LUDO WALTMAN

Computational and Game-Theoretic Approaches for Modeling Bounded Rationality
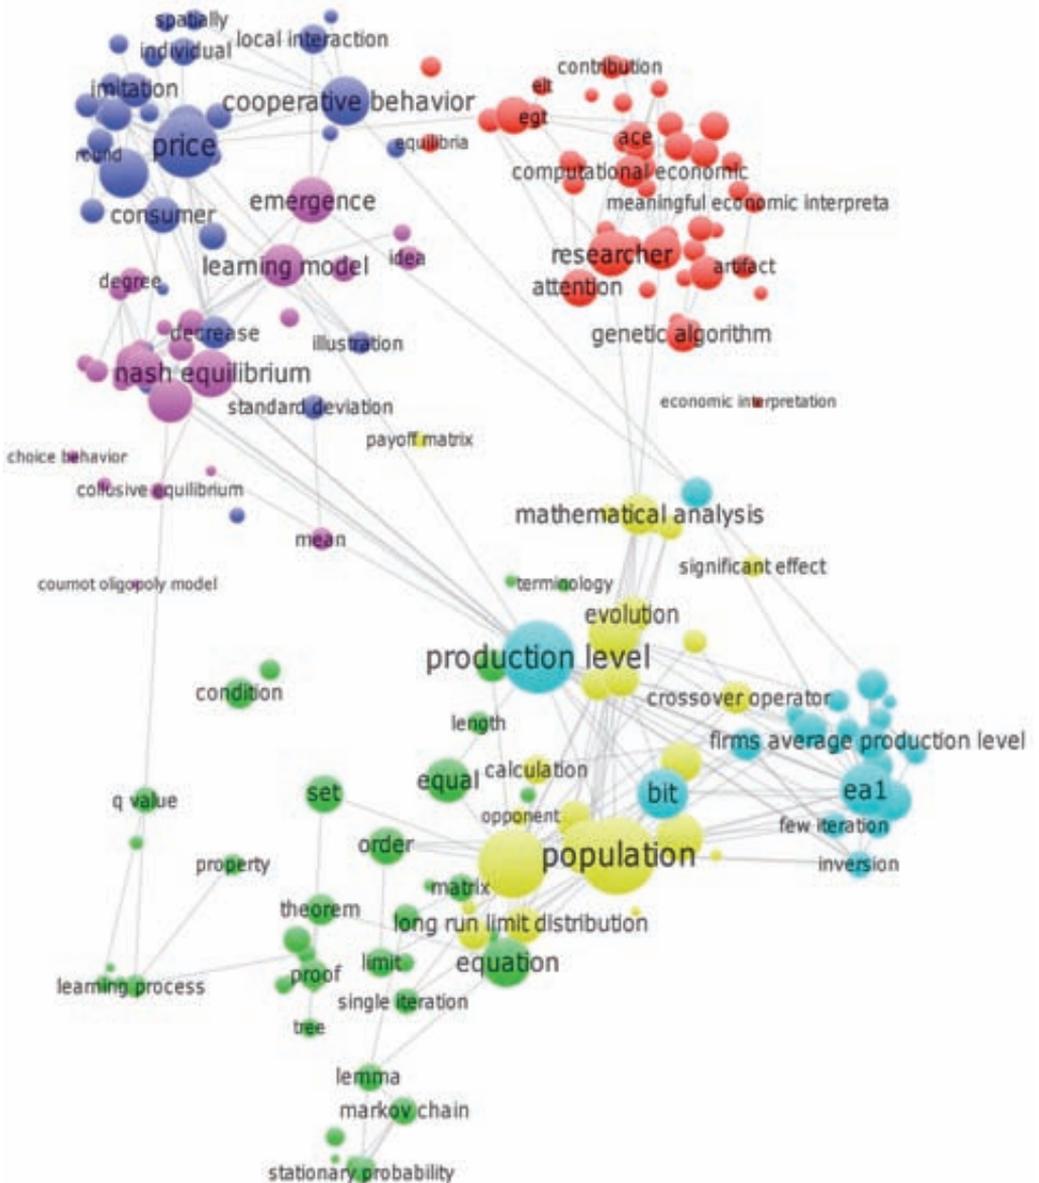
Computational and Game-Theoretic Approaches for Modeling Bounded Rationality

Computational and Game-Theoretic Approaches for Modeling Bounded Rationality

Computationele en speltheoretische methodes
voor het modelleren van beperkte rationaliteit

Thesis

to obtain the degree of Doctor from the
Erasmus University Rotterdam
by command of the
rector magnificus

Prof.dr. H.G. Schmidt

and in accordance with the decision of the Doctorate Board.

The public defence shall be held on

Thursday, 13 October 2011 at 15.30 hrs

by

LUDO WALTMAN
place of birth, Dordrecht.

ERASMUS UNIVERSITEIT ROTTERDAM

Doctoral Committee

Promotors:	Prof.dr.ir. R. Dekker
	Prof.dr.ir. U. Kaymak
Other members:	Prof.dr. C.A. van Bochove
	Prof.dr. M.C.W. Janssen
	Prof.dr. N.J. Vriend

Acknowledgements

Doing scientific research is one of the greatest joys in my life. Nevertheless, working on this thesis has not always been easy. Some research projects failed (at least in my own opinion) and therefore did not become part of this thesis. Other projects have been finished, but only after I managed to overcome the doubts I had on these projects. There are a number of people whom I would like to acknowledge for their support during the work on this thesis. Without these people, my time as a PhD student would have been much less enjoyable, and finishing this thesis would probably have taken even more time!

First of all, I want to thank the supervisors of this thesis, Rommert Dekker and Uzay Kaymak. Despite both of them being extremely busy, they managed to have regular meetings with me in which we discussed the progress on the various research projects. What I perhaps liked most was the freedom they gave me in this PhD project. When we had different ideas on which direction to take, they gave me the full freedom to choose my own preferred approach.

I am grateful to my fellow PhD students at Erasmus School of Economics for the nice time we had. There are a few colleagues whom I would like to thank in particular. Joost van Rosmalen has been my roommate for four years. During this time, we have had lots of discussions on our research projects, and it is hard to think of a roommate from whom I could have learned more than from Joost. Martijn Kagie never refused to have a drink with me in the afternoon. I always enjoyed the way in which he told about his ideas on doing scientific research (which are quite different from my own ideas!).

Of course, I am especially thankful to Nees (or actually: Nelis) Jan van Eck. Not only did we collaborate on most of the research projects for this thesis, but he also allowed me to interfere with his own projects (which, to be honest, were usually much

more fun than the projects that I was working on myself). Apart from our joint scientific work, I keep especially good memories of the trips we made together, both conference visits and holidays.

Finally, working on a PhD thesis is so much easier with a stable and happy personal life. I want to thank my parents, Nico and Sandra, for laying the foundation for such a life. I want to thank Eline for all the fun we have together (and nowadays with Tim as well), for never complaining about me being occupied so much by my scientific work, and very importantly, for pushing me more than anyone else to finish this thesis.

Ludo Waltman

August 2011

Table of Contents

Chapter 1

Introduction

1.1 The gradually shifting paradigm of mainstream economics

Economic theory is frequently criticized. The assumption of economic agents being fully rational, well-informed, and driven only by selfish motives plays a central role in economic theory. This assumption, however, is often argued to be highly unrealistic. Economic theory is also criticized for its heavy reliance on static models and equilibrium analysis. Related to this, it is argued that economists pay too much attention to mathematical rigor and neglect the practical relevance of their work. According to Mark Blaug, a scholar of economic methodology and of the history of economic thought, "economics has increasingly become an intellectual game played for its own sake and not for its practical consequences" (Blaug, 2002, p. 36).

I believe that the above criticism, whether one in essence agrees with it or not, has somewhat less relevance for today's economic research than for the economic research that was done, say, twenty or thirty years ago. Especially during the last two decades, economists have become more and more interested in studying boundedly rational behavior. They do so not only theoretically but also empirically. Economists have also started analyzing all kinds of dynamic processes, which means that they are no longer focusing exclusively on static equilibrium analysis. Because of these developments, criticism on unrealistic rationality assumptions made by economists and on economists'

fixation with equilibrium analysis seems somewhat less relevant today than it was a few decades ago.

This view of the developments taking place in economic research fits well into the account offered by Colander, Holt, and Rosser (2004b).[1] Colander et al. assert that mainstream economics is moving away from, what they call, the holy trinity of rationality, selfishness, and equilibrium, and they interpret this movement as a paradigm shift. Unlike Thomas Kuhn's notion of a paradigm shift (Kuhn, 1996), Colander et al. argue that the paradigm shift in mainstream economics occurs gradually rather than suddenly. They also argue that the paradigm shift is partly due to forces from within the mainstream.

The studies presented in this thesis all involve economic agents that are boundedly rational and that dynamically adjust their behavior. Before introducing the studies in more detail, I will first discuss three areas of economic research in which the modeling of bounded rationality and dynamic adjustment processes plays a central role. My focus will be on game-theoretical settings, that is, settings with a limited number of agents interacting with each other. In Sections 1.2 and 1.3, the areas of evolutionary game theory and economic learning theory will be discussed. Nowadays, these areas seem to be well-accepted in mainstream economics. In Section 1.4, the area of agent-based computational economics will be discussed. Even though there is some overlap with the other two areas, the area of agent-based computational economics seems to have gained much less acceptance in mainstream economics (Lehtinen & Kuorikoski, 2007; Leombruni & Richiardi, 2005). After the three research areas have been discussed, the studies presented in this thesis will be introduced in Section 1.5. A brief outline of the thesis will be provided in Section 1.6.

Some general characteristics of the research areas of evolutionary game theory, economic learning theory, and agent-based computational economics are summarized in Table 1.1. These characteristics will be discussed in the next three sections. To the best of my knowledge, there is no literature in which the areas of evolutionary game theory, economic learning theory, and agent-based computational economics are systematically compared with each other. There also does not seem to exist a generally accepted view on the characteristics of the three areas and on the similarities and dissimilarities be-

---

[1]See also Colander, Holt, and Rosser (2004a) for interesting background reading.

Table 1.1: Some general characteristics of the areas of evolutionary game theory (EGT), economic learning theory (ELT), and agent-based computational economics (ACE). See Sections 1.2, 1.3, and 1.4 for an explanation of the table.

	EGT	ELT	ACE
Model specification	mathematical	mathematical	algorithmic
Model complexity	low	fairly low	medium/high
Methodology	mathematical	mathematical; econometric	simulation; agent-based modeling
Inspiration	evolutionary	ad hoc; psychological	ad hoc; computer science; evolutionary
Focus	long run	short/long run	short/long run
Empirical orientation	low	fairly high	moderate

tween the areas. The discussion below reflects my personal view on how the areas of evolutionary game theory, economic learning theory, and agent-based computational economics can be characterized and on how they relate to each other.

1.2   Evolutionary game theory

Much of mainstream micro-economic theory rests on a foundation in classical game theory. Classical game theory is built on the assumption of fully rational behavior, and it focuses exclusively on equilibria. No attention is paid to issues such as bounded rationality and lack of information. Also, the question how equilibria emerge is more or less ignored and definitely not explicitly modeled.

During the last two decades, a fundamental change of direction has taken place in the field of game theory (see also Sugden, 2001). Many game theorists have shifted their attention from classical game theory to evolutionary game theory (EGT). EGT assumes that behavior in games is determined by some evolutionary mechanism. Economists usually do not give a biological interpretation to the evolutionary mechanism. Instead, they interpret the mechanism in terms of agents that learn from each other or that imitate each other. Alternatively, the mechanism is interpreted in terms of, for example, successful firms taking over market share from non-successful ones. Unlike classical game

theory, EGT does not rely on the assumption of fully rational behavior. Instead, it assumes some kind of bounded rationality. EGT also pays more attention to the dynamic processes through which equilibria may emerge.

EGT seems to have been largely inspired by the work of the biologist Maynard Smith (1982). He introduced the well-known notion of an evolutionary stable strategy (ESS). Even though the notion of an ESS is in fact just a refinement of the Nash equilibrium concept, its underlying justification is very different. It relies on an evolutionary justification rather than on a justification based on full rationality and common knowledge. Although ESS is basically a static equilibrium concept, much of the current EGT literature explicitly models dynamic processes and examines whether such processes converge to a certain equilibrium.

Unlike the classical game theory literature, the EGT literature has not yet resulted in a single generally accepted modeling framework. Instead, the models studied by different authors are sometimes based on quite different ideas. A number of books have appeared in which various approaches to EGT are discussed, see Gintis (2009), Vega-Redondo (1996), Samuelson (1997), Weibull (1995), and Young (1998).

In some way, the difference between classical game theory and EGT (as studied by economists) is smaller than it might seem to be (see also Sugden, 2001). Although EGT has dropped the important assumptions of full rationality and common knowledge, the methodology of EGT is in many respects not very different from that of classical game theory. Like classical game theory, EGT relies almost exclusively on model building and rigorous mathematical analysis. Models are almost always studied mathematically, and computer simulations are rarely used. Also, researchers in EGT usually focus mainly on the long-run outcomes of their models. Hence, like in classical game theory, researchers are primarily interested in equilibria. Another thing to note is that researchers in EGT generally do not pay much attention to empirical validation (Silva & Teixeira, 2009; Sugden, 2001). At the moment, the practical relevance of much of the EGT literature therefore still seems unclear.

1.3   Economic learning theory

During the last decades, more and more economists and game theorists have started studying simple models of learning in games. There seems to be no commonly used name for this research area. I will here refer to it as the area of economic learning theory (ELT).

Research in ELT is partly theoretically oriented and partly empirically oriented. In theoretical research in ELT, the properties of learning models are analyzed mathematically (e.g., Fudenberg & Levine, 1998; Young, 2004). Like in EGT, the focus is mainly on long-run outcomes and on convergence to equilibria. Empirical research in ELT is usually done by experimental and behavioral economists. These researchers try to build models that describe experimental data obtained from human participants learning to play games in laboratory settings (e.g., Chapter 6 in Camerer, 2003). In empirical ELT research, the focus is much more on the short-run outcomes of a learning model. These outcomes are important to determine how well a model fits experimental data.

Well-known learning models studied in ELT include Bayesian learning, best response learning, fictitious play, reinforcement learning, and experience-weighted attraction learning. Overviews of various models proposed in the literature are given by Brenner (1999, 2006) and Fudenberg and Levine (1998). ELT models typically assume some kind of bounded rationality. Unlike the models studied in EGT, ELT models usually do not rely on an evolutionary mechanism. Instead, more ad hoc mechanisms are employed, sometimes with some psychological justification.

1.4   Agent-based computational economics

Agent-based computational economics (ACE) is a research area in which economic phenomena are studied using agent-based modeling techniques. One leading researcher defines ACE as "the computational study of economic processes modeled as dynamic systems of interacting agents" (Tesfatsion, 2006, p. 835). The label 'agent-based computational economics' seems to be used only by a relatively small number of researchers. Many researchers who are following an ACE approach in much of their work do not use the ACE label. This is especially the case for researchers with an interest not only in economics but in the social sciences in general. An example is Epstein (2006), who

refers to his work as generative social science. Miller and Page (2007) are another example. They refer to their work as research into complex adaptive (social) systems. A third example is Axelrod (1997a, 2006a), who refers to his work simply as simulation in the social sciences.[2] Overviews of the area of ACE can be found in a survey paper by Tesfatsion (2003) and a handbook edited by Tesfatsion and Judd (2006). Discussions of the ACE approach are also provided by Dawid (1996) and Dawid, La Poutré, and Yao (2008). General introductions to agent-based modeling in the social sciences can be found in the books by Epstein (2006), Epstein and Axtell (1996), and Miller and Page (2007). The value of simulation and agent-based modeling in social science research is discussed by Axelrod (1997a, 2006a) and Axtell (2000).

The well-known segregation model of Nobel laureate Thomas Schelling (1969, 1971, 1978) is sometimes seen as an early example of the ACE approach (e.g., Epstein & Axtell, 1996). However, the ACE approach really started to attract attention about two decades ago, when researchers began to use techniques from computer science and artificial intelligence for economic modeling purposes (early research includes Arthur, 1991, 1994; Axelrod, 1987; Holland & Miller, 1991; Miller, 1986). Nowadays, ACE is still a rather heterogeneous research area, with some researchers publishing their work mainly in computer science oriented journals and conference proceedings and other researchers targeting more at an audience in mainstream economics.

Like EGT and ELT, ACE assumes that agents are boundedly rational. Following Epstein (2006, p. xvi–xviii), heterogeneity, local interaction, and non-equilibrium dynamics may be seen as other characteristics of the ACE approach. Similar characteristics are mentioned by Miller and Page (2007, Chapter 6). The most important difference between on the one hand ACE and on the other hand ELT and even more EGT seems to be that in ACE more weight is given to realistic modeling while less attention is paid

---

[2]Axelrod notes that many different labels are being used for simulation research in the social sciences. Examples are artificial society, complex system, agent-based model, multi-agent model, individual-based model, bottom-up model, and adaptive system. Given this plethora of labels, Axelrod calls for convergence on a commonly accepted terminology.

to parsimonious modeling.[3] Hence, in general the models studied in ACE are more complex than those studied in EGT and ELT.

A major difference between ACE and mainstream economics is that researchers in ACE perform almost all their analyses using computer simulations.[4] They perform almost no mathematical analyses. In mainstream economics this is exactly the other way around.[5] A related difference concerns the way in which researchers formulate their models. In mainstream economics models are almost always formulated mathematically, which usually means that the specification of a model is completely clear. In ACE many models are formulated in terms of an algorithm to be run on a computer. This sometimes causes some ambiguity about the exact specification of a model (see also Axelrod, 1997a; Axtell, Axelrod, Epstein, & Cohen, 1996).

Many ACE researchers study models that are evolutionarily inspired. From this point of view, ACE is quite closely related to EGT. In both areas, researchers employ evolutionary mechanisms for modeling boundedly rational behavior and dynamic adjustment processes. Despite this similarity between ACE and EGT, there is not much interaction between the two research communities (for an example of work in which ACE and EGT approaches are brought together, see Dawid, 1996).

The ACE approach seems to be somewhat less empirically oriented than the ELT approach. During the last years, however, ACE researchers have started to pay more attention to the issue of empirical validation (e.g., Fagiolo, Birchenhall, & Windrum, 2007; Windrum, Fagiolo, & Moneta, 2007). For an overview of work that has been done on the validation of ACE (and ELT) models based on experimental data, see Duffy (2006).

---

[3]According to Tesfatsion (2006, p. 838), ACE researchers "seek causal explanations grounded in the repeated interactions of agents operating in realistically rendered worlds. Ideally, the agents should have the same flexibility of action in their worlds as their corresponding entities have in the real world". The distinction between realistic models and parsimonious models is somewhat related to the distinction that Miller and Page (2007, p. 78–80) make between flexible models and precise models. Miller and Page argue that compared with mainstream economic models ACE models are more flexible but less precise.

[4]Tesfatsion (2006) seems to regard the use of a computer as an intrinsic element of the ACE approach. However, Epstein (2006, p. xiii) and Miller and Page (2007, p. 64–65) state that in their modeling approaches the use of a computer is not an essential element.

[5]For a discussion why computer simulation and agent-based modeling are not very popular in mainstream economics, see Lehtinen and Kuorikoski (2007) and Leombruni and Richiardi (2005). Miller and Page (2007, Chapter 5) list some common objections against agent-based modeling and try to refute these objections. Axelrod (2006b) discusses some personal experiences with resistance against agent-based modeling.

1.5   Contribution of the thesis

This thesis consists of four studies. These studies are presented in Chapters 2, 3, 4, and 5 of the thesis. Each study is an independent piece of research that makes a separate contribution to the literature.

A common theme of all four studies consists of bringing together elements from on the one hand the ACE modeling approach and on the other hand the EGT (Chapters 3, 4, and 5) and ELT (Chapter 2) modeling approaches. As discussed above, these three modeling approaches have a number of fundamental ideas in common, but they also differ from each other in important ways. Moreover, the three approaches have been developed relatively independently from each other, with only a limited degree of inter-action between the different scientific communities that were involved. In this thesis, relations between the three modeling approaches are pointed out and elements from the different approaches are employed in a complementary fashion. In this way, the thesis aims to bring the three modeling approaches closer together.

A second common theme of the four studies presented in this thesis consists of studying the emergence of cooperative (or altruistic) behavior. Cooperative behavior has been fascinating biologists and social scientists for quite some time already (e.g., Axelrod, 1984; Dawkins, 1989; Fehr & Fischbacher, 2003; Nowak, 2006; Ostrom, 2000). The emergence of cooperative behavior is puzzling because in general it seems to be at odds with basic evolutionary mechanisms. Cooperative behavior is also difficult to reconcile with one of the cornerstones of mainstream economics, namely the idea of fully rational self-interested economic agents. Chapters 2 and 3 of this thesis study the emergence of cooperative behavior in models in which firms compete with each other based on either quantity (Chapter 2) or price (Chapter 3). Chapters 4 and 5 of the thesis focus on methodological issues. The methodological insights from these chapters are il-lustrated using well-known models for studying the emergence of cooperative behavior. A prisoner's dilemma model is used in Chapter 4, and a model of quantity competition among firms is used in Chapter 5.

I will now discuss the contribution of each of the four studies presented in this thesis. As already mentioned, the studies in Chapters 4 and 5 focus on methodological issues. The studies in Chapters 2 and 3 are concerned with substantive economic issues. A purely theoretical approach is taken in these studies. No empirical or experimental

analyses are performed. Hence, the thesis contributes to methodological advancement as well as to increased theoretical understanding.

Chapter 2 of the thesis is concerned with a simple model of the learning behavior of boundedly rational agents. The model has been adopted from the computer science literature, where it is referred to as Q-learning (Watkins, 1989; Watkins & Dayan, 1992). The Q-learning model makes only very limited assumptions about the information available to agents and the cognitive abilities of agents. In Chapter 2, the Q-learning model is studied in the setting of a Cournot oligopoly market. The analysis of the model is performed partly mathematically and partly using computer simulations. The main contribution of Chapter 2 consists of showing that the Q-learning model is able to explain the emergence of cooperative behavior. Many similar learning models studied in the economic literature are unable to explain this phenomenon.

Chapter 3 of the thesis is concerned with a relatively recent explanation for the emergence of cooperative behavior. According to this explanation, cooperative behavior can be a consequence of evolutionary dynamics combined with local interaction among spatially distributed agents. This explanation was first proposed by the biologists Nowak and May (1992) and was introduced in the economic literature by Bergstrom and Stark (1993) and Eshel, Samuelson, and Shaked (1998). In the economic literature, the explanation has been studied mainly for agents that are located in a one-dimensional world and that can choose from only two actions (i.e., cooperation and defection). The advantage of this highly stylized setting is that it can be analyzed mathematically. The disadvantage is that it is unclear to what extent results derived in this setting generalize to other more complex settings. The contribution of Chapter 3 consists of studying some of these more complex settings, in particular settings in which agents are located in a two-dimensional world and settings in which agents can choose from more than two actions. The agents in the models studied in Chapter 3 are firms that compete with each other based on price. The models studied in Chapter 3 appear to be mathematically intractable, and most of the analysis is therefore performed using computer simulations. The main finding of Chapter 3 is that the emergence of cooperative behavior depends strongly on the amount of information available to agents. Agents tend to behave most cooperatively if they have only a very limited amount of information about their neigh-

bors. Contrary to earlier research, it is found that in some cases agents behave even less cooperatively than they would do according to the Nash equilibrium prediction.

The methodological issues studied in Chapters 4 and 5 of the thesis relate to the application of so-called evolutionary algorithms for modeling purposes in economic research. The application of evolutionary algorithms, and especially of genetic algorithms, is quite popular in the ACE approach. For a general introduction to genetic algorithms, see for example Mitchell (1996). For an extensive discussion of the application of genetic algorithms for economic modeling purposes, see Dawid (1996).

ACE researchers almost always use computer simulations to analyze genetic algorithm models (a notable exception is Dawid, 1996). In Chapter 4 of the thesis, it is shown how genetic algorithm models can be analyzed mathematically rather than using computer simulations. The proposed approach for mathematically analyzing genetic algorithm models relies on a mathematical technique that is frequently used in the EGT literature (e.g., Foster & Young, 1990; Kandori, Mailath, & Rob, 1993; Young, 1993). The main contribution of Chapter 4 consists of exploring the consequences of the similarities between evolutionary models studied in the ACE literature and evolutionary models studied in the EGT literature. As discussed earlier, the ACE and EGT research communities are quite separated from each other. The similarities between the models studied by the two communities suggest that the separation of the communities is somewhat artificial and may impede the further development of evolutionary modeling approaches. Chapter 4 illustrates this point by showing how the ACE community can benefit from a mathematical technique that is well-known in the EGT community.

Chapter 5 of the thesis is concerned with a problematic aspect of genetic algorithm models. These models typically rely on a so-called binary encoding of strategies. In Chapter 5, it is argued that such an encoding usually does not have a meaningful economic interpretation. It is also shown that the use of a binary encoding may lead to artifacts in the results of an analysis. Hence, the contribution of Chapter 5 consists of pointing out that in general it is not appropriate to use genetic algorithm models with a binary encoding of strategies for economic modeling purposes. The more general contribution of Chapter 5 consists of illustrating how the fairly high complexity of models in the ACE approach combined with a somewhat ad hoc justification may lead to results that have no economic significance and that are merely artifacts of the model. The anal-

ysis presented in Chapter 5 can be seen as an argument for simplifying the evolutionary models studied by ACE researchers. Simplifying these models would also illuminate the similarities with models studied by EGT researchers. Hence, it may help to reduce the somewhat artificial separation between the ACE and EGT research communities.

1.6   Outline of the thesis

This thesis consists of six chapters and is organized as follows. After this introductory chapter, there are four chapters that each present a separate study. As discussed above, Chapters 2 and 3 focus on the issue of the emergence of cooperative behavior. Chapter 2 presents an analysis of the Q-learning model, and Chapter 3 studies an evolutionary model with local interaction among spatially distributed agents. Chapters 4 and 5 focus on methodological issues concerning the application of evolutionary algorithms for economic modeling purposes. Chapter 4 points out how genetic algorithm models can be analyzed mathematically, and Chapter 5 criticizes the use of a binary encoding of strategies in genetic algorithm models. Chapter 6 provides a summary of the thesis.

Of the four studies presented in this thesis, two have been published in the peer-reviewed scientific literature. Chapter 2 has been published in the *Journal of Economic Dynamics and Control*, and Chapter 5 has been published in the *Journal of Evolutionary Economics*. Chapter 5 partly builds on a short contribution published in *Computational Economics* (Waltman & Van Eck, 2009). This contribution is not part of the thesis. Chapters 3 and 4 are currently under submission.

Chapter 2

Q-learning agents in a Cournot oligopoly model*

Abstract

Q-learning is a reinforcement learning model from the field of artificial intelligence. We study the use of Q-learning for modeling the learning behavior of firms in repeated Cournot oligopoly games. Based on computer simulations, we show that Q-learning firms generally learn to collude with each other, although full collusion usually does not emerge. We also present some analytical results. These results provide insight into the underlying mechanism that causes collusive behavior to emerge. Q-learning is one of the few learning models available that can explain the emergence of collusive behavior in settings in which there is no punishment mechanism and no possibility for explicit communication between firms.

2.1   Introduction

In this chapter, we model the learning behavior of firms in repeated Cournot oligopoly games using Q-learning. Q-learning is a reinforcement learning model of agent behavior originally developed in the field of artificial intelligence (Watkins, 1989). The model is based on two assumptions. First, for each possible strategy an agent is assumed

---

to remember some value indicating that strategy's performance. This value, referred to as a Q-value, is determined based on the agent's past experience with the strategy. Basically, the Q-value of a strategy is calculated as a weighted average of the payoffs obtained from the strategy in the past, where more recent payoffs are given greater weight. The second assumption of Q-learning states that, based on the Q-values, an agent probabilistically chooses which action to play. A logit model is used to describe the agent's choice behavior. The assumptions made by Q-learning can also be found in other reinforcement learning models. The models of Sarin and Vahid (1999, 2001) and Kirman and Vriend (2001) use ideas similar to Q-values, while the models of, for example, Mookherjee and Sopher (1997) and Camerer and Ho (1999) use a logit model to describe the way in which an agent chooses an action. Q-learning distinguishes itself from other reinforcement learning models in that it combines these two elements in a single model. In the economic literature, the combination of these elements has, to our knowledge, not been studied before.

