choices of *all possible* interactions. The measure is 1 if all actors are connected *and* cooperate in all interactions, and it is 0 if there is no cooperation in any interaction or if there are no interactions at all (i.e., the network is empty). This measure is not very comparable across different network structures because its maximal value is restricted by network density, but it does offer a better indication of overall welfare.

We use the proportion of cooperation per interaction (or cooperation per *tie*) to study emerging cooperation in fixed networks. To look at total cooperation (that is, the second measure) in fixed networks would not make much sense, as this measure is restricted by the density of the network. Networks with higher density show more cooperation, simply because there are more interactions. Conversely, for dynamic networks, total cooperation is a more suitable measure than cooperation per interaction, because the number of interactions is itself an outcome of the process.

Although these two measures may look very different at first sight, note that when one is applied to fixed networks and the other to dynamic networks, they both measure the *proportion of maximally-attainable cooperation*. Thus, we believe it is justified to use these two measures to compare cooperation between fixed and dynamic networks.

Table 5.2: Parameter values of the simulation

Parameter	Description	Values
T	Temptation payoff	5
R	Reward payoff	3
P	Punishment payoff	1
S	Sucker payoff	0
ω	Extent of reputation diffusion	0, $\frac{1}{4}$, $\frac{1}{2}$, $\frac{3}{4}$, 1
η	Speed of network formation	0, 300
N	Number of actors	30
α	Linear tie costs	0.9, 1.9, 2.9

5.4.2. *Parameters of the simulation*

We make a distinction between the *parameters* of the simulation, which govern the mechanisms of the process, and the *initial conditions*, which comprise the input for the process. The parameters are the payoffs of the game, that is, ω, η, α, and the number of actors in the network. We ran simulations with 30 actors. For the payoffs of the game, we chose $T = 5, R = 3, P = 1$, and $S = 0$. We varied the extent to which reputation traveled through the network ω between 0 and 1. The case in which $\omega = 0$ represents the situation in which there are no reputation effects and interactions are completely independent across dyads. We varied the speed of network formation η to be zero ($\eta = 0$; the network is fixed), relatively slow ($\eta = 30$, the number of actors), and relatively fast ($\eta = 435$, the number of dyads). $\eta = 0$ refers to the situation in which networks are fixed. The linear cost of ties α is chosen such that $S < \alpha < P$, $P < \alpha < (T+S)/2$, or $(T+S)/2 < \alpha < R$. These values are chosen such that, in the first case, the cost of tie formation allows relationships with mutual defection, mutual alternation and mutual cooperation. In the second case, the cost of ties allows only mutual alternation and mutual cooperation, and in the third case, it allows only mutual cooperation.

5.4.3. *Initial conditions of the simulation*

The initial conditions of the simulation consist of the initial beliefs of the actors and the initial network structure. As for beliefs, recall from Section 5.2.2 that actors base their beliefs on the future behavior of other actors using four quantities. We obtain initial values for these quantities in the following way. For each directed belief (i, j), we start with a random "assumed history" in which i has cooperated with j T_i^c times and defected T_i^d times, where T_i^d and T_i^c are random integers between 0 and $-t_{max}$. In a sense, $-t_{max}$ can be interpreted as the number of periods the process had been running before we began observing it. We have set this value to $-t_{max} = 5$. The remaining two quantities C_{ij}^c and C_{ij}^d for $t = 0$ are chosen as random proportions of T_i^c and T_i^d. In effect, this means that we choose $p(C_j \mid C_i)$ and

Table 5.3: Initial conditions in the simulation

	Mean	S.d.	Min	Max
Density	0.50	0.31	0.00	1.00
Centralization	0.06	0.08	0.00	0.33
λ	2.50	1.12	1.00	4.00
$-t_{max}$	5	0	5	5

$p(C_j \mid D_i)$ randomly for t_0. We introduce one extra parameter, λ, to govern the relative values of $p(C_j \mid C_i)_{t_0}$ and $p(C_j \mid D_i)_{t_0}$. We choose $p(C_j \mid c_i)_{t_0}$ and $p(C_j \mid d_i)_{t_0}$ such that

$$p(C_j \mid D_i)_{t_0} = \frac{x_{ij}}{\lambda} \tag{5.10}$$

and

$$p(C_j \mid C_i)_{t_0} = 1 - \frac{y_{ij}}{\lambda} \tag{5.11}$$

where x_{ij} and y_{ij} are random variables between 0 and 1, and $\lambda > 0$. The higher λ is, the larger is $p(C_j \mid C_i)_{t_0}$ relative to $p(C_j \mid D_i)_{t_0}$, and the higher is the average expectation that opponents will be reactive. Therefore, a higher λ is associated with a higher overall tendency toward cooperation.

For the initial network structure, we draw from a set of artificially-generated network structures. To construct these networks, we use various well-known network models, including the Erdős-Renyi random graph model, the small-world model by Watts and Strogatz (1998) and the preferential attachment model by Barabási and Albert (1999). These models have been shown to reproduce some key characteristics of empirical networks, and the resulting networks therefore provide a reasonably plausible set of initial networks for the simulation.[5] We vary the parameters of the algorithms in such a way as to obtain a reasonable variance in network density and network centralization. Table 5.3 summarizes the initial conditions used in the simulation.

5.4.4. Convergence of the simulation

As we described in Section 5.2.5, the process is theoretically stable when the beliefs of the actors are stable. In a numerical simulation, however, such convergence might take a very long time, and it is also possible that beliefs will never become stable at all, even though they do converge to a certain value. As an example, suppose that j defected once after i cooperated, but thereafter, j always cooperated. In that case, $C^c_{ij} = T^c_{ij} - 1$, and as t tends toward infinity,

[5] The results do not differ between different network-generating algorithms, which we take as an indication that the precise method used does not matter much and that studying additional methods is not likely to yield new insights.

the fraction $p(C_j \mid D_i) = \frac{C^c_{ij}}{T^c_{ij}}$ approaches 1, even though it never becomes perfectly stable. For this reason, we introduce a pragmatic convergence criterion, which states that the process is assumed to be converged when the largest change in the beliefs of all actors is smaller than 5%.

5.5. Simulation results

The results reported here are based on a total of 7200 simulation runs. In the fixed-network runs ($\eta = 0$), the process always converged within 1000 rounds. In the dynamic-network ($\eta > 0$), 99.8% converged within 1000 rounds. In the results that follow, we include only runs that converged within 1000 rounds.

5.5.1. Results for fixed networks

Before we turn to the results for dynamic networks, we discuss some of the results for fixed networks ($\eta = 0$). These results constitute the "baseline" to which we compare our results for dynamic networks. Figure 5.3 shows how cooperation (here represented by the average of cooperation per tie; see Section 5.4.1) depends on the density of the initial network for different combinations of λ ("optimism") and ω ("reputation diffusion"). To indicate the general trends, median splines where added to the scatter plots.

First, cooperation depends heavily on λ. If $\lambda = 1$ (which means that all beliefs are initially completely random), virtually no cooperation occurs, while if $\lambda = 4$, cooperation levels are rather high. This result indicates that the initial conditions of the process have a strong impact on the outcomes. Nevertheless, we also see effects of reputation diffusion on the outcomes. The simulation results are in line with the analytical result that the presence of reputation effects generally increases the range of stable states. While the outcomes lie close together when reputation effects are absent ($\omega = 0$), the stable states tend to "spread out" when ω increases. This is most visible when $\lambda = 3$. This implies that the presence of reputation effects allows for more extreme levels of cooperation, both high and low.

The exception to this trend is the case in which $\lambda = 4$, which is the most favorable condition for cooperation. Here, strong reputation effects *decrease* the variance in cooperation outcomes, especially in dense networks. If density is smaller than .75, a small reputation effect ($\omega = .25$) allows for lower levels of cooperation. Generally, however, it seems safe to say that the spread of reputations helps to maintain high levels when $\lambda = 4$.

5.5.2. Results for dynamic networks

Table 5.4 shows average levels of cooperation and network density for different combinations of values of α (cost of ties) and ω (reputation). Cooperation at the macrolevel is measured in two different ways (see Section 5.4.1): the average level of cooperation per *dyad* and the average level of cooperation per *actual interaction*. For $\alpha = .9$, all stable networks are complete, because the cost of maintaining a tie is lower than P, the minimally expected payoff (cf. Theorem 5.3, cases (ii) and (iv)). The net effects of reputation on cooperation are marginal at this cost level.

If $\alpha = 1.9$, the cost of maintaining a tie is higher than P, which makes mutually-defective relationships unstable in the absence of reputation effects (cf. Theorem 5.3, case (i)). Indeed, cooperation per interaction is almost 100% if $\omega = 0$.[6] The cooperation per dyad, however, is comparable to the lower cost regime. When $\omega > 0$, we find that the density of stable networks increases, while the level of cooperation per interaction remains more or less constant. There are more interactions, but these are not cooperative interactions. Thus, in line with the analytical results (Theorem 5.3), when reputations are allowed to spread through the network, it allows for defection to continue, even if the cost of maintaining a relationship is higher than the expected payoff from a mutually defective relationship. As a consequence, the level of cooperation across all dyads decreases if $\omega > 0$. These patterns are, however, not linear in ω; density jumps when ω increases from 0 to .25 but then decreases. Similarly, the drop in cooperation per dyad is largest between $\omega = 0$ and $\omega = .25$ and much smaller between higher levels of ω. Finally, if the cost of a tie is only slightly lower than the expected payoff from a mutually cooperative relationship ($\alpha = 2.9$), we see that in the absence of reputation effects ($\omega = 0$), cooperation is maximal in all existing relationships, but density is much lower than when the cost is lower. With $\omega > 0$, we again see that cooperation per interaction decreases. At the same time, density decreases with ω until almost 0, when $\omega = 1$. In this case, cooperation per interaction is still high, but there are very few interactions so that the network is extremely sparse.

The averages in Table 5.4 are informative, but given the strong effects of the initial tendency for cooperation that we found for fixed networks, we should also compare the outcomes for different values of λ. Figure 5.4 shows how the average cooperation *per dyad* depends on the strength of the reputation mechanism (ω) for different values of α and λ. This figure can be compared to Figure 5.3, but instead of initial density (which is not so informative for dynamic networks), we now look at different values of α. As compared to the case of fixed networks, we generally find a somewhat stronger effect of ω. The direction of this effect depends heavily on α and λ. For the two lower values of α, the effect of ω is negative for lower values of λ and positive for higher values of λ. Thus, on average, the spread of reputation "catalyzes" the tendency the process already had at the start. This does not mean, however, that a strong reputation effect leads to high cooperation when there is a high λ and low cooperation when there is low λ. As in fixed networks, the range of stable outcomes also

[6]The proportion of cooperation is not exactly 1 because in the simulation, convergence may be imperfect (see Section 5.4).

Table 5.4: Average density and cooperation in stable networks by α and ω

		ω				
α		0	.25	.5	.75	1
0.9	Coop. per interaction	0.45	0.48	0.47	0.46	0.45
	Coop. per dyad	0.45	0.48	0.47	0.46	0.45
	Density	1.00	1.00	1.00	1.00	1.00
1.9	Coop. per interaction	0.92	0.64	0.59	0.59	0.53
	Coop. per dyad	0.43	0.46	0.42	0.43	0.43
	Density	0.44	0.68	0.63	0.61	0.62
2.9	Coop. per interaction	1.00	0.99	0.97	0.95	0.89
	Coop. per dyad	0.18	0.11	0.10	0.04	0.02
	Density	0.18	0.11	0.10	0.04	0.02

increases with ω. For high costs ($\alpha = 2.9$), the result is different in a number of respects. First, we see that the effect of ω on cooperation is nearly always negative or zero. This is even the case for the highest value of λ, which mostly leads to full cooperation with lower costs. Second, we see an interesting *divergence* of outcomes when $\lambda = 4$ and $\omega = .25$. Here, a number of simulation runs converged on an average lower level of cooperation compared to the situation where $\omega = 0$, while another group of runs converged on significantly *higher* levels of cooperation. Closer analyses of these latter cases reveal that they are characterized by a relatively lower initial network density. Among all the runs with $\lambda = 4$, $\alpha = 2.9$, and $\omega = .25$, the correlation between initial density and cooperation is .59. An explanation of these results might be that if the network is initially sparse, limited diffusion of reputations helps to form a network of cooperative relations. If the network is dense from the start, in contrast, the diffusion of reputation mostly serves to "spread bad news," which prevents the further build-up of a cooperative network.

To conclude the discussion of simulation results, we look at the effects of *initial* density in dynamic networks. For this purpose, we again draw scatter-plots of average cooperation per dyad by ω, λ, and α, as depicted in Figure 5.4, but we make a further distinction between different initial network densities. To reduce the number of graphs, we round density by multiples of 0.25 and show only values of λ for which we find interesting differences. The result is shown in Figure 5.5. This figure can be read as an extension of Figure 5.4 in which different values of initial density are separated in the rows, while different values of cost α have different subfigures. For the two lower values of α, we only show the results for $\lambda = 2$ and $\lambda = 3$, while for the highest value of α, we only show results for $\lambda = 3$ and $\lambda = 4$. For the remaining values of λ in each case, results do not differ for the different values of initial density.

Figure 5.5 shows an interesting interaction effect between initial density and ω, which is the weight of third-party information. For lower values of α, the effect of ω is clearly stronger for lower initial density. Moreover, we see that the variance of stable states increases more strongly with ω, especially when $\lambda = 3$. This means that when the initial density is higher, higher levels of cooperation can be reached with lower levels of ω.

For the highest value of α, the effects are less clear. When $\lambda = 3$, the effect of ω becomes stronger for *higher initial density*, while for $\lambda = 4$, there is no clear interaction effect. We can, however, identify "special cases" as mentioned in the previous section in which an exceptionally high level of cooperation is reached when $\omega = \frac{1}{4}$.

The implication of these results is that the combination of a high initial density and the presence of reputation effects is *not* the best recipe for cooperation if the network is dynamic. On the contrary, if there is a positive effect of reputation, this effect is most pronounced when the initial density is low.

5.6. Conclusions and discussion

In this chapter, we developed a model to study the co-evolution of cooperation in dyadic relationships and social networks. Given that the spread of reputations is generally thought to improve the chances of cooperation if networks are given exogenously, we were specifically interested in studying how reputation effects change if the network *co-evolves* with behavior in strategic interactions.

For this purpose, we constructed a model in which boundedly rational actors play two-person Prisoner's Dilemmas and, at the same time, choose their relationships. Through the network, actors learn about the past behavior of other actors and adjust their behavior accordingly. We derived a number of formal propositions regarding theoretically stable states in this model. In addition, we applied computer simulation techniques to examine which stable states are more or less likely, depending on the parameters of the model and the initial conditions of the process. The most important model parameter was the *strength of the reputation effect*, that is, the extent to which information on the past behavior of others can diffuse through the network.

The overall results can be summarized as follows. First, we find that if networks are exogenously determined, the range of possible stable states increases with the extent of reputation diffusion and the density of the network. States with higher overall cooperation levels emerge as compared to situations with less reputation diffusion, but also states with *lower* cooperation rates. Thus, we do *not* find that reputation effects always lead to more cooperation, as is most commonly assumed in the sociological literature. Rather, we find that relatively higher cooperation is a possible consequence of reputation effects, but so is lower cooperation. These findings are in line with previously formulated arguments, such as Burt and Knez (1995), who argue that reputation effects generally lead to more extreme outcomes.

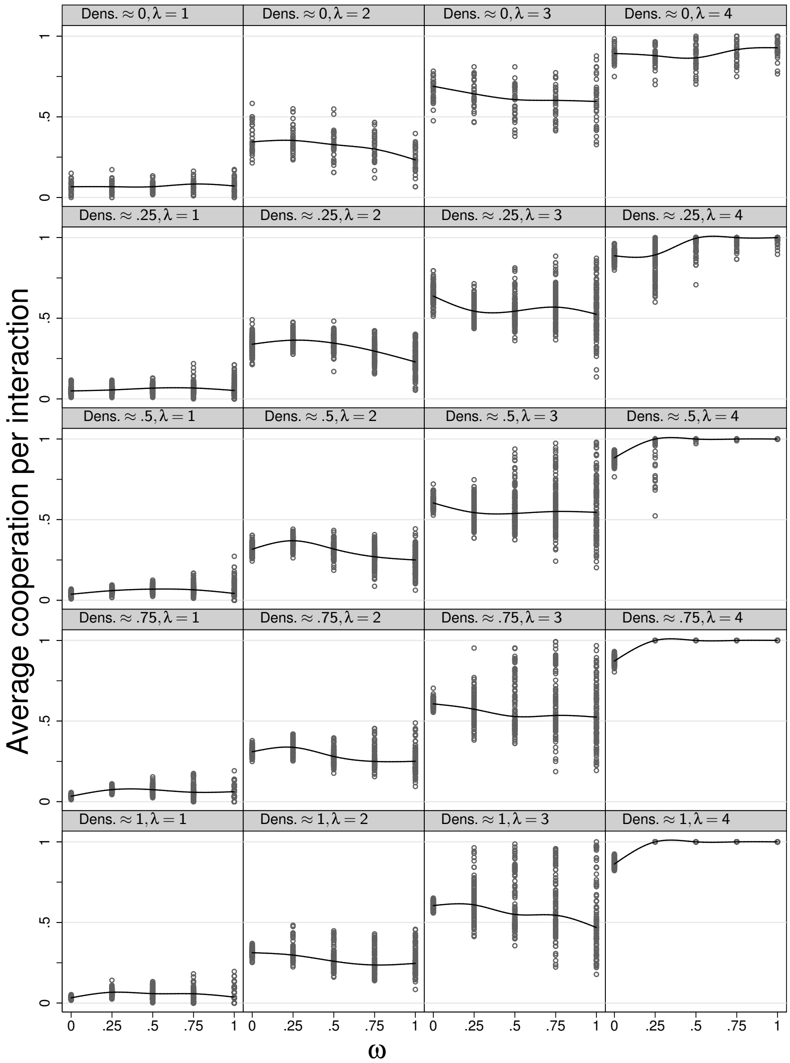

Figure 5.3: Average cooperation per interaction by ω in fixed networks with median splines added: graphs by λ (rows) and density (columns); density rounded to the nearest multiple of 0.25.

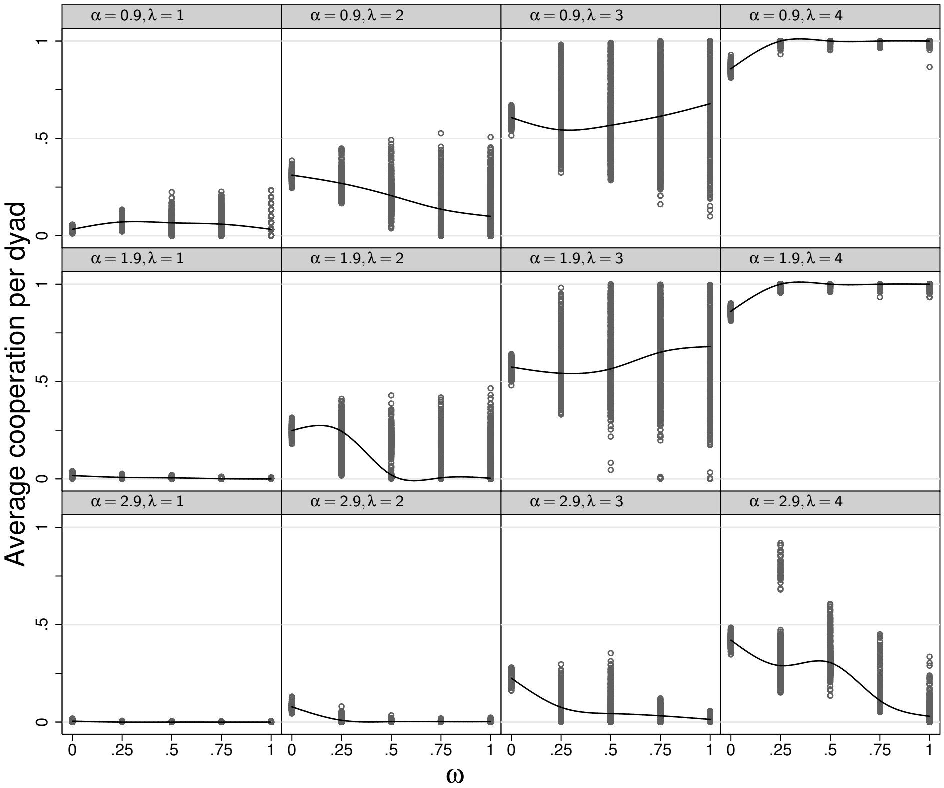

Figure 5.4: Average cooperation per dyad by ω in dynamic networks with median splines added: graphs by λ (rows) and α (columns)