In this chapter, we show that the use of Q-learning for modeling the learning behavior of firms in repeated Cournot oligopoly games generally leads to collusive behavior.[1] This is quite a remarkable result, since most Q-learning firms that we study do not have the ability to remember what happened in previous stage games. The firms therefore cannot use trigger strategies, that is, they cannot threaten to punish each other in case of non-collusive behavior. There is also no possibility for explicit communication between firms. However, despite the absence of punishment and communication mechanisms, collusive behavior prevails among firms. Apart from Q-learning, there are almost no models of the learning behavior of individual economic agents that predict collusive behavior in Cournot games. The only model of which we are aware is the so-called trial-and-error model studied by Huck, Normann, and Oechssler (2004a). Yet, experimental results (for an overview, see Huck, Normann, & Oechssler, 2004b) indicate that with two firms collusive behavior is quite common in Cournot games. Q-learning is one of the few models that does indeed predict this kind of behavior.

Models of the learning behavior of economic agents are studied both in agent-based computational economics (e.g., Tesfatsion, 2003, 2006) and in game theory (e.g., Fu-

---

[1]We refer to all firm behavior that results in a joint profit above the joint profit in the Nash equilibrium as collusive behavior. So, collusive behavior does not always mean that firms make the highest possible joint profit.

denberg & Levine, 1998). In agent-based computational economics the methodology of computer simulation is typically adopted, whereas in game theory the analytical methodology is predominant. It seems rather difficult to obtain analytical results for the behavior of multiple Q-learning agents interacting with each other in a strategic setting. In the field of artificial intelligence, it has been proven that under certain conditions a single Q-learning agent operating in a fixed environment learns to behave optimally (Watkins & Dayan, 1992). However, for settings with multiple agents learning simultaneously almost no analytical results are available. Given the difficulty of obtaining analytical results, most of the results that we present in this chapter are based on computer simulations. Analytical results are provided only for the special case in which Q-learning firms in a Cournot duopoly game can choose between exactly two production levels, the production level of the Nash equilibrium and some other, lower production level. The analytical results turn out to be useful for obtaining some basic intuition why Q-learning firms may learn to collude with each other.

The remainder of this chapter is organized as follows. First, in Sections 2.2 and 2.3, we provide an overview of related research and we introduce Q-learning. Then, in Section 2.4, we discuss the Cournot oligopoly model with which we are concerned throughout the chapter. We consider our computer simulations in Sections 2.5 and 2.6, in which we discuss the simulation setup and present the simulation results. We provide some analytical results in Section 2.7. Finally, in Section 2.8, we draw conclusions.

2.2   Related research

The literature on modeling the learning behavior of economic agents is quite large. Overviews of this literature are provided by Brenner (2006) and Duffy (2006). One can distinguish between individual learning models and social learning models (Vriend, 2000). In individual learning models an agent learns exclusively from its own experience, whereas in social learning models an agent also learns from the experience of other agents. Below, we first discuss the modeling of individual learning behavior, and we then consider the modeling of social learning behavior.

The two most important approaches to modeling individual learning behavior are belief-based learning and reinforcement learning. Examples of belief-based learning

models are Cournot adjustment and fictitious play (Fudenberg & Levine, 1998). These two models assume that an agent has the ability both to observe its opponents' action choices and to calculate best responses. In a Cournot oligopoly game, the models predict that firm behavior can converge only to the Nash equilibrium.

Reinforcement learning is based on a very simple idea: the higher the payoffs obtained from a strategy in the past, the more likely the strategy is to be played. Compared with belief-based learning models, reinforcement learning models make few assumptions about both the information available to an agent and the cognitive abilities of an agent. For example, in reinforcement learning an agent needs no information about its opponents' action choices or about the payoffs of the game. An agent is only assumed to have knowledge of the strategies that it can play and, after playing a strategy, of the payoff that it has obtained from that strategy. Reinforcement learning models are studied both in the economic literature and in the artificial intelligence literature (for an overview of the artificial intelligence literature on reinforcement learning, see Kaelbling, Littman, & Moore, 1996; Sutton & Barto, 1998). Q-learning is a reinforcement learning model that has been studied extensively by artificial intelligence researchers (e.g., Watkins, 1989; Watkins & Dayan, 1992) but that has received almost no attention from economists. In the economic literature, the reinforcement learning model studied by Roth and Erev (1995) and Erev and Roth (1998) is well-known. Bell (2001) performs a simulation study in which this model is compared with Q-learning. Some other reinforcement learning models have been proposed in the economic literature by Mookherjee and Sopher (1997), Sarin and Vahid (1999, 2001), and Kirman and Vriend (2001). These models are all in some way similar to Q-learning. We discuss their relationship with Q-learning in Section 2.3.

Some preliminary results on Q-learning behavior in a Cournot oligopoly game are reported by Kimbrough and Lu (2003). In their simulation study, the authors find a small tendency towards collusive behavior. In the present study, we extend the work of Kimbrough and Lu by analyzing Q-learning behavior in a Cournot game in more detail and by providing an explanation for the emergence of collusive behavior. Furthermore, in the artificial intelligence literature there are some papers in which Q-learning behavior in iterated prisoner's dilemmas or generalizations thereof is studied (Sandholm & Crites, 1996; Littman & Stone, 2001; Stimpson & Goodrich, 2003; Waltman & Kay-

mak, 2007). Whether Q-learning agents in an iterated prisoner's dilemma learn to cooperate with each other turns out to depend on the specific values of the prisoner's dilemma payoffs (Waltman & Kaymak, 2007).

The trial-and-error learning model studied by Huck et al. (2004a) also models individual learning behavior. Like reinforcement learning models, the trial-and-error model makes few assumptions about both the availability of information and the cognitive abilities of an agent. However, the underlying idea of the model is different. In a Cournot oligopoly game, the model assumes that a firm keeps increasing (decreasing) its production level as long as this results in a higher profit. As soon as profit falls, the firm starts decreasing (increasing) its production level. Like Q-learning, the trial-and-error model predicts collusive behavior in Cournot games.

A number of studies have investigated social learning behavior in Cournot oligopoly games. In a well-known study by Vega-Redondo (1997), an evolutionary model of firm behavior is analyzed. Vega-Redondo shows that the model predicts convergence of firm behavior to the Walrasian equilibrium. Alós-Ferrer (2004) and Bergin and Bernhardt (2005) extend the model of Vega-Redondo by providing firms with a memory. When firms have a memory, convergence to any outcome between the Walrasian equilibrium and the Nash equilibrium becomes possible (Alós-Ferrer, 2004) and even collusive behavior may emerge (Bergin & Bernhardt, 2005). Social learning behavior in Cournot games has also been investigated using models based on genetic algorithms (e.g., Arifovic, 1994; Vriend, 2000). Depending on the way in which genetic algorithms are applied, such models predict convergence of firm behavior to either the Walrasian equilibrium or the Nash equilibrium or some outcome in between. We further mention the work of Droste, Hommes, and Tuinstra (2002), in which social learning behavior in Cournot games is investigated using a model based on replicator dynamics.

Finally, Dixon (2000) and Oechssler (2002) use models with aspiration levels to investigate learning behavior in Cournot oligopoly games. In their models, a firm changes its production level only if its profit is below some aspiration level. The models predict collusive behavior in Cournot games.

2.3   Q-learning

In this chapter, Q-learning is applied as follows. An agent plays a repeated game. At the beginning of the stage game in period t, the agent's memory is in some state s_t. This state may be determined by, for example, the actions played by the agent and its opponents in the stage game in period $t - 1$. Taking into account the state of its memory, the agent chooses to play some action a_t. The choice of an action is made probabilistically based on the so-called Q-values of the agent. Playing action a_t results in some stage game payoff π_t that is obtained by the agent and in a transition of the state of the agent's memory from the old state s_t to some new state s_{t+1}. The agent uses the experience gained during the stage game to update its Q-values, thereby modifying the way in which it chooses actions in stage games in future periods.

For a formal discussion of Q-learning, let $Q_t(s, a)$ denote an agent's Q-value for state $s \in S$ and action $a \in A$ at the beginning of period t. The state space S and the action space A are assumed to be finite. The probability that in period t the agent chooses to play action a is given by

$$\Pr(a) = \frac{\exp(Q_t(s_t, a)/\beta)}{\sum_{a' \in A} \exp(Q_t(s_t, a')/\beta)}, \tag{2.1}$$

where s_t denotes the state of the agent's memory at the beginning of period t and the parameter $\beta > 0$ denotes the experimentation tendency. The larger the value of β, the higher the probability that the agent chooses to experiment, that is, chooses to play an action that does not have the highest Q-value. As β approaches zero, the probability that the agent chooses to experiment approaches zero too. In the artificial intelligence literature, action choice according to probabilities given by (2.1) is known as the Boltzmann exploration strategy (e.g., Kaelbling et al., 1996; Sandholm & Crites, 1996). Various other approaches to choosing actions have also been studied in the artificial intelligence literature. We model action choice behavior using (2.1) because this corresponds to a logit model, which is a quite commonly used model of choice behavior in the economic literature (e.g., McKelvey & Palfrey, 1995; Brock & Hommes, 1997; Mookherjee & So-

pher, 1997; Fudenberg & Levine, 1998; Camerer & Ho, 1999; Hofbauer & Sandholm, 2007)[2].

After the agent has played some action a_t in period t, the agent's Q-values are updated according to

$$Q_{t+1}(s,a) = \begin{cases} (1-\alpha)Q_t(s,a) + \alpha(\pi_t + \gamma \max_{a' \in A} Q_t(s_{t+1}, a')) & \text{if } s = s_t \text{ and } a = a_t, \\ Q_t(s,a) & \text{otherwise,} \end{cases}$$

(2.2)

where π_t denotes the stage game payoff obtained by the agent and the parameters $0 < \alpha \le 1$ and $0 \le \gamma < 1$ denote, respectively, the learning rate and the discount factor. The value of α determines the relative weight that is given to recent experience compared to older experience, while the value of γ indicates the time preference of the agent. The update rule in (2.2) has the appealing property that when there is only one learning agent (either because there is only one agent or because all other agents use fixed strategies), the update rule allows the agent, under certain conditions, to learn to behave optimally. This property has been proven by Watkins and Dayan (1992).

Unlike Q-learning, most reinforcement learning models studied in the economic literature (e.g., Roth & Erev, 1995; Mookherjee & Sopher, 1997; Erev & Roth, 1998; Sarin & Vahid, 1999, 2001) do not consider the possibility that an agent has a memory for remembering past events. In these models, it is not possible for an agent to learn a strategy in which the choice of an action in the current stage game depends on what happened in previous stage games. In this chapter, we consider both agents with a memory and agents without a memory. For an agent without a memory, (2.1) and (2.2) simplify to, respectively,

$$\Pr(a) = \frac{\exp(Q_t(a)/\beta)}{\sum_{a' \in A} \exp(Q_t(a')/\beta)},$$

(2.3)

---

[2]Note, however, that in each of these papers logit models are used in a somewhat different context. McKelvey and Palfrey, for example, use logit models as the basis of the logit equilibrium concept that they introduce in their paper. Brock and Hommes use logit models to model the way in which agents choose between different predictors based on publicly available information on each predictor's past performance. Hofbauer and Sandholm study settings in which there are a large number of agents and in which each of the agents occasionally changes its action according to, for example, a logit model.

and

$$Q_{t+1}(a) = \begin{cases} (1-\alpha)Q_t(a) + \alpha\pi_t & \text{if } a = a_t, \\ Q_t(a) & \text{otherwise.} \end{cases} \qquad (2.4)$$

Note that an agent without a memory cannot take into account the consequences of the action it plays in the current stage game on the payoffs it obtains in future stage games. For such an agent, the discount factor γ in (2.2) must therefore equal zero.

Sarin and Vahid (1999) and Kirman and Vriend (2001) propose reinforcement learning models that use the same update rule as Q-learning without a memory. The difference between Q-learning and the model of Sarin and Vahid is that in the latter model there is no experimentation, so that an agent always chooses the action from which it expects to obtain the highest payoff. Sarin and Vahid also propose a variant of their model in which an agent experiments depending on its 'state of mind or mood'. However, they do not seem to study this variant further in other papers (e.g., Sarin & Vahid, 2001). In the model of Kirman and Vriend, agents do experiment, but the probabilities with which the various actions are chosen are not the same as in Q-learning. Q-learning without a memory is also related to the learning model proposed by Mookherjee and Sopher (1997). In a similar way as in Q-learning without a memory, Mookherjee and Sopher use a logit model to describe an agent's choice behavior. Since Mookherjee and Sopher do not specify what kind of update rule to use, Q-learning without a memory can in fact be regarded as a special case of their model.

2.4   Cournot oligopoly model

We consider a simple Cournot oligopoly model with the following characteristics: the number of firms is fixed, firms produce perfect substitutes, the demand function is linear, firms have identical cost functions, and marginal cost is constant. The inverse demand function is given by

$$p = \max\left(u - v\sum_{i=1}^{n} q_i, 0\right), \qquad (2.5)$$

where n denotes the number of firms, p denotes the market price, q_i denotes firm i's production level, and $u > 0$ and $v > 0$ denote two parameters. Firm i's total cost equals

$$c_i = wq_i \qquad \text{for } i = 1, \ldots, n, \tag{2.6}$$

where the parameter w denotes a firm's constant marginal cost and satisfies $0 \leq w < u$. It follows from (2.5) and (2.6) that firm i's profit is given by

$$\pi_i = pq_i - c_i = q_i \max\left(u - w - v\sum_{i'=1}^{n} q_{i'}, -w\right) \qquad \text{for } i = 1, \ldots, n. \tag{2.7}$$

The Nash equilibrium of a Cournot model is obtained if each firm chooses the production level that maximizes its profit given the production levels of its competitors. Hence, in the Nash equilibrium $\partial\pi_i/\partial q_i = 0$ for $i = 1, \ldots, n$. In the above Cournot model, this implies that firms' joint production level in the Nash equilibrium is given by

$$q^* = \frac{(u-w)n}{v(n+1)}. \tag{2.8}$$

Consequently, firms' joint profit in the Nash equilibrium equals

$$\pi^* = \frac{(u-w)^2 n}{v(n+1)^2}. \tag{2.9}$$

Although in the Nash equilibrium firms individually maximize their profit, they do not maximize their joint profit. Firms maximize their joint profit in the collusive equilibrium, in which they collectively behave as a single monopolist. In the above Cournot model, firms jointly produce a quantity of $(u-w)/2v$ in the collusive equilibrium, which results in a joint profit of $(u-w)^2/4v$. In the collusive equilibrium of a Cournot model, firms produce a smaller quantity than the quantity that maximizes their individual profit. Firms therefore have an incentive to increase their production level. For this reason, a collusive equilibrium is unstable and is not a Nash equilibrium. However, things are different in repeated Cournot games, in which firms may be interested in maximizing their long-term profits. If firms play a Cournot game repeatedly and remember what happened in previous stage games, it may be possible to sustain collusion. Trigger

strategies may then be used to support collusive behavior, and collusive behavior may constitute a Nash equilibrium of the repeated game.

In addition to the Nash equilibrium and the collusive equilibrium, another outcome that may be obtained in a Cournot model is the Walrasian equilibrium. As discussed in Section 2.2, this equilibrium is sometimes encountered in studies on social learning behavior in Cournot models. The Walrasian equilibrium is obtained if firms are not aware of their influence on the market price and therefore behave as price takers. In the above Cournot model, firms' joint production level in the Walrasian equilibrium equals $(u - w)/v$. This production level results in zero profits for all firms.

2.5   Setup of the computer simulations

In this chapter, we focus on the long-run behavior of Q-learning agents when the probability of experimentation approaches zero. In this respect, the approach that we take is similar to the approach that is typically taken to analyze evolutionary game-theoretic learning models (e.g., Vega-Redondo, 1997; Alós-Ferrer, 2004; Bergin & Bernhardt, 2005). We further focus on settings in which the learning behavior of all agents is modeled using Q-learning. An alternative would be to consider settings in which the learning behavior of only one agent is modeled using Q-learning and in which all other agents use fixed strategies. However, such settings are less interesting to study. This is because a single Q-learning agent operating in a fixed environment is, under certain conditions, guaranteed to learn to behave optimally (Watkins & Dayan, 1992). This means that when a Q-learning agent competes with agents that use fixed strategies, the Q-learning agent will simply learn a best response to the strategies of the other agents. The settings on which we focus in this chapter, that is, settings with multiple Q-learning agents learning simultaneously, are more interesting to study, because for such settings analytical results are generally not available. Because of the difficulty of obtaining analytical results, most of the results that we present are based on computer simulations. We now discuss the setup of these simulations.

In the simulations, the Cournot oligopoly model introduced in the previous section was used. The values of the parameters u and v in the inverse demand function were, respectively, 40 and 1. The parameter w, which denotes a firm's constant marginal cost,

had a value of 4. Simulations were performed for various values for the number of firms n.

In each simulation run, firms played a repeated Cournot oligopoly game that lasted for one million periods. The learning behavior of firms was modeled using Q-learning. Three types of firms were considered in the simulations: firms without a memory, myopic firms with a memory, and non-myopic firms with a memory. All three types of firms had to choose their production level between 0 and 40. Only integer quantities were allowed. Simulations were performed for various values for the learning rate α. During a simulation run, the experimentation tendency β was gradually decreased over time according to

$$\beta(t) = 1000 \cdot 0.99999^t. \tag{2.10}$$

In this way, the probability of experimentation was almost one at the beginning of a simulation run and almost zero at the end. In other studies on Q-learning (e.g., Sandholm & Crites, 1996), β is decreased in a similar way. At the beginning of a simulation run, firms' Q-values were initialized to zero. Firms with a memory were able to remember their own production level in the previous period as well as their competitors' joint production level in the previous period. For myopic firms with a memory the discount factor γ had a value of zero, while for non-myopic firms with a memory it had a value of 0.9.

2.6   Results of the computer simulations

In this section, we present the results of the computer simulations that we performed. We first consider the simulations with firms that did not have a memory, and we then consider the simulations with firms that did have a memory.

Simulations with firms that did not have a memory were performed for various values for both the number of firms n and the learning rate α. For each combination of values for n and α, Table 2.1 shows firms' joint quantity produced and joint profit. Since we focus on the long-run behavior of firms when the probability of experimentation approaches zero, the quantities and profits in Table 2.1 were calculated by averaging firms' joint quantity produced and joint profit over the last 100 periods of a simulation run. Moreover, because the outcomes of a simulation run depend on the random num-

Table 2.1: Results of computer simulations with firms that did not have a memory.

		Nash	$\alpha = 0.05$		$\alpha = 0.25$		$\alpha = 0.50$		$\alpha = 1.00$	
$n = 2$	quantity	24.0	22.8	(1.3)	21.2	(1.4)	20.8	(1.2)	20.8	(1.4)
	profit	288.0	299.1	(11.6)	312.0	(10.2)	314.7	(6.2)	314.3	(7.0)
$n = 3$	quantity	27.0	25.1	(1.6)	22.0	(1.8)	21.5	(1.9)	22.1	(1.9)
	profit	243.0	270.7	(22.9)	304.6	(14.5)	307.8	(14.3)	303.7	(16.6)
$n = 4$	quantity	28.8	26.3	(1.8)	22.6	(1.9)	22.1	(2.4)	22.9	(2.6)
	profit	207.4	252.1	(29.8)	299.0	(18.7)	301.4	(19.2)	293.2	(25.8)
$n = 5$	quantity	30.0	27.6	(1.6)	23.2	(1.8)	22.2	(2.2)	23.3	(2.5)
	profit	180.0	229.3	(30.0)	294.1	(17.3)	301.1	(19.2)	290.2	(28.7)
$n = 6$	quantity	30.9	28.3	(1.5)	23.3	(2.2)	22.6	(2.6)	23.1	(3.1)
	profit	158.7	215.4	(32.2)	290.7	(23.5)	296.3	(27.1)	289.1	(34.8)

bers that are used, 100 simulation runs with different random numbers were carried out for each combination of values for n and α. Table 2.1 shows the mean of the outcomes of these 100 simulation runs. The corresponding standard deviation is reported within parentheses. For each value for n, firms' joint quantity produced and joint profit in the Nash equilibrium, calculated using (2.8) and (2.9), are also reported in Table 2.1. Firms' joint quantity produced and joint profit in the collusive equilibrium do not depend on n and are equal to, respectively, 18 and 324 (see Section 2.4).

Table 2.1 shows that for all combinations of values for n and α the average outcome that emerged in the simulation runs was somewhere in between the Nash equilibrium and the collusive equilibrium. For each combination of values for n and α, the mean of firms' joint quantity produced was significantly lower than firms' joint quantity produced in the Nash equilibrium ($p < 0.0001$), while the mean of firms' joint profit was significantly higher than firms' joint profit in the Nash equilibrium ($p < 0.0001$). So, on average firms learned to collude with each other. Full collusion usually did not emerge, since firms usually did not learn to make the highest possible joint profit. It can further be seen in Table 2.1 that on average firms' joint quantity produced increased as the number of firms n increased. However, since firms' joint quantity produced in the Nash equilibrium also increases as n increases, a substantial degree of collusion remained even for larger values for n. The same observation can be made if firms' joint profit rather than firms' joint quantity produced is considered. Now consider the effect of the learning rate α, which determined the relative weight that firms gave to recent experi-

Table 2.2: Results of computer simulations with firms that had a memory.

		Nash	myopic ($\gamma = 0.0$)		non-myopic ($\gamma = 0.9$)	
$n = 2$	quantity	24.0	20.8	(0.9)	19.6	(1.2)
	profit	288.0	314.2	(7.3)	318.0	(6.8)
$n = 3$	quantity	27.0	22.9	(1.3)	21.5	(1.7)
	profit	243.0	297.3	(15.8)	304.8	(18.5)
$n = 4$	quantity	28.8	24.1	(1.4)	23.8	(1.7)
	profit	207.4	284.2	(19.8)	277.5	(41.3)
$n = 5$	quantity	30.0	24.4	(1.4)	23.6	(2.1)
	profit	180.0	280.1	(20.5)	271.6	(75.8)
$n = 6$	quantity	30.9	24.7	(1.6)	21.9	(2.0)
	profit	158.7	274.5	(29.2)	288.5	(51.0)

ence compared to older experience. As can be seen in Table 2.1, a value of 0.05 for α resulted in a significantly lower degree of collusion than a value of 0.25 or higher. However, even for $\alpha = 0.05$ the degree of collusion was significant. For α equal to 0.25, 0.50, and 1.00, the differences in the degree of collusion were quite small. So, the relative weight that firms gave to recent experience compared to older experience did not have a very large effect on the degree of collusion. A substantial negative effect on the degree of collusion was found only when firms gave a rather low weight to recent experience compared to older experience.

Simulations with firms that had a memory were performed for various values for the number of firms n. In addition, both myopic and non-myopic firms were considered in the simulations. A value of 0.50 was used for the learning rate α. Table 2.2 shows the results of the simulations. The results in Table 2.2 were calculated in the same way as the results in Table 2.1. Comparing the results in the two tables, it can be seen that the results obtained for firms with a memory are quite similar to the results obtained for firms without a memory. Like firms without a memory, firms with a memory on average learned to collude with each other. Full collusion usually did not emerge. The degree of collusion seems to be somewhat lower for firms with a memory, both for myopic and for non-myopic firms, than for firms without a memory, but the difference is not very large. Between myopic and non-myopic firms with a memory, no clear difference can be observed in the degree of collusion.

2.7   Analytical results

In the previous section, we presented simulation results showing that the use of Q-learning for modeling the learning behavior of firms in a Cournot oligopoly game generally leads to collusive behavior. This turned out to be the case not only for firms with a memory but also for firms without a memory. This is quite remarkable, since firms without a memory cannot use trigger strategies, that is, they cannot threaten to punish each other in case of non-collusive behavior. So, collusive behavior prevails among Q-learning firms without a memory even though there is no punishment mechanism and no possibility for explicit communication between the firms. Interestingly, apart from Q-learning, there are very few learning models that predict collusive behavior in such a setting. In this section, we analyze why Q-learning results in collusive behavior. To do so, we make two simplifying assumptions. First, we assume that there are only two firms in the market. And second, we assume that firms can choose between only two production levels, the production level of the Nash equilibrium and some other, lower production level. Under these two assumptions, a Cournot game reduces to a prisoner's dilemma game. Moreover, the behavior of Q-learning firms becomes analytically tractable, and it becomes clear why Q-learning firms may learn to collude with each other.

Consider the Cournot oligopoly model introduced in Section 2.4. Let there be two firms in the market, that is, let $n = 2$. It follows from the results in Section 2.4 that in the Nash equilibrium each firm produces a quantity of $(u - w)/3v$ while in the symmetric collusive equilibrium each firm produces a quantity of $(u - w)/4v$. The following theorem provides sufficient conditions for the emergence of collusive behavior among Q-learning firms without a memory.

Theorem 2.1. *Consider an infinitely repeated Cournot duopoly game based on the Cournot model introduced in Section 2.4 with the number of firms n equal to 2. Let the firms' learning behavior be described by Q-learning, and assume that the firms do not have a memory. Assume that the firms can choose between two production levels, denoted by q_C and q_N, that satisfy*

$$(u - w)/4v < q_C < q_N = (u - w)/3v. \tag{2.11}$$

Let π_{CN}, π_{NN}, π_{CC}, and π_{NC} denote a firm's profit, respectively, if the firm produces q_{C} and its competitor produces q_{N}, if both the firm and its competitor produce q_{N}, if both the firm and its competitor produce q_{C}, and if the firm produces q_{N} and its competitor produces q_{C}. Let the learning rate α satisfy

$$\frac{\pi_{\mathrm{NN}} - \pi_{\mathrm{CN}}}{\pi_{\mathrm{CC}} - \pi_{\mathrm{CN}}} < \alpha < 1. \tag{2.12}$$

Let each firm's Q-value of producing q_{C} be initialized to a value strictly between π_{CN} and π_{CC}, and let each firm's Q-value of producing q_{N} be initialized to a value strictly between π_{NN} and π_{NC}. Then, in the limit as the experimentation tendency β approaches zero, the proportion of time in which both firms produce q_{C} equals one.

The proof of the theorem is provided in Appendix 2.A. The basic intuition of the proof is as follows. Consider first what happens when the firms never experiment, that is, when from the production levels q_{C} and q_{N} the firms always choose the one with the higher Q-value. Two special states can be distinguished: a 'collusive state' in which under the assumption of no experimentation both firms keep producing q_{C} forever and a 'Nash state' in which under the assumption of no experimentation both firms keep producing q_{N} forever. It can be shown that without experimentation the firms will always end up in one of these states, in which they will then remain forever. Consider now what happens when experimentation is introduced. Let the experimentation tendency β have a value close to zero, so that there is only a very low probability of experimentation. With experimentation, the firms will no longer keep producing the same quantity forever, and two kinds of transitions will start taking place: transitions from the collusive state to the Nash state and transitions the other way around. Since the probability of experimentation is very low, it will usually take a long time before a transition from one state to the other occurs. In which of the two states, the collusive state or the Nash state, the firms will spend most of their time depends on the relative likelihood of the two kinds of transitions. If transitions from the collusive state to the Nash state are more likely than transitions the other way around, the firms will spend most of their time in the Nash state. Conversely, if transitions from the collusive state to the Nash state are less likely than transitions the other way around, the firms will spend most of their time in the collusive state. It turns out that for β close to zero transitions from the collusive state to the Nash state are less likely than transitions the other way around. This can be

seen as follows. A transition from the collusive state to the Nash state can be shown to require at least one period in which one of the firms experiments (that is, produces q_N rather than q_C). Assuming that the learning rate α satisfies (2.12), it can also be shown that for a transition from the Nash state to the collusive state one period in which both firms experiment (that is, produce q_C rather than q_N) will usually be sufficient. The probability with which a firm experiments is given by (2.1) and depends on the difference between the firm's Q-values of producing q_C and q_N. The larger this difference, the lower the probability of experimentation. It can be shown that in the collusive state the difference between a firm's Q-values will usually be equal to approximately $\pi_{CC} - \pi_{NN}$, while in the Nash state the difference will usually be no larger than $\pi_{NN} - \pi_{CN}$. Because it can also be shown that $\pi_{CC} - \pi_{NN} > 2(\pi_{NN} - \pi_{CN})$, the probability of experimentation will usually be lower in the collusive state than in the Nash state. It even follows from (2.1) that the probability of one of the firms experimenting in the collusive state will usually be lower than the probability of both firms experimenting simultaneously in the Nash state. For that reason, transitions from the collusive state to the Nash state are less likely than transitions the other way around. As a consequence, the firms will spend most of their time in the collusive state.

The above informal argument provides the basic intuition why Q-learning firms may learn to collude with each other. There turn out to be two opposing forces at work. One force is due to firms trying to optimize their behavior in a given situation. This force is directed towards the Nash equilibrium. The other force is due to the possibility that firms experiment simultaneously and in that way discover the advantages of collusion. This force is directed towards collusive behavior. Somewhat surprisingly, the second force is typically stronger than the first one. This then leads to firms spending most of their time colluding with each other.

It is interesting to note that the argument depends crucially on the specific way in which the probability of experimentation is determined in Q-learning. In a setting with only two actions, the probability of experimentation depends on the difference between the actions' Q-values. The larger this difference, the lower the probability of experimentation. In many other learning models, for example in evolutionary models (e.g., Vega-Redondo, 1997) and in genetic algorithm models (e.g., Vriend, 2000), the probability of experimentation (sometimes referred to as the probability of mutation) has a

fixed value and does not depend on past experience. With a fixed probability of experimentation, the informal argument provided above no longer holds. (This is because with a fixed probability of experimentation the probability of one of the firms experimenting in the collusive state is higher than the probability of both firms experimenting simultaneously in the Nash state.) Computer simulations indicate that a fixed probability of experimentation typically leads to firms spending most of their time in the Nash state rather than in the collusive state (see also Waltman & Kaymak, 2007). Apparently, the way in which the probability of experimentation is determined can have a large effect on the learning behavior of economic agents. According to Brenner (2006), psychological research indicates that people take into account past experience when choosing their actions. A learning model like Q-learning, in which the probability of experimentation depends on past experience, therefore seems more in line with psychological findings than learning models with a fixed probability of experimentation.

Finally, we note that the informal argument for the emergence of collusive behavior that we provided above holds not only for two firms but for any number of firms. Although we do not have a formal proof, this suggests that the emergence of collusive behavior is always possible, regardless of the number of firms. This would be a somewhat counterintuitive result, since collusion is generally believed to become much more difficult, if not impossible, when the number of firms increases. However, it would be in line with the simulation results discussed in Section 2.6, which indicate a substantial degree of collusion for various numbers of firms.

2.8   Conclusions

We have studied the use of Q-learning for modeling the learning behavior of firms in repeated Cournot oligopoly games. Q-learning, which belongs to the family of reinforcement learning models, combines two elements that, individually, can also be found in other models of the reinforcement learning type. On the one hand, the way in which the performance of a strategy is measured is similar to the way in which this is done in the models of Sarin and Vahid (1999, 2001) and Kirman and Vriend (2001). On the other hand, the use of a logit model to describe an agent's choice behavior is fairly common and can also be found in the models of, for example, Mookherjee and Sopher

(1997) and Camerer and Ho (1999). Q-learning combines both elements in a single model.

Based on computer simulations, we have shown that Q-learning firms generally learn to collude with each other in Cournot oligopoly games, although full collusion usually does not emerge, that is, firms usually do not learn to make the highest possible joint profit. Interestingly, our results hold not only for firms with a memory but also for firms without a memory. The latter firms do not have the ability to remember the quantities produced by their competitors in past periods. Although these firms cannot use trigger strategies to sustain collusion, they still learn to collude with each other. Apart from Q-learning, there are very few learning models that predict collusive behavior among firms without a memory. The analytical results that we have obtained for Cournot duopoly games with two production levels provide some insight into why Q-learning firms may learn to collude with each other. The emergence of collusive behavior seems to depend crucially on the specific way in which the probability of experimentation is determined in Q-learning. More specifically, it seems crucial that in Q-learning the probability of experimentation does not have a fixed value, as is the case in many other learning models, but depends on an agent's past experience.

Whether Q-learning provides a good description of the learning behavior of economic agents is, of course, an empirical question. We have not considered this question in the present study. However, there is at least some correspondence between Q-learning behavior in Cournot oligopoly games and results from laboratory experiments (for an overview, see Huck et al., 2004b). Experimental results indicate that collusive behavior is quite common in Cournot duopoly games. Unlike most other learning models, Q-learning does indeed predict collusive behavior in these games. However, Q-learning also predicts a substantial degree of collusion in Cournot games with more than two firms. This does not match experimental results. In experimental studies, firm behavior usually turns out to be quite close to the Nash equilibrium when the number of firms is larger than two.

2.A Proof of Theorem 2.1

In this appendix, we provide a proof of Theorem 2.1. As a shorthand expression, below we sometimes write that the firms produce, for example, (q_C, q_N). With this we mean that one firm, firm 1, produces q_C while the other firm, firm 2, produces q_N.

Proof. Proof of Theorem 2.1 It follows from (2.7) and (2.11) that

$$\pi_{CN} = (-3vq_C + 2u - 2w)q_C/3, \tag{2.13}$$

$$\pi_{NN} = (u - w)^2/9v, \tag{2.14}$$

$$\pi_{CC} = (-2vq_C + u - w)q_C, \tag{2.15}$$

$$\pi_{NC} = (u - w)(-3vq_C + 2u - 2w)/9v, \tag{2.16}$$

and that

$$\pi_{CN} < \pi_{NN} < \pi_{CC} < \pi_{NC}. \tag{2.17}$$

For $i = 1, 2$ and $t = 0, 1, \ldots$, let $Q_{i,t}^C$ and $Q_{i,t}^N$ denote, respectively, firm i's Q-value of producing q_C in period t and firm i's Q-value of producing q_N in period t. The theorem assumes that $Q_{1,0}^C, Q_{2,0}^C \in (\pi_{CN}, \pi_{CC})$ and $Q_{1,0}^N, Q_{2,0}^N \in (\pi_{NN}, \pi_{NC})$. It then follows from (2.4) that $Q_{1,t}^C, Q_{2,t}^C \in (\pi_{CN}, \pi_{CC})$ and $Q_{1,t}^N, Q_{2,t}^N \in (\pi_{NN}, \pi_{NC})$ also holds for $t = 1, 2, \ldots$. Firm i is said to experiment in period t if it produces q_C while $Q_{i,t}^C < Q_{i,t}^N$ or if it produces q_N while $Q_{i,t}^N < Q_{i,t}^C$. For $t = 0, 1, \ldots$, let $X_t \in \{0, 1, 2, 3\}$ denote the state of the learning process in period t. Consider the following three conditions on the firms' Q-values:

$$Q_{1,t}^C, Q_{2,t}^C \leq \pi_{NN}, \tag{2.18}$$

$$Q_{1,t}^N, Q_{2,t}^N < \pi_{NN} + \varepsilon < \pi_{CC} - \varepsilon < Q_{1,t}^C, Q_{2,t}^C, \tag{2.19}$$

$$\max\left(Q_{1,t}^N, \pi_{CC} - \varepsilon\right) < Q_{1,t}^C \text{ and } \max\left(Q_{2,t}^N, \pi_{CC} - \varepsilon\right) < Q_{2,t}^C$$
$$\text{and } \max\left(Q_{1,t}^N, Q_{2,t}^N\right) \geq \pi_{NN} + \varepsilon. \tag{2.20}$$

In these conditions, ε denotes a constant that satisfies

$$0 < \varepsilon < \min((2\pi_{CN} - 3\pi_{NN} + \pi_{CC})/4, (1 - \alpha)\pi_{CN} - \pi_{NN} + \alpha\pi_{CC}), \tag{2.21}$$

where the positivity of the first argument of the min function follows from (2.11), (2.13), (2.14), and (2.15) and the positivity of the second argument of the min function follows from (2.12) and (2.17). Let the state of the learning process be determined by the above three conditions in the following way: $X_t = 1$ if (2.18) is satisfied, $X_t = 2$ if (2.19) is satisfied, $X_t = 3$ if (2.20) is satisfied, and $X_t = 0$ otherwise.

To prove the theorem, three properties of the learning process will be used. Each property is proven separately below.

Property 2.1. As the experimentation tendency β approaches zero, the proportion of time in which the learning process is in state 0 approaches zero.

Due to this property, state 0 need not be considered further. Consequently, a transition from state j to state k, where $j, k \in \{1, 2, 3\}$ and $j \neq k$, is said to occur between periods t and t', where $t < t'$, if $X_t = j$, $X_{t+1} = \ldots = X_{t'-1} = 0$, and $X_{t'} = k$.

Property 2.2. As the experimentation tendency β approaches zero, the probability that a transition from state 1 leads to state 2 approaches one.

Property 2.3. Consider the ratio between the average time it takes in state 1 before a transition to another state occurs and the average time it takes in state 2 before a transition to another state occurs. This ratio approaches zero as the experimentation tendency β approaches zero.

It follows from the last two properties that the ratio between the time in which the learning process is in state 1 and the time in which the learning process is in state 2 approaches zero as β approaches zero. Together with Property 2.1, this implies that as β approaches zero the proportion of time in which the learning process is in state 0 or 1 approaches zero. If the learning process is not in state 0 or 1, it will be in state 2 or 3. If the learning process is in one of the latter states, the probability that the firms produce (q_C, q_C) approaches one as β approaches zero. It follows that as β approaches zero the proportion of time in which the firms produce (q_C, q_C) approaches one. $\qquad\square$

Proof. Proof of Property 2.1 It will first be shown that if the learning process is in state 0 and the firms never experiment, the learning process will leave state 0 within a finite number of periods. Under the assumption that the firms never experiment, the following three observations can be made:

(1) *As long as the learning process is in state* 0*, the firms will not keep producing* (q_N, q_N) *forever.*

This can be seen as follows. Each firm's Q-value of producing q_N will decrease if the firms produce (q_N, q_N). Moreover, as long as the learning process is in state 0, at least one firm, say firm 1, will have a Q-value of producing q_C that is larger than π_{NN} (otherwise the learning process would be in state 1). If the firms produce (q_N, q_N) for a certain finite number of consecutive periods while the learning process is in state 0, firm 1's Q-value of producing q_N will become smaller than its Q-value of producing q_C. Due to the assumption of no experimentation, firm 1 will then produce q_C in the next period.

(2) *If the learning process is in state* 0 *and the firms produce* (q_C, q_C)*, the learning process will leave state* 0 *within a finite number of periods.*

This can be seen as follows. Due to the assumption of no experimentation, in some period t the firms will produce (q_C, q_C) only if $Q_{1,t}^C \geq Q_{1,t}^N$ and $Q_{2,t}^C \geq Q_{2,t}^N$. If the firms produce (q_C, q_C) in period t, it follows that $Q_{1,t+1}^C > Q_{1,t+1}^N$ and $Q_{2,t+1}^C > Q_{2,t+1}^N$. As a consequence, the firms will produce (q_C, q_C) another time in period $t + 1$ and will in fact keep producing it forever. After a finite number of periods, each firm's Q-value of producing q_C will then be larger than $\pi_{CC} - \varepsilon$ and the learning process will have reached state 2 or 3.

(3) *Regardless of the state of the learning process, the firms will produce* (q_C, q_N) *and* (q_N, q_C) *at most a finite number of times.*

To see this, consider one of the two quantity pairs, say (q_C, q_N). Firm 1's Q-value of producing q_C will decrease if the firms produce (q_C, q_N) and increase if the firms produce (q_C, q_C). If firm 1's Q-value of producing q_C decreases a certain finite number of times, denoted by m, and does not increase in between, it will no longer be larger than π_{NN}. It can be seen that the firms will produce (q_C, q_N) at most m times. Two cases have to be distinguished. In the first case, the firms produce (q_C, q_C) before they have produced (q_C, q_N) m times. As shown above, due to the assumption of no experimentation, the firms will then keep producing (q_C, q_C) forever. In the second case, the firms produce (q_C, q_N) m times and do not produce (q_C, q_C) in between. After the firms have produced (q_C, q_N) m times, firm 1's Q-value of producing q_C will no longer be larger than π_{NN}. Firm 1 will then

keep producing q_N forever. This is because firm 1's Q-value of producing q_N is always larger than π_{NN} and because, by assumption, firm 1 will never experiment. In both the first and the second case, the firms produce (q_C, q_N) no more than m times.

So, under the assumption of no experimentation, as long as the learning process is in state 0, the firms will not keep producing (q_N, q_N) forever and they will produce the other three quantity pairs at most a finite number of times. It follows that if the learning process is in state 0 and the firms never experiment, the learning process will leave state 0 within a finite number of periods. The other three states do not have this property. If the learning process is in state 1 and the firms never experiment, the firms will keep producing (q_N, q_N) forever and the learning process will never leave its current state. Similarly, if the learning process is in state 2 or 3 and the firms never experiment, the firms will keep producing (q_C, q_C) forever. Again, the learning process will never leave its current state.

As β approaches zero, the probability that a firm experiments approaches zero. Using the results obtained above, it can be seen that if the probability of experimentation approaches zero and the learning process is in state 0, the probability that within a finite number of periods another state is reached approaches one. Similarly, it can be seen that if the probability of experimentation approaches zero and the learning process is in state 1, 2, or 3, the probability that within a finite number of periods another state is reached approaches zero. It follows that as β approaches zero, the proportion of time in which the learning process is in state 0 approaches zero. □

Proof. Proof of Property 2.2 When the learning process is in state 1, two cases can be distinguished. If $Q_{1,t}^N, Q_{2,t}^N < \pi_{NN} + \varepsilon$, the learning process is said to be in state 1a, otherwise it is said to be in state 1b. The following three observations can now be made:

(1) *If the learning process is in state* 1*, it will leave that state only if the firms produce* (q_C, q_C).
    This is because in order to leave state 1, for at least one of the firms the Q-value of producing q_C must increase. Producing (q_C, q_C) is the only way in which a firm's Q-value of producing q_C can increase.

(2) *If the learning process is in state* $1a$ *and the firms produce* $(q_{\mathrm{C}}, q_{\mathrm{C}})$*, the probability that a transition to state* 2 *occurs approaches one as* β *approaches zero.*

This can be seen as follows. Due to (2.21), $\varepsilon < (1 - \alpha)\pi_{\mathrm{CN}} - \pi_{\mathrm{NN}} + \alpha\pi_{\mathrm{CC}}$. Therefore, if in some period t learning process is in state 1a and the firms produce $(q_{\mathrm{C}}, q_{\mathrm{C}})$, it follows that $Q^{\mathrm{C}}_{1,t+1}, Q^{\mathrm{C}}_{2,t+1} > \pi_{\mathrm{NN}} + \varepsilon$ and hence that $Q^{\mathrm{C}}_{1,t+1} > Q^{\mathrm{N}}_{1,t+1}$ and $Q^{\mathrm{C}}_{2,t+1} > Q^{\mathrm{N}}_{2,t+1}$. Consequently, the probability that the firms produce $(q_{\mathrm{C}}, q_{\mathrm{C}})$ another time in period $t + 1$ approaches one as β approaches zero. As long as the firms keep producing $(q_{\mathrm{C}}, q_{\mathrm{C}})$, the probability that they produce it again in the next period approaches one as β approaches zero. After the firms have produced $(q_{\mathrm{C}}, q_{\mathrm{C}})$ for some finite number of consecutive periods, each firm's Q-value of producing q_{C} will be larger than $\pi_{\mathrm{CC}} - \varepsilon$ and the learning process will have reached state 2.

(3) *Suppose that in some period* t *the learning process is in state* 1 *and the firms produce* $(q_{\mathrm{C}}, q_{\mathrm{C}})$*. The probability that the learning process was in state* $1a$ *in period* t *then approaches one as* β *approaches zero.*

This can be seen as follows. Let Δt denote the number of consecutive periods in which $(q_{\mathrm{N}}, q_{\mathrm{N}})$ must be produced so that a firm's Q-value of producing q_{N} decreases from π_{NC} to a value smaller than $\pi_{\mathrm{NN}} + \varepsilon$. Let $t' = t - \Delta t$. For the moment, assume that the learning process has been in state 1 all the time between periods t' and t. This assumption implies that the firms have never produced $(q_{\mathrm{C}}, q_{\mathrm{C}})$ between periods t' and $t - 1$. If the firms have produced $(q_{\mathrm{N}}, q_{\mathrm{N}})$ all the time between periods t' and $t - 1$, the learning process would have been in state 1a in period t. If, on the other hand, the firms have produced either $(q_{\mathrm{C}}, q_{\mathrm{N}})$ or $(q_{\mathrm{N}}, q_{\mathrm{C}})$ at least once between periods t' and $t - 1$, the learning process could have been in either state 1a or state 1b in period t. Given that the learning process was in state 1 in period t', the probability that the firms have produced $(q_{\mathrm{N}}, q_{\mathrm{N}})$ all the time between periods t' and $t - 1$ approaches one as β approaches zero. Furthermore, given the firms' Q-values in period t', the probability that the firms produce $(q_{\mathrm{C}}, q_{\mathrm{C}})$ in period t is higher if the firms have produced $(q_{\mathrm{N}}, q_{\mathrm{N}})$ all the time between periods t' and $t - 1$ than if they have produced either $(q_{\mathrm{C}}, q_{\mathrm{N}})$ or $(q_{\mathrm{N}}, q_{\mathrm{C}})$ at least once between these periods. It now follows that the probability that the learning process was in state 1a in period t approaches one as β approaches zero. This result

relies on the assumption that the learning process has been in state 1 all the time between periods t' and t. Since producing (q_C, q_C) in state 1 requires experimentation and since $t - t'$ is a finite number, it can be seen that the probability that this assumption is true approaches one as β approaches zero. Consequently, the result also holds without making the assumption.

As a consequence of the above three observations, the probability that a transition from state 1 leads to state 2 approaches one as β approaches zero. □

Proof. Proof of Property 2.3 First consider state 1. When the learning process is in this state, two cases can be distinguished. If $Q^N_{1,t}, Q^N_{2,t} < \pi_{NN} + \varepsilon$, the learning process is said to be in state 1a, otherwise it is said to be in state 1b. If in some period t the learning process is in state 1a, it will be in state 1b in period $t + 1$ only if the firms produce either (q_C, q_N) or (q_N, q_C) in period t. The probability that this happens approaches zero as β approaches zero. If in some period t the learning process is in state 1b, it will reach state 1a if from period t onwards the firms produce (q_N, q_N) for some finite number of consecutive periods. The probability that this happens approaches one as β approaches zero. So, in the limit as β approaches zero, it takes an infinite number of periods to reach state 1b from state 1a, whereas it takes a finite number of periods to reach state 1a from state 1b. Furthermore, if the learning process is in state 1, the number of periods it takes to leave that state approaches infinity as β approaches zero. It now follows that the conditional probability that the learning process is in state 1a given that it is in state 1 approaches one as β approaches zero. If in some period t the learning process is in state 1a, the probability that the firms produce (q_C, q_C) approaches zero as β approaches zero and is of order $\exp(-(Q^N_{1,t} - Q^C_{1,t} + Q^N_{2,t} - Q^C_{2,t})/\beta)$. (To see this, note that the firms choose their production levels independently according to probabilities given by (2.3).) A lower bound for this order is $\exp(-2(\pi_{NN} - \pi_{CN} + \varepsilon)/\beta)$. Furthermore, if the learning process is in state 1a and the firms produce (q_C, q_C), the probability that a transition from state 1 to another state occurs approaches one as β approaches zero. (This has been shown in the proof of Property 2.2.)

Now consider state 2. If in some period t the learning process is in this state, a transition to another state can occur only if the firms produce either (q_C, q_N) or (q_N, q_C). The probability that this happens approaches zero as β approaches zero and is of order

$\exp(-\min(Q_{1,t}^{\mathrm{C}} - Q_{1,t}^{\mathrm{N}}, Q_{2,t}^{\mathrm{C}} - Q_{2,t}^{\mathrm{N}})/\beta)$. An upper bound for this order is $\exp(-(\pi_{\mathrm{CC}} - \pi_{\mathrm{NN}} - 2\varepsilon)/\beta)$.

In summary, if the learning process is in state 1, then in the limit as β approaches zero the rate at which the probability of a state transition approaches zero equals, with probability one, at most $2(\pi_{\mathrm{NN}} - \pi_{\mathrm{CN}} + \varepsilon)$. Furthermore, if the learning process is in state 2, the rate at which the probability of a state transition approaches zero equals at least $\pi_{\mathrm{CC}} - \pi_{\mathrm{NN}} - 2\varepsilon$. From $\varepsilon < (2\pi_{\mathrm{CN}} - 3\pi_{\mathrm{NN}} + \pi_{\mathrm{CC}})/4$, which is due to (2.21), it follows that $2(\pi_{\mathrm{NN}} - \pi_{\mathrm{CN}} + \varepsilon) < \pi_{\mathrm{CC}} - \pi_{\mathrm{NN}} - 2\varepsilon$. Consequently, the ratio between the average time it takes in state 1 before a state transition occurs and the average time it takes in state 2 before a state transition occurs approaches zero as β approaches zero. $\qquad\square$

Chapter 3

An evolutionary model of price competition among spatially distributed firms*

Abstract

Various studies have shown the emergence of cooperative behavior in evolutionary models with spatially distributed agents. We investigate to what extent these findings generalize to evolutionary models of price competition among spatially distributed firms. We consider both one- and two-dimensional models, and we vary the amount of information firms have about competitors in their neighborhood. Our computer simulations show that the emergence of cooperative behavior depends strongly on the amount of information available to firms. Firms tend to behave most cooperatively if they have only a very limited amount of information about their competitors. We provide an intuitive explanation for this phenomenon. Our simulations further indicate that three other factors in our models, namely the accuracy of firms' information, the probability of experimentation, and the spatial distribution of consumers, have little effect on the emergence of cooperative behavior.

3.1 Introduction

The phenomenon of cooperative behavior among individuals in social, economic, and biological systems has been fascinating researchers for quite some time already. An important topic in the economic and biological literature is the emergence of cooperative behavior among individuals who are pursuing their self-interest. Researchers aim to identify the conditions under which the emergence of cooperative behavior among such individuals is possible.

In an economic context, the best-known explanation of cooperative behavior is probably the one based on the idea of reciprocity in repeated encounters. When individuals interact with each other repeatedly, they may choose to behave cooperatively, even though this has a negative effect on their short-term interests. Individuals may choose to behave cooperatively because they realize that if they do not behave this way, others won't do either. They also realize that in the long run they are better off in a cooperative world than in a non-cooperative one. Hence, although cooperative behavior harms one's short-term interests, it is likely to be beneficial to one's interests in the long run.

Explaining cooperative behavior in terms of reciprocity assumes that individuals interact with each other repeatedly and that they remember what happened in the past. These assumptions seem reasonable in some contexts but not in others. Because of this, a number of alternative explanations of cooperative behavior have been proposed in the literature. In this chapter, we focus on one such explanation. This is the explanation that cooperative behavior is a consequence of the spatial distribution of individuals and the local interaction among them. In the biological literature, this explanation was proposed in a well-known paper by Nowak and May (1992). Many biologists have built on this work, which has resulted in a substantial body of literature.

Inspired by the work done in biology, economists have also attempted to explain cooperative behavior in terms of local interaction among spatially distributed individuals. An evolutionary perspective is typically taken, in which individuals are assumed to imitate each other and to randomly experiment with new actions. The first work in the economic literature was done by Bergstrom and Stark (1993) and Eshel et al. (1998). In this work, cooperative behavior was shown to emerge in models in which individuals are organized in a circular structure. A large number of studies have expanded on

this early work.[1] Studies in the economic literature often focus on rather abstract models. Many studies for example assume that individuals are located in a one-dimensional world. Also, many studies assume a situation similar to a classical prisoners' dilemma, in which individuals can choose from only two actions (i.e., cooperation and defection). For examples of studies that make these assumptions, we refer to Bergstrom and Stark (1993), Eshel et al. (1998), Jun and Sethi (2007), Mengel (2009), and Stark and Behrens (2010).

In this chapter, we consider a somewhat less abstract level of modeling. We aim to determine to what extent the findings from earlier studies generalize to models of price competition among spatially distributed firms. In particular, we want to find out whether imitation and experimentation may cause cooperative behavior to emerge in spatial price competition models. Compared with the frequently studied prisoners' dilemma models, our model is of a more complex nature. There is no simple binary decision between cooperative and non-cooperative behavior in our models. Firms can cooperate by jointly increasing their price, and different price levels correspond with different levels of cooperation. Also, interactions in our models may involve more than two individuals. In one of our models, each consumer has four different firms from which he may choose to buy. Hence, firms in this model always have multiple competitors with which they fight for the same market share. Like in the literature mentioned above, we take an evolutionary perspective in our models. We assume that the behavior of firms is determined by imitation and experimentation. More specifically, we assume that firms change their price either by imitating successful competitors in their neighborhood or by experimenting with small price increases or decreases.

We study a variety of conditions under which firms may or may not start to cooperate. We consider both a model in which firms are organized in a one-dimensional space and a model in which firms are organized in a two-dimensional space. Our two-dimensional model has two variants, which differ in the way in which consumers are located. We also look at the effect of the information firms have about competitors in their neighborhood. In doing so, we distinguish between on the one hand the number of

---

[1]See Barr and Tassier (2010), Bilancini and Boncinelli (2009), Chen and Chow (2009), Eshel, Herreiner, Samuelson, Sansone, and Shaked (2000), Eshel, Sansone, and Shaked (1999), Fosco and Mengel (2011), Jun and Sethi (2007, 2009), Kirchkamp (1999, 2000), Mengel (2009), Noailly, Van den Bergh, and Withagen (2009), Noailly, Withagen, and Van den Bergh (2007), Outkin (2003), Stark and Behrens (2010), Tieman, Houba, and Van der Laan (2000), and Wilhite (2006).

competitors about which firms have information and on the other hand the accuracy of the information firms have. Another effect that we look at is the effect of firms' experimentation probability, that is, the probability with which firms experiment with small price increases or decreases. Due to the complexity of the models that we study, we perform our analyses mainly using computer simulations.

This chapter is organized as follows. The models that we study are introduced in Section 3.2. The analysis of the models is presented in Section 3.3. The main conclusions of our research are summarized in Section 3.4.

3.2 Models

We consider two closely related models. We refer to these models as the one-dimensional model and the two-dimensional model. The way in which firms and consumers are located is different in each model. Apart from this difference, the models are essentially identical. We discuss the one-dimensional model in Subsection 3.2.1 and the two-dimensional model in Subsection 3.2.2.

3.2.1 One-dimensional model

There are n firms, denoted by $1, \ldots, n$. In the one-dimensional model, firms are located equally spaced on a circle (see Figure 3.1).[2] The distance, measured over the circumference of the circle, between any two neighboring firms equals one. Consumers are uniformly distributed on the circle. They are modeled as a continuum. The density of consumers equals 1 everywhere on the circle. Firms all produce the same product, they all have an unlimited production capacity, and they all have the same constant marginal cost. Without loss of generality, firms' constant marginal cost is set to zero. The price at which firm i sells one unit of its product is denoted by p_i. Firms choose their prices simultaneously. A consumer's total cost of buying a unit from a firm equals the price charged by the firm plus transportation cost. Transportation cost equals the distance, measured over the circumference of the circle, between the consumer and the firm. Each consumer needs exactly one unit of the product produced by the firms. A consumer buys

---

[2]Firms are located on a circle rather than on a line in order to avoid boundary effects.
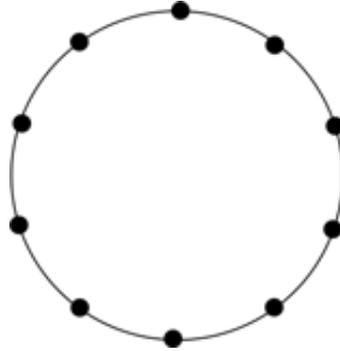
Figure 3.1: One-dimensional model with $n = 10$ firms. A firm is indicated by a black dot. Consumers are located everywhere on the circle.

this unit from the firm for which the consumer's total cost is lowest. This implies that the circle can be partitioned into n segments in such a way that all consumers located on the ith circle segment buy from firm i. Firm i's quantity demanded, denoted by q_i, then equals the length of the ith circle segment, and firm i's profit is given by $\pi_i = p_i q_i$.