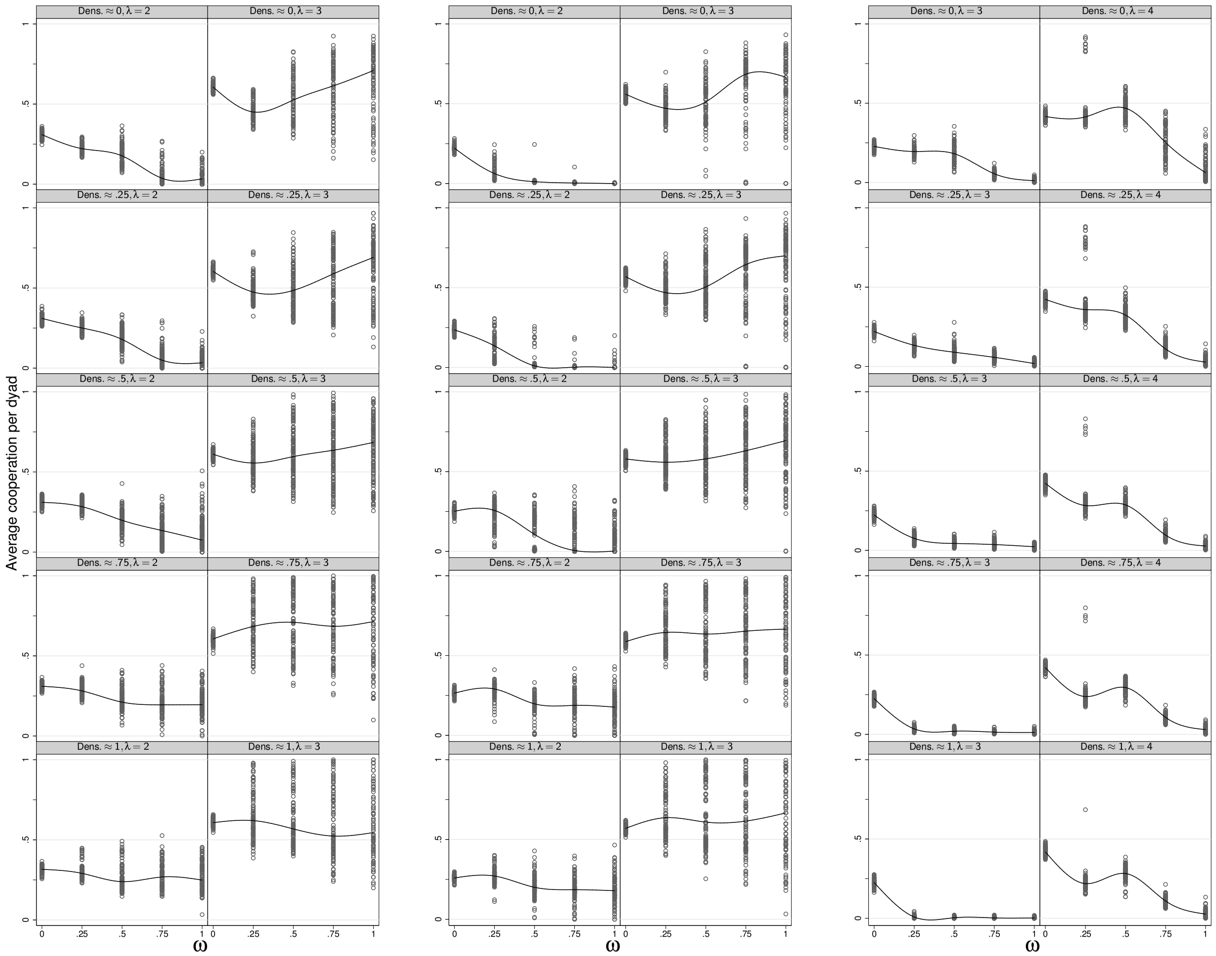

(a) $\alpha = 0.9$ **(b)** $\alpha = 1.9$ **(c)** $\alpha = 2.9$

Figure 5.5: Effects of reputation on cooperation by initial density, λ and α

While Burt and Knez (1995) rely on psychological mechanisms to explain this phenomenon, we have shown that it can also emerge from a simple learning model.

Second, we find that if the network is dynamic, the spread of reputation tends to “catalyze” the initial tendency of the process toward higher or lower levels of cooperation. Moreover, we find strong interaction effects with regard to the cost of maintaining ties and reputation effects in dynamic networks. When the cost of maintaining a tie becomes very high, maintaining a network of cooperative relationships becomes difficult; the addition of reputation effects makes this even worse, leading to empty networks in many cases.

Third, we find no indications that in a context in which the network structure is endogenous, high cooperation levels are likely to be the result of reputation effects in an initially dense network. Instead, we find that in dynamic networks, the effect of the spread of reputations tends to be stronger if the network is *less* dense. That is, the diffusion of reputations is most likely to lead to high cooperation rates if initial beliefs are “optimistic” (that is, λ is high) and the network is sparse. *As a result* of high cooperation, a dense network emerges. An interpretation of this effect could be that if the network is initially sparse, actors have the opportunity to initiate interactions only with those partners whom they expect to act cooperatively. The diffusion of reputation then helps in the further build-up of this “cooperative network.” If, in contrast, the network is dense from the beginning, there will also be some relationships in which actors do not behave cooperatively. In this case, the diffusion of reputations only helps to spread to “bad news,” which in turn hinders the development of cooperation.

While we believe that our analysis adds new insights to the study of cooperation in networks, the model also has some limitations. First, we modeled reputation only as learning, and we did not take into account that actors may care about their *future* reputation. Adding such “control” mechanisms to the model would not only make the model much more complex to analyze, but it would also put considerably higher demands on the rationality of actors.

Second, we did not model actor expectations with regard to the *linking* behavior of their interaction partners. In effect, actors in our model assume that their opponents will never unilaterally end the interaction. We decided not to do so because we wanted to focus strictly reputation effects, but in principle, the model could be extended in this direction by including expectations about the opponent’s linking behavior as conditional probabilities assigned to the beliefs of the actor.

Third, we did not consider information diffusion as a strategic choice. For simplicity, we assumed that the transfer of information is automatic and that actors do not make an explicit decision to pass along specific information. In many empirical applications, this assumption is unrealistic. As we can learn from studying gossip (Burt, 2001; Gambetta, 1994), people often have reasons to think strategically about what to tell to whom. For instance, actors may selectively relay information about their own interactions (only passing positive information about themselves) or even strategically spread negative information about others in order to sanction them (Ellickson, 1991).

Fourth, and related to the previous point, we treated the weight of reputation as a property of the *process* rather than as a property of *individual strategies*. In most applications, it would probably be more appropriate to assume that actors decide to what extent they want to make use of available third-party information. Considering this, one interpretation of our current model could be that all actors choose the same strategy with regard to the use of information. An alternative interpretation could be that all actors by default make maximal use of available information but that the availability of this information is determined by macro-level circumstances, such as the ease of communication (i.e., the value of ω in our model). A relatively straightforward way to check the robustness of the results to relaxation of this assumption would be to let ω vary between actors and thereby make actors heterogeneous in their susceptibility to third-party information. Alternatively, one might try to model the circumstances under which actors will take third-party information into account.

Fifth, the learning model applied here is rather simple. Actors have very simplified expectations about their opponent's behavior and update those beliefs using very basic methods. More complex learning models are conceivable in which actors, for example, assume that the behavior of their partners is conditional on a longer history than the previous round or take the reliability of their own estimations into account.

Finally, it would be interesting to look at the *stochastic* stability of the process. At present, our model is basically deterministic; the only stochastic element is the order in which actors update their behavior. Adding some "noise" to the process—in the sense that actors can make small "mistakes" every now and then—could have two advantages. First, it might help to resolve the problem that the model predicts many different stable states for any given set of parameters by showing which stable states recur in a perturbed process. This approach is adopted by, for example, Jackson and Watts (2002). Second, a "noisy" model would also be of substantive interest to investigate how reputation affects cooperation and network formation in a volatile environment. First intuition is that the introduction of noise would make cooperation even more fragile, but the spread of reputations may help to "protect" group cohesion and cooperation from small mistakes, as they can be compensated for by the good reputation of an actor.

Broadly speaking, we see two ways to develop the model further. The first is to address some of the theoretical issues mentioned above as extensions of the model. Another question, of course, concerns the *empirical* validity of the model. Eventually, we aim to explain and predict empirical processes related to cooperation and network formation with our model. In the present state, however, with so many theoretical issues still unresolved, directly testing predictions from the model with real-world data is likely to be problematic. For instance, if discrepancies between predictions and the data are found, it will be difficult to determine whether the discrepancy is caused by overly simplified assumptions about actor decision-making or by a misspecification of the underlying game. A more fruitful approach to test this model empirically would involve conducting controlled experiments. Using the methods developed in experimental economics and social psychology, one can study human behavior in complex strategic interactions “at close range” while also controlling properties of the larger environment. Such an approach would be most useful to assess the extent to which the model’s assumptions about actor decision-making are sufficient or are in need of modification.

CHAPTER 6

Summary and Suggestions for Further Research

6.1. Summary of the findings

It is well established in sociology that social networks influence many social and economic outcomes. Networks play an important role in, for instance, the diffusion of opinions and innovations, and can under certain conditions foster cooperation. The majority of past research has considered social networks as as exogenously fixed, and has studied the consequences of network embeddedness for individual behavior. However, it is increasingly recognized that social networks are dynamic, and are often the results of individual decisions about relationships. Thus, the question how networks evolve becomes salient. It is likely that individual decisions on relationships are related to the same individual traits that are also influenced by social networks. This implies that social networks and the traits of individuals within these networks develop interdependently, or *co-evolve*. For instance, we may be influenced by our close friends with regard to our political views, but our political views may also (partly) determine whom we pick as our friends. It is not trivial to predict what network structures would emerge from such a process, and what the emerging distribution of behavior would be.

The research presented here fits into this general theme, but focuses on a specific type of behavior, namely, behavior in social dilemmas. We focus on two types of social dilemmas in particular: coordination problems and cooperation problems. Coordination problems emerge when actors face individual incentives to align their behavior but are at risk of ending up in Pareto-suboptimal conventions. Cooperation problems emerge when actors have incentives to take advantage of each other, even though they could jointly gain more by cooperation (as in the Prisoner's Dilemma). Past research has shown that social networks influence behavior in these two dilemmas. This dissertation extends this research by relaxing the assumption that networks are fixed and by assuming that social networks and behavior are interdependent. We study how social networks and behavior in social dilemmas *co-evolve* and under which conditions optimal or suboptimal outcomes are more or less likely.

These question are addressed in four chapters, each from a different angle. Three chapters are concerned with coordination problems, while the fourth chapter studies cooperation problems. The core assumption of the theoretical approach is that both social networks and outcomes of social dilemmas are driven by individual *goal-directed* action. Hypotheses are derived through formal modelling of the interaction process, using both analytical methods and computer simulation. Empirically, we use two different strategies: laboratory experiments and field research.

Chapter 2 presents a theoretical study on how the emergence of different conventions in coordination games in dynamic networks depends on initial conditions in terms of behavior and social network structures. First, we specify a formal theoretical model for this co-evolution process. We then analytically characterize stable states of the process induced by the model. Because a multitude of stable states are possible, including efficient states, inefficient states, and states with multiple conventions, we use computer simulations to study which of these stable states are more or less likely to occur, given the initial conditions of the process. We focus on two features of stable states: *heterogeneity* and *efficiency* of emerging conventions. Emerging *networks* structures turn out to be completely determined by the constellation of behavioral choices in the coordination game. The results of the simulations show that although heterogeneous states are possible, they are unlikely to occur in a dynamic process. Moreover, we show that the initial density of the network strongly catalyzes the initial tendency of the process, while centralization has the opposite effect. When starting from a very heterogeneous situation, a higher density of the initial network makes coordination on the risk-dominant convention more likely, while higher centralization makes coordination on the risk-dominated convention more likely.

In Chapter 3, a laboratory experiment on coordination in dynamic networks is presented. The experiments tests predictions derived from the model developed in Chapter 2 on the effects of initial network density on efficiency and heterogeneity of conventions. In addition, we develop and test new hypotheses on the effect of information availability on the emergence of different conventions. Two information conditions are compared. Under global information, actors are informed about the actions of all other actors in the population. Under local information, actors are informed only about the actions of their direct neighbors. To model how actors choose new relations under local information, we assume that actors use the average behavior of current neighbors as an approximation of the expected behavior of a potential (unobserved) neighbor. The main theoretical result is that limiting information on behavior of others makes it *less likely* that a population will develop into subgroups supporting different conventions, because limited information makes it more difficult to avoid others who support a different convention.

In order to derive informative predictions for the experiment, we conduct new computer simulations that closely mimic the experimental conditions in terms of payoffs, initial networks, and information availability. We find only limited support for the hypotheses on emerging conventions at the macrolevel. In particular, we find no significant effects of the initial network structure or information availability on emergent conventions. However, if we examine individual behavior of the subjects, we find that the microfoundations of the model are largely supported. That is, we find clear indications that subjects' behavior resembles myopic best-reply play against direct neighbors. The way subjects choose new partners when they have limited information, however, differs from what we assumed in the model. The results are more consistent with a model in which actors assume that potential neighbors are likely to behave differently than their current neighbors.

Chapter 4 tests hypotheses on coordination in dynamic networks again, but in a "real-world" setting rather than in the laboratory. We analyze alcohol use among adolescents as a coordination game, arguing that using alcohol can be modeled as risk-dominant but inefficient behavior in a coordination game, given that adolescents face incentives to align their behavior with that of their friends. We test predictions from the model developed in previous chapters using data on alcohol use and friendships in Dutch high schools. Whereas most previous research on this topic studied only effects of *personal* networks of adolescents, our theoretical approach allows for predictions of effects of the *macro*level social network structure in a class on average alcohol use.

In the empirical analysis, we are able to replicate the "catalyzing" effect of initial network density on the development of alcohol use: the denser the initial network, the more likely the process will move further in the direction of the initial tendency. This confirmation, which we did not firmly establish in the experiment of Chapter 3, adds further support to the theoretical model. However, the predicted opposing effect of centralization could not be confirmed.

Chapter 5 studies cooperation problems, here conceptualized as two-person Prisoner's Dilemmas played in a co-evolving network. We develop a theoretical model in which actors play repeated Prisoner's Dilemmas and learn about their interaction partners using information from third parties (i.e., reputation). At the same time, actors can choose with whom they interact, thus changing the network endogenously. The model builds upon previous research that shows that cooperation is facilitated by cohesive networks. The model extends this research by relaxing the assumption that networks are fixed exogenously. We analyze this model using analytical methods and computer simulations. The analytical results show that in both fixed and dynamic networks, reputation effects allow for a broader range of stable states as compared to isolated encounters. However, cooperation levels in stable states are not necessarily higher with stronger reputation effects. The analytical results on dynamic networks show that stronger reputation effects lead to more homogeneous network structures.

Computer simulations of this model confirm the analytical results, and highlight that reputation effects in dense networks do not foster cooperation under all conditions. In fixed networks, reputation effects make *both* high and low cooperation rates more likely. In dynamic networks, stronger reputation effects lead to a higher variance in outcomes, but lead to lower cooperation on average. Lastly, the simulation results indicate that cohesive networks are not likely to emerge without a pre-existing tendency for cooperation.

6.2. Suggestions for further research

In each of the chapters, limitations and possible extensions are discussed that are specific to the research presented there. In this final section, we will not go again into the details of each study, but instead discuss a number of wider implications for further research that emerge from the collection of studies presented in this book. We first discuss theoretical extensions, then proceed with suggestions for empirical research.

6.2.1. Theoretical extensions

Alternative actor models

We developed theoretical models for coordination problems in dynamic networks in Chapters 2 and 3, and for cooperation problems in Chapter 5. While the models share a common theoretical approach, an important difference exists in the way actors are modelled at the microlevel. In the models for coordination, actors are modeled as simple myopic optimizers: they play a best response to what their interaction partners played in the round before.

In the model for cooperation problems of Chapter 5, the actor model is considerably more complex: actors are assumed to use the complete history of play with each interaction partner to form beliefs about these partners, to combine these beliefs with information from third parties, and to translate these beliefs into strategies for a repeated game as well as into linking decisions.

The reason for this difference, or discrepancy if you want, is relatively straightforward. Because defection is the dominant strategy in the Prisoner's Dilemma, myopic best-reply behavior would immediately lead to full defection by all actors. Thus, if we want to explain the variation in cooperation levels that we observe in real life, we need to use different decision models. The learning model that we used in Chapter 5 is a still relatively simple solution to this problem. The results from the experiments on coordination (Chapter 3) also suggest that the myopic best-reply model is an oversimplification. Subjects seem to use more complex strategies, such as strategies that are to some extent forward-looking.

In principle, of course, it would be desirable to have *one* consistent actor model that can be used to model behavior in various social dilemmas. After all, it seems hard to defend that people would behave according to simple best-reply rules when confronted with coordination problems, and switch to higher levels of rationality when facing cooperation problems as the models presented in this dissertation suggest. Rather, we would like to be able to explain behavior in *different* strategic situations using the *same*, preferably parsimonious, microlevel model.

Concerns about consistency are precisely one of the main reasons why many economists insist on assuming perfect rationality, even when it is clear that this assumption is unrealistic. In the context of network dynamics, however, we see two reasons to deviate from this principle. First, the complexity inherent in large-population network processes makes it extremely implausible that people are capable of determining the optimal choice in every situation. Second, the same complexity also makes that models that assume perfect rationality become too complex for the *modeler*. Using relatively simple microlevel models, then, can be a way to reduce complexity at the microlevel in order to derive macrolevel predictions.

Thus, we see a tension between, on the one hand, the desire to arrive at a model that is complex enough to be applied in various strategic situations, and, on the other hand, reasons to keep actor models relatively simple. A major task for future research is therefore to investigate the implications of different microlevel assumptions for macrolevel consequences in network formation processes. Building upon the models developed in this dissertation, one direction could be to apply learning models as in Chapter 5 to coordination problems. That is: actors would use the experiences of past actions of other actors to form beliefs about their typical behavior, and optimize their own behavior accordingly. Corbae and Duffy (2008) develop ideas along these lines.

Another important extension is to consider *forward-looking* behavior. A limitation of the models we developed is that actors are assumed to be mostly *backward looking*, in the sense that they base their decisions on what happened in past interactions, but do not—or to a limited extent—consider the long-term consequences of their decisions. In Chapters 2 through 4, actors are myopic: they simply try to maximize their payoff in the present period by choosing a best reply to the choices of interaction partners in the previous period. Actors do not anticipate reactions by their partners on these actions. In Chapter 5, actors are assumed to be somewhat more sophisticated and take into account that their interaction partners react to their actions in the next round. However, they do *not* anticipate reactions by third parties, who might sanction by defection or ostracism. Taking into account that actors anticipate reactions by other actors allows for studying control effects of reputation on cooperation, in addition to the learning effects studied in Chapter 5. Moreover, a model based on forward-looking behavior might also offer better explanations for the results of the experiments on coordination. At present, however, forward-looking behavior in the context of network formation is an area that is hardly explored. Some possible directions that might be taken by such research are sketched by Jackson (2008, Ch. 11). However, for the reasons sketched above, assuming that actors can perfectly anticipate *all* possible reactions by other actors is not desirable either. A model that considers forward-looking behavior in a network formation context without assuming perfect rationality is discussed by Berninghaus et al. (2008).

Other social dilemmas

We discussed two social dilemmas: cooperation and coordination problems. A natural extension of this research is to broaden the analysis to other social dilemmas, such as the Chicken Game, or multi-person public good games. Some other dilemmas have been studied in the networks literature; see for example Bramoullé et al. (2004) on anti-coordination games and Ule (2005) on multi-person Prisoner's Dilemmas. These models, however, are rather heterogeneous in terms of assumptions and analytical methods. Developing a more general theoretical framework that can explain behavior in various social dilemmas with co-evolving networks remains a major task for future research.

Heterogeneity

A simplifying assumption throughout the dissertation is that actors are *homogeneous* in terms of abilities, action alternatives, and payoffs associated with these alternatives. Differences between actors exist only as arbitrary initial conditions in a dynamic process (e.g., initial network positions, initial behavior, or initial beliefs) or emerge endogenously as a result of this process (e.g., different behaviors in a stable state). While clearly unrealistic, the homogeneity assumption was used with the aim to investigate to what extent individual differences can be *explained* as the result of a social process, without assuming differences a priori.

It is, however, possible that in reality, *structural* individual differences (as opposed to endogenous differences) exist that impact network evolution processes. For instance, it is likely that people differ in sociability, that is, the ability to maintain social relations. In terms of the models developed here, such differences would be expressed as individual differences in costs of ties. A first intuition is that such variances would cause a "natural" tendency towards centralization of the network, because some actors can more easily form ties than others. The results on centralization in Chapter 2 suggest that this, in turn, would influence the outcomes of co-evolution processes.

Another form of heterogeneity is that actors differ in terms of preferences for certain outcomes. For instance, in the context of coordination problems, some actors may prefer one convention while other actors prefer another. This approach is explored by Bojanowski and Buskens (2008), who show that such heterogeneity may lead to rather different network structures as compared to the homogeneous model studied in Chapter 2.

Differentiation of actions

Another difference between the models of Chapters 2 through 4 on the one hand, and Chapter 5 on the other hand, is that the models on coordination assume that actors can only choose *one* action against all their interaction partners, while in the models on cooperation actors play *dyadic* games in which they can cooperate with one partner and defect with another.

This difference in itself is not problematic, as the models can be considered as different answers to different types of problems. In many situations, differentiating behavior between interaction partners is simply impossible; the models of Chapters 2 through 4 apply to those situations, at least. In other situations, differentiating behavior is possible, but costly, for example, in choosing languages or computer operating systems. For such situations, our models might be considered as reasonable simplifications.

It would, however, be interesting to study the conditions under which people would be willing to differentiate their behavior between different interaction partners, even if this comes at a price. Under what conditions, for example, would people be willing to invest in learning an additional language, rather than adapt to their social environment or to change their relations?

6.2.2. Suggestions for empirical studies

With the studies in Chapter 3 and 4, we hope to have made some contributions to the empirical validation of theories on network dynamics. Much more empirical work is needed, however. Research on network dynamics is currently dominated by theoretical studies. This is not necessarily a problem, as the complexity of the problem requires that the theoretical implications of different assumptions and modeling approaches are carefully explored. However, without systematic empirical tests of the implications of theoretical models, it is difficult to judge which assumptions might be problematic, and to decide on how future theoretical research should be developed. In this final section, we sketch the outlines of two empirical studies that follow more or less directly from the research presented in the dissertation. Finally, we discuss some more general directions that empirical research on network dynamics might take.

Most attention in this dissertation went to coordination problems (three chapters), and less to cooperation problems (one chapter). Moreover, the chapter on cooperation cooperation problems does not include empirical work. Thus, an obvious way to proceed with this line of research would be to “complete the picture,” and conduct experimental studies as well as field research on cooperation problems in dynamic networks. As with coordination problems, experiments could be used as a first test of the predictions of the model of Chapter 5, and to examine to what extent individual behavior deviates from what is assumed in the model. Of particular interest for experimental studies on cooperation is the question whether a model that assumes only learning—and no control—can sufficiently explain individual choices in cooperation problems in dynamic networks.

Another interesting issue that can be addressed in experimental studies is the interplay of various sanctioning mechanisms. If cooperation problems are embedded in networks, reputation effects can lead to sanctioning of defectors by third parties. In dynamic networks, defectors may also be sanctioned by ostracism. These mechanisms are not necessarily compatible. The use of ostracism might prevent the formation of dense networks that facilitate the emergence of reputation. Experiments can be designed that manipulate the environment such that these mechanisms can be compared. Such an experiment might, for instance, compare cooperation rates in four conditions in a two-by-two setup: conditions with and without reputation effects, and with fixed and dynamic networks.

Next to experimental research, theory on co-evolution of cooperation and networks should be tested in field research. An interesting application is *inter-firm relations*. Strategic alliances, in particular, can be analyzed as dyadic cooperation problems. In these alliances, firms undertake collaborative R&D projects, in which both partners can benefit from the sharing of knowledge. However, firms also face incentives to act opportunistically and take advantage of the efforts of their partner. In this sense, strategic alliances share features of the Prisoner's Dilemma (Kogut, 1989; Parkhe, 1993). Economic sociologists have long emphasized the important role of networks in the success of inter-firm relations (Granovetter, 1985; Powell, 1990). Empirically, records of such collaborations can be used to reconstruct inter-organizational networks in a longitudinal fashion (e.g., Powell et al., 2005).

A possible cause of the relative lack of "real-life" empirical research on network dynamics is the shortage of suitable data. To test hypotheses on network dynamics, one typically needs detailed measurements of all relevant relations in a population, measured at various points in time. In order to study *co-evolution* of behavior and networks, one needs, in addition, longitudinal measures of behavior. Collecting such data takes considerable time and effort, and consequently, datasets that meet these requirements are rare, although some datasets exist (e.g., the data analyzed in Chapter 4, see Knecht, 2006).

Of course, data do not always need to be collected in the field. Sometimes, existing data sources can be used, as exemplified by the research on strategic alliances mentioned above. Another potentially interesting source of data is presented by the recent emergence of so-called "online social networks;" web-based services that allow users to maintain social relations. In the last few years, these services have become immensely popular.[1] The fact that relations via these services are mediated by technology allows for detailed longitudinal observations of networks and certain types of behavior. Although data collection using these services will have to deal with concerns over privacy, first steps in this direction (Lewis et al., 2008) prove this to be a promising development.

Data from such administrative data sources, however, are necessarily limited, because the researcher has little control over the measurements and can often not choose the specific variables needed to test particular hypotheses. Moreover, for many topics (e.g., alcohol use among adolescents as studied in Chapter 4) data records are simply not available. Therefore,

[1] In January 2009, the Dutch service *Hyves* claimed over 8,000,000 members at her website. This would amount to roughly 50% of the Dutch population.

own data collection in the field remains indispensable. To advance research on social network dynamics, it is essential that future research efforts and funds focus on the collection of new longitudinal network data, which can be used to test hypotheses from existing theoretical models and can suggest new research problems.

APPENDIX A

Instructions Used in the Experiment

This appendix contains the written instructions that were distributed among participants of the experiment described in Chapter 3, translated into English. The experiment was conducted in Dutch; the original Dutch instructions are available from the author upon request. In the experiment, we used four different versions of the written instructions: two for the two payoff conditions (high risk and low risk) and two for the different information conditions. We present the instructions for the high risk, local information condition. The instructions for the low-risk condition differ only from those in the high-risk condition with regard to the payoffs shown in Figure A.1. In the instructions for the global information condition, some sentences were formulated somewhat differently as compared to the local information condition. In the text below, these differences are indicated by footnotes. The numbering of tables and figures in the instructions has been adapted from the original version to be in line with the rest of the dissertation; otherwise, the layout has been preserved as much as possible.

—INSTRUCTIONS START HERE—

INSTRUCTION

PART 1

Welcome to this experiment, and thank you for coming. Please read the instructions until the bottom of page 7, marked "End of Instruction Part 1." If you at that point have any questions, feel free to ask for additional assistance by raising your hand. Also during the experiment, you can always ask for assistance by raising your hand in case something is not clear to you. However, please only do so if absolutely necessary; the instruction should speak for itself. All other participants receive the same instructions.

In this experiment, you can earn money depending on the choices you make. These earnings will be paid at the end of the experiment, without the other participants knowing how much you earned or which choices you made. From now on, you are no longer allowed to speak with anyone this room, and you should not look at the computer screens of the other participants. It is important that we know your choices based on your own considerations, and that you make your choices based only on this instruction and on what you see at your screen. There are no "right" or "wrong" choices.

During the experiment you will be part of a group with seven other participants. Some of these participants will be your "neighbors," others won't (these are not necessarily the participants sitting next to you; we get back to this issue later). The experiment consists of two parts. Each part consists of 15 rounds. In each round, you and the other participants in your group make choices. The number of points that you earn depends on both your choices and the choices of your neighbors in your group. For each point, you will be paid € 0.01. We now explain the choice alternatives in detail.

Choices

In each round, you and the others in your group choose simultaneously between "Left" or "Right." If you choose *Left*, you earn 20 points for each neighbor who also chooses Left; you earn 0 points for every neighbor who chooses Right. If you choose *Right*, you earn 16 points for every neighbor in the group who also chooses Right; you earn 14 points for every neighbor who chooses Left. The choices of participants who are not neighbors do not affect your score.

The number of points that you earn with each of your neighbors is summarized in the figure below:

		Neighbor	
		Left	Right
You	Left	20	0
	Right	14	16

Figure A.1: Choices and points

We will clarify this later using an example.

Relations and neighbors

Who of the other participants in your group are your neighbors, is decided by you. Per round, you can change one relation by proposing a new relation or by ending the relation with a current neighbor. *You can become an other participants' neighbor only if this participant also wants to become your neighbor; to end a current relation you do not need the permission of the other.* You can become neighbors with an other participant if one of you proposes this and the other accepts the proposal, or if both of you propose simultaneously. You can lose a neighbor if this neighbor does no longer want to be your neighbor, regardless of whether you wanted to remain a neighbor of this participant yourself. It will be explained below how this works in practice.

Lastly, you cannot only *earn* points in relations, but you also have to *pay* “maintenance costs” for each of your relations. The more relations you have, the more expansive these will be. The costs that you have to pay for each number of neighbors are shown in Table A.1.

Table A.1: Costs of relations

Number of neighbors	Total costs of your relations
1	7
2	16
3	27
4	40
5	55
6	72
7	91

Thus, the total profit in one round depends on:

- The number of points as depending the choices you and your neighbors make;
- The total costs of your relations.

Example
Suppose that, it some round, you have three neighbors. Your choices and the choices of your neighbors are as follows:
You choose Left;
Neighbor 1 chooses Left;
Neighbor 2 chooses Right;
Neighbor 3 chooses Left;

You earn 20 points from the relation with neighbor 1, 0 points from the relation with neighbor 2, and 20 points from the relation with neighbor 3. Taken together, your (gross) payoff from the relations with your neighbors is $20+0+20=40$ points. The costs for maintaining three relations is 27 (see Table A.1). Your total (net)payoff in this round in therefore $40-27=13$.

Note: because in every round, you do not only earn points but also have to pay costs, it is possible that you make a loss in some rounds.

Making choices using the computer

Shortly, you will be matched to seven other participants by the computer. You make your choices using the computer screen. The experiment consists of 15 rounds. Each round consists of the following steps:

Step 1: You have the opportunity to change one of your relations, i.e., end a relation with a neighbor or do a proposal to someone to become neighbors;

Step 2: If other participants want to start a new relation with you, you have the opportunity to decide whether you want to accept these proposals. If two participants do proposals to each other, a relation is created immediately without the need for acceptance.

Step 3: You and your neighbors each choose between Left and Right.

Step 4: Your payoffs and the choices of your neighbors are displayed.

The only exception to this is the first round, in which the computer determines who your neighbors are.

In each phase, it is displayed on screen who your neighbors are and what the choices of all other participants in the previous round were. This might, for example, look as in Figure A.2:

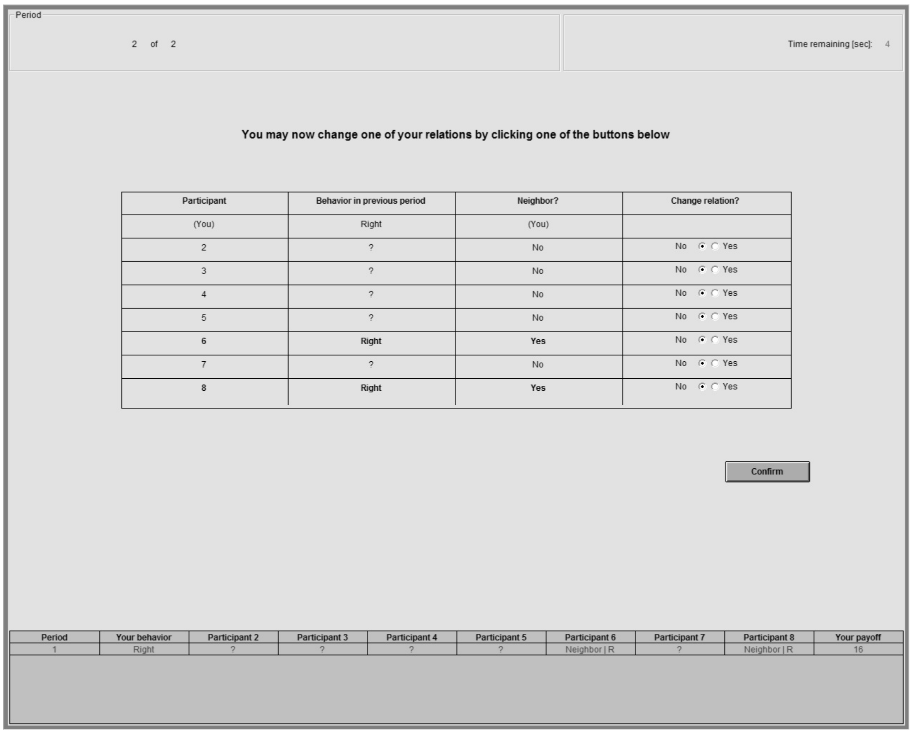

Figure A.2: Screenshot of the experimental interface: choosing relations with local information (translated from the original Dutch)

This is the screen in which you can change relations (Step 3). The first column shows the numbers identifying the other subjects, with yourself (number 1) at the top. Note: the numbers 2, 3, etc. refer to the second and third neighbor in *your* table. Number 2 in your table is not necessarily the same participant as number 2 in the table of one of the other participants. All participants really are people in this room.

The second column displays the behavioral choices of your neighbors (and yourself) in the previous round.[1] At moments at which you can change relations (as in the example screen above) this is displayed in the fourth column. To make it a bit easier to see who your neighbors are, all information about your neighbors is printed in **bold face**.

At the bottom of the screen, we summarize what happened in previous rounds. For each other participant, it is shown whether this participant was a neighbor, and if so, what his or her behavior was.[2] Your own behavioral choices and payoffs are displayed as well.

Before we start with the actual experiments, we will play a number of practice rounds. In these practice rounds, you will not win or lose any points. Also, you will not be matched in a group with other participants, but instead you will be playing against the computer. The computer's choices were pre-programmed. After the practice rounds, we will start with Part 1 of the experiment.

This concludes the instructions for Part 1; you will read the instructions for Part 2 only *after* Part 2 of the experiment has finished. If you have any questions at this point, please raise your hand now. If all is clear to you, please click "Continue" on the computer screen. As soon as all participants have finished reading the instructions, we will start the practice rounds.

On the next page, we again show Figure A.1 and Table A.1. During the experiment, you may use this page for reference.

[1] In the global information condition, this sentence read "The second column displays the behavioral choices of each participant (and yourself) in the previous round."

[2] In the instructions for the global information condition, this sentence read: "For each other participant, it is shown whether this participant was a neighbor, and what his or her behavior was."

		Neighbor Left	Neighbor Right
You	Left	20	0
	Right	14	16

Figure A.1: Choices and Points

Table A.1: Costs of relations

Number of neighbors	**Total costs of your relations**
1	7
2	16
3	27
4	40
5	55
6	72
7	91

— END OF THE INSTRUCTIONS FOR PART 1 —

— PART 2 —

Continue reading only if you have been instructed on the screen to do so!

INSTRUCTION PART 2

In Part 2, you will again be placed in a group with seven other people in this room; in part, this will be different people than in Part 1. Otherwise, Part 2 will proceed as Part 1. However, there is one difference. In Part 1, you were shown the behavior of your neighbors in the previous round. *In Part 2, you will be able to see the behavior of all other participants in the previous round*.[3] Otherwise, Part 2 is identical to Part 1; you may again use the scheme on page 7 for the scoring of points and the costs of relations. Just like in Part 1, the computer will determine who your neighbors will be in the first round.

When you have finished reading, please click "Continue".

—END OF THE INSTRUCTIONS—

[3] In the instructions for the global information condition, this line read: "In Part 2, you will be able to see only your neighbors' behavior in the previous round ."

APPENDIX B

The Computer Interface Used for the Experiment

This appendix describes the details of the computer interface used during the experiment described in Chapter 3. The interface was created using the z-Tree software (Fischbacher, 2007). Via this interface, subjects received information on the computer screen, and communicated their decisions by clicking on-screen buttons using the mouse.

We present the most important screens of the computer interface. For this purpose, all text on the screens has been translated from the original Dutch. The screens shown here are taken from the practice periods, which preceded the actual experiment in each session. The screens from the practice periods look identical to the screens used during the later periods.

At the start of each session, when subjects entered the laboratory, all computers were running the interface showing Screen 1. This screen was used to assign subjects to a computer, and instructed subjects to start reading the instructions once they were seated. When subjects had finished reading the instructions, they were asked to click "continue", and were subsequently presented with Screen 2. This screen announces the start of the practice periods, and informs subjects that their participant number is "1." The participant number was "1" for *all* participants, in order to ensure that the computer interface looked the same to all subjects.

As soon as all subjects had indicated that they were ready to start by clicking "continue," the first practice period was started. In this first period, subjects were connected in one of the three initial networks, and were asked to choose between "Left" and "Right" via Screen 3. In this screen, a table provides relevant information about the other participants in the subject's group. In this table, each row represents a participant. The left column in the table indicates the participants' number, while the right column indicates whether each of the other participants is a neighbor or not. Information on participants who are currently neighbors are printed in bold face. The first row of the table represents the subject herself. Subjects could choose between "Left" or "Right" by clicking one of the buttons at the bottom of the screen.

After this first behavior choice, Screen 4 informed informed about the results of their actions and the actions of other participants. Again, the information about other participants is communicated by a table, which on this screen includes a third column in the middle containing information on the behavior choices of the other participants. Screen 4 was taken from the practice periods of the global information condition, and shows the behavior of *all* other participants in the group. In the local information, the table showed only the behavior choices of current neighbors, and showed question marks in this column for non-neighbors.

From the second period onwards, subjects could choose relations at the start of each period via Screens 5 and 6. Thus, after Screen 4, subjects were presented with Screen 5, in which they could make a change to at most one relation. A change could be made by setting one of the radio buttons in the rightmost column of the table, and subsequently clicking "confirm." If subjects tried to change more than one relation, they were notified in a pop-up screen that this is not possible. At the bottom of Screen 5, subjects could review the history of play. After Screen 5, subjects had the opportunity to confirm or reject incoming proposals if there were any. Again, this choice was made by choosing "yes" or "no" in the rightmost column of the table, and then clicking confirm.

Next, Screen 3 ask for a behavior choice. This screen differs from Screen 3 in that also the behavior of other participants in the previous period is displayed, depending on the information condition. At the end of the period, subjects are informed about the results of the period via Screen 4. Thus, a typical period of the experiment consisted of a succession of Screen 5, Screen 6, Screen 7, and Screen 4.

After the three practice periods, subjects played 15 "real" periods in the same information condition as the practice periods, and then played another 15 periods in the alternative information condition. The interface used during those periods was the same as that of the practice periods shown here. At the end of each set of periods, subjects were informed about their total profits from that part of the experiment, and after the second set of 15 periods, also about their total profit from the whole experiment. The experiment concluded with an on-screen questionnaire.