The model has a symmetric pure-strategy Nash equilibrium in which $p_1 = \cdots = p_n = 1$. This can be seen as follows. Suppose that $p_1 = \cdots = p_n = 1$. We will show that a firm cannot increase its profit by unilaterally changing its price. Consider an arbitrary firm i, and suppose that this firm changes its price p_i. The other firms do not change their price. Next, consider a consumer located somewhere in between firm i and firm j, where firm j is one of the two neighboring firms of firm i. Let the distance between the consumer and firm i be denoted by d. The distance between the consumer and firm j is then given by $1 - d$. The consumer's total cost of buying from firm i equals $p_i + d$, while the consumer's total cost of buying from firm j equals $p_j + 1 - d = 2 - d$. Hence, the consumer will buy from firm i if $p_i + d < 2 - d$ or, equivalently, if $d < 1 - p_i/2$. This means that firm i's quantity demanded equals $q_i = 2(1 - p_i/2) = 2 - p_i$ and that firm i makes a profit of $\pi_i = p_i(2 - p_i)$. Clearly, firm i maximizes its profit by choosing a price of $p_i = 1$. In other words, if firm i changes its price to a value different from 1, its profit will decrease. This implies that $p_1 = \cdots = p_n = 1$ is a Nash equilibrium. In this equilibrium, each firm makes a profit of 1.

It is straightforward to see that firms find themselves in a situation that is somewhat

Table 3.1: Payoff matrix that illustrates the situation of firms in the one-dimensional model. For the purpose of illustration, price is treated as a discrete variable that can take four different values. The row player represents an arbitrary firm i. The column player represents firms $i-1$ and $i+1$, which are the two neighbors of firm i. (The two neighbors are assumed to choose the same price.) The payoffs represent the profits of firm i.[3]

	$p_{i-1} = p_{i+1} = 1.50$	$p_{i-1} = p_{i+1} = 1.25$	$p_{i-1} = p_{i+1} = 1.00$	$p_{i-1} = p_{i+1} = 0.75$
$p_i = 1.50$	1.50	1.13	0.75	0.38
$p_i = 1.25$	1.56	1.25	0.94	0.63
$p_i = 1.00$	1.50	1.25	1.00	0.75
$p_i = 0.75$	1.31	1.13	0.94	0.75

similar to a prisoners' dilemma. Choosing the Nash equilibrium price can be seen as defection, while choosing a price above the Nash equilibrium level can be seen as co-operation. If a firm cooperates while its neighbors defect, the firm will make a lower profit than in the Nash equilibrium. However, if a firm cooperates and its neighbors do so as well, the firm will make a higher profit than in the Nash equilibrium. The payoff matrix shown in Table 3.1 illustrates the situation in which firms find themselves. Notice that there is one important difference with a prisoners' dilemma. This is because a firm may also choose a price below the Nash equilibrium level. For this action, there is no analogous action in a prisoners' dilemma.

In our model, we assume that firms are boundedly rational. Hence, firms need not use Nash equilibrium strategies. We take an evolutionary game theory approach and assume that the behavior of firms is determined by imitation and experimentation. The stage game described above is played repeatedly for a large number of rounds. After each round, firms may change their price. Firms change their price by imitating successful neighbors or by experimenting with a small price increase or decrease. Price is modeled as a discrete variable. That is, firms choose their price from a finite set of price levels.

---

[3]In the payoff matrix shown in Table 3.1, there are three Nash equilibria, namely a strict Nash equilibrium in which each firm charges a price of 1 and two weak Nash equilibria, one in which each firm charges a price of 1.25 and one in which each firm charges a price of 1.5. The weak Nash equilibria are due to the treatment of price as a discrete variable. We will come back to this issue later on in this chapter (see Footnote 5).

Imitation is modeled as follows. At the end of each round, a firm is randomly selected. The selected firm knows its own price in the most recent round and the prices of its ρ closest neighbors, where ρ is an even number that indicates the size of the information neighborhood of a firm. The selected firm also observes its own profit in the most recent round and the profits of its ρ closest neighbors. In the case of its neighbors, however, the firm does not observe their true profits but rather their true profits perturbed by some noise. For each neighbor, the noise is modeled by adding a normally distributed random variable to the neighbor's true profit. The random variable has mean 0 and standard deviation σ, where we refer to σ as the noise level. The selected firm chooses a new price by copying the price that appears to have been most profitable in the most recent round. More specifically, the firm first averages the observed profits of firms that used the same price in the most recent round. The firm then chooses the price associated with the highest observed profit as its new price. (Ties are broken randomly.) An illustration of the imitation mechanism is provided in Figure 3.2.

Experimentation takes place after imitation and is modeled as follows. At the end of each round, each firm independently decides whether to experiment with a new price or not. The probability that a firm chooses to experiment is given by the parameter μ. This probability is typically very small. If a firm chooses to experiment, there is a 50% probability of a price increase and a 50% probability of a price decrease. The firm will set its new price to the closest price level above or below its current price. If the firm's current price equals the highest (lowest) price level that can be chosen in the model, no price increase (decrease) is possible.

One-dimensional models have been frequently studied in the literature (Barr & Tassier, 2010; Bergstrom & Stark, 1993; Chen & Chow, 2009; Eshel et al., 1998, 1999, 2000; Jun & Sethi, 2007; Mengel, 2009; Noailly et al., 2007; Outkin, 2003; Stark & Behrens, 2010; Wilhite, 2006).[4] Our model is somewhat similar to the model of Eshel et al. (1998). Eshel et al. study a population of altruists and egoists located on a circle. Like in our model, agents imitate the strategies of successful neighbors. It turns out that, even though being an egoist is a dominant strategy, altruism can still prevail in the long run. Altruism can prevail if altruists are grouped together on the circle, so that they benefit from each other's altruism. An important difference between the model of Eshel

---

[4]We focus on the theoretical literature. In the experimental literature, a model similar to our one-dimensional model is considered by Selten and Apesteguia (2005).

	Firm									
	1	2	3	4	5	6	7	8	9	10
Price	1.00	1.10	1.10	1.10	1.00	0.90	1.00	1.00	0.90	0.80
Quantity	0.95	0.95	1.00	0.95	1.00	1.10	0.95	0.95	1.00	1.15
Profit	0.95	1.05	1.10	1.05	1.00	0.99	0.95	0.95	0.90	0.92
Obs. profit			1.17	1.23	1.00	0.98	0.86			
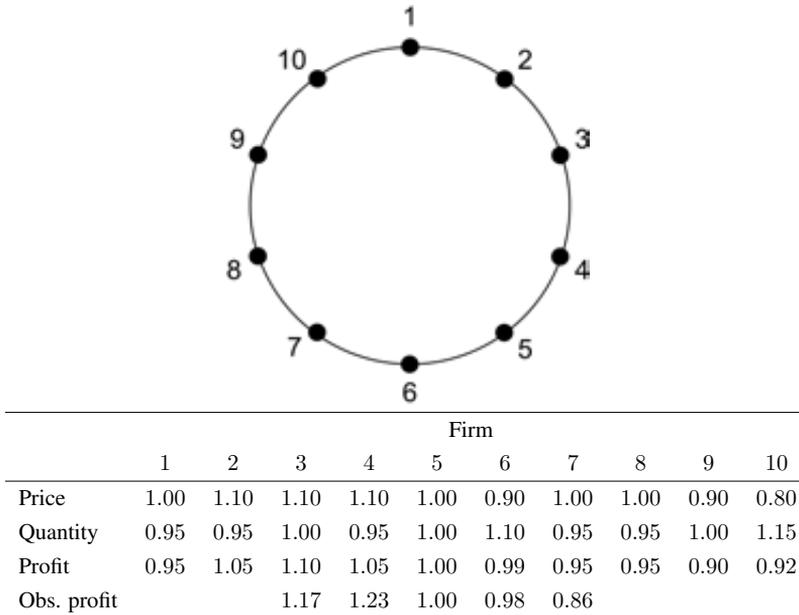
Figure 3.2: Illustration of the imitation mechanism. There are $n = 10$ firms. Firms have an information neighborhood of size $\rho = 4$, and the noise level equals $\sigma = 0.1$. For each firm, the price, the quantity demanded, and the profit are listed in the table. Firm 5 is randomly selected to change its price. Firm 5 does not observe the true profits of its neighbors but rather their true profits perturbed by some noise. The profits observed by firm 5 are listed in the table as well. The average observed profit equals 0.98 for a price of 0.9, $(1.00 + 0.86)/2 = 0.93$ for a price of 1.0, and $(1.17 + 1.23)/2 = 1.20$ for a price of 1.1. Hence, firm 5 will increase its price from 1.0 to 1.1.

et al. and our model is that in our model agents can choose from more than two actions (i.e., firms can choose from more than two price levels). Another difference is that in our model agents do not always have noise-free information about their neighbors' payoffs.

3.2.2 Two-dimensional model

The two-dimensional model is very similar to the one-dimensional model except that firms and consumers are located differently. In the two-dimensional model, we start
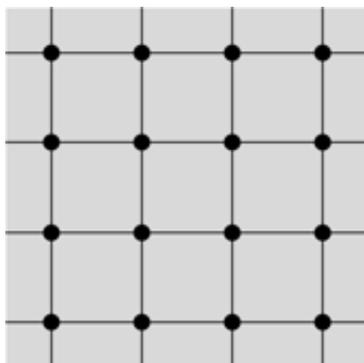
Figure 3.3: Two-dimensional model with $n = 16$ firms. A firm is indicated by a black dot. In variant A of the model, consumers are located only on the black lines. In variant B of the model, consumers are located everywhere in the gray area.

with a square lattice of $m \times m$ points. There are $n = m^2$ firms, which are located on the points of the lattice (see Figure 3.3). The distance between firms that are direct neighbors equals one. The model has two variants. These variants differ in the way in which consumers are located (see Figure 3.3). In one variant, referred to as variant A, consumers are located only on line segments between firms that are direct neighbors. In the other variant, referred to as variant B, consumers are located everywhere in the two-dimensional space in between the firms. Consumers are modeled as a continuum in both variants of the model. Also, in both variants, the distribution of consumers is uniform, with a density of one everywhere. All distances in the model are calculated using the city block or Manhattan distance measure rather than using the Euclidean distance measure (see Figure 3.4). The use of the city block measure is mathematically convenient, and it also is a natural choice if we interpret the model in terms of firms and consumers located in a city with a block design.

In our one-dimensional model discussed in Subsection 3.2.1, firms and consumers are located on a circle rather than on a line. This simplifies the analysis of the model, because there are no boundary effects that need to be taken into account. In a similar way, we also want to avoid boundary effects in our two-dimensional model. We therefore treat firms at opposite edges of the two-dimensional space as direct neighbors. Hence, each of the leftmost firms has a direct neighbor among the rightmost firms. Similarly,
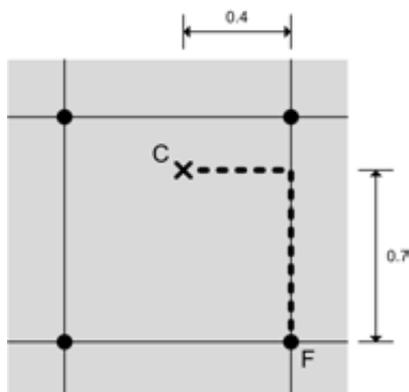
Figure 3.4: Illustration of the calculation of distances using the city block measure. The distance between firm F and consumer C equals $0.7 + 0.4 = 1.1$.

each of the topmost firms has a direct neighbor among the bottommost firms. In this way, there are no boundary effects in the model. That is, each firm finds itself in exactly the same position, with the same number of neighboring firms and an equally-sized consumer market.

Both variants of the model have a symmetric pure-strategy Nash equilibrium. Variant A has an equilibrium in which $p_1 = \cdots = p_n = 1$. This can be shown using an argument analogous to the argument used in the case of the one-dimensional model. Variant B has an equilibrium in which $p_1 = \cdots = p_n = 1/2$. This can be seen as follows. Suppose that $p_1 = \cdots = p_n = 1/2$. We will show that a firm cannot increase its profit by unilaterally changing its price. Consider an arbitrary firm i, and suppose that this firm changes its price p_i. The other firms do not change their price. We distinguish between two cases, namely the case of a price increase and the case of a price decrease. We first analyze the case of a price increase. If firm i increases its price p_i to a value above $1/2$, its quantity demanded will become $q_i = \max(3/2 - p_i, 0)^2$ (see Figure 3.5(a) for an illustration for $p_i = 0.7$). Consequently, firm i will make a profit of $\pi_i = p_i \max(3/2 - p_i, 0)^2$. Clearly, for $p_i \geq 1/2$, this profit function is monotonically decreasing. Hence, if firm i increases its price to a value above $1/2$, its profit will decrease. We now analyze the case of a price decrease. In the case of a price decrease, firm i's quantity demanded is given by $q_i = (1/2)p_i^2 - (5/2)p_i + 17/8$

(a)   (b)
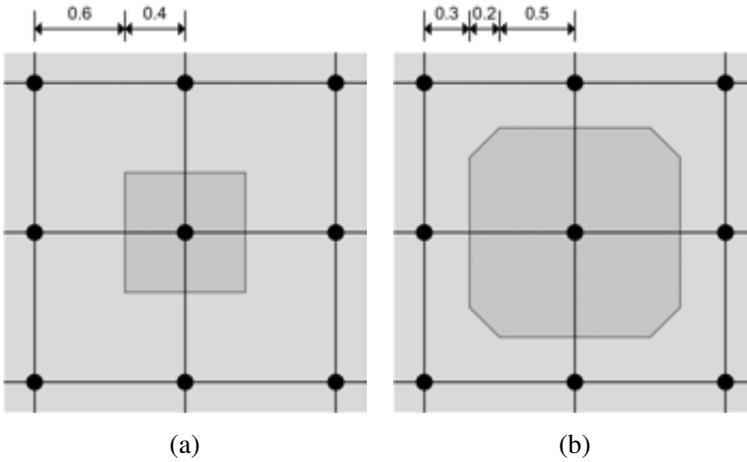
Figure 3.5: Illustration of the calculation of a firm's quantity demanded. The firm in the center of panel (a) charges a price of 0.7. The firm in the center of panel (b) charges a price of 0.1. The surrounding firms all charge a price of 0.5. In both panels, the shaded area marks the consumers that buy from the firm in the center. In panel (a), the quantity demanded of the firm in the center equals $0.8 \times 0.8 = 0.64$. In panel (b), the quantity demanded of the firm in the center equals $1.4 \times 1.4 - 4 \times 0.5 \times 0.2 \times 0.2 = 1.88$.

(see Figure 3.5(b) for an illustration for $p_i = 0.1$). This results in a profit function of $\pi_i = (1/2)p_i^3 - (5/2)p_i^2 + (17/8)p_i$. For $p_i \leq 1/2$, this function is monotonically increasing. This implies that a decrease of firm i's price to a value below $1/2$ will lead to a decrease of firm i's profit. Hence, both a price increase and a price decrease will lead to a decrease in profit. This shows that $p_1 = \cdots = p_n = 1/2$ is a Nash equilibrium of variant B of the model.

Both in variant A and in variant B of the model, firms find themselves in a situation that resembles a prisoners' dilemma. If a single firm unilaterally deviates from the Nash equilibrium by increasing its price, the firm will make a lower profit than in the equilibrium. However, if all firms jointly deviate from the Nash equilibrium by increasing their price, they will all make a higher profit than in the equilibrium.

Like in our one-dimensional model, we assume that firms are boundedly rational in our two-dimensional model. Firms' behavior is assumed to be determined by imitation and experimentation in the same way as in the one-dimensional model. In the

one-dimensional model, imitation takes place by looking at the prices and profits of $\rho + 1$ firms, namely a randomly selected firm and its ρ closest neighbors. In the two-dimensional model, we focus on two scenarios for the way in which firms imitate each other. In the first scenario, imitation takes place based on the prices and profits of five firms, namely a randomly selected firm and its direct neighbors in horizontal and vertical direction. In this scenario, firms have an information neighborhood of size $\rho = 4$. In the second scenario, imitation takes place based on the prices and profits of nine firms. In this scenario, the selected firm's direct neighbors in diagonal direction are included as well. The second scenario results in an information neighborhood of size $\rho = 8$.

A number of two-dimensional models have been studied in the economic literature (Barr & Tassier, 2010; Kirchkamp, 1999, 2000; Noailly et al., 2009; Outkin, 2003; Tieman et al., 2000; Wilhite, 2006). The models of Kirchkamp (2000) and Tieman et al. (2000) are the ones that are most closely related to our model. Kirchkamp studies a two-dimensional model in which agents play prisoners' dilemma games with their neighbors. He shows that under certain conditions cooperation can prevail in the long run. Important differences between Kirchkamp's model and our model are that in our model agents can choose from more than two actions and that in our model agents do not always have noise free information about their neighbors' payoffs. Tieman et al. study a local interaction model in which agents play generalized prisoners' dilemma games, that is, prisoners' dilemma games in which there can be more than two actions. They find that with a high probability a moderate level of cooperation emerges in their model. An essential difference between the model of Tieman et al. and our model is that in the model of Tieman et al. agents do not imitate each others' strategies. Instead, agents increase or decrease their cooperativeness depending on whether their average payoff is higher or lower than the average payoff of their neighbors.

3.3   Analysis

We are interested in the long-run behavior of firms in our one-dimensional and two-dimensional models. In particular, we want to find out whether in the long run firms behave cooperatively by charging prices above the Nash equilibrium level. Because our models do not seem analytically tractable, we use computer simulations to perform our

analysis. In Subsection 3.3.1, the setup of the simulations is discussed. The results obtained using the simulations are presented in Subsections 3.3.2 (one-dimensional model) and 3.3.4 (two-dimensional model). Some intuitive insight into the one-dimensional model is provided in Subsection 3.3.3.

3.3.1   Simulation setup

In our simulations, there are $n = 400$ firms. This means that in the two-dimensional model firms are located on the points of a 20×20 square lattice. Firms choose their price from a set of 21 price levels. These price levels are uniformly distributed between $0.5p_N$ and $1.5p_N$, where p_N denotes the Nash equilibrium price. Hence, in the one-dimensional model and in variant A of the two-dimensional model, the price levels that can be chosen are $0.50, 0.55, \ldots, 1.50$. In variant B of the two-dimensional model, the price levels that can be chosen are $0.250, 0.275, \ldots, 0.750$.[5] In the case of the one-dimensional model, simulations are run for six different values of the information neighborhood size ρ, namely 2, 4, 6, 8, 10, and 20. In the case of the two-dimensional model, simulations are run for an information neighborhood of size $\rho = 4$ and for an information neighborhood of size $\rho = 8$. Furthermore, both in the case of the one-dimensional model and in the case of the two-dimensional model, simulations are run for four different noise levels σ, namely 0, $0.1p_N$, $0.2p_N$, and $0.5p_N$, and for four different experimentation probabilities μ, namely 0, 0.00001, 0.0001, and 0.001.

At the beginning of a simulation run, each firm's price is initialized by randomly drawing a price from a uniform distribution over all price levels. A simulation run lasts for one million rounds. One million rounds turns out to be sufficient for studying firms' long-run behavior. We performed some tests which indicate that after one million rounds the results of our simulations are insensitive to the way in which firms' prices were initialized. The tests that we performed also indicate that a larger number of rounds would yield essentially the same simulation results.

---

[5]Modeling price as a discrete rather than a continuous variable may introduce additional Nash equilibria. This turns out to be the case in the one-dimensional model and in variant A of the two-dimensional model. In addition to a strict Nash equilibrium in which each firm charges a price of 1.00, there are two weak Nash equilibria in these models, one in which each firm charges a price of 0.95 and one in which each firm charges a price of 1.05. There turn out to be no additional Nash equilibria in variant B of the two-dimensional model. In the rest of this chapter, when we refer to a Nash equilibrium of a model, we always mean a strict Nash equilibrium.

Table 3.2: Simulation results for the one-dimensional model for $\rho = 2$ and for different values of σ and μ. The table shows the mean price at the end of the simulation runs.

	$\sigma = 0.0$	$\sigma = 0.1$	$\sigma = 0.2$	$\sigma = 0.5$
$\mu = 0$	1.28	1.33	1.31	1.31
$\mu = 0.00001$	1.27	1.25	1.27	1.28
$\mu = 0.0001$	1.25	1.20	1.22	1.23
$\mu = 0.001$	1.14	1.11	1.13	1.15

Table 3.3: Simulation results for the one-dimensional model for different values of ρ, σ, and μ. The table shows the mean price at the end of the simulation runs.

	$\sigma = 0.0; \mu = 0$	$\sigma = 0.0; \mu = 0.0001$	$\sigma = 0.2; \mu = 0$	$\sigma = 0.2; \mu = 0.0001$
$\rho = 2$	1.28	1.25	1.31	1.22
$\rho = 4$	0.90	0.90	0.93	0.95
$\rho = 6$	0.90	0.90	0.95	0.95
$\rho = 8$	0.90	0.93	0.95	0.96
$\rho = 10$	0.93	0.95	0.96	0.96
$\rho = 20$	0.98	1.00	0.98	0.98

The source code of the simulations is available online at www.ludowaltman.nl/ price_competition/. The source code runs in MATLAB and has been written partly in the MATLAB language and partly in the C language.

3.3.2   Simulation results for the one-dimensional model

The results of the simulations for the one-dimensional model are reported in Tables 3.2 and 3.3. For each combination of an information neighborhood size ρ, a noise level σ, and an experimentation probability μ, 500 simulation runs were performed. For each simulation run, we calculated the mean price of the firms at the end of the last round (i.e., at the end of the one millionth round). In Tables 3.2 and 3.3, this mean price is averaged over the 500 simulations runs that were performed. Standard deviations over the 500 simulation runs (not reported in the tables) are always less than 0.05. The relatively small standard deviations indicate that there is little variation between simulation runs.

In Table 3.2, results are reported of simulations in which the information neighborhood has a size of $\rho = 2$. In these simulations, firms can imitate only their direct

neighbors. As can be seen in the table, the simulations yield prices that are substantially above the Nash equilibrium level of 1. The prices are not very sensitive to the noise level σ. They are somewhat more sensitive to the experimentation probability μ. A higher experimentation probability clearly leads to a lower price. The results in Table 3.2 are in line with the findings of earlier studies in which somewhat similar models were analyzed (e.g., Eshel et al., 1998).

We now turn to the effect of the information neighborhood size ρ. Simulation results for different values of the size of the information neighborhood are reported in Table 3.3. If the size of the information neighborhood is larger than 2, firms can imitate not only their direct neighbors but also some of their more distant neighbors. The results in Table 3.3 are quite remarkable. It turns out that prices are no longer above the Nash equilibrium level if the size of the information neighborhood is larger than 2. On the contrary, if the size of the information neighborhood is not too large, prices turn out to be below the Nash equilibrium level. This is especially the case if there is no noise and no experimentation (i.e., $\sigma = 0$ and $\mu = 0$). In earlier studies (Hoffmann, 1999; Ifti, Killingback, & Doebeli, 2004; Mengel, 2009; Stark & Behrens, 2010), it was found that cooperative behavior (i.e., prices above the equilibrium level in our context) tends to be more difficult to sustain if the size of the information neighborhood is increased. This is consistent with our findings, but our findings go one step further. If the size of the information neighborhood is increased, firms not only stop behaving cooperatively but they in fact start behaving in exactly the opposite way, that is, they decrease their prices to values below the equilibrium level. Hence, our results show that in some cases the combination of local interaction and imitation of neighboring individuals has a negative rather than a positive effect on the degree to which individuals cooperate.

3.3.3 Further analysis of the one-dimensional model

Why does our one-dimensional model yield completely opposite simulation results for an information neighborhood of size $\rho = 2$ on the one hand and for an information neighborhood of size $\rho \in \{4, 6, 8, 10\}$ on the other hand? To provide some intuitive insight, we first focus on the case of an information neighborhood of size 2 and we then consider the case of an information neighborhood of size 4. To simplify the analysis, we assume that in both cases firms can choose from only two price levels. We also assume

that there is no noise and no experimentation (i.e., $\sigma = 0$ and $\mu = 0$). In other words, the only way in which a firm can change its price is by means of imitation, and if a firm imitates, it does so based on noise-free information about the profits of its neighbors.

In the case of an information neighborhood of size 2, we assume that firms charge a price of either 1.0 (i.e., the Nash equilibrium price) or 1.1.[6] We refer to these prices as, respectively, the low price and the high price, and we refer to firms charging the low price as low-price firms and to firms charging the high price as high-price firms. Suppose that we have a cluster of low-price firms and a cluster of high-price firms. By a cluster of low-price (high-price) firms, we mean a number of low-price (high-price) firms that are direct neighbors of each other. Suppose further that the cluster of low-price firms and the cluster of high-price firms are located next to each other in the one-dimensional space of our model. This is illustrated in Figure 3.6(a). The figure also shows the profit made by each firm. Based on Figure 3.6(a), let us look what will happen. A low-price firm that is surrounded by two other low-price firms cannot change its price. The same holds for a high-price firm that is surrounded by two other high-price firms. We therefore focus on firms 5 and 6 in Figure 3.6(a). Firm 6 will not change its price. This is because, based on the information available to this firm, the average profit resulting from the high price (i.e., $\pi_6/2 + \pi_7/2 \approx 1.073$) exceeds the average profit resulting from the low price (i.e., $\pi_5 = 1.050$). Hence, firm 6 will stick to the high price. Unlike firm 6, firm 5 will change its price. Firm 5 is currently a low-price firm, but based on the information available to the firm, the high price appears to be more profitable than the low price (since $\pi_6 = 1.045 > \pi_4/2 + \pi_5/2 = 1.025$). As a consequence, firm 5 will change to the high price. This will lead to a new situation, which is illustrated in Figure 3.6(b). Looking at Figure 3.6(b), it is clear that the next step will be firm 4 changing from the low price to the high price. Hence, the general pattern is that the cluster of high-price firms is growing more and more while the cluster of low-price firms is shrinking. This

---

[6]The choice of these two prices is fairly arbitrary. However, our analysis is valid for many other prices as well.
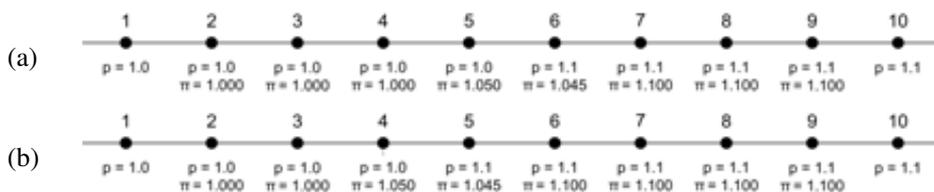
Figure 3.6: Illustration of the effect of imitation in the one-dimensional model with an information neighborhood of size $\rho = 2$.

is the basic intuition underlying our simulation results for an information neighborhood of size 2.[7]

We now consider the case of an information neighborhood of size 4.[8] In this case, we assume that firms charge either a low price of 0.9 or a high price of 1.0 (i.e., the Nash equilibrium price). Notice that these prices are different from the prices used in the above analysis for an information neighborhood of size 2. This is because we now want to explain why firms charge prices below the Nash equilibrium level, while in the analysis presented above we wanted to explain why firms charge prices above the Nash equilibrium level. We again start from a situation with a cluster of low-price firms and a cluster of high-price firms. This situation is illustrated in Figure 3.7(a). Based on Figure 3.7(a), it can be seen that there are two firms for which a price change is possible, namely firm 4 and firm 5. In both cases, there would be a change from the low price to the high price. Firm 4 and firm 5 cannot both change their price at the same time. Instead, one of the two firms will be randomly selected to change its price. If firm 5 is selected, the effect will be that the cluster of high-price firms grows while the cluster of low-price firms shrinks. This is similar to what happens in the case of

---

[7]The full story is more complicated. In particular, it can be shown that low-price firms will not disappear altogether. Suppose we have a cluster of low-price firms surrounded on both sides by a cluster of high-price firms. As explained above, the cluster of low-price firms will shrink more and more. However, when there are just two low-price firms left, the cluster of low-price firms will not shrink any further. Hence, in the end there will be mostly high-price firms, but in between these firms there will also be some small islands of low-price firms. We refer to Eshel et al. (1998) for an extensive discussion of this kind of dynamics.