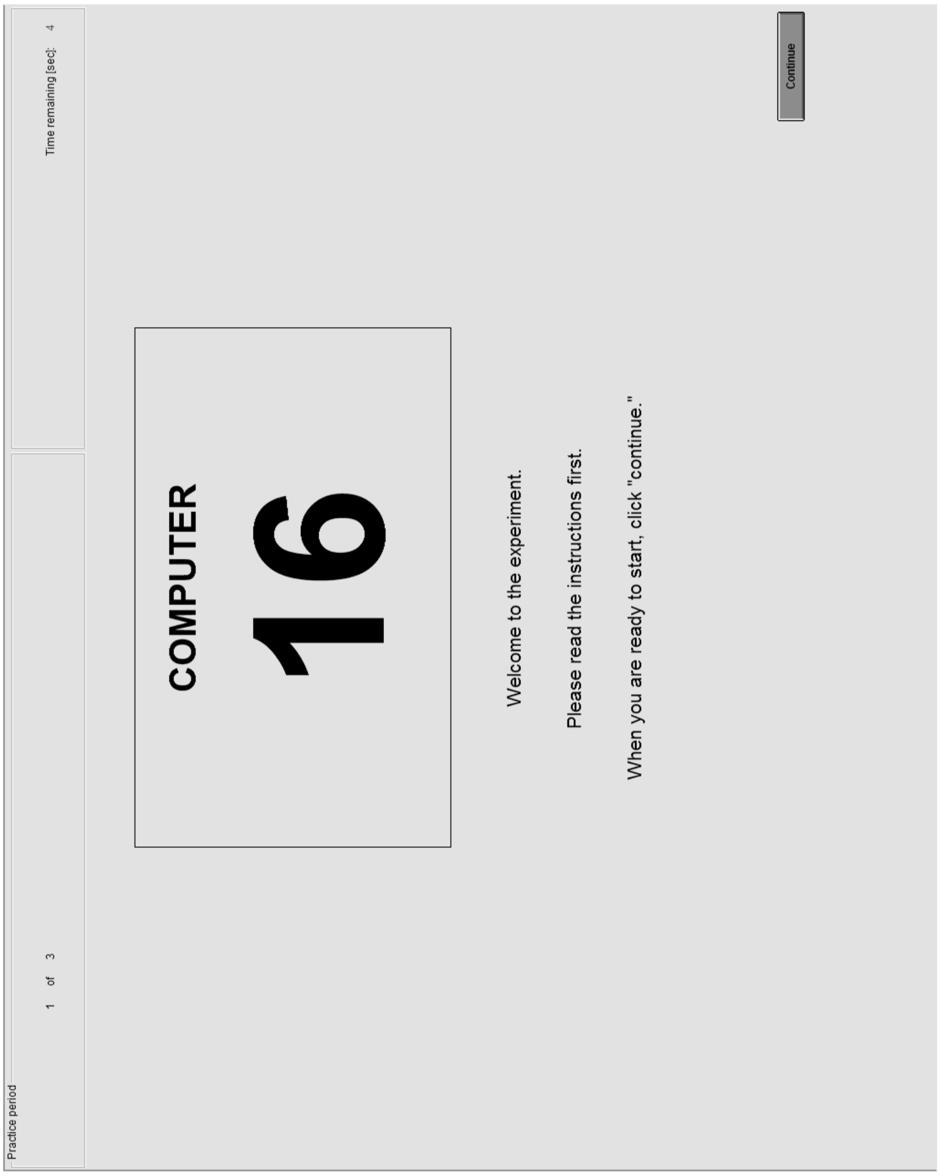

Screen 1: Welcome

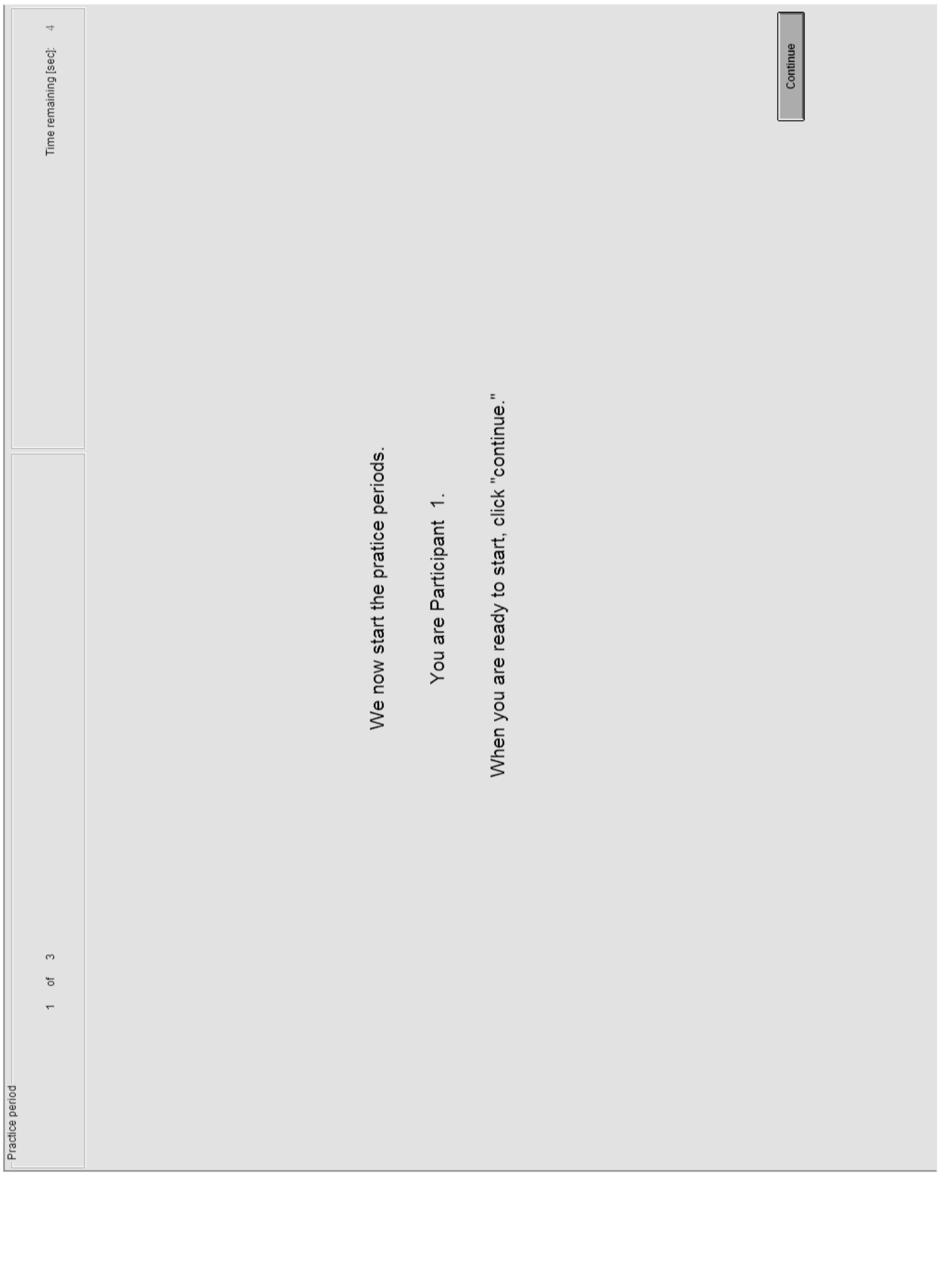

Screen 2: Start of practice periods

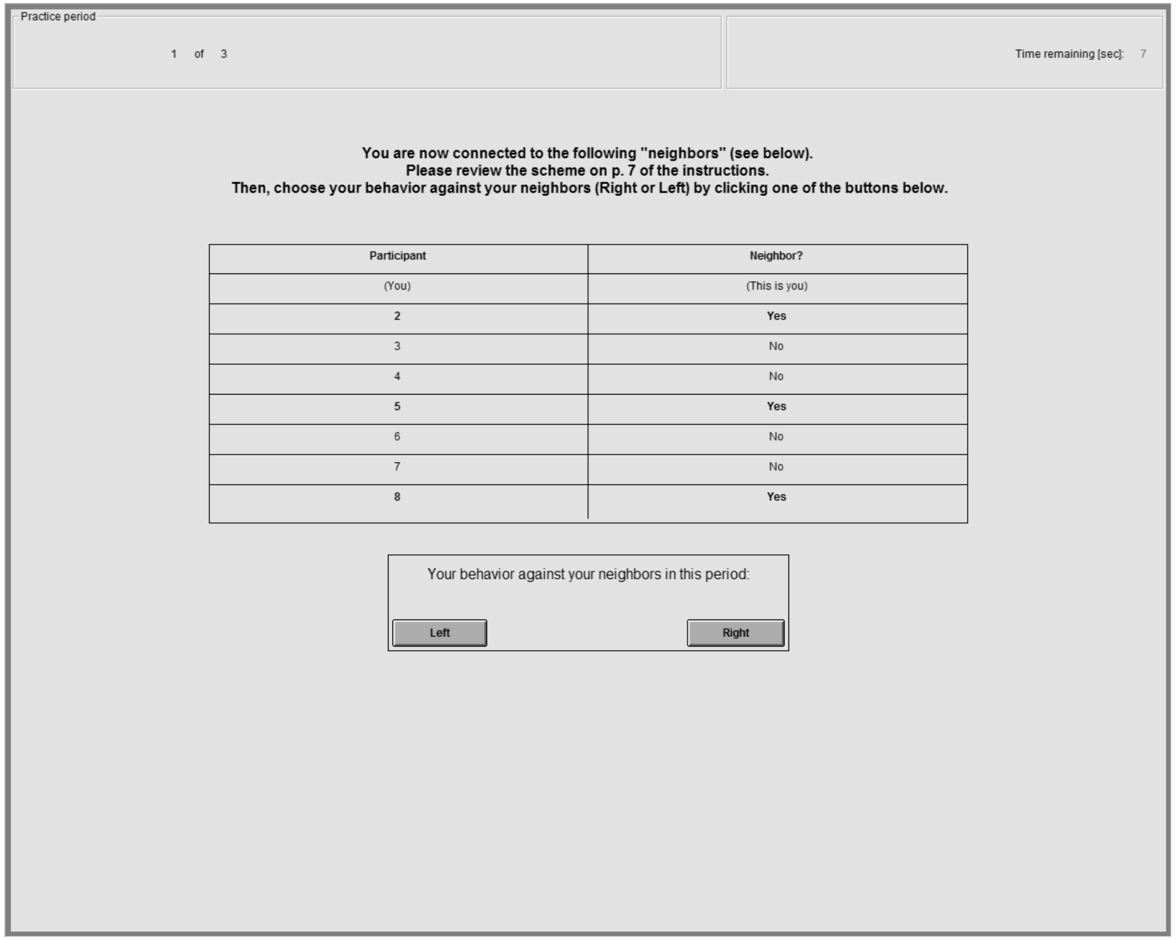

Screen 3: First behavior choice

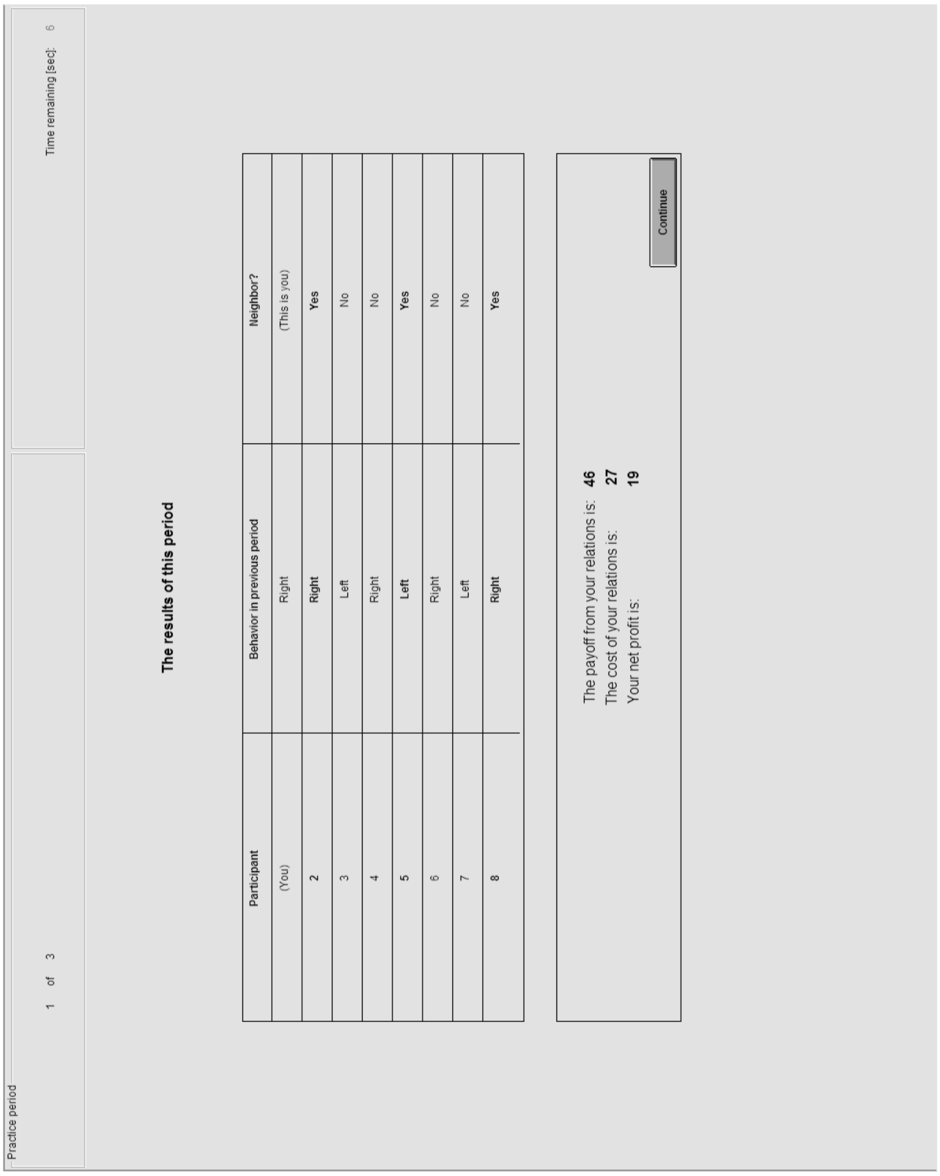

Screen 4: Results

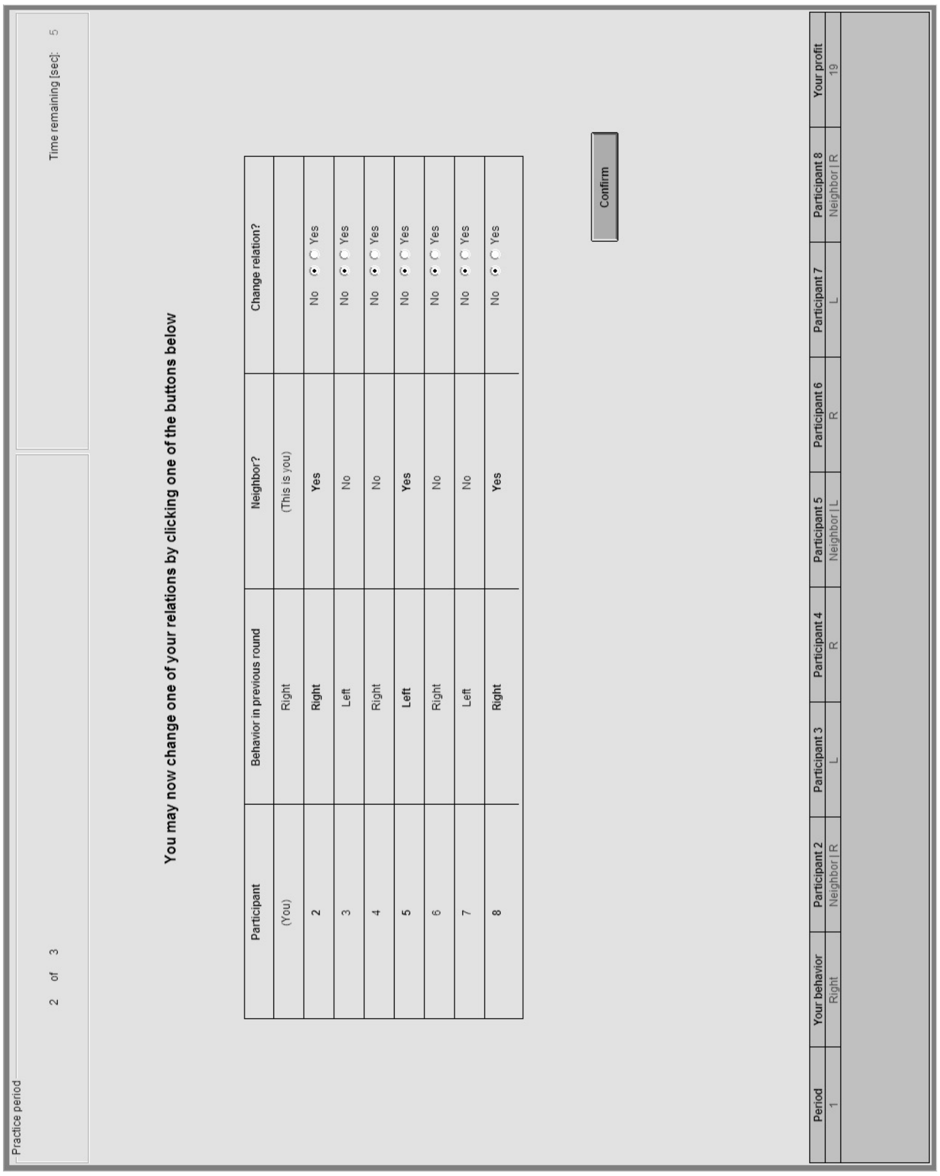

Screen 5: Tie proposals

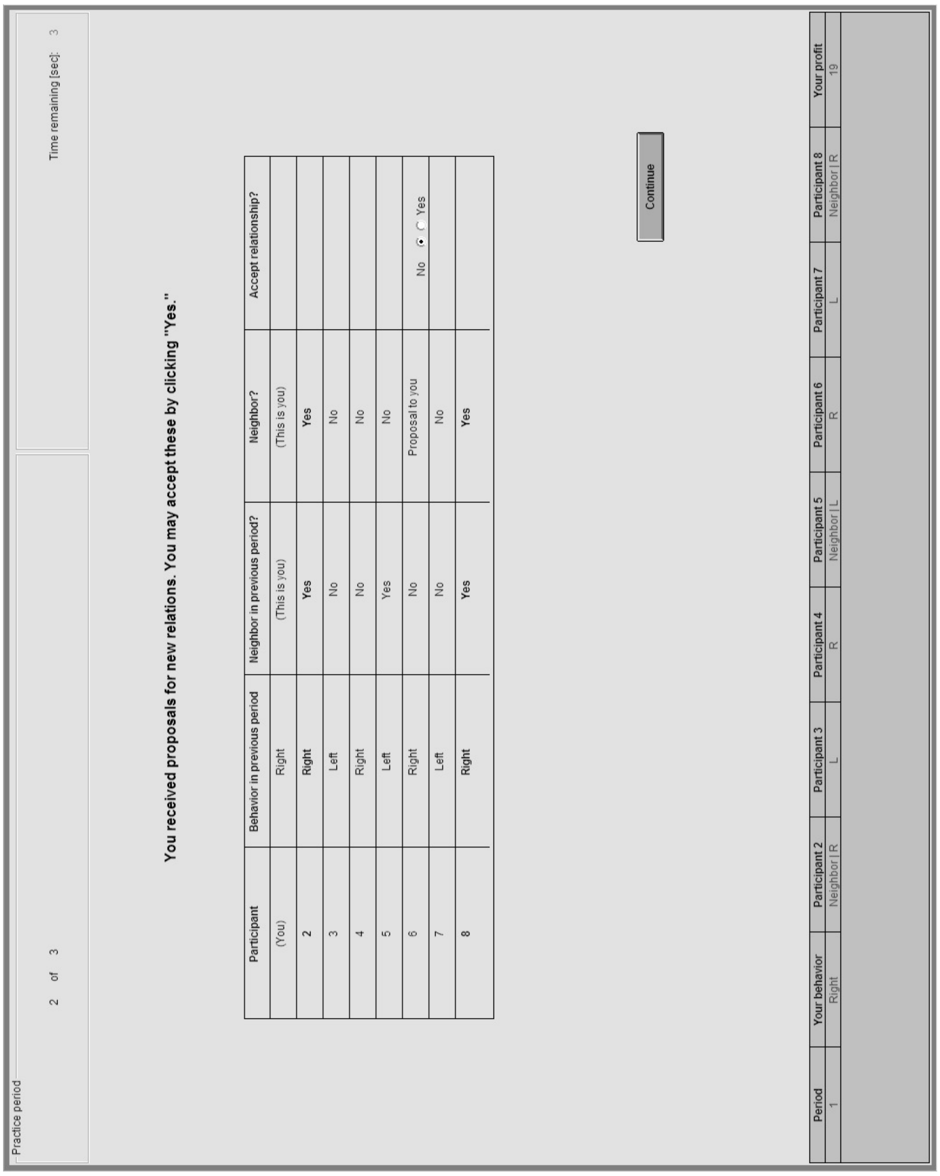

Screen 6: Incoming proposals

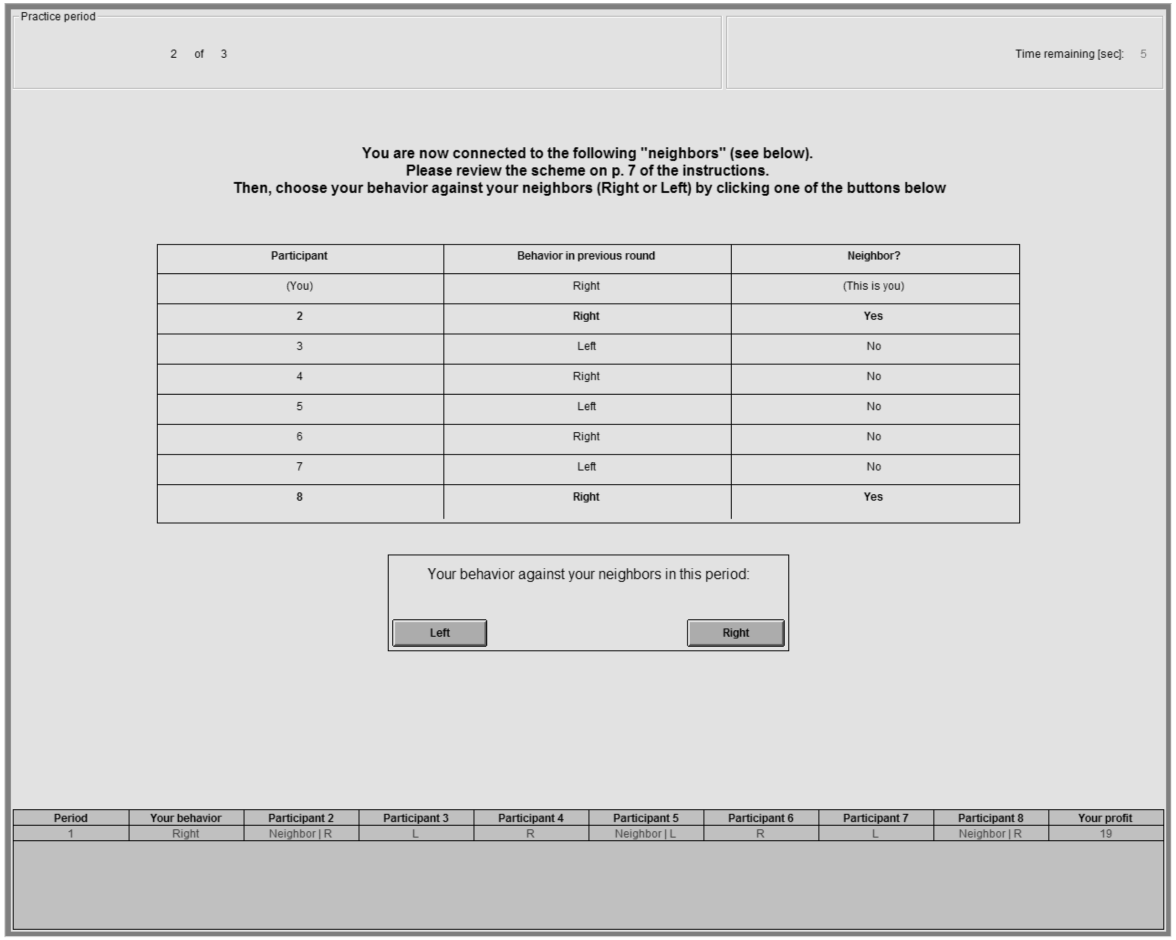

Screen 7: Behavior choice

References

Abrahamson, Eric and Lori Rosenkopf. 1997. "Social Network Effects on the Extent of Innovation Diffusion: A Computer Simulation." *Organization Science* 8:289–309.

Agresti, Alan and Barbara F. Agresti. 1978. "Statistical Analysis of Qualitative Variation." *Sociological Methodology* 9:204–237.

Albert, Réka and Albert-László Barabási. 1999. "Emergence of Scaling in Random Networks." *Science* 286:509–512.

Axelrod, Robert. 1984. *The Evolution of Cooperation*. New York: Basic Books.