[8]We refer to Mengel (2009) for a somewhat similar analysis. One of the differences between our analysis and the analysis of Mengel is that we do not consider the effect of experimentation while Mengel focuses on the limit case in which the probability of experimentation (referred to as trembling by Mengel) approaches zero.

	1	2	3	4	5	6	7	8	9	10
(a)	p = 0.9	p = 0.9 π = 0.900	p = 0.9 π = 0.900	p = 0.9 π = 0.945	p = 0.9 π = 0.950	p = 1.0 π = 1.000	p = 1.0 π = 1.000	p = 1.0 π = 1.000	p = 1.0 π = 1.000	p = 1.0
(b)	p = 0.9	p = 0.9 π = 0.900	p = 0.9 π = 0.945	p = 1.0 π = 0.900	p = 0.9 π = 0.990	p = 1.0 π = 0.950	p = 1.0 π = 1.000	p = 1.0 π = 1.000	p = 1.0 π = 1.000	p = 1.0
(c)	p = 0.9	p = 0.9 π = 0.900	p = 0.9 π = 0.945	p = 1.0 π = 0.900	p = 0.9 π = 0.990	p = 1.0 π = 0.950	p = 1.0 π = 1.000	p = 1.0 π = 1.000	p = 1.0 π = 1.000	p = 1.0
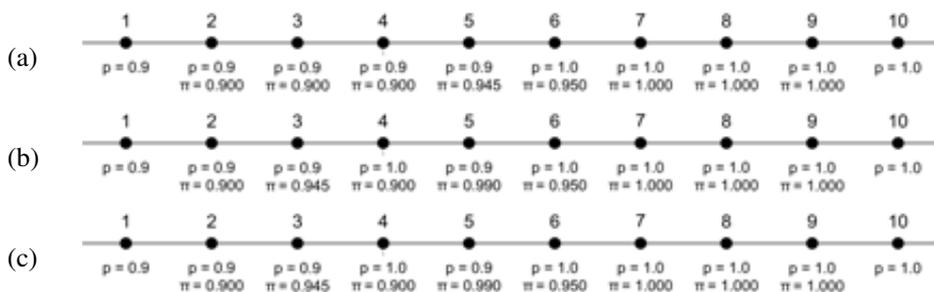
Figure 3.7: Illustration of the effect of imitation in the one-dimensional model with an information neighborhood of size $\rho = 4$.

an information neighborhood of size 2. If firm 4 is selected, the effect will be quite different. We will then end up in the situation illustrated in Figure 3.7(b). As can be seen in the figure, there will no longer be a perfect separation of low-price and high-price firms. In this new situation, there turn out to be three firms for which a price change is possible, namely firm 4, firm 6, and firm 7. In all three cases, the price change would be a movement from the high price to the low price. At this point, a comprehensive analysis of the various possibilities becomes cumbersome. Let us therefore focus on the most interesting possibility. This is the possibility of a price change by firm 7. If firm 7 changes its price, the cluster of high-price firms will shrink, as is illustrated in Figure 3.7(c). A next step could then be that firm 8 or firm 9 also changes its price, which would mean that the cluster of high-price firms will shrink even further. Going back to the initial situation illustrated in Figure 3.7(a), it is now clear that there are two counteracting forces at work. On the one hand the cluster of high-price firms may grow, while on the other hand this cluster may shrink. In the situation illustrated in Figure 3.7(a), the cluster of high-price firms will grow if firm 5 changes its price. On the other hand, if firm 4 changes its price, this may cause the cluster of high-price firms to shrink. It is not immediately clear which of these two counteracting forces is stronger. However, based on the simulation results reported in Table 3.3, it can be concluded that the force working against the high-price firms must be the stronger one.

We have now looked at our one-dimensional model both in the case of an information neighborhood of size 2 and in the case of an information neighborhood of size 4.

What is the essential difference between these two cases? In both cases, a cluster of high-price firms may take over a neighboring low-price firm. However, the difference is that in the case of an information neighborhood of size 2 the high-price firms will always remain organized in a single cluster (see Figure 3.6) while in the case of an information neighborhood of size 4 the high-price firms may become separated from each other (see Figure 3.7). When high-price firms are separated from each other, they become vulnerable. This is because an isolated high-price firm makes a relatively low profit while an isolated low-price firm makes a relatively high profit. The result may therefore be that low-price firms start to take over high-price firms. This is the basic mechanism that explains why in our one-dimensional model prices are lower in the case of an information neighborhood of size 4 than in the case of an information neighborhood of size 2.

3.3.4 Simulation results for the two-dimensional model

The results of the simulations for the two-dimensional model are reported in Tables 3.4 to 3.7. Tables 3.4 and 3.5 relate to variant A of the model. This is the variant in which consumers are located on line segments between neighboring firms. Tables 3.6 and 3.7 relate to variant B of the model. In this variant, consumers are located everywhere in the two-dimensional space in between the firms. The results in Tables 3.4 to 3.7 were obtained in the same way as the results in Tables 3.2 and 3.3. Hence, the results are averages over 500 simulation runs. Standard deviations over the 500 simulation runs (not reported in the tables) are always less than 0.03, indicating that there is little variation between simulation runs.

We first focus on variant A of the two-dimensional model. As can be seen in Tables 3.4 and 3.5, prices tend to be relatively close to the Nash equilibrium level of 1. They do not exceed the equilibrium level by more than 18%. This is much less than in the one-dimensional model, in which prices exceed the equilibrium level by at most 33% (see Table 3.2). Also, in variant A of the two-dimensional model, prices do not fall below the equilibrium level (except for $\sigma = 0.5$ and $\mu = 0$, where the price is marginally below the equilibrium level). This is another difference with the one-dimensional model. In the one-dimensional model, prices can be up to 10% below the equilibrium level (see Table 3.3). Comparing Tables 3.4 and 3.5, it can be seen that increasing the size of the information neighborhood from $\rho = 4$ to $\rho = 8$ leads to sub-

Table 3.4: Simulation results for variant A of the two-dimensional model for $\rho = 4$ and for different values of σ and μ. The table shows the mean price at the end of the simulation runs.

	$\sigma = 0.0$	$\sigma = 0.1$	$\sigma = 0.2$	$\sigma = 0.5$
$\mu = 0$	1.15	1.05	1.06	1.06
$\mu = 0.00001$	1.17	1.07	1.06	1.07
$\mu = 0.0001$	1.18	1.10	1.08	1.08
$\mu = 0.001$	1.16	1.12	1.11	1.11

Table 3.5: Simulation results for variant A of the two-dimensional model for $\rho = 8$ and for different values of σ and μ. The table shows the mean price at the end of the simulation runs.

	$\sigma = 0.0$	$\sigma = 0.1$	$\sigma = 0.2$	$\sigma = 0.5$
$\mu = 0$	1.04	1.01	1.00	0.99
$\mu = 0.00001$	1.04	1.00	1.00	1.00
$\mu = 0.0001$	1.05	1.01	1.01	1.01
$\mu = 0.001$	1.07	1.05	1.04	1.03

stantially lower prices. This is similar to what was observed for the one-dimensional model, and it is also somewhat similar to earlier findings reported in the literature (Ifti et al., 2004). The effect of the noise level σ and the experimentation probability μ is different than in the one-dimensional model. The noise level turns out to have a negative effect on prices, while the experimentation probability turns out to have a positive effect. Notice, however, that especially the effect of the experimentation probability is not very strong.

We now consider variant B of the two-dimensional model. As discussed in Subsection 3.2.2, variant B has a Nash equilibrium price of 0.5, which is only half of the Nash equilibrium price of variant A. This explains why the prices in Tables 3.6 and 3.7 are much lower than the prices in Tables 3.4 and 3.5. When looking at prices relative to the equilibrium price, it can be seen that the results in Tables 3.6 and 3.7 are in fact very similar to the results in Tables 3.4 and 3.5. The effects of the information neighborhood size ρ, the noise level σ, and the experimentation probability μ are also very similar.

Table 3.6: Simulation results for variant B of the two-dimensional model for $\rho = 4$ and for different values of σ and μ. The table shows the mean price at the end of the simulation runs.

	$\sigma = 0.00$	$\sigma = 0.05$	$\sigma = 0.10$	$\sigma = 0.25$
$\mu = 0$	0.58	0.53	0.53	0.54
$\mu = 0.00001$	0.59	0.54	0.54	0.54
$\mu = 0.0001$	0.59	0.55	0.55	0.55
$\mu = 0.001$	0.59	0.56	0.56	0.57

Table 3.7: Simulation results for variant B of the two-dimensional model for $\rho = 8$ and for different values of σ and μ. The table shows the mean price at the end of the simulation runs.

	$\sigma = 0.00$	$\sigma = 0.05$	$\sigma = 0.10$	$\sigma = 0.25$
$\mu = 0$	0.52	0.50	0.51	0.51
$\mu = 0.00001$	0.52	0.51	0.51	0.51
$\mu = 0.0001$	0.53	0.51	0.51	0.52
$\mu = 0.001$	0.54	0.53	0.53	0.53

Hence, it turns out that the way in which firms behave is very similar in the two variants of the two-dimensional model.

3.4   Conclusions

We have studied evolutionary models of price competition among spatially distributed firms. In our models, firms are organized either in a one-dimensional space or in a two-dimensional space. The behavior of firms is determined by imitation and experimentation. Imitation means that firms copy the price of one or more successful competitors in their neighborhood. Experimentation means that firms randomly increase or decrease their price by a small amount.

In earlier studies (e.g., Bergstrom & Stark, 1993; Eshel et al., 1998; Nowak & May, 1992), often in the context of prisoners' dilemma games, it was found that spatially distributed individuals that interact locally and that imitate successful neighbors tend to behave cooperatively in many cases. In this chapter, our aim has been to investigate

whether a similar tendency towards cooperative behavior can be found in the context of price competition among spatially distributed firms. In this context, cooperative behavior would mean that firms have prices and profits above the ordinary equilibrium level.

We have performed our analyses mainly using computer simulations. The results of the simulations provide a mixed picture. The emergence of cooperative behavior turns out to depend strongly on the amount of information available to firms. In the one-dimensional model, firms behave cooperatively only if the information they have about the prices and profits of other firms is restricted to their two direct neighbors. In the two-dimensional model, firms behave more cooperatively if they have information about four neighbors than if they have information about eight neighbors. Hence, the general pattern seems to be that having too much information may hurt cooperation (for similar results, see Hoffmann, 1999; Ifti et al., 2004; Mengel, 2009; Stark & Behrens, 2010). We have shown that in the one-dimensional model this is because having too much information may cause cooperative firms to become separated from each other, which weakens their position relative to non-cooperative firms. The two-dimensional model is more difficult to analyze, but the mechanism at work in this model may well be similar. A remarkable finding is that in the one-dimensional model having too much information may even lead to prices and profits below the ordinary equilibrium level. This shows that the combination of local interaction and imitation of neighboring individuals can have both a positive and a negative effect on the degree to which individuals cooperate. To the best of our knowledge, negative effects have not been reported before in the literature. We have also investigated a number of other factors that may affect the degree of cooperative behavior among firms. One of these factors is the accuracy of the information firms have about the profits of their neighbors. Another factor is the probability with which firms experiment with small price increases or decreases. The effect of these two factors turns out to be relatively small. In the case of the two-dimensional model, we have also looked at the effect of the way in which consumers are located in the two-dimensional space. There turn out to be no substantial differences between the two variants that we have considered.

Chapter 4

A mathematical analysis of the long-run behavior of genetic algorithms for economic modeling[*]

Abstract

We present a mathematical analysis of the long-run behavior of genetic algorithms that are used for modeling economic phenomena. The analysis relies on commonly used mathematical techniques in evolutionary game theory. Assuming a positive but infinitely small mutation rate, we derive results that can be used to calculate the exact long-run behavior of a genetic algorithm. Using these results, the need to rely on computer simulations can be avoided. We also show that if the mutation rate is infinitely small the crossover rate has no effect on the long-run behavior of a genetic algorithm. To demonstrate the usefulness of our mathematical analysis, we replicate a well-known study by Axelrod in which a genetic algorithm is used to model the evolution of strategies in iterated prisoner's dilemmas. The theoretically predicted long-run behavior of the genetic algorithm turns out to be in perfect agreement with the long-run behavior observed in computer simulations. Also, in line with our theoretically informed expectations, computer simulations indicate that the crossover rate has virtually no long-run effect. Some general new insights into the behavior of genetic algorithms in the prisoner's dilemma context are provided as well.

---

[*]This chapter is joint work with Nees Jan van Eck. The chapter is currently under submission.

4.1   Introduction

The field of evolutionary computation is concerned with the study of all kinds of evolutionary algorithms. These algorithms can be used for various purposes. Perhaps the most popular purpose for which they can be used is function optimization (e.g., Goldberg, 1989; Michalewicz, 1996; Gen & Cheng, 2000). In the function optimization context, evolutionary algorithms can be seen as heuristics that serve as alternatives to more traditional techniques from the fields of combinatorial optimization and mathematical programming. Another important purpose for which evolutionary algorithms can be used is the modeling of biological, social, and economic phenomena (e.g., Mitchell, 1996). This is the topic with which we are concerned in this chapter. Our focus is in particular on the use of evolutionary algorithms for modeling economic phenomena.

When using evolutionary algorithms in the economic modeling context, one of the assumptions one makes is that the agents whose behavior is being modeled are boundedly rational. This basically means that the agents are assumed not to behave in a utility maximizing manner. There are numerous ways in which boundedly rational behavior can be modeled (e.g., Fudenberg & Levine, 1998; Brenner, 2006). A popular approach is to rely on an evolutionary metaphor. This is the approach that is taken by evolutionary algorithms. In its simplest form, the evolutionary approach assumes that there is a population of agents and that for each agent in the population the strategy it uses depends on the population-wide past performance of strategies. The better the past performance of a strategy, the more likely the strategy is to be used again. The evolutionary approach also assumes that there always is a small probability that an agent experiments with a new strategy.

The evolutionary approach to modeling boundedly rational behavior has attracted a lot of attention, not only from researchers in the field of evolutionary computation but also from researchers in the social sciences, in particular from economists. Traditionally, economists have typically relied on game-theoretic models to analyze interactions between agents. These models usually assume agents to behave in a fully rational way. Nowadays, however, the limitations of game-theoretic models based on full rationality are well recognized and many economists have started to study evolutionary models of agent behavior. These models are based on the assumption that the behavior of agents

can best be described using some evolutionary mechanism rather than using the idea of full rationality.

In the field of economics, there are two quite separate streams of research that are both concerned with the evolutionary approach to modeling boundedly rational behavior. One stream of research, which is usually referred to as agent-based computational economics (e.g., Tesfatsion, 2006), makes use of techniques from the field of evolutionary computation. Especially genetic algorithms (GAs) are frequently used. Early work in this stream of research includes Miller (1986); Holland and Miller (1991); Marks (1992); Arifovic (1994); Andreoni and Miller (1995); Arifovic (1996); Dawid (1996); Miller (1996), and examples of more recent work are Vriend (2000); Lux and Schornstein (2005); Alkemade, La Poutré, and Amman (2006); Georges (2006); Haruvy, Roth, and Utku Ünver (2006); Alkemade, La Poutré, and Amman (2009); Waltman and Van Eck (2009). The other stream of research is more closely related to traditional game theory and is referred to as evolutionary game theory (e.g., Maynard Smith, 1982; Weibull, 1995; Vega-Redondo, 1996; Gintis, 2000). Like the traditional game-theoretic approach, the evolutionary game-theoretic approach is model-based and relies heavily on mathematical analysis. The use of computer simulations is not very common in evolutionary game theory.

In this chapter, it is not our aim to argue in favor of either the agent-based computational economics approach, which emphasizes algorithms and computer simulations, or the evolutionary game-theoretic approach, which emphasizes models and mathematical analysis. Instead, we want to show how the former approach can benefit from the mathematical techniques used in the latter approach. More specifically, we want to show how evolutionary algorithms that are used for modeling economic phenomena can be analyzed mathematically using techniques that are popular in evolutionary game theory. Our focus in this chapter is on one particular type of evolutionary algorithm, namely GAs with a binary encoding. However, we emphasize that the approach that we take can be applied to other types of evolutionary algorithms as well. The reason for focusing on GAs with a binary encoding is that this seems to be the type of evolutionary algorithm that is used most frequently for modeling economic phenomena (e.g., Miller, 1986; Axelrod, 1987; Marks, 1992; Arifovic, 1994; Yao & Darwen, 1994; Andreoni & Miller, 1995; Arifovic, 1996; Ashlock, Smucker, Stanley, & Tesfatsion, 1996; Crowley

et al., 1996; Dawid, 1996; Miller, 1996; Vriend, 2000; Van Bragt, Van Kemenade, & La Poutré, 2001; Alkemade, Van Bragt, & La Poutré, 2005; Ishibuchi & Namikawa, 2005; Lux & Schornstein, 2005; Alkemade et al., 2006; Georges, 2006; Alkemade, La Poutré, & Amman, 2007; Alkemade et al., 2009; Waltman & Van Eck, 2009).

The mathematical analysis that we present in this chapter deals with the long-run behavior of GAs with a binary encoding. The GAs are assumed to be used in the economic modeling context (for theoretical work on GAs in the function optimization context, see e.g. Nix & Vose, 1992; Rudolph, 1994; Mitchell, 1996; Rudolph, 1998; Vose, 1999). In the terminology of Vriend (2000), we are concerned with GAs that are used for modeling social learning (as opposed to individual learning). Our work can be seen as an extension of the work of Dawid (1996), who derived a number of important mathematical results on the behavior of GAs. For small and moderate population sizes, the results of Dawid do not provide a full characterization of the long-run behavior of GAs. We extend the work of Dawid by deriving results that do provide a full characterization of the long-run behavior of GAs for small and moderate population sizes. Using our results, the long-run behavior of a GA can be calculated exactly and needs not be estimated using computer simulations. This means that it is no longer necessary to run a GA a large number of times for a large number of iterations in order to get insight into its long-run behavior. The use of our mathematical results has at least three advantages over the use of computer simulations:

(1) Our mathematical results can be used to calculate the long-run behavior of a GA exactly, while computer simulations can only be used to estimate the long-run behavior of a GA.

(2) When using computer simulations, it can be difficult to determine how many iterations of a GA are required to approximate the long-run behavior of the GA reasonably closely. Our mathematical results do not have this problem.

(3) Calculating the exact long-run behavior of a GA using our mathematical results requires less computing time than obtaining a reasonably accurate estimate of the long-run behavior of a GA using computer simulations.

Our mathematical results have one important limitation, which is that on most of today's computers they can only be used if the chromosome length is not greater than about 24

bits. If the chromosome length is greater than about 24 bits, the use of our mathematical results to calculate the long-run behavior of a GA most likely requires a prohibitive amount of computer memory.

Like in Dawid (1996), the mathematical analysis presented in this chapter relies on the assumption that the mutation rate is positive but infinitely small. (In other words, the analysis is concerned with the limit case in which the mutation rate approaches zero.) In simulation studies with GAs, researchers typically work with values between 0.001 and 0.01 for the mutation rate. This seems to be a rather pragmatic choice (cf. Dawid, 1996). On the one hand, lower values for the mutation rate would lead to very slow convergence and, consequently, very long simulation runs. On the other hand, higher values for the mutation rate would lead to convergence to unstable, difficult to interpret outcomes. We believe that our assumption of an infinitely small mutation rate is justified because an infinitely small mutation rate is less arbitrary than a mutation rate whose value is determined solely based on pragmatic grounds (cf. Foster & Young, 1990). The assumption of an infinitely small mutation rate is also in line with the common practice in evolutionary game theory, in which a similar assumption is almost always made. The advantage of assuming an infinitely small mutation rate is that it greatly simplifies the mathematical analysis of the long-run behavior of GAs (see also Dawid, 1996). In fact, GAs with an infinitely small mutation rate can be analyzed in a similar way as well-known models in evolutionary game theory (e.g., Foster & Young, 1990; Kandori et al., 1993; Young, 1993; Vega-Redondo, 1997). Like in evolutionary game theory, mathematical results provided by Freidlin and Wentzell (1998) are the key tool for analyzing the long-run behavior to which convergence will take place. We note that, in addition to the assumption of an infinitely small mutation rate, there are some other assumptions on which our mathematical analysis relies. However, these assumptions are quite mild. Most GAs will probably satisfy them, and if a GA does not satisfy them, a minor modification of the GA will usually be sufficient to meet the assumptions.

To demonstrate the usefulness of our mathematical analysis, we replicate a well-known study by Axelrod (1987, 1997b; see also Dawid, 1996; Mitchell, 1996). Axelrod used a GA to model the evolution of strategies in iterated prisoner's dilemmas. He showed that an evolutionary mechanism can lead to cooperative behavior. Axelrod's study has been one of the first and also one of the most influential studies on the use

of evolutionary algorithms for modeling social and economic phenomena. Directly or indirectly, his study seems to have inspired many researchers (e.g., Mühlenbein, 1991; Fogel, 1993; Yao & Darwen, 1994; Ashlock et al., 1996; Crowley et al., 1996; Van Bragt et al., 2001; Alkemade et al., 2005; Chong & Yao, 2005; Ishibuchi & Namikawa, 2005; Ashlock, Kim, & Leahy, 2006; Mittal & Deb, 2006; Chong & Yao, 2007; Thibert-Plante & Charbonneau, 2007; Ashlock & Kim, 2008). The results obtained by Axelrod are all based on computer simulations. In this chapter, we show that more or less the same results can be calculated exactly, with no need to rely on simulations. We also discuss some new insights that exact calculations provide.

The mathematical analysis that we present in this chapter also has an important implication for the choice of the parameters of a GA. The analysis indicates that if the mutation rate is infinitely small the crossover rate has no effect on the long-run behavior of a GA. This is a quite remarkable result that, to the best of our knowledge, has not been reported before in the theoretical literature on GAs. The result implies that when GAs are used for modeling economic phenomena the crossover rate is likely to be a rather insignificant parameter, at least when one is mainly interested in the behavior of GAs in the long run (for the short run, see Thibert-Plante & Charbonneau, 2007). This suggests that in many cases the crossover rate can simply be set to zero, in which case no crossover will take place at all. Simulation results that we report in this chapter indeed show no significant effect of the crossover rate on the long-run behavior of a GA.

The remainder of this chapter is organized as follows. In Section 4.2, we present a mathematical analysis of the long-run behavior of GAs that are used for modeling economic phenomena. Based on the analysis, we derive an algorithm for calculating the long-run behavior of GAs in Section 4.3. In Section 4.4, we demonstrate an application of the algorithm by replicating Axelrod's study (Axelrod, 1987). Finally, we discuss the conclusions of our research in Section 4.5. Proofs of our mathematical results are provided in the appendix.

4.2   Analysis

The general form of the GAs that we analyze in this chapter is shown in Figure 4.1. In this figure, and also in the rest of this chapter, the positive integers n and m and the

---

Input: $n, m, \gamma,$ and ε

1   Initialize the population by randomly setting nm bits to zero or one
2   **repeat**
3   Selection: Apply the selection operator to select n chromosomes from the population (a chromosome may be selected more than once), and use the selected chromosomes as the new population
4   Crossover: Randomly partition the population into $n/2$ pairs of two chromosomes, and apply the crossover operator to each pair of chromosomes with probability γ
5   Mutation: Mutate the population by inverting each bit with probability ε
6   **until** some stopping criterion is satisfied

Figure 4.1: General form of the genetic algorithms analyzed in this chapter.

probabilities γ and ε denote, respectively, the population size, the chromosome length, the crossover rate, and the mutation rate. For simplicity, we assume the population size n to be even. We further assume the crossover rate γ and the mutation rate ε to remain constant over time. We also assume ε to be positive. The GAs that we analyze are generalizations of the canonical GA discussed by, for example, Goldberg (1989) and Mitchell (1996). Like the canonical GA, we assume the use of a binary encoding, that is, chromosomes correspond to bit strings in our GAs. Unlike the canonical GA, we do not assume the use of specific selection and crossover operators. Instead, the GAs that we analyze may use almost any selection operator, such as roulette wheel selection (sometimes referred to as fitness-proportionate selection), tournament selection, or rank selection, and any crossover operator, such as single-point crossover, two-point crossover, or uniform crossover. Furthermore, in the GAs that we analyze, the fitness of a chromosome may depend, either deterministically or stochastically, on the entire population rather than only on the chromosome itself. When using GAs for economic modeling, the fitness of a chromosome typically depends on the entire population. This is referred to as state-dependent fitness by Dawid (1996). In most studies, GAs that are used for economic modeling have the same general form as the GAs that we analyze in this chapter.

We now introduce the terminology and the mathematical notation that we use in our analysis. We note that an overview of the mathematical notation is provided in Table 4.1. There are $\mu = 2^m$ different chromosomes, denoted by $0, \ldots, \mu - 1$. Each

chromosome has a unique binary encoding, which is given by a bit string of length m.[1] $\mathcal{C} = \{0, \ldots, \mu - 1\}$ denotes the set of all chromosomes. i and j denote typical chromosomes and take values in $\mathcal{C}$. The following definition introduces the notion of uniform and non-uniform populations.

Definition 4.1. A population is said to be *uniform* if and only if all n chromosomes in the population are identical. A population is said to be *non-uniform* if and only if some chromosomes in the population are different.

$\mathcal{U}$ denotes the set of all uniform populations. Obviously, since there are μ different chromosomes, there are also μ different uniform populations, that is, $|\mathcal{U}| = \mu$. $u(i) \in \mathcal{U}$ denotes the uniform population consisting of n times chromosome i. $\delta(i, j)$ denotes the Hamming distance between chromosomes i and j, that is, the number of corresponding bits in the binary encodings of i and j that are different. $\mathcal{G}(i)$ denotes the set of all chromosomes that have the same binary encoding as chromosome i except that one bit has been changed from one into zero. Conversely, $\mathcal{H}(i)$ denotes the set of all chromosomes that have the same binary encoding as chromosome i except that one bit has been changed from zero into one. In mathematical notation,

$$\mathcal{G}(i) = \{j \mid j < i \text{ and } \delta(i, j) = 1\}$$
$$\mathcal{H}(i) = \{j \mid j > i \text{ and } \delta(i, j) = 1\}.$$

Notice that $j \in \mathcal{G}(i)$ if and only if $i \in \mathcal{H}(j)$. There are

$$\nu = \mu m / 2 = m 2^{m-1}$$

combinations of two chromosomes i and j such that $\delta(i, j) = 1$, that is, such that the binary encodings of i and j differ by exactly one bit. k and k' denote indices that take values in $\{1, \ldots, \nu\}$. $\widetilde{\mathcal{V}}$ denotes the set of all populations in which there are exactly two different chromosomes and in which the binary encodings of these chromosomes differ

---

[1]In this chapter, we use a standard binary encoding. Hence, if $m = 2$, chromosomes 0, 1, 2, and 3 have binary encodings 00, 01, 10, and 11, respectively. We emphasize that the use of a standard binary encoding is by no means essential for our analysis. Other binary encoding schemes, such as Gray encoding, can be used as well. This does not require any significant changes in our analysis.

by exactly one bit. There are

$$\xi = |\widetilde{\mathcal{V}}| = \nu(n-1) = (n-1)m2^{m-1}$$

such populations. (The order of the chromosomes within a population has no effect on the behavior of a GA. Populations consisting of the same chromosomes in different orders are therefore considered identical.) $\mathcal{V}$ denotes the set that is obtained by adding the uniform populations to $\widetilde{\mathcal{V}}$, that is, $\mathcal{V} = \widetilde{\mathcal{V}} \cup \mathcal{U}$. For i and j such that $\delta(i,j) = 1$ and for $\lambda \in \{0, \ldots, n\}$, $v(i, j, \lambda) \in \mathcal{V}$ denotes the population consisting of λ times chromosome i and $n - \lambda$ times chromosome j. Notice that $v(i, j, \lambda) = v(j, i, n - \lambda)$ and that $v(i, j, 0) = u(j)$ and $v(i, j, n) = u(i)$. $\mathcal{W}$ denotes the set of all possible populations. As shown by Nix and Vose (1992, Lemma 1) and Dawid (1996), the number of possible populations equals

$$|\mathcal{W}| = \binom{n + \mu - 1}{\mu - 1} = \frac{(n + \mu - 1)!}{n!(\mu - 1)!}.$$

(Again, populations consisting of the same chromosomes in different orders are considered identical.) For $t \in \{0, 1, \ldots\}$, the random variable $W_t \in \mathcal{W}$ denotes the population at the beginning of iteration t of a GA. For i and j such that $\delta(i,j) = 1$ and for $\lambda \in \{1, \ldots, n-1\}$ and $\lambda' \in \{0, \ldots, n\}$, $\pi(i, j, \lambda, \lambda')$ denotes the limit as the mutation rate ε approaches zero of the probability that population $v(i, j, \lambda)$ is turned into population $v(i, j, \lambda')$ in a single iteration of a GA. In mathematical notation,

$$\pi(i, j, \lambda, \lambda') = \lim_{\varepsilon \to 0} \Pr(W_{t+1} = v(i, j, \lambda') \,|\, W_t = v(i, j, \lambda)) \tag{4.1}$$

where $t \in \{0, 1, \ldots\}$. Because the binary encodings of the chromosomes i and j differ by only one bit, the crossover operator has no effect on $\pi(i, j, \lambda, \lambda')$. Moreover, because ε approaches zero, the mutation operator has no effect on $\pi(i, j, \lambda, \lambda')$ either. $\pi(i, j, \lambda, \lambda')$ therefore equals the probability that the selection operator turns population $v(i, j, \lambda)$ into population $v(i, j, \lambda')$ in a single iteration of a GA.