Axelrod, Robert. 1997. "The Dissemination of Culture; a Model with Local Convergence and Global Polarization." *Journal of Conflict Resolution* 41:203–226.

Baerveldt, Chris and Tom A. B. Snijders. 1994. "Influences on and from the Segmentation of Networks: Hypotheses and Tests." *Social Networks* 16:213–232.

Bala, Venkatesh and Sanjeev Goyal. 2000. "A Noncooperative Model of Network Formation." *Econometrica* 68:1181–1229.

Barabási, Albert-László and Réka Albert. 1999. "The Emegence of Scaling in Random Networks." *Science* 286:509–512.

Bauman, Karl E. and Susan T. Ennett. 1996. "On the Importance of Peer Influence for Adolescent Drug Use: Commonly Neglected Considerations." *Addiction* 91:185–198.

Bearman, Peter S., James Moody, and Katherine Stovel. 2004. "Chains of Affection: The Structure of Adolescent Romantic and Sexual Networks." *American Journal of Sociology* 110:44–91.

Berninghaus, Siegfried, Karl-Martin Ehrhardt, and Claudia Keser. 2002. "Conventions and Local Interaction Structures: Experimental Evidence." *Games and Economic Behavior* 39:177–205.

Berninghaus, Siegfried and Ulrich Schwalbe. 1996. "Conventions, Local Interaction, and Automata Networks." *Journal of Evolutionary Economics* 6:231–324.

Berninghaus, Siegfried and Bodo Vogt. 2006. "Network Formation in Symmetric 2×2 Games." *Homo Oeconomicus* 23:421–466.

Berninghaus, Siegfried K. and Karl-Martin Ehrhardt. 1998. "Time Horizon and Equilibrium Selection in Tacit Coordination Games: Experimental Results." *Journal of Economic Behavior and Organisation* 37:231–248.

Berninghaus, Siegfried K., Karl-Martin Ehrhart, and Marion Ott. 2008. "Myopically Forward-Looking Agents in a Network Formation Game: Theory and Experimental Evidence." Working Paper, University of Karlsruhe.

Biely, Christoly, Klaus Dragostis, and Stefan Thurner. 2007. "The Prisoner's Dilemma on Co-evolving Networks under Perfect Rationality." *Physica D* 228:40–48.

Binmore, Ken. 2007. *Playing for Real: A Text on Game Theory*. Oxford: Oxford University Press.

Blume, Lawrence E. 1993. "The Statistical Mechanics of Strategic Interaction." *Games and Economic Behavior* 5:387–424.

Bojanowski, Michał and Vincent Buskens. 2008. "Coordination in Dynamic Social Networks under Heterogeneity." Working Paper, Utrecht University.

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

Bot, Sander M., Rutger C. M. E. Engels, Ronald A. Knibbe, and Wim H. J. Meeus. 2005. "Friend's Drinking Behaviour and Adolescent Alcohol Consumption: The Moderating Role of Friendship Characteristics." *Addictive Behaviors* 30:929–947.

Boudon, Raymond. 1981. *The Logic of Social Action. An Introduction to Sociological Analysis*. London: Routlege.

Bramoullé, Yann, Dunia López-Pintado, Sanjeev Goyal, and Fernando Vega-Redondo. 2004. "Network Formation and Anti-Coordination Games." *International Journal of Game Theory* 33:1–19.

Buchanan, Mark. 2002. *Nexus: Small Worlds and the Groundbreaking Science of Networks*. New York: W.W. Norton & Company.

Burgess, Robert L. and Don Bushell (eds.). 1969. *Behavioral Sociology: The Experimental Analysis of Social Process*. New York: Columbia University Press.

Burt, Ronald S. 1992. *Structural Holes: The Social Structure of Competition*. Cambridge, MA: Harvard University Press.

Burt, Ronald S. 2001. "Bandwidth and Echo: Trust, Information, and Gossip in Social Networks." In *Networks and Markets: Contributions from Economics and Sociology*, edited by Alessandra Casella and James E. Rauch, pp. 30–74. New York: Russell Sage Foundation.

Burt, Ronald S. 2005. *Brokerage and Closure: An Introduction to Social Capital*. New York: Oxford University Press.

Burt, Ronald S. and Marc Knez. 1995. "Kinds of Third-Party Effects on Trust." *Rationality and Society* 7:255–292.

Burt, Ronald S. and Marc Knez. 1996. "Trust and Third-Party Gossip." In *Trust in Organizations: Frontiers of Theory and Research*, edited by Roderick M. Kramer and Tom R. Tyler, pp. 68–89. Thousand Oaks, CA: Sage.

Buskens, Vincent. 2002. *Social Networks and Trust*. Boston, MA: Kluwer Academic Publishers.

Buskens, Vincent. 2003. "Trust in Triads: Effects of Exit, Control, and Learning." *Games and Economic Behavior* 42:235–252.

Buskens, Vincent, Rense Corten, and Jeroen Weesie. 2008. "Consent or Conflict: Coevolution of Coordination and Networks." *Journal of Peace Research* 45:205–222.

Buskens, Vincent and Werner Raub. 2002. "Embedded Trust: Control and Learning." In *Group Cohesion, Trust and Solidarity*, edited by Shane R. Thye and Edward J. Lawler, volume 19 of *Advances in Group Processes*, pp. 167–202. Amsterdam: Elsevier Science.

Buskens, Vincent and Arnout van de Rijt. 2008. "Dynamics of Networks If Everyone Strives for Structural Holes." *American Journal of Sociology* 114:371–407.

Buskens, Vincent and Chris Snijders. 2008. "Effects of Network Characteristics on Reaching the Payoff-Dominant Equilibrium in Coordination Games: A Simulation Study." ISCORE paper 232, Utrecht University.

Camerer, Colin F. 2003. *Behavioral Game Theory: Experiments in Strategic Interaction*. Princeton, NJ: Princeton University Press.

Camerer, Colin F. and Teck-Hua Ho. 1999. "Experience-Weighted Attraction Learning in Games: Estimates from Weak-Link Games." In *Games and Human Behavior*, edited by David W. Budescu, Ido Erev, and Rami Zwick, pp. 31–51. Mahwah, NJ: Lawrence Erlbaum.

Cassar, Allesandra. 2007. "Coordination and Cooperation in Local, Random and Small World Networks: Experimental Evidence." *Games and Economic Behavior* 58:209–230.

Cederman, Lars-Erik. 2005. "Computational Models of Social Forms: Advancing Generative Process Theory." *American Journal of Sociology* 110:864–893.

Centola, Damon and Michael W Macy. 2007. "Complex Contagions and the Weakness of Long Ties." *American Journal of Sociology* 113:702–734.

Chwe, Michael Suk-Young. 2001. *Rational Ritual; Culture, Coordinaton, and Common Knowledge*. Princeton, NJ: Princeton University Press.

Cohen, Jere M. 1977. "Sources of Peer Group Homogeneity." *Sociology of Education* 50:227–241.

Cohen, Michael D., Rick L. Riolo, and Robert Axelrod. 2001. "The Role of Social Structure in the Maintenance of Cooperative Regimes." *Rationality and Society* 13:5–32.

Coleman, James S. 1961. *The Adolescent Society*. New York: The Free Press.

Coleman, James S. 1964. *Introduction to Mathematical Sociology*. New York: Free Press.

Coleman, James S. 1984. "Introducing Social Structure into Economic Analysis." *American Economic Review* 74:84–88.

Coleman, James S. 1987. "Microfoundations and Macrobehavior." In *The Micro-Macro Link*, edited by Jeffrey C. Alexander, Bernhard Giesen, Richard Münch, and Neil J. Smelser, pp. 155–173. Berkeley, CA: University of California Press.

Coleman, James S. 1988. "Social Capital in the Creation of Human Capital." *American Journal of Sociology* 94:95–120.

Coleman, James S. 1990. *Foundations of Social Theory*. Cambridge, MA: Belknap.

Coleman, James S., Elihu Katz, and Herbert Menzel. 1957. "The Diffusion of an Innovation among Physicians." *Sociometry* 20:253–270.

Conlisk, John. 1996. "Why Bounded Rationality?" *Journal of Economic Literature* 34:669–700.

Cooper, Russel, Douglas V. DeJong, Robert Forsythe, and Thomas W. Ross. 1990. "Selection Criteria in Coordination Games: Some Experimental Results." *American Economic Review* 80:218–233.

Corbae, Dean and John Duffy. 2008. "Experiments with Network Formation." *Games and Economic Behavior* 64:81–120.

Corten, Rense and Vincent Buskens. 2008. "Co-evolution of Conventions and Networks: An Experimental Study." ISCORE Paper 252, Utrecht University.

Corten, Rense and Karen S. Cook. 2009. "Cooperation and Reputation in Dynamic Networks." In *Proceedings of the First International Conference on Reputation: Theory and Technology - ICORE 09*, edited by Mario Paolucci, pp. 20–34. Gargonza, Italy.

Crawford, Vincent P. 1997. "Theory and Experiment in the Analysis of Strategic Interaction." In *Advances in Economics and Econometrics: Theory and Applications*, edited by David M. Kreps and Kenneth F. Wallis, volume I, pp. 206–242. Cambridge: Cambridge University Press.

Davies, Mark and Denise B. Kandel. 1981. "Parental and Peer Influences on Adolescents' Educational Plans: Some Further Evidence." *American Journal of Sociology* 87:363–387.

Dawes, Robyn N. 1980. "Social Dilemmas." *Annual Review of Psychology* 31:169–93.

Diekmann, Andreas. 2008. "Soziologie und Ökonomie: Der Beitrag experimenteller Wirtschaftsforschung zur Sozialtheorie." *Kölner Zeitschrift für Soziologie und Sozialpsychologie* 3:528–550.

Dogan, Gönül, Marcel A. L. M. van Assen, Arnout van de Rijt, and Vincent Buskens. 2009. "The Stability of Exchange Networks." *Social Networks* Forthcoming.

Doreian, Patrick and Frans N. Stokman. 1997. "The Dynamics and Evolution of Social Networks." In *Evolution of Social Networks*, edited by Patrick Doreian and Frans N. Stokman, pp. 1–17. Amsterdam: Gordon and Breach Publishers.

Downing, Leslie. 1975. "The Prisoner's Dilemma Game as a Problem-Solving Phenomenon." *Simulation & Games* 6:366–391.

Duesenberry, James. 1960. "An Economic Analysis of Fertility: Comment." In *Demographic and Economic Change in Developed Countries*, volume 11 of *Universities-National Bureau Conference Series*, pp. 231–234. Princeton, NJ: Princeton Univesity Press.

Durkheim, Emile. [1897] 2002. *Suicide*. London: Routledge, 2nd edition.

Dutta, Bashkar and Matthew O. Jackson (eds.). 2003. *Networks and Groups. Models of Strategic Formation*. Heidelberg: Springer-Verlag.

EdK-Group. 2000. "Exit, Anonimity and the Chances of Egoistical Cooperation." *Analyse und Kritik* 22:114–129.

Eguíluz, Victor M., Martin G. Zimmerman, and Camilo J. Cela-Conde. 2005. "Cooperation and the Emergence of Role Differentiation in the Dynamics of Social Networks." *American Journal of Sociology* 110:977–1008.

Ehrhardt, George, Matteo Marsili, and Fernando Vega-Redondo. 2006. "Phenomenological Models of Socioeconomic Networks." *Physical Review E* 74.

Elias, Norbert. 1969. *The Civilizing Process, Vol.I. The History of Manners*. Oxford: Blackwell.

Ellickson, Robert C. 1991. *Order without Law. How Neighbors Settle Disputes*. Cambridge, MA: Harvard University Press.

Ellison, Glenn. 1993. "Learning, Local Interaction, and Coordination." *Econometrica* 61:1047–1071.

Elster, Jon. 1989. *Nuts and Bolts for the Social Sciences*. Cambridge: Cambridge University Press.

Ennett, Susan T. and Karl E. Bauman. 1993. "Peer Group Structure and Adolescent Cigarette Smoking: A Social Network Analysis." *Journal of Health and Social Behavior* 34:226–236.

Epstein, Joyce Levy and Nancy Karweit (eds.). 1983. *Friends in School; Patters of Selection and Influence in Secondary Schools*. New York: Academic Press.

Epstein, Joshua M. (ed.). 2007. *Generative Social Science: Studies in Agent-Based Computational Modeling*. Princeton, NJ: Princeton University Press.

Erev, Ido and Alvin E. Roth. 1998. "Predicting How People Play Games: Reinforcement Learning in Experimental Games with Unique, Mixed Strategy Equilibria." *American Economic Review* 88:848–881.

Erickson, Bonnie H. 1988. "The Relational Basis of Attitudes." In *Social Structures: A Network Approach*, edited by Barry Wellman and Stephen D. Berkowitz, pp. 99–121. New

York: Cambridge University Press.

Esteban, Joan and Debraj Ray. 1999. "Conflict and Distribution." *Journal of Economic Theory* 87:379–415.

Esteban, Joan and Gerald Scheider. 2008. "Polarization and Conflict: Theoretical and Empirical Issues." *Journal of Peace Research* 45:131–141.

Fararo, Thomas J. 1973. *Mathematical Sociology. An Introduction to Fundamentals*. New York: Wiley.

Fehr, Ernst and Klaus M. Schmidt. 1999. "A Theory of Fairness, Competition, and Cooperation." *Quarterly Journal of Economics* 114:817–868.

Fischbacher, Urs. 2007. "Z-Tree: Zurich Toolbox for Readymade Economic Experiments - Experimenter's Manual." *Journal of Experimental Economics* 10:171–178.

Fisher, Claude. 1982. *To Dwell among Friends: Personal Networks in Town and City*. Chicago: University of Chicago Press.

Flap, Henk. 2002. "No Man Is an Island: The Research Programme of a Social Capital Theory." In *Conventions and Structures in Economic Organization: Markets, Networks, and Hierarchies*, edited by Olivier Favereau and Emmanuel Lazega, pp. 29–59. Cheltenham: Edward Elgar.

Flap, Henk. 2004. "Creation and Returns of Social Capital: A New Research Program." In *Creation and Returns of Social Capital: A New Research Program*, edited by Henk Flap and Beate Völker, pp. 3–23. London: Routledge.

Freeman, Linton C. 1978. "Segregation in Social Networks." *Sociological Methods and Research* 6:411–429.

Freeman, Linton C. 1979. "Centrality in Social Networks: Conceptual Clarification." *Social Networks* 1:215–239.

Fudenberg, Drew and David K. Levine. 1998. *The Theory of Learning in Games*. Cambridge, MA: MIT Press.

Gambetta, Diego. 1994. "Godfather's Gossip." *Archives Européennes de Sociologie* 35:199–223.

Gilbert, Nigel and Klaus G. Troitzsch. 2005. *Simulation for the Social Scientist*. Maidenhead: Open University Press.

Giordano, Peggy C. 2003. "Relationships in Adolescence." *Annual Review of Sociology* 29:257–281.

Gould, Roger V. 1991. "Multiple Networks and Mobilization in the Paris Commune, 1871." *American Sociological Review* 56:716–729.

Gould, Roger V. 1993. "Collective Action and Network Structure." *American Sociological Review* 58:182–196.

Gould, Roger V. 1995. *Insurgent Identities: Class, Community, and Protest in Paris from 1848 to the Commune*. Chicago, IL: University of Chicago Press.

Gould, Roger V. 1999. "Collective Violence and Group Solidarity: Evidence from a Feuding Society." *American Sociological Review* 64:356–380.

Goyal, Sanjeev. 2007. *Connections: An Introduction to the Economics of Networks*. Princeton, NJ: Princeton University Press.

Goyal, Sanjeev and Fernando Vega-Redondo. 2005. "Network Formation and Social Coordination." *Games and Economic Behavior* 50:178–207.

Graham, John W., Gary Marks, and William B. Hansen. 1991. "Social Influence Processes Affecting Adolescent Substance Use." *Journal of Applied Psychology* 76:291–298.

Granovetter, Mark S. 1973. "The Strength of Weak Ties." *American Journal of Sociology* 78:1360–1380.

Granovetter, Mark S. 1974. *Getting a Job. A Study of Contacts and Careers*. Cambridge, MA: Harvard University Press.

Granovetter, Mark S. 1978. "Threshold Models of Collective Behavior." *American Journal of Sociology* 83:1420–1443.

Granovetter, Mark S. 1979. "The Theory-Gap in Social Network Analysis." In *Perspectives on Social Network Research*, edited by Paul W. Holland and Samuel Leinhardt, pp. 501–518. New York: Academic Press.

Granovetter, Mark S. 1985. "Economic Action and Social Structure: The Problem of Embeddedness." *American Journal of Sociology* 91:481–510.

Greif, Avner. 1989. "Reputation and Coalitions in Medieval Trade: Evidence on the Maghribi Traders." *Journal of Economic History* 49:857–882.

Greif, Avner. 1994. "Cultural Beliefs and the Organization of Society: A Historical and Theoretical Reflection on Collectivist and Individualist Societies." *The Journal of Political Economy* 102:912–950.

Greiner, Ben. 2004. "The Online Recruitment System Orsee 2.0 - a Guide for the Organization of Experiments in Economics." Working Paper Series in Economics, University of Cologne.

Halpin, Brendan. 1999. "Simulation in Sociology." *American Behavioral Scientist* 42:1488–1508.

Hardin, Russell. 1995. *One for All: The Logic of Group Conflict*. Princeton, NJ: Princeton University Press.

Hardin, Russell. 2007. *David Hume: Moral and Political Theorist*. Oxford: Oxford University Press.

Harsanyi, John C. 1977. *Rational Behavior and Bargaining Equilibrium in Games and Social*

Situations. Cambridge: Cambridge University Press.

Harsanyi, John C. and Reinhard Selten. 1988. *A General Theory of Equilibrium Selection in Games*. Cambridge, MA: MIT Press.

Hauk, Esther. 2001. "Leaving the Prison: Permitting Partner Choice and Refusal in Prisoner's Dilemma Games." *Computational Economics* 18:65–87.

Hawkins, J. David, Richard F. Catalano, and Janet Y. Miller. 1996. "Risk and Protective Factors for Alcohol and Other Drug Problems in Adolescence and Early Adulthood: Implications for Substance Abuse Prevention." *Psychological Bulletin* 112:64–105.

Heckathorn, Douglas D. 1996. "The Dynamics and Dilemmas of Collective Action." *American Sociological Review* 61:250–277.

Hedström, Peter. 2005. *Dissecting the Social: On the Principles of Analytical Sociology*. Cambridge: Cambridge University Press.

Hobbes, Thomas. [1651] 1988. *Leviathan*. New York: Prometheus Books.

Homans, George C. 1951. *The Human Group*. London: Routledge.

Hume, David. [1739-40] 1978. *A Treatise of Human Nature*. Oxford: Clarendon Press.

Jackson, Matthew O. 2008. *Social and Economic Networks*. Princeton, NJ: Princeton University Press.

Jackson, Matthew O. and Alison Watts. 2002. "On the Formation of Interaction Networks in Social Coordination Games." *Games and Economic Behavior* 41:265–291.

Jackson, Matthew O. and Asher Wolinsky. 1996. "A Strategic Model of Social and Economic Networks." *Journal of Economic Theory* 71:44–74.

Kagel, John H. and Alvin E. Roth (eds.). 1995. *The Handbook of Experimental Economics*. Princeton, NJ: Princeton University Press.

Kandel, Denise B. 1978. "Homophily, Selection, and Socialization in Adolescent Friendships." *American Journal of Sociology* 84:427–436.

Kandori, Michihiro, George J. Mailath, and Rafael Rob. 1993. "Learning, Mutation and Long Run Equilibria in Games." *Econometrica* 61:29–56.

Keser, Claudia, Karl-Martin Ehrhardt, and Siegfried Berninghaus. 1998. "Coordination and Local Interaction: Experimental Evidence." *Economic Letters* 58:269–275.

Kirke, Deirdre M. 2004. "Chain Reactions in Adolescents' Cigarette, Alcohol and Drug Use: Similarity through Peer Influence or the Patterning of Ties in Peer Networks?" *Social Networks* 26:3–28.

Klandermans, Bert. 1988. "Union Action and the Free-Rider Dilemma." In *Social Movenemts as a Factor of Change in the Contemporary World*, edited by Louis Kriesberg and Bronislaw Misztal, volume 10 of *Research in Social Movements, Conflicts and Change*, pp. 77–92. Greenwhich, CT: JAI Press.

Knecht, Andrea B. 2006. *The Dynamics of Networks and Behavior in Early Adolescence [2003/04]*. Utrecht: Utrecht University.

Knecht, Andrea B. 2008. *Friendship Selection and Friends' Influence: Dynamics of Networks and Actor Attributes in Early Adolescence*. ICS dissertation series. Utrecht: Utrecht University.

Knecht, Andrea B. and Thomas N. Friemel. 2008. "Praktikable vs. tatsächliche Grenzen von sozialen Netzwerken: Eine Diskussion zur Validität von Schulklassen als komplette Netzwerke." Working Paper, University of Erlangen-Nürnberg.

Kogut, Bruce. 1989. "The Stability of Joint Ventures: Reciprocity and Competitive Rivalry." *Journal of Industrial Economics* 38:183–198.

Kollock, Peter. 1998. "Social Dilemmas: The Anatomy of Cooperation." *Annual Review of Sociology* 24:183–214.

Kolm, Serge-Christophe and Jean Mercier Ythier (eds.). 2006. *Handbook of the Economics of Giving, Altrusim and Reciprocity*, volume 1 of *Handbooks in Economics*. Amsterdam: Elsevier.