The following definition introduces the notion of almost uniform populations.

Definition 4.2. A non-uniform population $w \in \mathcal{W} \setminus \mathcal{U}$ is said to be *almost uniform* if and only if

$$\lim_{\varepsilon \to 0} \Pr(W_{t+N} = u \,|\, W_t = w) > 0$$

Table 4.1: Overview of the mathematical notation.

$\mathcal{C}$	Set of all chromosomes
$\mathcal{G}(i)$	Set of all chromosomes that have the same binary encoding as chromosome i except that one bit has been changed from one into zero
$\mathcal{H}(i)$	Set of all chromosomes that have the same binary encoding as chromosome i except that one bit has been changed from zero into one
m	Chromosome length
n	Population size
$\bar{q}(w)$	Long-run probability of population w
$\hat{q}(w)$	Long-run limit probability of population w
$\hat{\mathbf{q}}$	Long-run limit distribution
$\mathcal{U}$	Set of all uniform populations
$u(i)$	Uniform population consisting of n times chromosome i
$\mathcal{V}$	Set of all populations in which there are at most two different chromosomes and in which the binary encodings of chromosomes differ by at most one bit
$v(i,j,\lambda)$	Population consisting of λ times chromosome i and $n - \lambda$ times chromosome j
$\mathcal{W}$	Set of all populations
W_t	Population at the beginning of iteration t of a GA
γ	Crossover rate
$\delta(i,j)$	Hamming distance between chromosomes i and j
ε	Mutation rate
μ	Number of different chromosomes Number of uniform populations
ν	Number of combinations of two chromosomes whose binary encodings differ by exactly one bit
ξ	Number of populations in which there are exactly two different chromosomes and in which the binary encodings of chromosomes differ by at most one bit
$\pi(i,j,\lambda,\lambda')$	Probability that the selection operator turns population $v(i,j,\lambda)$ into population $v(i,j,\lambda')$ in a single iteration of a GA

for all $t \in \{0, 1, \ldots\}$, some finite positive integer N, and some $u \in \mathcal{U}$.

Hence, a non-uniform population is almost uniform if and only if no mutation is required to go from the non-uniform population to some uniform population. We note that in many cases all non-uniform populations are almost uniform. For example, if a GA uses roulette wheel selection or tournament selection, the selection operator can turn any non-uniform population into a uniform population in a single iteration and, consequently, all non-uniform populations are almost uniform.

The following two definitions introduce the notion of a connection from one chromosome to another.

Definition 4.3. A *direct connection* from chromosome i to chromosome j is said to exist if and only if $\delta(i, j) = 1$ and

$$\lim_{\varepsilon \to 0} \Pr(W_{t+N} = u(j) \,|\, W_t = v(i, j, n - 1)) > 0$$

for all $t \in \{0, 1, \ldots\}$ and some finite positive integer N.

Definition 4.4. A *connection* from chromosome i to chromosome j is said to exist if and only if there exists a sequence $(i_1, \ldots, i_N)$ such that $i_1, \ldots, i_N \in \mathcal{C}$, $i_1 = i$, $i_N = j$, and i_M is directly connected to i_{M+1} for all $M \in \{1, \ldots, N - 1\}$.

Definition 4.3 states that there is a direct connection from chromosome i to chromosome j if and only if the minimum number of mutations required to go from uniform population $u(i)$ to uniform population $u(j)$ is one. We note that in many cases all chromosomes i and j such that $\delta(i, j) = 1$ have mutual direct connections. This is for example the case if a GA uses roulette wheel selection and the fitness of a chromosome is always positive. Definition 4.4 states that there is a connection from chromosome i to chromosome j if and only if there is a sequence of chromosomes starting at i and ending at j such that each chromosome in the sequence is directly connected to its successor. Clearly, if all chromosomes i and j such that $\delta(i, j) = 1$ have mutual direct connections, then each chromosome is connected to all other chromosomes.

It is well-known that the population in the current iteration of a GA has no effect on the behavior of the GA in the long run (e.g., Nix & Vose, 1992; Dawid, 1996). More specifically, the population an infinite number of iterations in the future is statistically

independent of the population in the current iteration. The following lemma states this result in a formal way.

Lemma 1. *For each population $w \in \mathcal{W}$, there exists a* long-run probability $\bar{q}(w)$ *such that*

$$\lim_{N \to \infty} \Pr(W_{t+N} = w \,|\, W_t = w_t) = \bar{q}(w) \tag{4.2}$$

for all $t \in \{0, 1, \ldots\}$ and all $w_t \in \mathcal{W}$.

Proof. See the appendix.

In our analysis, we are concerned with the long-run behavior of GAs in the limit as the mutation rate ε approaches zero. We therefore use the following definition.

Definition 4.5. For $w \in \mathcal{W}$, $\hat{q}(w) = \lim_{\varepsilon \to 0} \bar{q}(w)$ is called the *long-run limit probability* of population w.

We now introduce the vectors and matrices that we need to state the main result of our analysis. We first note that throughout this chapter vectors and matrices are represented by, respectively, bold lowercase and bold uppercase letters and that the transpose of a matrix $\mathbf{X}$ is written as $\mathbf{X}^{\mathsf{T}}$. $\mathbf{I}_N$ denotes an identity matrix of order $N \times N$, and $\mathbf{0}_{M \times N}$ and $\mathbf{1}_{M \times N}$ denote matrices of order $M \times N$ in which all elements are equal to, respectively, zero and one. We simply write $\mathbf{I}$, $\mathbf{0}$, or $\mathbf{1}$ when the order of a matrix is clear from the context. $\mathbf{g} = [g_k]$ and $\mathbf{h} = [h_k]$ denote vectors of length ν that satisfy

$$\forall k : g_k, h_k \in \mathcal{C}$$
$$\forall k : h_k \in \mathcal{H}(g_k)$$
$$\forall k, k' : k \neq k' \Rightarrow (g_k, h_k) \neq (g_{k'}, h_{k'}).$$

Hence, for each k, (g_k, h_k) denotes a combination of two chromosomes such that the binary encodings of the chromosomes differ by exactly one bit. $\mathbf{g}$ and $\mathbf{h}$ together contain all such combinations of two chromosomes. $\mathbf{A}$, $\mathbf{B}$, $\mathbf{C}$, and $\mathbf{D}$ denote matrices of order

$\mu \times \xi$, $\xi \times \mu$, $\xi \times \xi$, and $\mu \times \mu$, respectively. Matrix $\mathbf{A}$ is given by

$$\mathbf{A} = \begin{bmatrix} \mathbf{a}(0,1) & \cdots & \mathbf{a}(0,\nu) \\ \vdots & \ddots & \vdots \\ \mathbf{a}(\mu-1,1) & \cdots & \mathbf{a}(\mu-1,\nu) \end{bmatrix} \tag{4.3}$$

where

$$\mathbf{a}(i,k) = \begin{cases} \tilde{\mathbf{a}}_1, & \text{if } g_k = i \\ \tilde{\mathbf{a}}_2, & \text{if } h_k = i \\ \mathbf{0}_{1 \times (n-1)}, & \text{otherwise} \end{cases} \tag{4.4}$$

and

$$\tilde{\mathbf{a}}_1 = \begin{bmatrix} \mathbf{0}_{1 \times (n-2)} & 1 \end{bmatrix} \qquad\qquad \tilde{\mathbf{a}}_2 = \begin{bmatrix} 1 & \mathbf{0}_{1 \times (n-2)} \end{bmatrix}. \tag{4.5}$$

Matrix $\mathbf{B}$ is given by

$$\mathbf{B} = \begin{bmatrix} \mathbf{b}(1,0) & \cdots & \mathbf{b}(1,\mu-1) \\ \vdots & \ddots & \vdots \\ \mathbf{b}(\nu,0) & \cdots & \mathbf{b}(\nu,\mu-1) \end{bmatrix} \tag{4.6}$$

where

$$\mathbf{b}(k,i) = \begin{cases} \tilde{\mathbf{b}}(g_k, h_k, n), & \text{if } g_k = i \\ \tilde{\mathbf{b}}(g_k, h_k, 0), & \text{if } h_k = i \\ \mathbf{0}_{(n-1) \times 1}, & \text{otherwise} \end{cases} \tag{4.7}$$

and

$$\tilde{\mathbf{b}}(i,j,\lambda) = \begin{bmatrix} \pi(i,j,1,\lambda) \\ \vdots \\ \pi(i,j,n-1,\lambda) \end{bmatrix}. \tag{4.8}$$

Matrix $\mathbf{C}$ is given by

$$\mathbf{C} = \begin{bmatrix} \mathbf{C}(1,1) & \cdots & \mathbf{C}(1,\nu) \\ \vdots & \ddots & \vdots \\ \mathbf{C}(\nu,1) & \cdots & \mathbf{C}(\nu,\nu) \end{bmatrix} \tag{4.9}$$

where

$$\mathbf{C}(k,k') = \begin{cases} \widetilde{\mathbf{C}}(g_k, h_k), & \text{if } k = k' \\ \mathbf{0}_{(n-1)\times(n-1)}, & \text{otherwise} \end{cases} \tag{4.10}$$

and

$$\widetilde{\mathbf{C}}(i,j) = \begin{bmatrix} \pi(i,j,1,1) & \cdots & \pi(i,j,1,n-1) \\ \vdots & \ddots & \vdots \\ \pi(i,j,n-1,1) & \cdots & \pi(i,j,n-1,n-1) \end{bmatrix}. \tag{4.11}$$

Matrix $\mathbf{D}$ is obtained from $\mathbf{A}$, $\mathbf{B}$, and $\mathbf{C}$ and is given by

$$\mathbf{D} = \mathbf{A}(\mathbf{I} - \mathbf{C})^{-1}\mathbf{B} - m\mathbf{I}. \tag{4.12}$$

The following theorem states the main result of our analysis.

Theorem 4.1. *Let all non-uniform populations be almost uniform, and let each chromosome in $\mathcal{C}$ be connected to all other chromosomes in $\mathcal{C}$. Then, (i) all non-uniform populations have a long-run limit probability of zero, that is, $\hat{q}(w) = 0$ for all $w \in \mathcal{W} \setminus \mathcal{U}$, and (ii) the long-run limit distribution $\hat{\mathbf{q}} = \begin{bmatrix} \hat{q}(u(0)) & \cdots & \hat{q}(u(\mu-1)) \end{bmatrix}$ satisfies*

$$\hat{\mathbf{q}}\mathbf{D} = \mathbf{0} \tag{4.13}$$

$$\hat{\mathbf{q}}\mathbf{1} = 1 \tag{4.14}$$

which has a unique solution.

Proof. See the appendix.

There are three comments that we would like to make on the above theorem. First, the result that under certain assumptions non-uniform populations have a long-run limit probability of zero is not new. A similar result can be found in Dawid (1996, Proposition 4.2.1). Second, under the assumptions of the theorem, the long-run limit probability

of a population does not depend on the crossover rate γ. This is a quite remarkable result that, to the best of our knowledge, has not been reported before in the theoretical literature on GAs. It indicates that in the limit as the mutation rate ε approaches zero γ has no effect on the long-run behavior of a GA. Third, the theorem can be used to calculate the long-run limit distribution $\hat{q}$ only if the probabilities $\pi(i, j, \lambda, \lambda')$ defined in (4.1) can be calculated for all i and all j such that $\delta(i, j) = 1$ and for all $\lambda \in \{1, \ldots, n-1\}$ and all $\lambda' \in \{0, \ldots, n\}$. Whether this is possible depends on the way in which the fitness of a chromosome is determined and on the selection operator that is used. This in turn depends heavily on the specific problem that one wants to model using a GA. Because of the dependence on the problem to be modeled, we cannot provide any general results for the calculation of the probabilities $\pi(i, j, \lambda, \lambda')$. In Section 4.4, however, we demonstrate how the probabilities $\pi(i, j, \lambda, \lambda')$ can be calculated for a GA that is similar to the GA used by Axelrod in his seminal paper on GA modeling (Axelrod, 1987).

4.3   Algorithm

In this section, we present an algorithm for calculating the long-run limit distribution $\hat{q}$. The algorithm is based on Theorem 4.1. Like Theorem 4.1, it assumes that all non-uniform populations are almost uniform and that each chromosome in $\mathcal{C}$ is connected to all other chromosomes in $\mathcal{C}$. It also assumes that the probabilities $\pi(i, j, \lambda, \lambda')$ defined in (4.1) can be calculated for all i and all j such that $\delta(i, j) = 1$ and for all $\lambda \in \{1, \ldots, n-1\}$ and all $\lambda' \in \{0, \ldots, n\}$.

The most straightforward approach to calculating the long-run limit distribution $\hat{q}$ would be to start with calculating the matrices $\mathbf{A}$, $\mathbf{B}$, and $\mathbf{C}$ using (4.3)–(4.11). Matrix $\mathbf{D}$ would then be calculated using (4.12), which would require solving a linear system. Finally, $\hat{q}$ would be obtained by solving the linear system given by (4.13) and (4.14). Unfortunately, this approach to calculating $\hat{q}$ requires a lot of computer memory and is therefore infeasible even for problems of only moderate size. Most memory is required for storing matrix $\mathbf{C}$. This matrix has (at most) $\nu(n-1)^2 = (n-1)^2 m 2^{m-1}$ non-zero elements. Clearly, as the population size n and the chromosome length m increase, storing the non-zero elements of $\mathbf{C}$ in a computer's main memory soon becomes infeasible. The algorithm that we propose for calculating $\hat{q}$ exploits the sparsity of the matrices $\mathbf{A}$,

B, and C in order to calculate matrix D in a memory-efficient way. The algorithm does not require the entire matrices A, B, and C to be stored in memory. The algorithm also solves the linear system given by (4.13) and (4.14) in a memory-efficient way. This is achieved by exploiting the sparsity of D. The algorithm is shown in Figure 4.2. We now discuss it in more detail.

We first consider the efficient calculation of matrix D. Let $\widehat{C} = (I - C)^{-1}$. Because C is a block diagonal matrix, $\widehat{C}$ can be written as

$$\widehat{C} = \begin{bmatrix} \widehat{C}(1,1) & \cdots & \widehat{C}(1,\nu) \\ \vdots & \ddots & \vdots \\ \widehat{C}(\nu,1) & \cdots & \widehat{C}(\nu,\nu) \end{bmatrix}$$

where

$$\widehat{C}(k,k') = \begin{cases} (I - \widetilde{C}(g_k, h_k))^{-1}, & \text{if } k = k' \\ 0_{(n-1)\times(n-1)}, & \text{otherwise.} \end{cases}$$

Hence, $\widehat{C}$ is a block diagonal matrix too. Let D be written as

$$D = \begin{bmatrix} d(0,0) & \cdots & d(0,\mu-1) \\ \vdots & \ddots & \vdots \\ d(\mu-1,0) & \cdots & d(\mu-1,\mu-1) \end{bmatrix}.$$

Taking into account the sparsity of A, B, and $\widehat{C}$, it can be seen that

$$d(i,j) = \begin{cases} \sum_{i' \in \mathcal{G}(i)} \tilde{a}_2 \tilde{e}(i',i,0) + \sum_{i' \in \mathcal{H}(i)} \tilde{a}_1 \tilde{e}(i,i',n) - m, & \text{if } i = j \\ \tilde{a}_2 \tilde{e}(j,i,n), & \text{if } j \in \mathcal{G}(i) \\ \tilde{a}_1 \tilde{e}(i,j,0), & \text{if } j \in \mathcal{H}(i) \\ 0, & \text{otherwise} \end{cases} \quad (4.15)$$

where

$$\tilde{e}(i,j,\lambda) = (I - \widetilde{C}(i,j))^{-1} \tilde{b}(i,j,\lambda).$$

This result shows that each non-zero element of D can be calculated by solving one or

---

Input: n, m, $\hat{\mathbf{q}}_0$, and ω

Output: $\hat{\mathbf{q}}$

1 // Calculation of $\mathbf{D} = \Big[d(i,j)\Big]$

2 // Only non-zero elements of $\mathbf{D}$ should be stored

3 $\mu \leftarrow 2^m$

4 $\mathbf{D} \leftarrow -m\mathbf{I}_\mu$

5 $\tilde{\mathbf{a}}_1 \leftarrow \tilde{\mathbf{a}}_1$ given by (4.5)

6 $\tilde{\mathbf{a}}_2 \leftarrow \tilde{\mathbf{a}}_2$ given by (4.5)

7 **for** $i \leftarrow 0$ **to** $\mu - 1$ **do**

8   **for all** $j \in \mathcal{H}(i)$ **do**

9          $\tilde{\mathbf{b}} \leftarrow \tilde{\mathbf{b}}(i,j,n)$ given by (4.8)

10          $\widetilde{\mathbf{C}} \leftarrow \widetilde{\mathbf{C}}(i,j)$ given by (4.11)

11          $\tilde{\mathbf{e}} \leftarrow (\mathbf{I} - \widetilde{\mathbf{C}})^{-1}\tilde{\mathbf{b}}$ // Use, e.g., Gaussian elimination

12          $d(i,i) \leftarrow d(i,i) + \tilde{\mathbf{a}}_1\tilde{\mathbf{e}}$

13          $d(j,j) \leftarrow d(j,j) + 1 - \tilde{\mathbf{a}}_2\tilde{\mathbf{e}}$

14          $d(i,j) \leftarrow 1 - \tilde{\mathbf{a}}_1\tilde{\mathbf{e}}$

15          $d(j,i) \leftarrow \tilde{\mathbf{a}}_2\tilde{\mathbf{e}}$

16   **end for**

17 **end for**

18 // Calculation of $\hat{\mathbf{q}} = \Big[\hat{q}(u(i))\Big]$

19 // The linear system given by (4.13) and (4.14) will be solved using successive overrelaxation

20 $\hat{\mathbf{q}} \leftarrow \hat{\mathbf{q}}_0$

21 **repeat**

22   **for** $i \leftarrow 0$ **to** $\mu - 1$ **do**

23          $\sigma \leftarrow 0$

24   **for all** $j \in \mathcal{G}(i) \cup \mathcal{H}(i)$ **do**

25              $\sigma \leftarrow \sigma + \hat{q}(u(j))d(j,i)$

26   **end for**

27          $\sigma \leftarrow -\sigma/d(i,i)$

28          $\hat{q}(u(i)) \leftarrow (1 - \omega)\hat{q}(u(i)) + \omega\sigma$

29   **end for**

30 **until** some convergence criterion is satisfied

31 $\hat{\mathbf{q}} \leftarrow \hat{\mathbf{q}}/(\hat{\mathbf{q}}\mathbf{1})$

Figure 4.2: Algorithm for calculating the long-run limit distribution of a genetic algorithm.

more relatively small linear systems, that is, systems of $n-1$ equations and unknowns. Moreover, by calculating the elements of $\mathbf{D}$ one by one, there is no need to store the entire matrices $\mathbf{A}$, $\mathbf{B}$, and $\mathbf{C}$ in memory. Solving a linear system of $n-1$ equations and unknowns can be done using standard Gaussian elimination methods. Except for very large values for the population size n, today's computers have sufficient main memory to apply Gaussian elimination methods to such systems. We further note that the amount of computation required for obtaining $\mathbf{D}$ can be reduced by taking into account that

$$
\begin{aligned}
\tilde{\mathbf{e}}(i,j,0) &= (\mathbf{I} - \widetilde{\mathbf{C}}(i,j))^{-1}\tilde{\mathbf{b}}(i,j,0) \\
&= (\mathbf{I} - \widetilde{\mathbf{C}}(i,j))^{-1}(1 - \textstyle\sum_{\lambda=1}^{n}\tilde{\mathbf{b}}(i,j,\lambda)) \\
&= (\mathbf{I} - \widetilde{\mathbf{C}}(i,j))^{-1}(1 - \widetilde{\mathbf{C}}(i,j)\mathbf{1} - \tilde{\mathbf{b}}(i,j,n)) \\
&= (\mathbf{I} - \widetilde{\mathbf{C}}(i,j))^{-1}(\mathbf{I} - \widetilde{\mathbf{C}}(i,j))\mathbf{1} - \tilde{\mathbf{e}}(i,j,n) \\
&= \mathbf{1} - \tilde{\mathbf{e}}(i,j,n).
\end{aligned}
$$

Because of this, $d(i,j)$ can be written as

$$
d(i,j) = \begin{cases}
\sum_{i'\in\mathcal{G}(i)}(1 - \tilde{\mathbf{a}}_{\mathbf{2}}\tilde{\mathbf{e}}(i',i,n)) + \sum_{i'\in\mathcal{H}(i)}\tilde{\mathbf{a}}_{\mathbf{1}}\tilde{\mathbf{e}}(i,i',n) - m, & \text{if } i = j \\
\tilde{\mathbf{a}}_{\mathbf{2}}\tilde{\mathbf{e}}(j,i,n), & \text{if } j \in \mathcal{G}(i) \\
1 - \tilde{\mathbf{a}}_{\mathbf{1}}\tilde{\mathbf{e}}(i,j,n), & \text{if } j \in \mathcal{H}(i) \\
0, & \text{otherwise.}
\end{cases}
\tag{4.16}
$$

Using (4.16) rather than (4.15) to calculate $\mathbf{D}$ halves the number of linear systems that need to be solved. In the algorithm in Figure 4.2, the calculation of $\mathbf{D}$ based on (4.16) is performed between lines 1 and 17.

Matrix $\mathbf{D}$ has $\mu^2 = 2^{2m}$ elements. Consequently, storing all elements of $\mathbf{D}$ in a computer's main memory is possible only if the chromosome length m is not too large. It follows from (4.15) and (4.16) that the number of non-zero elements in $\mathbf{D}$ equals $\mu(m+1) = (m+1)2^m$. Hence, $\mathbf{D}$ is a rather sparse matrix and a lot of memory can be saved by storing only its non-zero elements.[2] In addition to the memory efficiency of

---

[2]The non-zero elements of $\mathbf{D}$ can be stored efficiently by using two arrays: a one-dimensional array of size μ for the diagonal elements of $\mathbf{D}$ and a two-dimensional array of size $m \times \mu$ for the non-zero off-diagonal elements of $\mathbf{D}$. The element in the κth row and the ith column of the latter array is used to store $d(j,i)$, where j has the same binary encoding as i except that the κth bit is inverted.

the way in which $\mathbf{D}$ is stored, one should also pay attention to the memory efficiency of the method that is used to solve the linear system given by (4.13) and (4.14). Gaussian elimination and other direct (i.e., non-iterative) methods for solving linear systems generally require that at least a large number of elements of the coefficient matrix, including zero elements, are stored in memory. Consequently, when using such a method to solve the linear system given by (4.13) and (4.14), it would not be possible to fully exploit the sparsity of $\mathbf{D}$. Linear systems can also be solved using iterative methods that require only the non-zero elements of the coefficient matrix to be stored in memory. One such method is the method of successive overrelaxation (e.g., Stewart, 1994; Tijms, 1994, 2003; Barrett et al., 2008). In the algorithm in Figure 4.2, this method is used to solve the linear system given by (4.13) and (4.14) (see lines 18–31 of the algorithm). In addition to an initial guess $\hat{\mathbf{q}}_0$ for the solution of the linear system, the method of successive overrelaxation also requires a value for the relaxation factor ω. The value of ω, which should be between 0 and 2, may have a large effect on the rate of convergence of the method, and for some values of ω the method may not converge at all. An appropriate value for ω has to be determined experimentally. For $\omega = 1$, the method of successive overrelaxation reduces to the Gauss-Seidel method, which is another iterative method for solving linear systems. We refer to Stewart (1994) for an in-depth discussion of both the method of successive overrelaxation and a number of alternative methods for solving linear systems similar to the one given by (4.13) and (4.14). We further note that the amount of main memory in most of today's computers allows the algorithm in Figure 4.2 to be run for chromosomes with length m up to about 24 bits.

4.4   Application

In this section, we demonstrate an application of the algorithm presented in the previous section. We study the use of a GA for modeling the evolution of strategies in iterated prisoner's dilemmas (IPDs). The use of GAs in this context was first studied by Axelrod (1987, 1997b; see also Dawid, 1996; Mitchell, 1996) and after him by many others (e.g., Mühlenbein, 1991; Yao & Darwen, 1994; Ashlock et al., 1996; Crowley et al., 1996; Miller, 1996; Van Bragt et al., 2001; Alkemade et al., 2005; Ishibuchi & Namikawa, 2005; Ashlock et al., 2006; Mittal & Deb, 2006; Thibert-Plante & Charbonneau, 2007;

Ashlock & Kim, 2008). The algorithm presented in the previous section is used to analyze the long-run behavior of our GA. The results of the analysis are compared with results obtained using computer simulations. We emphasize that our primary aim is merely to illustrate the usefulness of the mathematical analysis provided in Section 4.2 and of the algorithm derived from the analysis in Section 4.3. It is not our primary aim to provide new insights into the behavior of GAs in the context of IPDs.

4.4.1 Genetic algorithm modeling in iterated prisoner's dilemmas

The way in which we model the evolution of strategies in IPDs is similar to the way in which this was done by Axelrod (1987). However, Axelrod studied two approaches for modeling the evolution of strategies. In one approach, the fitness of a chromosome is determined by the performance of the chromosome in IPD games against a fixed set of opponents. In the other approach, the fitness of a chromosome is determined by the performance of the chromosome in IPD games against other chromosomes in the population. We restrict our attention to the second approach. This is the approach on which almost all studies after Axelrod's work have focused (an exception is Mittal & Deb, 2006).

We model the evolution of strategies in IPDs using a GA with a population size of $n = 20$ chromosomes. Each chromosome represents a strategy for playing IPD games. Players in IPD games are assumed to choose the action they play, that is, whether they cooperate or defect, based on their own actions and their opponent's actions in the previous τ periods of the game, where τ is referred to as players' memory length. Players are further assumed to play only pure strategies. We use the same binary encoding of strategies as was used by Axelrod (1987). For a description of this encoding, we refer to Dawid (1996); Mitchell (1996); Axelrod (1987, 1997b). Using Axelrod's encoding, the chromosome length m depends on the memory length τ. We consider three memory lengths, 1, 2, and 3 periods, which result in chromosome lengths of, respectively, 6, 20, and 70 bits. In each iteration of the GA, each chromosome in the population plays an IPD game of 151 periods against all other chromosomes. In addition, each chromosome also plays a game against itself. The payoff matrix for a single period of an IPD game

Table 4.2: Payoff matrix for a single period of an iterated prisoner's dilemma game. The payoff obtained by the row (column) player is reported first (second).

	Cooperate	Defect
Cooperate	R, R	S, T
Defect	T, S	P, P

is shown in Table 4.2. The payoffs in this matrix must satisfy

$$S < P < R < T$$

and

$$S + T < 2R.$$

The payoff obtained by a chromosome in an IPD game equals the mean payoff obtained by the chromosome in all periods of the game. The fitness f of a chromosome equals the mean payoff obtained by the chromosome in the IPD games that it has played in the current iteration of the GA. Like in Axelrod's work (Axelrod, 1987), we use sigma scaling (e.g., Mitchell, 1996) to normalize the fitness of a chromosome. The normalized fitness $\tilde{f}$ of a chromosome is given by

$$\tilde{f} = \begin{cases} \max\left(\dfrac{f - \mu_f}{\sigma_f} + 1, 0\right), & \text{if } \sigma_f > 0 \\ 1, & \text{otherwise} \end{cases} \tag{4.17}$$

where μ_f and σ_f denote, respectively, the mean and the standard deviation of the fitness of the chromosomes in the population. The selection operator that we use is roulette wheel selection. Selection is performed based on the normalized fitness of the chromosomes in the population. The crossover operator that we use is single-point crossover.

4.4.2 Calculation of the long-run limit distribution of the genetic algorithm

In this subsection, we are concerned with the calculation of the long-run limit distribution of the GA discussed in the previous subsection. To calculate the long-run limit

distribution of the GA, we use the algorithm presented in Section 4.3. This algorithm assumes that the probabilities $\pi(i, j, \lambda, \lambda')$ defined in (4.1) can be calculated for all i and all j such that $\delta(i, j) = 1$ and for all $\lambda \in \{1, \ldots, n-1\}$ and all $\lambda' \in \{0, \ldots, n\}$. We now discuss the calculation of the probabilities $\pi(i, j, \lambda, \lambda')$ for our GA. For $i', j' \in \mathcal{C}$, let $\varphi(i', j')$ denote the payoff obtained by chromosome i' in an IPD game against chromosome j'. Suppose that the population in the current iteration of our GA equals $v(i, j, \lambda)$, where i and j satisfy $\delta(i, j) = 1$ and where $\lambda \in \{1, \ldots, n-1\}$. That is, the population in the current iteration of our GA consists of λ times chromosome i and $n - \lambda$ times chromosome j. The fitness f_i of chromosome i is then given by

$$f_i = \frac{\lambda\varphi(i, i) + (n - \lambda)\varphi(i, j)}{n}.$$

Similarly, the fitness f_j of chromosome j is given by

$$f_j = \frac{\lambda\varphi(j, i) + (n - \lambda)\varphi(j, j)}{n}.$$

Furthermore, the mean μ_f and the standard deviation σ_f of the fitness of the chromosomes in the population are equal to, respectively,

$$\mu_f = \frac{\lambda f_i + (n - \lambda)f_j}{n}$$

and

$$\sigma_f = \sqrt{\frac{\lambda(f_i - \mu_f)^2 + (n - \lambda)(f_j - \mu_f)^2}{n}}.$$

The normalized fitness $\tilde{f}_i$ of chromosome i is obtained by substituting f_i, μ_f, and σ_f into (4.17). The normalized fitness $\tilde{f}_j$ of chromosome j is obtained in a similar way. Let $\tilde{\pi}_i$ and $\tilde{\pi}_j$ denote the probabilities that the roulette wheel selection operator selects, respectively, chromosome i and chromosome j. Obviously, $\tilde{\pi}_i$ and $\tilde{\pi}_j$ equal

$$\tilde{\pi}_i = \frac{\lambda\tilde{f}_i}{\lambda\tilde{f}_i + (n - \lambda)\tilde{f}_j} \qquad\qquad \tilde{\pi}_j = \frac{(n - \lambda)\tilde{f}_j}{\lambda\tilde{f}_i + (n - \lambda)\tilde{f}_j}.$$

$\pi(i, j, \lambda, \lambda')$, where $\lambda' \in \{0, \ldots, n\}$, equals the probability that the roulette wheel selection operator turns population $v(i, j, \lambda)$ into population $v(i, j, \lambda')$ in a single iteration of

our GA. Taking into account that the roulette wheel selection operator selects chromosomes independently of each other, it can be seen that $\pi(i, j, \lambda, \lambda')$ equals the probability mass function of a binomial distribution and is given by

$$\pi(i, j, \lambda, \lambda') = \binom{n}{\lambda'} \tilde{\pi}_i^{\lambda'} \tilde{\pi}_j^{n-\lambda'}$$

where the binomial coefficient $\binom{n}{\lambda'}$ is defined as

$$\binom{n}{\lambda'} = \frac{n!}{\lambda'!(n - \lambda')!}.$$

The algorithm presented in Section 4.3 also assumes that all non-uniform populations are almost uniform and that each chromosome in $\mathcal{C}$ is connected to all other chromosomes in $\mathcal{C}$. Because of the use of roulette wheel selection, the assumption that all non-uniform populations are almost uniform is satisfied. The assumption that each chromosome in $\mathcal{C}$ is connected to all other chromosomes in $\mathcal{C}$ is satisfied if and only if matrix $\mathbf{D}$ calculated in lines 1–17 of the algorithm in Figure 4.2 is irreducible. ($\mathbf{D} = \left[d(i, j) \right]$ is said to be irreducible if and only if there does not exist a non-empty set of chromosomes $\widetilde{\mathcal{C}} \subset \mathcal{C}$ such that $d(i, j) = 0$ for all $i \in \widetilde{\mathcal{C}}$ and all $j \in \mathcal{C} \setminus \widetilde{\mathcal{C}}$.) For the particular values that we use for the parameters S, P, R, T, and τ (see the next subsection), $\mathbf{D}$ turns out to be irreducible. Hence, the assumption that each chromosome in $\mathcal{C}$ is connected to all other chromosomes in $\mathcal{C}$ is satisfied.

4.4.3 Analysis of the long-run behavior of the genetic algorithm

In this subsection, we analyze the long-run behavior of our GA for the prisoner's dilemma payoffs $S = 0$, $P = 1$, $R = 3$, and $T = 5$. These are the same payoffs as were used by Axelrod (1987) (see also Axelrod, 1984) and by many others. The analysis is performed using the algorithm presented in Section 4.3. The use of this algorithm to analyze the long-run behavior of our GA was discussed in the previous subsection. We compare the results obtained using the algorithm with results obtained using computer simulations.[3]

---

[3]The software used to obtain the results reported in this subsection is available online at www.ludowaltman.nl/ga_analysis/. The software runs in MATLAB and has been written partly in the MATLAB programming language and partly in the C programming language.
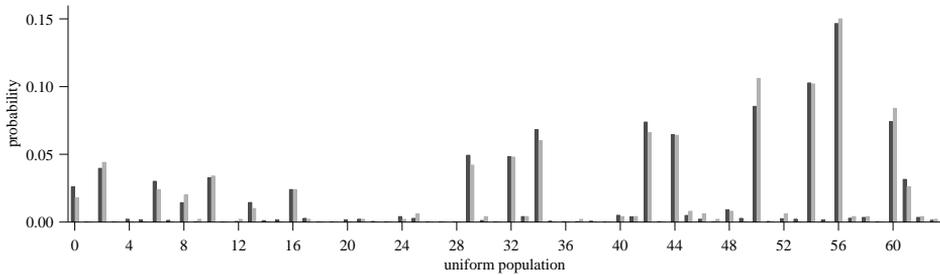
Figure 4.3: The long-run limit distribution calculated using the algorithm presented in Section 4.3 (in dark grey) and a probability distribution over the uniform populations estimated using computer simulations (in light grey). The memory length τ equals 1. On the horizontal axis, integers between 0 and 63 are used to represent the uniform populations. Integer i represents the uniform population consisting of 20 times chromosome i.

The long-run limit distribution for a memory length of $\tau = 1$ period is shown in Figure 4.3 (in dark grey). The distribution was calculated using the algorithm from Section 4.3. As mentioned before, $\tau = 1$ results in a chromosome length of $m = 6$ bits. This implies that there are $\mu = 2^m = 64$ different chromosomes and, as a consequence, that there are 64 different uniform populations. The long-run limit distribution is a probability distribution over these populations. As can be seen in Figure 4.3, the long-run limit distribution spreads most of its mass over approximately fifteen populations. It puts almost no mass on the remaining populations. Since all chromosomes in a uniform population are identical and represent the same strategy, the long-run limit distribution can be used to determine the long-run limit probability that a particular strategy is played. However, when doing so, it should be noted that there is some redundancy in the binary encoding of strategies that we use (as was already pointed out by Axelrod, 1987). Due to this redundancy, it is possible that different chromosomes represent the same strategy. Some strategies can be encoded in two or three different ways, and the strategies *always cooperate* and *always defect* can even be encoded in twelve different ways. Taking into account the redundancy in the encoding, we have calculated the long-run limit probabilities of all possible strategies. The six strategies with the highest long-run limit probability are reported in Table 4.3. Together, these strategies have a

Table 4.3: The six strategies with the highest long-run limit probability (reported in the first column). The memory length τ equals 1.

Prob.	Strategy	Chromosomes
0.430	Always defect	0, 2, 8, 10, 16, 24, 32, 34, 40, 42, 48, 50
0.147	Start cooperating; cooperate if and only if both you and your opponent cooperated in the previous period	56
0.139	Start cooperating; cooperate if and only if your opponent cooperated in the previous period (tit for tat)	44, 60
0.133	Start defecting; cooperate if and only if you and your opponent played different actions in the previous period	6, 54
0.051	Start cooperating; cooperate unless you cooperated in the previous period and your opponent did not	13, 45, 61
0.049	Start defecting; cooperate unless you cooperated in the previous period and your opponent did not	29

long-run limit probability of almost 0.95. The remaining strategies all have very low long-run limit probabilities. It is sometimes claimed (e.g., Axelrod, 1984, 1987) that a very effective strategy for playing IPD games is the *tit for tat* strategy, which is the strategy of cooperating in the first period and repeating the opponent's previous action thereafter. The results reported in Table 4.3 do not really support this claim. As can be seen in the table, the *always defect* strategy has by far the highest long-run limit probability. In the long run, this strategy is played about 43% of the time. The *tit for tat* strategy has a long-run limit probability of no more than 0.14. This is even slightly less than the long-run limit probability of another cooperative strategy, namely the strategy that keeps cooperating until the opponent defects and then keeps defecting forever.

In order to check the correctness of the algorithm presented in Section 4.3, we have also used computer simulations to analyze the long-run behavior of our GA. Like above, we first focus on the behavior of the GA for a memory length of $\tau = 1$ period. We performed 500 runs of the GA. The crossover rate was set to $\gamma = 1.0$, and the mutation rate was set to $\varepsilon = 10^{-5}$. Because of the very small value of ε, the simulation results should be similar to the results obtained using the algorithm from Section 4.3. (Recall that the latter results hold in the limit as ε approaches zero.) Each run of the GA lasted $2 \cdot 10^5$ iterations. This seemed sufficient for the GA to reach its steady state. After the last iteration of a GA run, we almost always observed that the population was uniform.

Based on the 500 GA runs that we had performed, we estimated for each uniform population the probability of observing that population at the end of a GA run. In this way, we obtained a probability distribution over the uniform populations. This distribution is shown in Figure 4.3 (in light grey). Figure 4.3 allows us to compare the distribution with the long-run limit distribution calculated using the algorithm from Section 4.3. It can be seen that the two distributions are very similar. This confirms the correctness of the algorithm presented in Section 4.3.

In order to examine to what extent our GA results in the evolution of cooperative strategies, we now focus on the long-run mean fitness, that is, the mean fitness of a chromosome after a large number of iterations of the GA. For various values of the memory length τ, the crossover rate γ, and the mutation rate ε, the long-run mean fitness estimated using computer simulations is reported in Table 4.4. The associated 95% confidence interval is also provided in the table. The simulation results for $\tau = 1$ are based on 500 runs of the GA, and the simulation results for $\tau = 2$ and $\tau = 3$ are based on 200 runs. Each run lasted $2 \cdot 10^5$ iterations. The long-run mean fitness was estimated by taking the average over all GA runs of the mean fitness of a chromosome at the end of a run. In the limit as ε approaches zero, the long-run mean fitness can be calculated exactly and does not depend on γ. The calculation of the long-run mean fitness is based on the long-run limit distribution of the GA, which can be obtained using the algorithm presented in Section 4.3. For $\tau = 1$ and $\tau = 2$, the long-run mean fitness in the limit as ε approaches zero is reported in Table 4.4. For $\tau = 3$, we cannot calculate the long-run limit distribution of the GA and we therefore do not know the long-run mean fitness in the limit as ε approaches zero. Calculating the long-run limit distribution of the GA is impossible for $\tau = 3$ because the chromosome length equals $m = 70$ bits and because for such a chromosome length storing the long-run limit distribution requires a prohibitive amount of computer memory.

Based on the results in Table 4.4, a number of observations can be made. First, for $\tau = 1$ and $\tau = 2$, the results obtained for $\varepsilon = 10^{-4}$ and $\varepsilon = 10^{-5}$ turn out to be very similar to the results obtained for $\varepsilon \to 0$. This again confirms the correctness of the algorithm presented in Section 4.3. Second, for $\tau = 1$, we find that the results are quite sensitive to the value of ε. Studies on GA modeling sometimes report that the long-run behavior of a GA is relatively insensitive to the value of ε. Our results demonstrate

that this need not always be the case. Third, for small values of ε, it can be seen that increasing τ leads to a higher long-run mean fitness and, hence, to more cooperation. The evolution of cooperative strategies in IPD games therefore seems more likely when players have longer memory lengths. Finally, it can be observed that the value of γ has no significant effect on our results. This is in line with the mathematical analysis provided in Section 4.2. The mathematical analysis implies that for $\varepsilon \to 0$ the long-run mean fitness is independent of γ. The results in Table 4.4 indicate that this is the case not only for $\varepsilon \to 0$ but more generally.

4.5   Conclusions

In this chapter, we have presented a mathematical analysis of the long-run behavior of GAs that are used for modeling economic phenomena. Under the assumption of a positive but infinitely small mutation rate, the analysis provides a full characterization of the long-run behavior of GAs with a binary encoding. Based on the analysis, we have derived an algorithm for calculating the long-run behavior of GAs. In an economic context, the algorithm can for example be used to determine whether convergence to an equilibrium will take place and, if so, what kind of equilibrium will emerge. Compared with computer simulations, the main advantage of the algorithm that we have derived is that it calculates the long-run behavior of GAs exactly. Computer simulations only estimate the long-run behavior of GAs.

To demonstrate the usefulness of our mathematical analysis, we have replicated a well-known study by Axelrod in which a GA is used to model the evolution of strategies in iterated prisoner's dilemmas (Axelrod, 1987). We have used both our exact algorithm and computer simulations to replicate Axelrod's study. By comparing the results of the two approaches, we have confirmed the correctness of our algorithm. We have also obtained some interesting new insights. For example, when players have a memory length of one period, the *tit for tat* strategy turns out to be less important than is sometimes claimed (e.g., Axelrod, 1984, 1987). In the long run, the strategy is played only 14% of the time. Another finding is that the long-run behavior of a GA can be quite sensitive to the value of the mutation rate. We regard this as a serious problem, since the value

Table 4.4: Estimated long-run mean fitness and associated 95% confidence interval for various values of the memory length τ, the crossover rate γ, and the mutation rate ε. For $\varepsilon \to 0$, the long-run mean fitness has been calculated exactly.

	$\tau = 1$			$\tau = 2$			$\tau = 3$		
	$\gamma = 0.0$	$\gamma = 0.5$	$\gamma = 1.0$	$\gamma = 0.0$	$\gamma = 0.5$	$\gamma = 1.0$	$\gamma = 0.0$	$\gamma = 0.5$	$\gamma = 1.0$
$\varepsilon = 10^{-2}$	2.76 ± 0.05	2.71 ± 0.05	2.79 ± 0.04	2.64 ± 0.08	2.72 ± 0.07	2.67 ± 0.07	2.64 ± 0.07	2.70 ± 0.06	2.67 ± 0.06
$\varepsilon = 10^{-3}$	2.23 ± 0.08	2.24 ± 0.08	2.25 ± 0.08	2.34 ± 0.12	2.41 ± 0.11	2.38 ± 0.11	2.60 ± 0.09	2.59 ± 0.08	2.55 ± 0.09
$\varepsilon = 10^{-4}$	1.93 ± 0.09	1.94 ± 0.09	1.90 ± 0.09	2.25 ± 0.12	2.24 ± 0.12	2.32 ± 0.12	2.57 ± 0.09	2.53 ± 0.09	2.50 ± 0.09
$\varepsilon = 10^{-5}$	1.85 ± 0.09	1.81 ± 0.09	1.85 ± 0.09	2.28 ± 0.12	2.31 ± 0.11	2.22 ± 0.12	2.58 ± 0.09	2.44 ± 0.10	2.44 ± 0.10
$\varepsilon \to 0$	1.84	1.84	1.84	2.29	2.29	2.29	?	?	?

of the mutation rate is typically chosen in a fairly arbitrary way without any empirical justification (see also Dawid, 1996).

The mathematical analysis that we have presented also reveals that if the mutation rate is infinitely small the crossover rate has no effect on the long-run behavior of a GA. This remarkable result is perfectly in line with the simulation results that we have reported in Section 4.4. For various values of the mutation rate, the simulation results show no significant effect of the crossover rate on the long-run behavior of a GA. Hence, when GAs are used for modeling economic phenomena, the crossover rate seems to be a rather unimportant parameter, at least when the focus is on the long run (for the short run, see Thibert-Plante & Charbonneau, 2007). It seems likely that in many cases leaving out the crossover operator altogether has no significant effect on the long-run behavior of a GA. Interestingly, leaving out the crossover operator brings GAs quite close to well-known models in evolutionary game theory, such as those studied by Kandori et al. (1993) and Vega-Redondo (1997).

Finally, we note that an analysis such as the one presented in this chapter can be performed not only for GAs with a binary encoding but also for other types of evolutionary algorithms. From a modeling point of view, a binary encoding in many cases has the disadvantage that it lacks a clear interpretation (e.g., Dawid, 1996). The use of a binary encoding can therefore be difficult to justify and may even lead to artifacts, as is shown in Chapter 5 of this thesis (see also Waltman & Van Eck, 2009). Probably for these reasons, some researchers use evolutionary algorithms without a binary encoding (e.g., Lux & Schornstein, 2005; Haruvy et al., 2006). The analysis presented in this chapter then does not directly apply. However, when the action space of agents is assumed discrete, the long-run behavior of evolutionary algorithms without a binary encoding can still be analyzed in a similar way as we have done in this chapter, namely by relying on mathematical results provided by Freidlin and Wentzell (1998). This indicates that our approach is quite general and can be adapted relatively easily to other types of evolutionary algorithms.

Appendix

In this appendix, we prove the mathematical results presented in Section 4.2. Before proving the results, we first provide some definitions and lemmas on Markov chains.

Definition 4.6. A collection of random variables $\{X_t\}$, where the index t takes values in $\{0, 1, \ldots\}$ and where $X_0, X_1, \ldots$ take values in a finite set $\mathcal{X}$, is called a *finite discrete-time Markov chain* if

$$\Pr(X_{t+1} = x_{t+1} \mid X_t = x_t) = \Pr(X_{t+1} = x_{t+1} \mid X_t = x_t, \ldots, X_0 = x_0)$$

for all t and all $x_0, \ldots, x_{t+1} \in \mathcal{X}$. The elements of $\mathcal{X}$ are called the *states* of the Markov chain. $\mathcal{X}$ is called the *state space* of the Markov chain.

Definition 4.7. A finite discrete-time Markov chain $\{X_t\}$ is said to be *time-homogeneous* if

$$\Pr(X_{t+1} = x_{t+1} \mid X_t = x_t) = p(x_t, x_{t+1})$$

for all t, all $x_t, x_{t+1} \in \mathcal{X}$, and some function $p : \mathcal{X}^2 \to [0, 1]$ that does not depend on t. For $x, x' \in \mathcal{X}$, the probability $p(x, x')$ is called the *transition probability* from state x to state x'. The matrix

$$\mathbf{P} = \Big[p(x, x') \Big]_{x, x' \in \mathcal{X}}$$

is called the *transition matrix* of the Markov chain.

In the remainder of this appendix, the term Markov chain always refers to a finite discrete-time Markov chain that is time-homogeneous.

Definition 4.8. Consider a Markov chain $\{X_t\}$. A row vector $\bar{\mathbf{p}} = [\bar{p}(x)]_{x \in \mathcal{X}}$ that satisfies

$$\bar{\mathbf{p}}\mathbf{P} = \bar{\mathbf{p}}$$
$$\bar{\mathbf{p}}\mathbf{1} = 1$$

is called a *stationary distribution* of the Markov chain. For $x \in \mathcal{X}$, the probability $\bar{p}(x)$ is called the *stationary probability* of state x.

Definition 4.9. A Markov chain $\{X_t\}$ is said to be *irreducible* if for each $x, x' \in \mathcal{X}$ there exists a positive integer N such that $\Pr(X_{t+N} = x' \mid X_t = x) > 0$.

Lemma 2. *If a Markov chain $\{X_t\}$ is irreducible, it has a unique stationary distribution $\bar{\mathbf{p}}$.*

Proof. See, for example, Tijms (1994, Theorem 2.3.3).

Definition 4.10. An irreducible Markov chain $\{X_t\}$ is said to be *aperiodic* if for each $x \in \mathcal{X}$ there exists a positive integer N such that $\Pr(X_{t+M} = x \mid X_t = x) > 0$ for all integers $M \geq N$.

Lemma 3. *If a Markov chain $\{X_t\}$ is irreducible and aperiodic, then*

$$\lim_{t \to \infty} \Pr(X_t = x \mid X_0 = x_0) = \bar{p}(x)$$

for all $x, x_0 \in \mathcal{X}$.

Proof. See, for example, Tijms (1994, Theorem 2.3.1 and Lemma 2.3.2).

Lemma 4. *Let a Markov chain $\{X_t\}$ be irreducible. Let $\mathcal{Y} \subset \mathcal{X}$ and $\mathcal{Y} \neq \varnothing$. Let*

$$\mathbf{T} = \Big[p(x, x')\Big]_{x, x' \in \mathcal{Y}} \qquad\qquad \mathbf{U} = \Big[p(x, x')\Big]_{x \in \mathcal{Y}, x' \in \mathcal{X} \setminus \mathcal{Y}}$$

$$\mathbf{V} = \Big[p(x, x')\Big]_{x \in \mathcal{X} \setminus \mathcal{Y}, x' \in \mathcal{Y}} \qquad\qquad \mathbf{W} = \Big[p(x, x')\Big]_{x, x' \in \mathcal{X} \setminus \mathcal{Y}}$$

and let

$$\mathbf{P}_Y = \mathbf{T} + \mathbf{U}(\mathbf{I} - \mathbf{W})^{-1}\mathbf{V}.$$

Let $\{Y_t\}$ denote a Markov chain with state space $\mathcal{Y}$ and transition matrix $\mathbf{P}_Y$. Markov chain $\{Y_t\}$ is then irreducible and has stationary probabilities $\bar{p}_Y(y)$ that are given by

$$\bar{p}_Y(y) = \frac{\bar{p}(y)}{\sum_{y' \in \mathcal{Y}} \bar{p}(y')}$$

where $y \in \mathcal{Y}$.

Proof. See Kemeny and Snell (1960, Theorem 6.1.1).[4]

---

[4] The terminology used by Kemeny and Snell (1960) differs from the terminology used in many other texts on Markov chains. In particular, an ergodic Markov chain in Kemeny and Snell (1960) corresponds to an irreducible Markov chain in this chapter.

Definition 4.11. Consider a set $\mathcal{X}$. For $x, x' \in \mathcal{X}$, the ordered pair (x, x') is called an *arrow* from x to x'. For $x_1, \ldots, x_N \in \mathcal{X}$, the sequence of arrows

$$((x_1, x_2), (x_2, x_3), \ldots, (x_{N-2}, x_{N-1}), (x_{N-1}, x_N))$$

is called a *path* from x_1 to x_N. For $x \in \mathcal{X}$, a set of arrows E is called an *x-tree* on $\mathcal{X}$ if it satisfies the following conditions:

(1) E contains no arrow that starts at x.

(2) For each $x' \in \mathcal{X} \setminus \{x\}$, E contains exactly one arrow that starts at x'.

(3) For each $x' \in \mathcal{X} \setminus \{x\}$, E contains a path from x' to x (or, formulated more accurately, for each $x' \in \mathcal{X} \setminus \{x\}$, there exists a path from x' to x such that E contains all arrows of the path).

Lemma 5. *Let a Markov chain $\{X_t\}$ be irreducible. For $x \in \mathcal{X}$, let $\mathcal{E}(x)$ denote the set of all x-trees on $\mathcal{X}$. The stationary probabilities $\bar{p}(x)$ of the Markov chain are then given by*

$$\bar{p}(x) = \frac{\tilde{p}(x)}{\sum_{x' \in \mathcal{X}} \tilde{p}(x')}$$

where $x \in \mathcal{X}$ and

$$\tilde{p}(x) = \sum_{E \in \mathcal{E}(x)} \prod_{(x,x') \in E} p(x, x').$$

Proof. A proof is provided by Freidlin and Wentzell (1998, Chapter 6, Lemma 3.1) (see also Dawid, 1996, Theorem 4.2.1).

Using the above definitions and lemmas, we now prove the mathematical results presented in Section 4.2.