Kosfeld, Michael. 1999. *Individual Decision-Making and Social Interaction*. Ph.D. thesis, Tilburg University.

Kosfeld, Michael. 2004. "Economic Networks in the Laboratory: A Survey." *Review of Network Economics* 3:20–41.

Kreps, David M. and R. Wilson. 1982. "Reputation and Imperfect Information." *Journal of Economic Theory* 27:253–279.

Kretzschmar, Mirjam and Jacco Wallinga. 2007. "Editorial: Networks in Epidemiology." *Mathematical Population Studies* 14:203–209.

Lazarsfeld, Paul F. and Robert K. Merton. 1954. "Friendship as as Social Process: A Substantive and Methodological Analysis." In *Freedom and Control in Modern Society*, edited by Morroe Berger, Theodore Abel, and Charles H. Page, pp. 18–66. Princeton, NJ: Van Nostrand.

Lewis, David K. 1969. *Convention*. Cambridge, MA: Harvard University Press.

Lewis, Kevin, Jason Kaufman, Marco Gonzalez, and Nicholas Christakis. 2008. "Tastes, Ties, and Time: A New Social Network Dataset Using Facebook.Com." *Social Networks* 30:330–342.

Lieberson, Stanley. 1969. "Measuring Population Diversity." *American Sociological Review* 34:850–862.

Light, John M. and Thomas J. Dishion. 2007. "Early Adolescent Antisocial Behavior and Peer Rejection: A Dynamic Test of a Developmental Process." *New Directions for Child* *and Adolescent Development* 118:77–89.

Lin, Nan. 2001. *Social Capital: A Theory of Social Structure and Action*. New York: Cambridge University Press.

Lippert, Steffen and Giancarlo Spagnolo. 2005. "Networks of Relations and Social Capital." SSE/EFI Working Paper Series in Economics and Finance, Stockholm School of Economics.

Lohmann, Susanne. 1994. "The Dynamics of Informational Cascades: The Monday Demonstrations in Leipzig, East Germany, 1989-1991." *World Politics* 47:42–101.

Macaulay, Steward. 1963. "Non-Contractual Relations in Business: A Preliminary Study." *American Sociological Review* 28:55–67.

Macy, Michael W. 1991. "Chains of Cooperation: Threshold Effects in Collective Action." *American Sociological Review* 56:730–747.

Macy, Michael W. and Andreas Flache. 2002. "Learning Dynamics in Social Dilemmas." *Proceedings of the National Academy of Sciences* 99:7229–7236.

Marsden, Peter V. and Noah E. Friedkin. 1993. "Network Studies of Social Influence." *Sociological Methods and Research* 22:127–151.

Marwell, Gerald and Pamela Oliver. 1993. *The Critical Mass in Collective Action*. Cambridge: Cambridge University Press.

McPherson, Miller, Lynn Smith-Lovin, and James M. Cook. 2001. "Birds of a Feather: Homophily in Social Networks." *Annual Review of Sociology* 27:415–444.

Merton, Robert K. 1936. "The Unanticipated Consequences of Purposive Social Action." *American Sociological Review* 1:894–904.

Merton, Robert K. 1968. *Social Theory and Social Structure*. New York: Free Press, enlarged edition.

Moffitt, Terrie E. 1993. "Adolescence-Limited and Life-Course-Persistent Antisocial Behavior: A Developmental Taxonomy." *Psychological Review* 100:674–701.

Mollenhorst, Gerald, Beate Völker, and Henk Flap. 2008. "Social Contexts and Core Discussion Networks: Using a Choice-Constraint Approach to Study Similarity in Intimate Relationships." *Social Forces* 56:937–965.

Montalvo, José G. and Marta Reynal-Querol. 2005. "Etnic Polarization, Potential Conflict, and Civil Wars." *American Economic Review* 95:796–816.

Morris, Martina, J. Zavisca, and L. Dean. 1995. "Social and Sexual Networks: Their Role in the Spread of HIV/AIDS among Young Gay Men." *AIDS education and prevention* 7:24–35.

Mueller, John H. and Karl F. Schuessler. 1961. *Statistical Reasoning in Sociology*. Boston, MA: Houghton Mifflin.

Newcomb, Michael D. and Peter M. Bentler. 1989. "Substance Use and Abuse among Children and Teenagers." *American Psychologist* 44:242–248.

Newman, Mark, Albert-László Barabási, and Duncan Watts (eds.). 2006. *The Structure and Dynamics of Networks*. Princeton, NJ: Princeton University Press.

Ohtsuki, Hisashi, Christoph Hauert, Erez Lieberman, and Martin A. Nowak. 2006. "A Simple Rule for the Evolution of Cooperation on Graphs and Social Networks." *Nature* 441:502–505.

Olson, Mancur. 1971. *The Logic of Collective Action*. Cambridge, MA: Harvard University Press, second edition.

Opp, Karl-Dieter and Christiane Gern. 1993. "Dissident Groups, Personal Networks, and Spontaneous Cooperation: The East German Revolution of 1989." *American Sociological Review* 58:659–680.

Parkhe, Arvind. 1993. "Strategic Alliance Structuring: A Game Theoretic and Transaction Cost Examination of Interfirm Cooperation." *Academy of Management Journal* 36:794–829.

Parsons, Talcott. 1937. *The Structure of Social Action. Volume I: Marshall, Pareto, Durkheim*. New York: Free Press.

Plott, Charles R. and Vernon L. Smith (eds.). 2008. *Handbook of Experimental Economics Results, Volume 1*. Amsterdam: North Holland.

Powell, Walter W. 1990. "Neither Market nor Hierarchy: Network Forms of Organization." *Research in Organizational Behavior* 12:295–336.

Powell, Walter W., Douglas R. White, Kenneth W. Koput, and Jason Owen-Smith. 2005. "Network Dynamics and Field Evolution: The Growth of Interorganizational Collaboration in the Life Sciences." *American Journal of Sociology* 110:1132–1205.

Pujol, Josep M., Andreas Flache, Jordi Delgado, and Ramon Sangüesa. 2005. "How Can Social Networks Ever Become Complex? Modelling the Emergence of Complex Networks from Local Social Exchanges." *Journal of Artificial Societies and Social Simulation* 8.

Rabin, Matthew. 1993. "Incorporating Fairness into Game Theory and Economics." *American Economic Review* 83:1281–1302.

Rapoport, Anatol. 1974. "Prisoner's Dilemma. Recollections and Observations." In *Game Theory as a Theory of Conflict Resolution*, edited by Anatol Rapoport, pp. 17–34. Dordrecht: Reidel.

Rapoport, Anatol, Andreas Diekmann, and Axel Franzen. 1995. "Experiments with Social Traps IV: Reputation Effects in the Evolution of Cooperation." *Rationality and Society* 7:431–441.

Rasmusen, Eric. 2001. *Games and Information*. Oxford: Wiley-Blackwell, 4th edition.

Raub, Werner and Thomas Voss. 1986. "Conditions for Cooperation in Problematic Social Situations." In *Paradoxal Effects of Social Behavior. Essays in Honor of Anatol Rapoport*, edited by Andreas Diekmann and Peter Mitter, pp. 85–103. Wien: Physica-Verlag.

Raub, Werner and Jeroen Weesie. 1990. "Reputation and Efficiency in Social Interactions: An Example of Network Effects." *American Journal of Sociology* 96:626–654.

Reynal-Querol, Marta. 2002. "Ethnicity, Political Systems and Civil War." *Journal of Conflict Resolution* 46:29–54.

Rijt, Arnout van de and Michael W. Macy. 2008. "The Problem of Social Order: Egoism or Autonomy?" Working Paper, Cornell University.

Rogers, Everett M. 1995. *Diffusion of Innovations*. New York: Free Press, fourth edition.

Rogers, William M. 1993. "Regression Standard Errors in Clustered Samples." *Stata Technical Bulletin* 13:19–23.

Rubinstein, Ariel. 1998. *Modeling Bounded Rationality*. Cambridge, MA: MIT Press.

Schelling, Thomas C. 1978. *Micromotives and Macrobehavior*. New York: Norton.

Schuessler, Rudolf. 1989. "Exit Threats and Cooperation under Anonymity." *Journal of Conflict Resolution* 33:728–749.

Scott, John. 2000. *Social Network Analysis: A Handbook*. London: Sage, second edition.

Scott, James C. 1990. *Domination and the Art of Resistance: Hidden Transcripts*. New Haven, CT: Yale University Press.

Sherif, Muzafer and Carolyn W. Sherif. 1964. *Reference Groups: Exploration into Conformity and Deviation of Adolescents*. New York: Harper and Row.

Simon, Herbert A. 1956. "Rational Choice and the Structure of the Environment." *Psychological Review* 63:129–138.

Skyrms, Brian and Robin Pemantle. 2000. "A Dynamic Model of Social Network Formation." *Proceedings of the National Academy of Sciences U.S.A.* 97:9340–9346.

Snijders, Chris and Richard Zijdeman. 2004. "Reputation and Internet Auctions: eBay and Beyond." *Analyse und Kritik* 26:158–184.

Snijders, Tom A. B. 1981. "The Degree Variance: An Index of Graph Heterogeneity." *Social Networks* 3:163–174.

Snijders, Tom A. B. 2001. "The Statistical Evaluation of Social Network Dynamics." In *Sociological Methodology*, edited by Michael E. Sobel and Mark P. Becker, volume 31, pp. 361–395. Boston, MA: Blackwell.

Snijders, Tom A. B. 2006. "Statistical Methods for Network Dynamics." In *Proceedings of the XLIII Scientific Meeting, Italian Statistical Society*, edited by S.R. Luchini, pp. 281–296. Padova: CLEUP.

Snijders, Tom A. B. and Roel J. Bosker. 1999. *Multilevel Analysis. An Introduction to Basic and Advanced Multilevel Modeling*. London: Sage.

Snijders, Tom A. B., Christian Steglich, and Michael Schweinberger. 2007. "Modelling the Co-evolution of Networks and Behavior." In *Longitudinal Models in the Behavioral and Related Sciences*, edited by Kees v. Montfort, Han Oud, and Albert Sattora, pp. 41–71. Mahwah, NJ: Lawrence Erlbaum.

Stanley, E. Anne, Dan Ashlock, and Leigh Tesfatsion. 1994. "Iterated Prisoner's Dilemma with Choice and Refusal of Partners." In *Artificial Life III*, edited by Christopher G. Langton, volume XVII of *Santa Fe Institute Studies in the Sciences of Complexity*. Reading, MA: Addison-Wesley.

StataCorp. 2005. "Stata Statistical Software: Release 9."

Steele, Claude M. and Robert A. Josephs. 1990. "Alcohol Myopia: Its Prized and Dangerous Effects." *American Psychologist* 45:921–933.

Steglich, Christian, Tom A. B. Snijders, and Patrick West. 2006. "Applying Siena: An Illustrative Analysis of the Co-evolution of Adolescents' Friendship Networks, Taste in Music, and Alcohol Consumption." *Methodology* 2:48–56.

Steinberg, Laurence and Amanda Sheffield Morris. 2001. "Adolescent Development." *Annual Review of Psychology* 52:83–110.

Stokman, Frans N. and Evelien P. H. Zeggelink. 1996. "Is Politics Power or Policy Orientated? A Comperative Analysis of Dynamic Access Models in Policy Networks." *Journal of Mathematical Sociology* 21:77–111.

Straub, Paul S. 1995. "Risk Dominance and Coordination Failure in Static Games." *The Quarterly Review of Economics and Finance* 35:339–363.

Swadi, Harith. 1999. "Individual Risk Factors for Adolescent Substance Use." *Drug and Alcohol Dependence* 55:209–224.

Takács, Károly. 2001. "Structural Embeddedness and Group Conflict." *Journal of Conflict Resolution* 45:743–769.

Takács, Károly, Béla Janky, and Andreas Flache. 2008. "Collective Action and Network Change." *Social Networks* 30:177–189.

Taylor, Michael. 1976. *Anarchy and Cooperation*. London: Wiley.

Taylor, Michael. 1987. *The Possibility of Cooperation*. Cambridge: Cambridge University Press.

Ule, Aljaž. 2005. *Exclusion and Cooperation in Networks*. Amsterdam: Thela Thesis.

Ullmann-Margalit, Edna. 1977. *The Emergence of Norms*. Oxford: Clarendon Press.

Uzzi, Brian. 1996. "The Sources and Consequences of Embeddedness for the Economic Performance of Organizations: The Network Effect." *American Sociological Review* 61:674–

698.

Uzzi, Brian. 1997. "Social Structure and Competition in Interfirm Networks: The Paradox of Embeddedness." *Administrative Science Quarterly* 42:35–67.

Valente, Tom. 1996. "Social Network Thresholds in the Diffusion of Innovations." *Social Networks* 18:69–89.

Valente, Tom. 2005. "Network Models and Methods for Studying the Diffusion of Innovations." In *Models and Methods in Social Network Analysis*, edited by Peter J. Carrington, John Scott, and Stanley Wasserman, pp. 98–116. New York: Cambridge University Press.

Van Huyck, John B., Raymond C. Battalio, and Richard O. Beil. 1990. "Tacit Coordination Games, Strategic Uncertainty and Coordination Failure." *American Economic Review* 80:234–248.

Vanberg, Viktor J. and Roger D. Congleton. 1992. "Rationality, Morality, and Exit." *American Political Science Review* 86:418–431.

Vega-Redondo, Fernando. 2006. "Building up Social Capital in a Changing World." *Journal of Economic Dynamics and Control* 30:2305–2338.

Vega-Redondo, Fernando. 2007. *Complex Social Networks*. Econometric Society Monographs. New York: Cambridge University Press.

Verbrugge, Lois. 1977. "The Structure of Adult Friendship Choices." *Social Forces* 56:576–597.

Voss, Thomas. 2001. "Game-Theoretical Perspectives on the Emergence of Social Norms." In *Social Norms*, edited by Michael Hechter and Karl-Dieter Opp, pp. 105–136. New York: Russell Sage Foundation.

Wasserman, Stanley and Katherine Faust. 1994. *Social Network Analysis: Methods and Applications*. Cambridge: Cambridge University Press.

Watts, Duncan J. 1999. "Networks, Dynamics, and the Small-World Phenomenon." *American Journal of Sociology* 105:493–527.

Watts, Duncan J. 2002. "A Simple Model of Global Cascades on Random Networks." *Proceedings of the National Academy of Sciences of the United States of America* 99:5766–5771.

Watts, Duncan J. and Steven H. Strogatz. 1998. "Collective Dynamics of 'Small World' Networks." *Nature* 393:440–442.

Weber, Max. [1921] 1976. *Wirtschaft Und Gesellschaft*. Tübingen: Mohr, fifth edition.

Weesie, Jeroen. 1996. "Disciplining via Exit and Voice." ISCORE Paper, Utrecht University.

Weesie, Jeroen and Henk Flap (eds.). 1990. *Social Networks through Time*. Utrecht: ISOR.

Willer, David. 1999. *Network Exchange Theory*. Westport, CT: Praeger.

Wilson, Roger. 1985. "Reputations in Games and Markets." In *Game-Theoretic Models of Bargaining*, edited by Alvin E. Roth, pp. 27–62. Cambridge: Cambridge University Press.

Wippler, Reinhard and Siegwart Lindenberg. 1987. "Collective Phenomena and Rational Choice." In *The Micro-Macro Link*, edited by Jeffrey C. Alexander, Bernhard Giesen, Richard Münch, and Neil J. Smelser, pp. 135–151. Berkeley, CA: University of California Press.

Yamagishi, Toshio, Nahoko Hayashi, and Nobuhito Jin. 1994. "Prisoner's Dilemma Networks: Selection Strategy versus Action Strategy." In *Social Dilemmas and Cooperation*, edited by Ulrich Schulz, Wulf Albers, and Ulrich Müller, pp. 233–250. Berlin: Springer-Verlag.

Young, H. Peyton. 1993. "The Evolution of Conventions." *Econometrica* 61:57–84.

Young, H. Peyton. 1998. *Individual Strategy and Social Structure. An Evolutionary Theory of Institutions*. Princeton, NJ: Princeton University Press.

Zeggelink, Evelien P. H., Frans N. Stokman, and Gerhard G. van de Bunt. 1996. "The Emergence of Groups in the Evolution of Friendship Networks." *Journal of Mathematical Sociology* 21:29–55.

Author index

SUMMARY IN DUTCH

Samenvatting: Co-evolutie van Sociale Netwerken en Gedrag in Sociale Dilemma's

Inleiding

Sociale netwerken spelen een belangrijke rol in de verklaring van sociale verschijnselen, zoals bijvoorbeeld de verspreiding van ideeën of gedrag en het ontstaan van samenwerking en vertrouwen. In veel van het voorgaande onderzoek naar deze effecten worden sociale netwerken beschouwd als vaststaande structuren die 'van buitenaf' bepaald worden. De aandacht gaat dan met name uit naar de gevolgen van deze structuren op individueel gedrag. Ook sociale netwerken zelf zijn echter sociale verschijnselen en gegeven het belang van sociale netwerken in sociologische verklaringen ligt de vraag voor de hand hoe netwerken zelf ontstaan. Deze vraag heeft aanzienlijk minder aandacht gekregen.

In veel gevallen zijn sociale netwerken het gevolg van individuele keuzes van relaties, en het ligt voor de hand dat deze keuzes tenminste ten dele samenhangen met hetzelfde gedrag dat ook door deze sociale relaties beïnvloed wordt. Dit suggereert dat individueel gedrag en sociale netwerken zich in onderlinge afhankelijkheid ontwikkelen, ofwel *co-evolueren.*

Sociale dilemma's en sociale netwerken

Dit proefschrift richt zich op de co-evolutie van sociale netwerken en gedrag in een specifiek soort situaties, namelijk in *sociale dilemma's*. Een sociaal dilemma is een situatie waarin individueel rationeel gedrag kan leiden tot een uitkomst die sub-optimaal is voor het collectief. Een voorbeeld van een sociaal dilemma is het Gevangenendilemma, waarin twee actoren kunnen profiteren van samenwerking, maar tegelijkertijd verleid worden om misbruik van elkaar te maken. De standaard speltheoretische voorspelling is dat er geen samenwerking zal plaatsvinden.

Een andere, minder bekende, maar daarom niet minder belangrijke klasse van sociale dilemma's wordt gevormd door *coördinatieproblemen*. Hierin zijn actoren beter af wanneer zij *hetzelfde* gedrag kiezen. Dit spel staat model voor het ontstaan van *conventies*. Voorbeelden van alledaagse conventies zijn verkeersregels (links of rechts rijden), technologische standaarden, en etiquette. In al deze voorbeelden is het belangrijker dat we het *hetzelfde* kiezen als anderen dan *wat* we kiezen, hoewel in veel gevallen de ene conventie gunstiger is voor de betrokkenen dan de andere. Toch is het niet vanzelfsprekend dat de meest gunstige conventie altijd gekozen wordt. Een conventie kan weliswaar een hoger nut opleveren, maar kan ook meer risicovol zijn in de zin dat een afwijking van de ene conventie schadelijker is dan afwijking van de andere conventie. Een belangrijke vraag in het onderzoek naar coördinatieproblemen is hoe collectief gunstige conventies tot stand kunnen komen, óók wanneer deze meer risicovol zijn.

Sociale dilemma's spelen een belangrijke rol in sociaal-wetenschappelijk onderzoek naar sociale interacties, met name naar de vraag onder welke omstandigheden sociale orde kan ontstaan uit individueel doelgericht handelen. Deze vraag staat ook bekend als het *ordeprobleem*. Uit eerder onderzoek blijkt dat sociale netwerken een belangrijke rol kunnen spelen bij het verklaren van uitkomsten van sociale dilemma's. Wanneer samenwerkingsproblemen zoals het Gevangenendilemma in een hecht netwerk van sociale relaties ingebed zijn, kan samenwerking bevorderd worden door reputatie-effecten. Bij coördinatieproblemen kunnen sociale netwerken beïnvloeden hoe conventies zich verspreiden.

Dynamische netwerken

Een gangbare aanname in de theorievorming over en onderzoek naar netwerkeffecten op uitkomsten van sociale dilemma's is dat sociale netwerken vaststaande, veelal 'van buitenaf' bepaalde structuren zijn. Inmiddels is echter duidelijk dat ook sociale netwerken het resultaat zijn van individuele keuzes van relaties. Gegeven dat de structuur van een netwerk de uitkomst van sociale dilemma's mede bepaalt, ligt het voor de hand dat actoren zullen proberen deze structuren in hun voordeel te beïnvloeden. Daarnaast is het aannemelijk dat de keuze van relaties samenhangt met de uitkomsten van sociale dilemma's; zo is te verwachten dat actoren in samenwerkingsrelaties partners zullen vermijden waarmee ze in het verleden slechte ervaringen hadden, en dat actoren in coördinatieproblemen actoren die kiezen voor een andere conventie uit de weg zullen gaan.

Het is echter *niet* vanzelfsprekend dat de netwerkstructuren die gunstige uitkomsten van sociale dilemma's bevorden, ook als vanzelf zullen ontstaan als gevolg van individuele relatiekeuzes. Individuen hebben weliswaar (tot op zekere hoogte) controle over hun eigen sociale relaties, maar kunnen niet de structuur van het netwerk als geheel bepalen: deze grotere structuur is het gevolg van de *gecombineerde* keuzes van de individuen in het netwerk. De relatiekeuzes van de ene actor beïnvloeden indirect de relaties van andere actoren; zo kan het verbreken van een enkele relatie tussen twee actoren vele indirecte verbindingen tussen derden verbreken. In die zin kunnen sociale netwerkstructuren beschouwd worden als onbedoelde collectieve gevolgen van individueel gedrag. Dit proefschrift richt zich op de vraag hoe sociale netwerken en gedrag in sociale dilemma's co-evolueren, en onder welke omstandigheden optimale of juist suboptimale situaties ontstaan.

Deze vragen worden onderzocht in vier onafhankelijke studies, waarin zowel aandacht is voor theorievorming als voor empirische toetsing van hypothesen. De theoretische benadering is gestoeld op de aanname dat zowel uitkomsten van sociale dilemma's als de vorming van netwerken voortkomen uit doelgericht gedrag door individuen. Omdat individuele keuzes hierin onderling afhankelijk zijn, speelt speltheorie een belangrijke rol in de theorievorming. Daarnaast wordt gebruik gemaakt van computersimulatie als hulpmiddel om op systematische wijze hypothesen af te leiden over complexe processen. cryptisch Om hypothesen te toetsen wordt gebruik gemaakt van verschillende onderzoeksstrategieën: laboratoriumexperimenten en vragenlijstonderzoek. Experimenten hebben als voordeel dat de onderzoeker een grote mate van controle heeft over de omstandigheden van het onderzoek en daardoor beter in staat is causale verbanden te leggen dan in niet-experimenteel empirisch onderzoek. Daarnaast bieden experimenten de mogelijkheid individueel gedrag in complexe co-evolutieprocessen in detail te observeren. Daar tegenover staat dat laboratoriumexperimenten altijd tot op zekere hoogte kunstmatig blijven. De pretentie van sociaal-wetenschappelijk onderzoek is om daadwerkelijke sociale verschijnselen te verklaren; vragenlijstonderzoek is een methode om te onderzoeken in hoeverre dat gelukt is.

Theoretische en empirische bevindingen

Coördinatie in dynamische netwerken: theorie

Hoofdstuk 2 richt zich op coördinatieproblemen in dynamische netwerken. De aandacht beperkt zich daarbij tot coördinatieproblemen met twee evenwichten ('conventies'), waarvan het ene een hogere opbrengst oplevert voor de betrokken actoren, terwijl het andere minder risico met zich meebrengt. Een belangrijke aanname is dat actoren niet alleen hun gedrag in coördinatieproblemen kunnen kiezen, maar ook *met wie* ze interacties willen hebben. De centrale vragen in dit hoofdstuk zijn:

- Onder welke omstandigheden is het waarschijnlijker dat een efficiënte conventie tot stand komt?
- Onder welke omstandigheden is het waarschijnlijker dat verschillende conventies naast elkaar voortbestaan?

Om deze vragen te beantwoorden wordt een speltheoretisch model opgesteld waarbij actoren in een populatie zowel hun gedrag in een coördinatiespel kunnen kiezen als hun interactiepartners in dit spel. We nemen bovendien aan dat sociale relaties niet alleen opbrengsten, maar ook kosten hebben, zodat er een maximum is aan het aantal interactiepartners dat een actor kan hebben. Verder nemen we aan dat een relatie alleen tot stand kan komen met instemming van beide betrokken actoren. Het coördinatiespel wordt vervolgens gespeeld in ronden. Deze opzet resulteert in een dynamisch proces, waarbij gedrag en het sociale netwerk van interacties co-evolueren. Dit proces is stabiel als geen van de actoren een reden heeft om zijn gedrag te veranderen of een relatie te verbreken, en als er geen paren van actoren zijn die een nieuwe relatie willen beginnen.

De eerste stap in de analyse van dit model bestaat uit het formeel beschrijven van de stabiele toestanden van het co-evolutieproces. Deze formele resultaten laten zien dat vele constellaties van gedrag en netwerken stabiel kunnen zijn, waaronder toestanden waarin er één conventie bestaat en alle actoren met elkaar verbonden zijn, toestanden met één conventie maar meerdere niet met elkaar verbonden componenten in het netwerk, en toestanden waarin het netwerk uiteen valt in meerdere componenten waarin verschillende conventies bestaan.

Om uitspraken te kunnen doen over welke van deze toestanden meer waarschijnlijk zijn als uitkomsten van een dynamisch proces, gebruiken we computersimulaties. We laten het proces lopen vanuit een groot aantal begintoestanden, waarbij we zowel de beginverdeling van gedrag als het beginnetwerk variëren. Door middel van statistische analyse van de gesimuleerde data kan vervolgens onderzocht worden hoe het ontstaan van specifieke stabiele toestanden afhangt van specifieke begintoestanden. De resultaten laten zien dat er een belangrijke interactie bestaat tussen de beginverdeling van gedrag en kenmerken van de aanvankelijke netwerkstructuur. Als het beginnetwerk hechter is, heeft de beginverdeling van gedrag een grotere invloed op de eindverdeling van gedrag. Centralisatie van het beginnetwerk heeft het tegenovergestelde effect.

Coördinatie in dynamische netwerken: een experiment

In Hoofdstuk 3 worden de implicaties van het model uit Hoofdstuk 2 aan een eerste empirische toets onderworpen door middel van een laboratoriumexperiment. Daarnaast worden er nieuwe hypothesen afgeleid over het effect van beperkte beschikbaarheid van informatie op het ontstaan van conventies. We vergelijken twee verschillende scenario's voor de beschikbaarheid van informatie. In het *globale informatie*-scenario hebben actoren de beschikking over informatie over het gedrag van alle andere actoren in de populatie. In het *locale informatie*-scenario hebben actoren alleen informatie over het gedrag van hun directe 'buren' (interactiepartners) in het netwerk. Om te voorspellen hoe actoren nieuwe relaties kiezen met locale informatie, nemen we aan dat zij het gemiddelde gedrag van hun huidige buren gebruiken als een benadering van het gedrag van onbekende actoren in de populatie.

Het belangrijkste theoretische resultaat is dat het beperken van informatie het *minder* waarschijnlijk maakt dat de populatie uiteen zal vallen in subgroepen waarin verschillende conventies gesteund worden. De reden is dat beperkte informatie het moeilijker maakt contact te vermijden met actoren die een andere conventie steunen. Om gerichte voorspellingen af te leiden voor het experiment gebruiken we computersimulaties die de experimentele condities precies nabootsen wat betreft de opbrengsten, de beginnetwerken en de beschikbaarheid van informatie.

De resultaten van het experiment bieden beperkte steun voor de hypothesen over het ontstaan van conventies op het macroniveau. De voorspelling dat inefficiënte conventies vaker ontstaan hechtere netwerken wordt bevestigd, maar we vinden geen significante resultaten over interactie-effecten tussen het beginnetwerk en de beginverdeling van gedrag zoals voorspeld in Hoofdstuk 2. Wel vinden we ondersteuning voor de voorspellingen van *individueel* gedrag van proefpersonen. Dat wil zeggen: we vinden duidelijke aanwijzingen dat proefpersonen hun gedrag vooral op korte termijn aanpassen aan dat van hun directe buren. De manier waarop proefpersonen nieuwe interactiepartners kiezen met locale informatie komt echter niet overeen met de voorspelling. De resultaten suggereren eerder dat proefpersonen aannemen dat onbekende actoren *verschillen* van hun huidige buren.

Alcoholgebruik onder adolescenten als een coördinatieprobleem

Ook in Hoofdstuk 4 worden hypothesen over coördinatie in dynamische netwerken getoetst, maar nu in een natuurlijke omgeving in plaats van in het laboratorium. We analyseren alcoholgebruik onder adolescenten als een coördinatieprobleem, onder de aanname dat alcoholgebruik gezien kan worden als inefficiënt maar risicodominant gedrag in een coördinatieprobleem, gegeven dat adolescenten redenen hebben hun gedrag in overeenstemming te brengen met dat van hun vrienden. We toetsen voorspellingen afgeleid uit het model uit de eerdere hoofdstukken met data over alcoholgebruik en vriendschappen onder scholieren van Nederlandse middelbare scholen. Waar het meeste onderzoek over dit onderwerp beperkt blijft tot effecten van de *persoonlijke* netwerken van adolescenten, biedt onze theoretische aanpak de mogelijkheid voorspellingen te doen over effecten van de netwerkstructuur op *macroniveau* op alcoholgebruik.

De statistische analyses bevestigen het voorspelde 'katalyserende' effect van de aanvankelijke netwerkdichtheid op de ontwikkeling van alcoholgebruik in een klas: hoe dichter het beginnetwerk, hoe waarschijnlijker het is dat een kleine meerderheid van alcoholgebruikers zich zal uitbreiden. De empirische bevestiging van dit effect, dat we niet vonden in het experiment in Hoofdstuk 3, biedt verdere ondersteuning voor het theoretisch model. Het tegengestelde effect van centralisatie kan echter niet bevestigd worden.

Samenwerking in dynamische netwerken

In Hoofdstuk 5 wordt de aandacht verlegd naar *samenwerkingsproblemen*, hier geconceptualiseerd als twee-personen Gevangenendilemma's in een dynamisch netwerk. We ontwikkelen een model waarin actoren herhaalde Gevangenendilemma's spelen en over hun interactiepartners kunnen leren via derden. Met andere woorden, er is sprake van *reputatie-effecten*. Tegelijkertijd kunnen actoren kiezen *met wie* ze interacties hebben, waardoor een dynamisch netwerk ontstaat. Dit model bouwt voort op eerder onderzoek dat laat zien dat samenwerking bevorderd wordt door hechte netwerken, maar breidt dit uit door het netwerk niet langer als een exogene factor te beschouwen. Het model wordt zowel formeel als met behulp van computersimulatie geanalyseerd. De formele analyses laten zien dat in zowel statische als dynamische netwerken meer verschillende stabiele toestanden mogelijk worden door de aanwezigheid van reputatie-effecten, maar dat er niet noodzakelijkerwijs meer samenwerking ontstaat. Daarnaast blijkt dat reputatie-effecten in dynamische netwerken leiden tot meer homogene netwerkstructuren.

Simulaties van het model bevestigen de analytische resultaten, en benadrukken dat reputatie-effecten in hechte netwerken niet altijd tot meer samenwerking leiden. In statische netwerken zorgen reputatie-effecten ervoor dat stabiele toestanden met zowel zeer veel als zeer weinig samenwerking waarschijnlijker worden. In dynamische netwerken leiden reputatie-effecten daarnaast tot (gemiddeld) lagere niveaus van samenwerking. Tot slot laten de simulaties zien dat het niet waarschijnlijk is dat hechte netwerken spontaan ontstaan zonder een bestaande neiging tot samenwerking.

Suggesties voor vervolgonderzoek

Hoofdstuk 6 sluit het proefschrift af met suggesties voor vervolgonderzoek. Deze suggesties richten zich op zowel theoretische uitbreidingen als op verdere empirische toetsingen van de in dit proefschrift ontwikkelde modellen.

Een belangrijke theoretische aanbeveling is het onderzoeken van alternatieve modellen voor individueel gedrag. De modellen voor gedrag in coördinatie- en samenwerkingsproblemen verschillen op een aantal belangrijke punten, met name wat betreft aannames over de mate waarin actoren vooruit denken. Het is wenselijk meer consistente modellen te ontwikkelen, die in staat zijn gedrag in *beide* types sociale dilemma's te verklaren.

De empirische studies in dit proefschrift gaan uitsluitend over coördinatieproblemen. De aanbevelingen voor verder empirische onderzoek richten zich dan ook voornamelijk op toetsingen van het model voor samenwerkingsproblemen in dynamische netwerken uit Hoofdstuk 5. Een eerste stap zou kunnen bestaan uit experimenteel onderzoek, waarin de effecten van partnerkeuze en verspreiding van reputaties systematisch vergeleken worden. Gegevens over strategische allianties tussen bedrijven bieden mogelijkheden voor niet-experimenteel onderzoek naar samenwerking in dynamische netwerken.

Acknowledgments

In this dissertation, I study what happens when actors are influenced by their social network, but can also choose their relationships in this network. While writing the dissertation, I benefited much from my own social environment. Although I did not choose all the people in that environment, I don't think I could have made a better choice if given the chance. I take this opportunity to thank a number of these people.

First of all, I owe much to my daily supervisor Vincent Buskens. His all-round expertise and merciless eye for detail have been crucial for the overall quality of the dissertation. More in particular, Chapters 2 and 3 also benefited much from his skills as a programmer and an experimentalist. However, it was especially his cheerful and informal personal style that made working with him a great pleasure. Stephanie Rosenkranz, as my second supervisor, has improved the dissertation in many places by her sharp insights, and as a relative outsider from economics, made sure that the multidisciplinary perspective didn't get lost. My promotor Werner Raub always made sure that I used the right words in the right place, so that what I wrote was what I meant, even when I didn't know yet that that was what I meant. His strong analytical perspective has much improved the theoretical consistency of the dissertation. Moreover, he always took care that everything around my project was organized smoothly. Finally, I am grateful for his encouragement and confidence since I was an undergraduate, which strongly influenced my choice for an academic career.

I am grateful to my co-authors for their collaboration. Jeroen Weesie dropped out of my supervision team even before the project started, but has nevertheless contributed to the dissertation in direct (as a coauthor of Chapter 2) and indirect ways. He provided statistical advice on many of the analyses presented here, gave some crucial advice on Chapter 5, but was above all an inspiring teacher and colleague. The study presented in Chapter 4 would have been impossible without the collaboration of Andrea Knecht, who not only provided the data but also shared her expertise on the data and the substantive application. Karen Cook has been a great host during my stay at Stanford University, where I developed the early versions of Chapter 5, which she also coauthored.

I thank Michał Bojanowski, Richard Zijdeman, Bastian Westbrock, and Sytske Weidema for proof-reading various parts of the manuscript and providing highly useful comments on the text.

I thank Joris van der Veer and Dennie van Dolder for their assistance in running the experiments reported in Chapter 3.

I am indebted to many people who contributed to the writing of my dissertation in less direct ways.

It has been a great pleasure to work with my colleagues in the research group "Dynamics of cooperations, networks and institutions" (led by Vincent Buskens and Stephanie Rosenkranz): Bastian Westbrock, Ines Lindner, Michał Bojanowski, Kris De Jaegher, Jesse Kaijser, Dominik Morbitzer, and Britta Hoyer. The interdisciplinary character of this group has contributed tremendously to my academic development. At the same time, the friendly and informal atmosphere (fueled by the various dinners and borrels at Stephanie and Utz's place, and elsewhere) has made working in this group all the more enjoyable. I thank in particular Bastian and Michał, my brothers in arms as "HiPo-AIOs" almost from the beginning. I have very much enjoyed the work on our joint paper, our never-ending discussions on network segregation measures, our conference visits together, and our numerous nightly gatherings. I am confident that both our academic collaboration and our friendship will last. I thank Michał in addition for his patience with all my questions on LaTeX, Vim and other software issues.

In the first years of my dissertation project, I learned a lot in the meetings of the "Network Reading Group," an informal interdisciplinary group of young researchers that met regularly for a number of years and included, among others, Aljaž Ule, Willemien Kets, Marco van der Leij, Ana Babus, and Jurjen Kamphorst.

I thank my colleagues at the Department of Sociology at Utrecht University and the ICS research school for maintaining a lively research community in which I've always felt at home. The administrative staff at the department has done a great job creating a comfortable, well-organized work environment.

I also thank my ICS yeargroup for all the good times and sharing the burden.

Richard Zijdeman has been a great friend ever since we met as undergraduate students in the sociology program. He is now also a wonderful colleague, and it is only fitting that we finally started our first research project together.

I thank my family for their love and support, and in particular my father, for always stimulating my curiosity and critical thinking. Finally, I thank Sytske for being by my side.

Rense Corten
Utrecht, April 2009

Curriculum Vitae

Rense Corten (1979) was born in Groningen, the Netherlands, where he graduated from the Praedinius Gymnasium in 1998. Between 1999 and 2004, he studied Sociology and Philosophy at Utrecht University. During this period, he also spend some months at the University of Cologne, Germany, and the European University Institute in Florence, Italy. In 2004, he obtained a Master's degree in Sociology from Utrecht University. In September 2004, he became a Ph. D. student at the Interuniversity Center for Social Science Theory and Methodology (ICS) in Utrecht, where he completed this dissertation. In 2007, he spent a research period at the Institute for Research in the Social Sciences (IRISS) at Stanford University, CA. As of March 2009, he is employed as a postdoctoral researcher (Dutch: *universitair docent*) at the Department of Sociology/ICS at Utrecht University.

ICS dissertation series

The ICS-series presents dissertations of the Interuniversity Center for Social Science Theory and Methodology. Each of these studies aims at integrating explicit theory formation with state-of-the-art empirical research or at the development of advanced methods for empirical research. The ICS was founded in 1986 as a cooperative effort of the universities of Groningen and Utrecht. Since 1992, the ICS expanded to the University of Nijmegen. Most of the projects are financed by the participating universities or by the Netherlands Organization for Scientific Research (NWO). The international composition of the ICS graduate students is mirrored in the increasing international orientation of the projects and thus of the ICS-series itself.

1. C. van Liere, (1990), *Lastige Leerlingen. Een empirisch onderzoek naar sociale oorzaken van probleemgedrag op basisscholen*, Amsterdam: Thesis Publishers.
2. Marco H.D. van Leeuwen, (1990), *Bijstand in Amsterdam, ca. 1800–1850. Armenzorg als beheersings- en overlevingsstrategie*, ICS-dissertation, Utrecht.
3. I. Maas, (1990), *Deelname aan podiumkunsten via de podia, de media en actieve beoefening. Substitutie of leereffecten?*, Amsterdam: Thesis Publishers.
4. M.I. Broese van Groenou, (1991), *Gescheiden Netwerken. De relaties met vrienden en verwanten na echtscheiding*, Amsterdam: Thesis Publishers.
5. Jan M.M. van den Bos, (1991), *Dutch EC Policy Making. A Model-Guided Approach to Coordination and Negotiation*, Amsterdam: Thesis Publishers.
6. Karin Sanders, (1991), *Vrouwelijke Pioniers. Vrouwen en mannen met een 'mannelijke' hogere beroepsopleiding aan het begin van hun loopbaan*, Amsterdam: Thesis Publishers.
7. Sjerp de Vries, (1991), *Egoism, Altruism, and Social Justice. Theory and Experiments on Cooperation in Social Dilemmas*, Amsterdam: Thesis Publishers.
8. Ronald S. Batenburg, (1991), *Automatisering in bedrijf*, Amsterdam: Thesis Publishers.
9. Rudi Wielers, (1991), *Selectie en allocatie op de arbeidsmarkt. Een uitwerking voor de informele en geïnstitutionaliseerde kinderopvang*, Amsterdam: Thesis Publishers.
10. Gert P. Westert, (1991), *Verschillen in ziekenhuisgebruik*, ICS-dissertation, Groningen.
11. Hanneke Hermsen, (1992), *Votes and Policy Preferences. Equilibria in Party Systems*, Amsterdam: Thesis Publishers.
12. Cora J.M. Maas, (1992), *Probleemleerlingen in het basisonderwijs*, Amsterdam: Thesis Publishers.
13. Ed A.W. Boxman, (1992), *Contacten en carrière. Een empirisch-theoretisch onderzoek naar de relatie tussen sociale netwerken en arbeidsmarktposities*, Amsterdam: Thesis Publishers.
14. Conny G.J. Taes, (1992), *Kijken naar banen. Een onderzoek naar de inschatting van arbeidsmarktkansen bij schoolverlaters uit het middelbaar beroepsonderwijs*, Amsterdam: Thesis Publishers.
15. Peter van Roozendaal, (1992), *Cabinets in Multi-Party Democracies. The Effect of Dominant and Central Parties on Cabinet Composition and Durability*, Amsterdam: Thesis Publishers.
16. Marcel van Dam, (1992), *Regio zonder regie. Verschillen in en effectiviteit van gemeentelijk arbeidsmarktbeleid*, Amsterdam: Thesis Publishers.
17. Tanja van der Lippe, (1993), *Arbeidsverdeling tussen mannen en vrouwen*, Amsterdam: Thesis Publishers.

18. Marc A. Jacobs, (1993), *Software: Kopen of Kopiëren? Een sociaal-wetenschappelijk onderzoek onder PC-gebruikers*, Amsterdam: Thesis Publishers.
19. Peter van der Meer, (1993), *Verdringing op de Nederlandse arbeidsmarkt. Sector- en sekseverschillen*, Amsterdam: Thesis Publishers.
20. Gerbert Kraaykamp, (1993), *Over lezen gesproken. Een studie naar sociale differentiatie in leesgedrag*, Amsterdam: Thesis Publishers.
21. Evelien Zeggelink, (1993), *Strangers into Friends. The Evolution of Friendship Networks Using an Individual Oriented Modeling Approach*, Amsterdam: Thesis Publishers.
22. Jaco Berveling, (1994), *Het stempel op de besluitvorming. Macht, invloed en besluitvorming op twee Amsterdamse beleidsterreinen*, Amsterdam: Thesis Publishers.
23. Wim Bernasco, (1994), *Coupled Careers. The Effects of Spouse's Resources on Success at Work*, Amsterdam: Thesis Publishers.
24. Liset van Dijk, (1994), *Choices in Child Care. The Distribution of Child Care Among Mothers, Fathers and Non-Parental Care Providers*, Amsterdam: Thesis Publishers.
25. Jos de Haan, (1994), *Research Groups in Dutch Sociology*, Amsterdam: Thesis Publishers.
26. K. Boahene, (1995), *Innovation Adoption as a Socio-Economic Process. The Case of the Ghanaian Cocoa Industry*, Amsterdam: Thesis Publishers.
27. Paul E.M. Ligthart, (1995), *Solidarity in Economic Transactions. An Experimental Study of Framing Effects in Bargaining and Contracting*, Amsterdam: Thesis Publishers.
28. Roger Th. A.J. Leenders, (1995), *Structure and Influence. Statistical Models for the Dynamics of Actor Attributes, Network Structure, and their Interdependence*, Amsterdam: Thesis Publishers.
29. Beate Völker, (1995), *Should Auld Acquaintance Be Forgot...? Institutions of Communism, the Transition to Capitalism and Personal Networks: the Case of East Germany*, Amsterdam: Thesis Publishers.
30. A. Cancrinus-Matthijsse, (1995), *Tussen hulpverlening en ondernemerschap. Beroepsuitoefening en taakopvattingen van openbare apothekers in een aantal West-Europese landen*, Amsterdam: Thesis Publishers.
31. Nardi Steverink, (1996), *Zo lang mogelijk zelfstandig. Naar een verklaring van verschillen in oriëntatie ten aanzien van opname in een verzorgingstehuis onder fysiek kwetsbare ouderen*, Amsterdam: Thesis Publishers.
32. Ellen Lindeman, (1996), *Participatie in vrijwilligerswerk*, Amsterdam: Thesis Publishers.
33. Chris Snijders, (1996), *Trust and Commitments*, Amsterdam: Thesis Publishers.
34. Koos Postma, (1996), *Changing Prejudice in Hungary. A Study on the Collapse of State Socialism and Its Impact on Prejudice Against Gypsies and Jews*, Amsterdam: Thesis Publishers.
35. Jooske T. van Busschbach, (1996), *Uit het oog, uit het hart? Stabiliteit en verandering in persoonlijke relaties*, Amsterdam: Thesis Publishers.
36. René Torenvlied, (1996), *Besluiten in uitvoering. Theorieën over beleidsuitvoering modelmatig getoetst op sociale vernieuwing in drie gemeenten*, Amsterdam: Thesis Publishers.
37. Andreas Flache, (1996), *The Double Edge of Networks. An Analysis of the Effect of Informal Networks on Cooperation in Social Dilemmas*, Amsterdam: Thesis Publishers.
38. Kees van Veen, (1997), *Inside an Internal Labor Market: Formal Rules, Flexibility and Career Lines in a Dutch Manufacturing Company*, Amsterdam: Thesis Publishers.
39. Lucienne van Eijk, (1997), *Activity and Well-being in the Elderly*, Amsterdam: Thesis Publishers.
40. Róbert Gál, (1997), *Unreliability. Contract Discipline and Contract Governance under Economic Transition*, Amsterdam: Thesis Publishers.
41. Anne-Geerte van de Goor, (1997), *Effects of Regulation on Disability Duration*, ICS-dissertation, Utrecht.
42. Boris Blumberg, (1997), *Das Management von Technologiekooperationen. Partnersuche und Verhandlungen mit dem Partner aus Empirisch-Theoretischer Perspektive*, ICS-dissertation, Utrecht.
43. Marijke von Bergh, (1997), *Loopbanen van oudere werknemers*, Amsterdam: Thesis Publishers.
44. Anna Petra Nieboer, (1997), *Life-Events and Well-Being: A Prospective Study on Changes in Well-Being of Elderly People Due to a Serious Illness Event or Death of the Spouse*, Amsterdam: Thesis Publishers.

45. Jacques Niehof, (1997), *Resources and Social Reproduction: The Effects of Cultural and Material Resources on Educational and Occupational Careers in Industrial Nations at the End of the Twentieth Century*, ICS-dissertation, Nijmegen.
46. Ariana Need, (1997), *The Kindred Vote. Individual and Family Effects of Social Class and Religion on Electoral Change in the Netherlands, 1956–1994*, ICS-dissertation, Nijmegen.
47. Jim Allen, (1997), *Sector Composition and the Effect of Education on Wages: an International Comparison*, Amsterdam: Thesis Publishers.
48. Jack B.F. Hutten, (1998), *Workload and Provision of Care in General Practice. An Empirical Study of the Relation Between Workload of Dutch General Practitioners and the Content and Quality of their Care*, ICS-dissertation, Utrecht.
49. Per B. Kropp, (1998), *Berufserfolg im Transformationsprozeß, Eine theoretisch-empirische Studie über die Gewinner und Verlierer der Wende in Ostdeutschland*, ICS-dissertation, Utrecht.
50. Maarten H.J. Wolbers, (1998), *Diploma-inflatie en verdringing op de arbeidsmarkt. Een studie naar ontwikkelingen in de opbrengsten van diploma's in Nederland*, ICS-dissertation, Nijmegen.
51. Wilma Smeenk, (1998), *Opportunity and Marriage. The Impact of Individual Resources and Marriage Market Structure on First Marriage Timing and Partner Choice in the Netherlands*, ICS-dissertation, Nijmegen.
52. Marinus Spreen, (1999), *Sampling Personal Network Structures: Statistical Inference in Ego-Graphs*, ICS-dissertation, Groningen.
53. Vincent Buskens, (1999), *Social Networks and Trust*, ICS-dissertation, Utrecht.
54. Susanne Rijken, (1999), *Educational Expansion and Status Attainment. A Cross-National and Over-Time Comparison*, ICS-dissertation, Utrecht.
55. Mérove Gijsberts, (1999), *The Legitimation of Inequality in State-Socialist and Market Societies, 1987–1996*, ICS-dissertation, Utrecht.
56. Gerhard G. Van de Bunt, (1999), *Friends by Choice. An Actor-Oriented Statistical Network Model for Friendship Networks Through Time*, ICS-dissertation, Groningen.
57. Robert Thomson, (1999), *The Party Mandate: Election Pledges and Government Actions in the Netherlands, 1986–1998*, Amsterdam: Thela Thesis.
58. Corine Baarda, (1999), *Politieke besluiten en boeren beslissingen. Het draagvlak van het mestbeleid tot 2000*, ICS-dissertation, Groningen.
59. Rafael Wittek, (1999), *Interdependence and Informal Control in Organizations*, ICS-dissertation, Groningen.
60. Diane Payne, (1999), *Policy Making in the European Union: an Analysis of the Impact of the Reform of the Structural Funds in Ireland*, ICS-dissertation, Groningen.
61. René Veenstra, (1999), *Leerlingen – Klassen – Scholen. Prestaties en vorderingen van leerlingen in het voortgezet onderwijs*, Amsterdam: Thela Thesis.
62. Marjolein Achterkamp, (1999), *Influence Strategies in Collective Decision Making. A Comparison of Two Models*, ICS-dissertation, Groningen.
63. Peter Mühlau, (2000), *The Governance of the Employment Relation. A Relational Signaling Perspective*, ICS-dissertation, Groningen.
64. Agnes Akkerman, (2000), *Verdeelde vakbeweging en stakingen. Concurrentie om leden*, ICS-dissertation, Groningen.
65. Sandra van Thiel, (2000), *Quangocratization: Trends, Causes and Consequences*, ICS-dissertation, Utrecht.
66. Rudi Turksema, (2000), *Supply of Day Care*, ICS-dissertation, Utrecht.
67. Sylvia E. Korupp (2000), *Mothers and the Process of Social Stratification*, ICS-dissertation, Utrecht.
68. Bernard A. Nijstad (2000), *How the Group Affects the Mind: Effects of Communication in Idea Generating Groups*, ICS-dissertation, Utrecht.
69. Inge F. de Wolf (2000), *Opleidingsspecialisatie en arbeidsmarktsucces van sociale wetenschappers*, ICS-dissertation, Utrecht.
70. Jan Kratzer (2001), *Communication and Performance: An Empirical Study in Innovation Teams*, ICS-dissertation, Groningen.

71. Madelon Kroneman (2001), *Healthcare Systems and Hospital Bed Use*, ICS/NIVEL-dissertation, Utrecht.
72. Herman van de Werfhorst (2001), *Field of Study and Social Inequality. Four Types of Educational Resources in the Process of Stratification in the Netherlands*, ICS-dissertation, Nijmegen.
73. Tamás Bartus (2001), *Social Capital and Earnings Inequalities. The Role of Informal Job Search in Hungary*, ICS-dissertation, Groningen.
74. Hester Moerbeek (2001), *Friends and Foes in the Occupational Career. The Influence of Sweet and Sour Social Capital on the Labour Market*, ICS-dissertation, Nijmegen.
75. Marcel van Assen (2001), *Essays on Actor Perspectives in Exchange Networks and Social Dilemmas*, ICS-dissertation, Groningen.
76. Inge Sieben (2001), *Sibling Similarities and Social Stratification. The Impact of Family Background across Countries and Cohorts*, ICS-dissertation, Nijmegen.
77. Alinda van Bruggen (2001), *Individual Production of Social Well-Being. An Exploratory Study*, ICS-dissertation, Groningen.
78. Marcel Coenders (2001), *Nationalistic Attitudes and Ethnic Exclusionism in a Comparative Perspective: An Empirical Study of Attitudes Toward the Country and Ethnic Immigrants in 22 Countries*, ICS-dissertation, Nijmegen.
79. Marcel Lubbers (2001), *Exclusionistic Electorates. Extreme Right-Wing Voting in Western Europe*, ICS-dissertation, Nijmegen.
80. Uwe Matzat (2001), *Social Networks and Cooperation in Electronic Communities. A theoretical-empirical Analysis of Academic Communication and Internet Discussion Groups*, ICS-dissertation, Groningen.
81. Jacques P.G. Janssen (2002), *Do Opposites Attract Divorce? Dimensions of Mixed Marriage and the Risk of Divorce in the Netherlands*, ICS-dissertation, Nijmegen.
82. Miranda Jansen (2002), *Waardenoriëntaties en partnerrelaties. Een panelstudie naar wederzijdse invloeden*, ICS-dissertation, Utrecht.
83. Anne Rigt Poortman (2002), *Socioeconomic Causes and Consequences of Divorce*, ICS-dissertation, Utrecht.
84. Alexander Gattig (2002), *Intertemporal Decision Making*, ICS-dissertation, Groningen.
85. Gerrit Rooks (2002), *Contract en Conflict: Strategisch Management van Inkooptransacties*, ICS-dissertation, Utrecht.
86. Károly Takács (2002), *Social Networks and Intergroup Conflict*, ICS-dissertation, Groningen.
87. Thomas Gautschi (2002), *Trust and Exchange, Effects of Temporal Embeddedness and Network Embeddedness on Providing and Dividing a Surplus*, ICS-dissertation, Utrecht.
88. Hilde Bras (2002), *Zeeuwse meiden. Dienen in de levensloop van vrouwen, ca. 1850–1950*, Amsterdam: Aksant Academic Publishers.
89. Merijn Rengers (2002), *Economic Lives of Artists. Studies into Careers and the Labour Market in the Cultural Sector*, ICS-dissertation, Utrecht.
90. Annelies Kassenberg (2002), *Wat scholieren bindt. Sociale gemeenschap in scholen*, ICS-dissertation, Groningen
91. Marc Verboord (2003), *Moet de meester dalen of de leerling klimmen? De invloed van literatuuronderwijs en ouders op het lezen van boeken tussen 1975 en 2000*, ICS-dissertation, Utrecht.
92. Marcel van Egmond (2003), *Rain Falls on All of Us (but Some Manage to Get More Wet than Others): Political Context and Electoral Participation*, ICS-dissertation, Nijmegen.
93. Justine Horgan (2003), *High Performance Human Resource Management in Ireland and the Netherlands: Adoption and Effectiveness*, ICS-dissertation, Groningen.
94. Corine Hoeben (2003), *LETS' Be a Community. Community in Local Exchange Trading Systems*, ICS-dissertation, Groningen.
95. Christian Steglich (2003), *The Framing of Decision Situations. Automatic Goal Selection and Rational Goal Pursuit*, ICS-dissertation, Groningen.

96. Johan van Wilsem (2003), *Crime and Context. The Impact of Individual, Neighborhood, City and Country Characteristics on Victimization*, ICS-dissertation, Nijmegen.
97. Christiaan Monden (2003), *Education, Inequality and Health. The Impact of Partners and Life Course*, ICS-dissertation, Nijmegen.
98. Evelyn Hello (2003), *Educational Attainment and Ethnic Attitudes. How to Explain their Relationship*, ICS-dissertation, Nijmegen.
99. Marnix Croes en Peter Tammes (2004), *Gif laten wij niet voortbestaan. Een onderzoek naar de overlevingskansen van joden in de Nederlandse gemeenten, 1940–1945*, Amsterdam: Aksant Academic Publishers.
100. Ineke Nagel (2004), *Cultuurdeelname in de levensloop*, ICS-dissertation, Utrecht.
101. Marieke van der Wal (2004), *Competencies to Participate in Life. Measurement and the Impact of School*, ICS-dissertation, Groningen.
102. Vivian Meertens (2004), *Depressive Symptoms in the General Population: a Multifactorial Social Approach*, ICS-dissertation, Nijmegen.
103. Hanneke Schuurmans (2004), *Promoting Well-Being in Frail Elderly People. Theory and Intervention*, ICS-dissertation, Groningen.
104. Javier Arregui (2004), *Negotiation in Legislative Decision-Making in the European Union*, ICS-dissertation, Groningen.
105. Tamar Fischer (2004), *Parental Divorce, Conflict and Resources. The Effects on Children's Behaviour Problems, Socioeconomic Attainment, and Transitions in the Demographic Career*, ICS-dissertation, Nijmegen.
106. René Bekkers (2004), *Giving and Volunteering in the Netherlands: Sociological and Psychological Perspectives*, ICS-dissertation, Utrecht.
107. Renée van der Hulst (2004), *Gender Differences in Workplace Authority: An Empirical Study on Social Networks*, ICS-dissertation, Groningen.
108. Rita Smaniotto (2004), *'You Scratch My Back and I Scratch Yours' Versus 'Love Thy Neighbour'. Two Proximate Mechanisms of Reciprocal Altruism*, ICS-dissertation, Groningen.
109. Maurice Gesthuizen (2004), *The Life-Course of the Low-Educated in the Netherlands: Social and Economic Risks*, ICS-dissertation, Nijmegen.
110. Carlijne Philips (2005), *Vakantiegemeenschappen. Kwalitatief en Kwantitatief Onderzoek naar Gelegenheid- en Refreshergemeenschap tijdens de Vakantie*, ICS-dissertation, Groningen.
111. Esther de Ruijter (2005), *Household Outsourcing*, ICS-dissertation, Utrecht.
112. Frank van Tubergen (2005), *The Integration of Immigrants in Cross-National Perspective: Origin, Destination, and Community Effects*, ICS-dissertation, Utrecht.
113. Ferry Koster (2005), *For the Time Being. Accounting for Inconclusive Findings Concerning the Effects of Temporary Employment Relationships on Solidary Behavior of Employees*, ICS-dissertation, Groningen.
114. Carolien Klein Haarhuis (2005), *Promoting Anti-Corruption Reforms. Evaluating the Implementation of a World Bank Anti-Corruption Program in Seven African Countries (1999–2001)*, ICS-dissertation, Utrecht.
115. Martin van der Gaag (2005), *Measurement of Individual Social Capital*, ICS-dissertation, Groningen.
116. Johan Hansen (2005), *Shaping Careers of Men and Women in Organizational Contexts*, ICS-dissertation, Utrecht.
117. Davide Barrera (2005), *Trust in Embedded Settings*, ICS-dissertation, Utrecht.
118. Mattijs Lambooij (2005), *Promoting Cooperation. Studies into the Effects of Long-Term and Short-Term Rewards on Cooperation of Employees*, ICS-dissertation, Utrecht.
119. Lotte Vermeij (2006), *What's Cooking? Cultural Boundaries among Dutch Teenagers of Different Ethnic Origins in the Context of School*, ICS-dissertation, Utrecht.
120. Mathilde Strating (2006), *Facing the Challenge of Rheumatoid Arthritis. A 13-year Prospective Study among Patients and Cross-Sectional Study among Their Partners*, ICS-dissertation, Groningen.
121. Jannes de Vries (2006), *Measurement Error in Family Background Variables: The Bias in the Intergenerational Transmission of Status, Cultural Consumption, Party Preference, and Religiosity*, ICS-dissertation, Nijmegen.

122. Stefan Thau (2006), *Workplace Deviance: Four Studies on Employee Motives and Self-Regulation*, ICS-dissertation, Groningen.
123. Mirjam Plantinga (2006), *Employee Motivation and Employee Performance in Child Care. The effects of the Introduction of Market Forces on Employees in the Dutch Child-Care Sector*, ICS-dissertation, Groningen.
124. Helga de Valk (2006), *Pathways into Adulthood. A Comparative Study on Family Life Transitions among Migrant and Dutch Youth*, ICS-dissertation, Utrecht.
125. Henrike Elzen (2006), *Self-Management for Chronically Ill Older People*, ICS-dissertation, Groningen.
126. Ayşe Güveli (2007), *New Social Classes within the Service Class in the Netherlands and Britain. Adjusting the EGP Class Schema for the Technocrats and the Social and Cultural Specialists*, ICS-dissertation, Nijmegen.
127. Willem-Jan Verhoeven (2007), *Income Attainment in Post-Communist Societies*, ICS-dissertation, Utrecht.
128. Marieke Voorpostel (2007), *Sibling support: The Exchange of Help among Brothers and Sisters in the Netherlands*, ICS-dissertation, Utrecht.
129. Jacob Dijkstra (2007), *The Effects of Externalities on Partner Choice and Payoffs in Exchange Networks*, ICS-dissertation, Groningen.
130. Patricia van Echtelt (2007), *Time-Greedy Employment Relationships: Four Studies on the Time Claims of Post-Fordist Work*, ICS-dissertation, Groningen.
131. Sonja Vogt (2007), *Heterogeneity in Social Dilemmas: The Case of Social Support*, ICS-dissertation, Utrecht.
132. Michael Schweinberger (2007), *Statistical Methods for Studying the Evolution of Networks and Behavior*, ICS-dissertation, Groningen.
133. István Back (2007), *Commitment and Evolution: Connecting Emotion and Reason in Long-term Relationships*, ICS-dissertation, Groningen.
134. Ruben van Gaalen (2007), *Solidarity and Ambivalence in Parent-Child Relationships*, ICS-dissertation, Utrecht.
135. Jan Reitsma (2007), *Religiosity and Solidarity - Dimensions and Relationships Disentangled and Tested*, ICS-dissertation, Nijmegen.
136. Jan Kornelis Dijkstra (2007) *Status and Affection among (Pre)Adolescents and Their Relation with Antisocial and Prosocial Behavior*, ICS-dissertation, Groingen.
137. Wouter van Gils (2007), *Full-time Working Couples in the Netherlands. Causes and Consequences*, ICS-dissertation, Nijmegen.
138. Djamila Schans (2007), *Ethnic Diversity in Intergenerational Solidarity*, ICS-dissertation, Utrecht.
139. Ruud van der Meulen (2007), *Brug over Woelig Water: Lidmaatschap van Sportverenigingen, Vriendschappen, Kennissenkringen en Veralgemeend Vertrouwen*, ICS-dissertation, Nijmegen.
140. Andrea Knecht (2008), *Friendship Selection and Friends' Influence. Dynamics of Networks and Actor Attributes in Early Adolescence*, ICS-dissertation, Utrecht.
141. Ingrid Doorten (2008), *The Division of Unpaid Work in the Household: A Stubborn Pattern?*, ICS-dissertation, Utrecht.
142. Stijn Ruiter (2008), *Association in Context and Association as Context: Causes and Consequences of Voluntary Association Involvement*, ICS-dissertation, Nijmegen.
143. Janneke Joly (2008), *People on Our Minds: When Humanized Contexts Activate Social Norms*, ICS-dissertation, Groningen.
144. Margreet Frieling (2008), *'Joint production' als motor voor actief burgerschap in de buurt*, ICS-dissertion, Groningen.
145. Ellen Verbakel (2008), *The Partner as Resource or Restriction? Labour Market Careers of Husbands and Wives and the Consequences for Inequality Between Couples*, ICS-dissertation, Nijmegen.
146. Gijs van Houten (2008), *Beleidsuitvoering in gelaagde stelsels. De doorwerking van aanbevelingen van de Stichting van de Arbeid in het CAO-overleg*, ICS-dissertation, Utrecht.

147. Eva Jaspers (2008), *Intolerance over Time. Macro and Micro Level Questions on Attitudes Towards Euthanasia, Homosexuality and Ethnic Minorities*, ICS-dissertation, Nijmegen.
148. Gijs Weijters (2008), *Youth delinquency in Dutch cities and schools: A multilevel approach*, ICS-dissertation, Nijmegen.
149. Jessica Pass (2009), *The Self in Social Rejection*, ICS-dissertation, Groningen.
150. Gerald Mollenhorst (2009), *Networks in Contexts. How Meeting Opportunities Affect Personal Relationships*, ICS-dissertation, Utrecht.
151. Tom van der Meer (2009), *States of freely associating citizens: comparative studies into the impact of state institutions on social, civic and political participation*, ICS-dissertation, Nijmegen.
152. Manuela Vieth (2009), *Commitments and Reciprocity in Trust Situations. Experimental Studies on Obligation, Indignation, and Self-Consistency*, ICS-dissertation, Utrecht.