Proof of Lemma 1. Notice that

$$\Pr(W_{t+1} = w_{t+1} \,|\, W_t = w_t) = \Pr(W_{t+1} = w_{t+1} \,|\, W_t = w_t, \ldots, W_0 = w_0)$$

for all $t \in \{0, 1, \ldots\}$ and all $w_0, \ldots, w_{t+1} \in \mathcal{W}$. That is, the population in iteration $t + 1$ of a GA depends only on the population in iteration t. Given the population in

iteration t, the population in iteration $t+1$ is independent of the populations in iterations $0, \ldots, t-1$. Notice further that

$$\Pr(W_{t+1} = w_{t+1} \mid W_t = w_t) = q(w_t, w_{t+1})$$

for all $t \in \{0, 1, \ldots\}$, all $w_t, w_{t+1} \in \mathcal{W}$, and some function $q : \mathcal{W}^2 \to [0,1]$ that does not depend on t. That is, the probability of going from one population to some other population remains constant over time. (Recall that the crossover rate γ and the mutation rate ε are assumed to remain constant over time.) It now follows from Definitions 4.6 and 4.7 that $\{W_t\}$, where the index t takes values in $\{0, 1, \ldots\}$, is a Markov chain with state space $\mathcal{W}$ and transition probabilities $q(w, w')$. Since the mutation rate ε is assumed to be positive, any population can be turned into any other population in a single iteration of a GA. Hence, $q(w, w') > 0$ for all $w, w' \in \mathcal{W}$. Consequently, it follows from Definitions 4.9 and 4.10 that Markov chain $\{W_t\}$ is irreducible and aperiodic. Lemma 3 then implies that for each population $w \in \mathcal{W}$ there exists a stationary probability $\bar{q}(w)$ such that

$$\lim_{t \to \infty} \Pr(W_t = w \mid W_0 = w_0) = \bar{q}(w) \tag{4.18}$$

for all $w_0 \in \mathcal{W}$. We refer to a stationary probability $\bar{q}(w)$ as the long-run probability of population w. Finally, (4.2) is obtained from (4.18) by taking into account the time-homogeneity of Markov chain $\{W_t\}$. This completes the proof of Lemma 1. $\qquad\square$

Proof of Theorem 4.1. As shown in the proof of Lemma 1, $\{W_t\}$, where the index t takes values in $\{0, 1, \ldots\}$, is an irreducible and aperiodic Markov chain with state space $\mathcal{W}$. Markov chain $\{W_t\}$ has stationary probabilities $\bar{q}(w)$, to which we refer as long-run probabilities. We now introduce some additional mathematical notation. Like in the proof of Lemma 1, the function $q : \mathcal{W}^2 \to [0,1]$ denotes the transition probabilities of Markov chain $\{W_t\}$. For $w, w' \in \mathcal{W}$, $q(w, w')$ is a polynomial in the mutation rate ε and can therefore be written as

$$q(w, w') = \sum_{l=0}^{\infty} \alpha(w, w', l)\varepsilon^l \tag{4.19}$$

where $\alpha(w, w', 0), \alpha(w, w', 1), \ldots$ denote the coefficients of the polynomial. $c(w, w')$ is

defined as

$$c(w, w') = \min\{l \mid \alpha(w, w', l) \neq 0\}. \tag{4.20}$$

That is, $c(w, w')$ is defined as the rate at which $q(w, w')$ approaches zero as ε approaches zero. It follows from this definition that $c(w, w')$ equals the minimum number of mutations required to go from population w to population w' in a single iteration of a GA. $\alpha(w, w')$ is defined as

$$\alpha(w, w') = \alpha(w, w', c(w, w')). \tag{4.21}$$

For $w \in \mathcal{W}$, $\tilde{q}(w)$ is defined as

$$\tilde{q}(w) = \sum_{E \in \mathcal{E}(w)} \prod_{(w,w') \in E} q(w, w') \tag{4.22}$$

where $\mathcal{E}(w)$ denotes the set of all w-trees on $\mathcal{W}$. Since the transition probabilities $q(w, w')$ are polynomials in ε, $\tilde{q}(w)$ is a polynomial in ε too. $\tilde{q}(w)$ can therefore be written as

$$\tilde{q}(w) = \sum_{l=0}^{\infty} \tilde{\alpha}(w, l) \varepsilon^l \tag{4.23}$$

where $\tilde{\alpha}(w, 0), \tilde{\alpha}(w, 1), \ldots$ denote the coefficients of the polynomial. $\tilde{c}(w)$ is defined as

$$\tilde{c}(w) = \min\{l \mid \tilde{\alpha}(w, l) \neq 0\}. \tag{4.24}$$

That is, $\tilde{c}(w)$ is defined as the rate at which $\tilde{q}(w)$ approaches zero as ε approaches zero. $\tilde{\alpha}(w)$ is defined as

$$\tilde{\alpha}(w) = \tilde{\alpha}(w, \tilde{c}(w)). \tag{4.25}$$

Using the mathematical notation introduced above, we first prove part (i) of Theorem 4.1. It follows from (4.19), (4.20), and (4.22)–(4.24) that $\tilde{c}(w)$ can be written as

$$\tilde{c}(w) = \min_{E \in \mathcal{E}(w)} \sum_{(w,w') \in E} c(w, w'). \tag{4.26}$$

At least one mutation is required to go from a uniform population $u \in \mathcal{U}$ to any other population $w \in \mathcal{W} \setminus \{u\}$. Hence, $c(u, w) \geq 1$ for all $u \in \mathcal{U}$ and all $w \in \mathcal{W}$ such that $u \neq w$. Consequently, it follows from (4.26) that $\tilde{c}(u) \geq \mu - 1$ for all $u \in \mathcal{U}$ and that

$\tilde{c}(w) \geq \mu$ for all $w \in \mathcal{W} \setminus \mathcal{U}$. We now show that for each chromosome i it is possible to construct a $u(i)$-tree E on $\mathcal{W}$ that satisfies

$$\sum_{(w,w') \in E} c(w, w') = \mu - 1. \tag{4.27}$$

Consider an arbitrary chromosome i. Let the function $\rho : \mathcal{C} \to \mathcal{C}$ satisfy the following conditions:

(1) For each $j \neq i$, chromosome j is directly connected to chromosome $\rho(j)$.

(2) For each $j \neq i$, $\rho^N(j) = i$ for some positive integer N.

In condition (2), $\rho^N(j)$ is defined as

$$\rho^N(j) = \begin{cases} \rho(j), & \text{if } N = 1 \\ \rho\big(\rho^{N-1}(j)\big), & \text{otherwise.} \end{cases}$$

Because Theorem 4.1 assumes that each chromosome is connected to all other chromosomes, a function ρ that satisfies the above two conditions is guaranteed to exist. In order to construct a $u(i)$-tree E on $\mathcal{W}$ that satisfies (4.27), we start with an empty set of arrows E. For each $j \neq i$, we then add an arrow to E that starts at $u(j)$ and ends at $v(j, \rho(j), n - 1)$. It follows from condition (1) that one mutation is required to go from $u(j)$ to $v(j, \rho(j), n - 1)$ in a single iteration of a GA. Hence, $c(u(j), v(j, \rho(j), n - 1)) = 1$. Next, for each $j \neq i$, we add a path to E that starts at $v(j, \rho(j), n - 1)$ and ends at $u(\rho(j))$. Each path that we add to E must contain no cycles, that is, it must contain no two arrows (w_1, w_1') and (w_2, w_2') such that either $w_1 = w_2$ or $w_1' = w_2'$. In addition, each path must only contain arrows (w, w') for which $c(w, w') = 0$. Condition (1) guarantees the existence of paths that satisfy the latter requirement. Due to condition (2), for each $u \in \mathcal{U} \setminus \{u(i)\}$, E now contains a path from u to $u(i)$. Finally, for each $w \in \mathcal{W} \setminus \mathcal{U}$, if E does not yet contain an arrow that starts at w, we add such an arrow to E. We choose the arrows that we add to E in such a way that, after adding the arrows, E contains, for each $w \in \mathcal{W} \setminus \mathcal{U}$, a path from w to some $u \in \mathcal{U}$ (which implies that E contains a path from w to $u(i)$). In addition, we only choose arrows (w, w') for which $c(w, w') = 0$. We can choose the arrows

in this way because Theorem 4.1 assumes that all non-uniform populations are almost uniform. Using Definition 4.11, it can be seen that the set of arrows E constructed as discussed above is a $u(i)$-tree on $\mathcal{W}$. Moreover, E satisfies (4.27). We have therefore shown that for each chromosome i a $u(i)$-tree E on $\mathcal{W}$ that satisfies (4.27) can be constructed. Consequently, it follows from (4.26) that $\tilde{c}(u) \leq \mu - 1$ for all $u \in \mathcal{U}$. Since it has been shown above that $\tilde{c}(u) \geq \mu - 1$ for all $u \in \mathcal{U}$, this implies that $\tilde{c}(u) = \mu - 1$ for all $u \in \mathcal{U}$. It has also been shown above that $\tilde{c}(w) \geq \mu$ for all $w \in \mathcal{W} \setminus \mathcal{U}$. Hence, as the mutation rate ε approaches zero, $\tilde{q}(w)$ approaches zero faster for $w \in \mathcal{W} \setminus \mathcal{U}$ than for $w \in \mathcal{U}$. It then follows from Lemma 5 that for all non-uniform populations $w \in \mathcal{W} \setminus \mathcal{U}$ the long-run probability $\bar{q}(w)$ approaches zero as ε approaches zero. In other words, the long-run limit probability $\hat{q}(w)$ equals zero for all non-uniform populations $w \in \mathcal{W} \setminus \mathcal{U}$. This completes the proof of part (i) of Theorem 4.1.

We now prove part (ii) of Theorem 4.1. It has been shown above that $\tilde{c}(u) = \mu - 1$ for all $u \in \mathcal{U}$. Consequently, as the mutation rate ε approaches zero, $\tilde{q}(u)$ approaches zero equally fast for all $u \in \mathcal{U}$. Using Lemma 5, it can therefore be seen that the long-run limit probability $\hat{q}(u)$ of a uniform population $u \in \mathcal{U}$ is given by

$$\hat{q}(u) = \lim_{\varepsilon \to 0} \bar{q}(u) = \frac{\tilde{\alpha}(u)}{\sum_{u' \in \mathcal{U}} \tilde{\alpha}(u')}. \tag{4.28}$$

For $u \in \mathcal{U}$, let $\widetilde{\mathcal{E}}(u)$ be defined as

$$\widetilde{\mathcal{E}}(u) = \left\{ E \in \mathcal{E}(u) \, \middle| \, \sum_{(w,w') \in E} c(w, w') = \mu - 1 \right\}. \tag{4.29}$$

It then follows from (4.19)–(4.25) that $\tilde{\alpha}(u)$ can be written as

$$\tilde{\alpha}(u) = \sum_{E \in \widetilde{\mathcal{E}}(u)} \prod_{(w,w') \in E} \alpha(w, w'). \tag{4.30}$$

Consider an arbitrary uniform population $u \in \mathcal{U}$ and an arbitrary u-tree E on $\mathcal{W}$, where $E \in \widetilde{\mathcal{E}}(u)$. Let E_1 and E_2 denote sets of arrows that are given by

$$E_1 = \{(w, w') \in E \,|\, w \in \mathcal{V}\}$$
$$E_2 = E \setminus E_1.$$

It is immediately clear that E_1 satisfies the following conditions:

(A1) E_1 contains no arrow that starts at u or at some $w \in \mathcal{W} \setminus \mathcal{V}$.

(A2) For each $v \in \mathcal{V} \setminus \{u\}$, E_1 contains exactly one arrow that starts at v.

Notice that $c(u', w) \geq 1$ for all $u' \in \mathcal{U}$ and all $w \in \mathcal{W}$ such that $u' \neq w$. Notice further that, due to (4.29), $\sum_{(w,w') \in E_1} c(w, w') \leq \mu - 1$. These observations imply that, for each $(w, w') \in E_1$, $c(w, w') = 1$ if $w \in \mathcal{U}$ and $c(w, w') = 0$ otherwise. They also imply that E_1 satisfies the following condition:

(A3) $\displaystyle\sum_{(w,w') \in E_1} c(w, w') = \mu - 1.$

It is easy to see that $c(v, w) \geq 1$ for all $v \in \mathcal{V}$ and all $w \in \mathcal{W} \setminus \mathcal{V}$ and that $c(u', w) \geq 2$ for all $u' \in \mathcal{U}$ and all $w \in \mathcal{W} \setminus \mathcal{V}$. Consequently, E_1 contains no arrows that end at some $w \in \mathcal{W} \setminus \mathcal{V}$. This implies the following condition on E_1:

(A4) For each $v \in \mathcal{V} \setminus \{u\}$, E_1 contains a path from v to u.

It is immediately clear that E_2 satisfies the following conditions:

(B1) E_2 contains no arrow that starts at some $v \in \mathcal{V}$.

(B2) For each $w \in \mathcal{W} \setminus \mathcal{V}$, E_2 contains exactly one arrow that starts at w.

(B3) For each $w \in \mathcal{W} \setminus \mathcal{V}$, E_2 contains a path from w to some $v \in \mathcal{V}$.

Furthermore, taking into account that E_1 satisfies condition (A3), (4.29) implies that E_2 satisfies the following condition:

(B4) $\displaystyle\sum_{(w,w') \in E_2} c(w, w') = 0.$

For $u \in \mathcal{U}$, let $\widetilde{\mathcal{E}}_1(u)$ denote a set that contains all sets of arrows E_1 satisfying conditions (A1)–(A4). Let $\widetilde{\mathcal{E}}_2$ denote a set that contains all sets of arrows E_2 satisfying conditions (B1)–(B4). Notice that $\widetilde{\mathcal{E}}_2$ does not depend on u. Clearly, for each $E \in \widetilde{\mathcal{E}}(u)$, there exist an $E_1 \in \widetilde{\mathcal{E}}_1(u)$ and an $E_2 \in \widetilde{\mathcal{E}}_2$ such that $E = E_1 \cup E_2$. Conversely, it can

be seen that for each $E_1 \in \widetilde{\mathcal{E}}_1(u)$ and each $E_2 \in \widetilde{\mathcal{E}}_2$ there exists an $E \in \widetilde{\mathcal{E}}(u)$ such that $E = E_1 \cup E_2$. Hence,

$$\widetilde{\mathcal{E}}(u) = \Big\{ E_1 \cup E_2 \,\Big|\, E_1 \in \widetilde{\mathcal{E}}_1(u),\, E_2 \in \widetilde{\mathcal{E}}_2 \Big\}.$$

Equation (4.30) can now be written as

$$\tilde{\alpha}(u) = \Big(\sum_{E_1 \in \widetilde{\mathcal{E}}_1(u)} \prod_{(w,w') \in E_1} \alpha(w,w') \Big) \Big(\sum_{E_2 \in \widetilde{\mathcal{E}}_2} \prod_{(w,w') \in E_2} \alpha(w,w') \Big).$$

Consequently, it follows from (4.28) that

$$\hat{q}(u) = \lim_{\varepsilon \to 0} \bar{q}(u) = \frac{\sum_{E_1 \in \widetilde{\mathcal{E}}_1(u)} \prod_{(w,w') \in E_1} \alpha(w,w')}{\sum_{u' \in \mathcal{U}} \sum_{E_1 \in \widetilde{\mathcal{E}}_1(u')} \prod_{(w,w') \in E_1} \alpha(w,w')}. \qquad (4.31)$$

Based on (4.31), the following observations can be made:

(1) For $w, w' \in \mathcal{W}$ such that $w \neq w'$ and such that there exists an $E_1 \in \bigcup_{u' \in \mathcal{U}} \widetilde{\mathcal{E}}_1(u')$ that contains an arrow (w, w'), $\lim_{\varepsilon \to 0} \bar{q}(u)$ depends on the term of lowest degree in the transition probability $q(w, w')$ and does not depend on other terms in $q(w, w')$.

(2) For $w, w' \in \mathcal{W}$ such that $w \neq w'$ and such that there does not exist an $E_1 \in \bigcup_{u' \in \mathcal{U}} \widetilde{\mathcal{E}}_1(u')$ that contains an arrow (w, w'), $\lim_{\varepsilon \to 0} \bar{q}(u)$ does not depend on any of the terms in the transition probability $q(w, w')$.

Let $\{V_t\}$, where the index t takes values in $\{0, 1, \ldots\}$, denote a Markov chain with state space $\mathcal{V}$, transition probabilities $r(v, v')$, and stationary probabilities $\bar{r}(v)$, where $v, v' \in \mathcal{V}$. For $v \neq v'$, let

$$r(v, v') = \begin{cases} \alpha(v, v')\varepsilon, & \text{if } v \in \mathcal{U} \text{ and } c(v, v') = 1 \\ \alpha(v, v'), & \text{if } v \notin \mathcal{U} \text{ and } c(v, v') = 0 \\ 0, & \text{otherwise.} \end{cases} \qquad (4.32)$$

Furthermore, let $r(v, v) = 1 - \sum_{v' \in \mathcal{V} \setminus \{v\}} r(v, v')$. Clearly, Markov chain $\{V_t\}$ is irreducible. Taking into account the two observations made above, it can be seen that

$\lim_{\varepsilon \to 0} \bar{r}(v) = \lim_{\varepsilon \to 0} \bar{q}(v)$ for all $v \in \mathcal{V}$. That is, in the limit as ε approaches zero, corresponding states of Markov chains $\{V_t\}$ and $\{W_t\}$ have the same stationary probability. It follows from this that $\lim_{\varepsilon \to 0} \bar{r}(v) = \hat{q}(v)$ for all $v \in \mathcal{V}$.

The following observations can be made:

(1) For $v \in \mathcal{U}$ and $v' \in \mathcal{V}$, $c(v, v') = 1$ if and only if $v = u(i)$ and $v' = v(i, j, n - 1)$ for some i and some j such that $\delta(i, j) = 1$.

(2) For $v \in \mathcal{U}$ and $v' \in \mathcal{V}$ such that $c(v, v') = 1$, $q(v, v')$ equals the probability that the mutation operator inverts one specific bit in the binary encoding of an arbitrarily chosen chromosome and that it does not invert any other bits in the binary encoding of the chosen chromosome or of any other chromosome in the population. This probability does not depend on v or v'. Consequently, for all $v_1, v_2 \in \mathcal{U}$ and all $v_1', v_2' \in \mathcal{V}$ such that $c(v_1, v_1') = c(v_2, v_2') = 1$, $q(v_1, v_1') = q(v_2, v_2')$ and hence $\alpha(v_1, v_1') = \alpha(v_2, v_2')$.

(3) For $v \in \mathcal{V} \setminus \mathcal{U}$ and $v' \in \mathcal{V}$, $c(v, v') = 0$ only if $v = v(i, j, \lambda)$ and $v' = v(i, j, \lambda')$ for some i and some j such that $\delta(i, j) = 1$ and for some $\lambda \in \{1, \dots, n - 1\}$ and some $\lambda' \in \{0, \dots, n\}$.

(4) For $v \in \mathcal{V} \setminus \mathcal{U}$ and $v' \in \mathcal{V}$ such that $c(v, v') = 0$, $\alpha(v, v') = \pi(i, j, \lambda, \lambda')$, where i, j, λ, and λ' satisfy $v = v(i, j, \lambda)$ and $v' = v(i, j, \lambda')$ and where $\pi(i, j, \lambda, \lambda')$ is defined in (4.1).

Let $\alpha = \alpha(v, v')$ for all $v \in \mathcal{U}$ and all $v' \in \mathcal{V}$ such that $c(v, v') = 1$. Using (4.32), it follows from the first two observations made above that $r(v, v') = \alpha\varepsilon$ if $v = u(i)$ and $v' = v(i, j, n - 1)$ for some i and some j such that $\delta(i, j) = 1$. It also follows that $r(v, v') = 1 - m\alpha\varepsilon$ if $v = v' \in \mathcal{U}$. Furthermore, taking into account the last two observations made above, it can be seen from (4.32) that $r(v, v') = \pi(i, j, \lambda, \lambda')$ if $v = v(i, j, \lambda)$ and $v' = v(i, j, \lambda')$ for some i and some j such that $\delta(i, j) = 1$ and for some $\lambda \in \{1, \dots, n - 1\}$ and some $\lambda' \in \{0, \dots, n\}$. Finally, (4.32) implies that $r(v, v') = 0$ if none of the above conditions is satisfied. Let the vector $\tilde{\mathbf{v}} = \begin{bmatrix} \tilde{v}_1 & \cdots & \tilde{v}_\xi \end{bmatrix}$ be given

by

$$
\tilde{\mathbf{v}}^{\mathrm{T}} =
\begin{bmatrix}
v(g_1, h_1, 1) \\
\vdots \\
v(g_1, h_1, n-1) \\
v(g_2, h_2, 1) \\
\vdots \\
\vdots \\
v(g_{\nu-1}, h_{\nu-1}, n-1) \\
v(g_\nu, h_\nu, 1) \\
\vdots \\
v(g_\nu, h_\nu, n-1)
\end{bmatrix}
$$

where $\mathbf{g} = [g_k]$ and $\mathbf{h} = [h_k]$ are defined in Section 4.2. Notice that $\tilde{\mathbf{v}}$ contains each population in $\mathcal{V} \setminus \mathcal{U}$ exactly once. It can be seen that

$$
(1 - m\alpha\varepsilon)\mathbf{I} =
\begin{bmatrix}
r(u(0), u(0)) & \cdots & r(u(0), u(\mu-1)) \\
\vdots & \ddots & \vdots \\
r(u(\mu-1), u(0)) & \cdots & r(u(\mu-1), u(\mu-1))
\end{bmatrix}
\tag{4.33}
$$

$$
\alpha\varepsilon\mathbf{A} =
\begin{bmatrix}
r(u(0), \tilde{v}_1) & \cdots & r(u(0), \tilde{v}_\xi) \\
\vdots & \ddots & \vdots \\
r(u(\mu-1), \tilde{v}_1) & \cdots & r(u(\mu-1), \tilde{v}_\xi)
\end{bmatrix}
\tag{4.34}
$$

$$
\mathbf{B} =
\begin{bmatrix}
r(\tilde{v}_1, u(0)) & \cdots & r(\tilde{v}_1, u(\mu-1)) \\
\vdots & \ddots & \vdots \\
r(\tilde{v}_\xi, u(0)) & \cdots & r(\tilde{v}_\xi, u(\mu-1))
\end{bmatrix}
\tag{4.35}
$$

$$\mathbf{C} = \begin{bmatrix} r(\tilde{v}_1, \tilde{v}_1) & \cdots & r(\tilde{v}_1, \tilde{v}_\xi) \\ \vdots & \ddots & \vdots \\ r(\tilde{v}_\xi, \tilde{v}_1) & \cdots & r(\tilde{v}_\xi, \tilde{v}_\xi) \end{bmatrix} \qquad (4.36)$$

where $\mathbf{A}$, $\mathbf{B}$, and $\mathbf{C}$ are defined in (4.3), (4.6), and (4.9). Let $\mathbf{S}$ denote a $\mu \times \mu$ matrix that is obtained from the matrices in (4.33)–(4.36) and that is given by

$$\mathbf{S} = (1 - m\alpha\varepsilon)\mathbf{I} + \alpha\varepsilon\mathbf{A}(\mathbf{I} - \mathbf{C})^{-1}\mathbf{B}. \qquad (4.37)$$

This can be written more simply as

$$\mathbf{S} = \mathbf{I} + \alpha\varepsilon\mathbf{D}$$

where $\mathbf{D}$ is defined in (4.12). Let $\{U_t\}$, where the index t takes values in $\{0, 1, \ldots\}$, denote a Markov chain with state space $\mathcal{U}$ and transition matrix $\mathbf{S}$. Using (4.33)–(4.37), it follows from Lemma 4 that Markov chain $\{U_t\}$ is irreducible and has stationary probabilities $\bar{s}(u)$ that are given by

$$\bar{s}(u) = \frac{\bar{r}(u)}{\sum_{u' \in \mathcal{U}} \bar{r}(u')} \qquad (4.38)$$

where $u \in \mathcal{U}$. Definition 4.8 states that the stationary distribution

$$\bar{\mathbf{s}} = \begin{bmatrix} \bar{s}(u(0)) & \cdots & \bar{s}(u(\mu - 1)) \end{bmatrix}$$

of Markov chain $\{U_t\}$ satisfies

$$\bar{\mathbf{s}}\mathbf{S} = \bar{\mathbf{s}} \qquad (4.39)$$

$$\bar{\mathbf{s}}\mathbf{1} = 1. \qquad (4.40)$$

Lemma 2 implies that this linear system has a unique solution. The equality in (4.39) can be written as

$$\bar{\mathbf{s}}(\mathbf{S} - \mathbf{I}) = \alpha\varepsilon\bar{\mathbf{s}}\mathbf{D} = \mathbf{0}.$$

Since $\alpha > 0$ and $\varepsilon > 0$, this can be simplified to

$$\bar{\mathbf{s}}\mathbf{D} = \mathbf{0}. \tag{4.41}$$

Notice that $\mathbf{D}$ does not depend on ε. $\bar{\mathbf{s}}$ therefore does not depend on ε either. Recall further that $\lim_{\varepsilon \to 0} \bar{r}(v) = \hat{q}(v)$ for all $v \in \mathcal{V}$ and that $\hat{q}(w) = 0$ for all $w \in \mathcal{W} \setminus \mathcal{U}$. Using (4.38), it now follows that

$$\bar{s}(u) = \lim_{\varepsilon \to 0} \bar{s}(u) = \lim_{\varepsilon \to 0} \frac{\bar{r}(u)}{\sum_{u' \in \mathcal{U}} \bar{r}(u')} = \frac{\hat{q}(u)}{\sum_{u' \in \mathcal{U}} \hat{q}(u')} = \hat{q}(u)$$

for all $u \in \mathcal{U}$. Hence, the stationary distribution $\bar{\mathbf{s}}$ of Markov chain $\{U_t\}$ equals the long-run limit distribution $\hat{\mathbf{q}}$. Consequently, (4.40) and (4.41) imply that $\hat{\mathbf{q}}$ satisfies (4.13) and (4.14). It also follows that the linear system given by (4.13) and (4.14) has a unique solution. This completes the proof of part (ii) of Theorem 4.1.   □

Chapter 5

Economic modeling using evolutionary algorithms: The effect of a binary encoding of strategies[*]

Abstract

We are concerned with evolutionary algorithms that are employed for economic modeling purposes. We focus in particular on evolutionary algorithms that use a binary encoding of strategies. These algorithms, commonly referred to as genetic algorithms, are popular in agent-based computational economics research. In many studies, however, there is no clear reason for the use of a binary encoding of strategies. We therefore examine to what extent the use of such an encoding may influence the results produced by an evolutionary algorithm. It turns out that the use of a binary encoding can have quite significant effects. Since these effects do not have a meaningful economic interpretation, they should be regarded as artifacts. Our findings indicate that in general the use of a binary encoding is undesirable. They also highlight the importance of employing evolutionary algorithms with a sensible economic interpretation.

5.1   Introduction

Evolutionary algorithms (EAs) are algorithms that are inspired by the process of natural evolution. EAs have their origins in the field of computer science, where they are mainly applied for optimization purposes. Nowadays, EAs are also regularly employed in the field of economics. In economic research, EAs frequently serve as a tool for modeling boundedly rational behavior. When EAs are applied as a modeling tool in economic research, a binary encoding of strategies is typically used. This means that strategies are represented by bit strings (i.e., strings of zeros and ones, often referred to as chromosomes) and that evolutionarily inspired operations such as crossover and mutation take place at the level of individual bits. EAs that use a binary encoding of strategies are commonly referred to as genetic algorithms. For early research in which genetic algorithms are employed, we refer to Miller (1986, 1996), Axelrod (1987), Marks (1992), Arifovic (1994, 1996), Andreoni and Miller (1995), and Dawid (1996). Examples of more recent research can be found in the work of, among others, Lux and Schornstein (2005), Alkemade et al. (2006, 2007, 2009), Arifovic and Maschek (2006), Wheeler, Bean, Gaffney, and Taylor (2006), Xu (2006), Casari (2008), and Maschek (2010).

Researchers who apply genetic algorithms as a tool for modeling boundedly rational behavior typically do not justify why they use a binary encoding of strategies. If the agents whose behavior is being modeled have to make decisions that are intrinsically binary, such as decisions between cooperation and defection in a prisoner's dilemma (e.g., Axelrod, 1987), the use of a binary encoding of strategies is a very natural choice. However, in the case of non-binary decisions, such as decisions by firms on their production level (e.g., Arifovic, 1994; Price, 1997; Dawid & Kopel, 1998; Franke, 1998; Vriend, 2000; Alkemade et al., 2006, 2007, 2009; Arifovic & Maschek, 2006; Wheeler et al., 2006; Casari, 2008; Maschek, 2010), there is no clear reason for the use of a binary encoding of strategies.

In this chapter, we examine to what extent the use of a binary encoding of strategies may influence the results of studies in which EAs are employed. It turns out that the use of a binary encoding can have quite significant effects. In general, these effects do not have a meaningful economic interpretation and should be regarded as artifacts. In order to avoid these artifacts, we argue that in most cases researchers should not use a binary encoding of strategies.

Our research is inspired by results reported by Alkemade et al. (2006, 2007, 2009; see also Waltman & Van Eck, 2009). Alkemade et al. show that under certain conditions an EA that is employed for modeling purposes may exhibit premature convergence. By premature convergence Alkemade et al. mean that different runs of the EA can lead to very different results. Alkemade et al. argue that premature convergence is caused by a too small population size. In this chapter, we report results that point in a different direction. We show that the observation of premature convergence by Alkemade et al. depends crucially on their use of a binary encoding of strategies. Using the same economic environment as Alkemade et al. (i.e., a Cournot oligopoly market), we demonstrate that premature convergence does occur in the case of EAs with a binary encoding while it does not occur in the case of EAs without a binary encoding.

We note that the consequences of the use of a binary encoding of strategies are also studied extensively by Dawid (1996; see also Dawid & Kopel, 1998). However, the approach taken by Dawid is quite different from the approach that we take in the present study. Dawid focuses on EAs with a large population size, and he is concerned with aggregate results, that is, results averaged over many EA runs. We do not assume the population size to be large, and we are specifically interested in comparing results of individual EA runs. Another difference is that the crossover operator plays a crucial role in Dawid's approach while in our approach the crossover operator is not important at all.

The remainder of this chapter is organized as follows. In Section 5.2, we introduce the Cournot market that we consider in this chapter. In Sections 5.3 and 5.4, we present the various EAs that we study and we discuss the economic interpretation of EAs. We report the results of the computer simulations that we have performed in Section 5.5. Based on these results, we provide an elaborate analysis of the effect of a binary encoding of strategies in Section 5.6. Finally, in Section 5.7, we discuss the conclusions of our research.

5.2   Cournot oligopoly market

To analyze the effect of a binary encoding of strategies, we study the behavior of firms in a Cournot oligopoly market. To facilitate comparison, we consider exactly the same

Cournot market as Alkemade et al. (2006, 2007, 2009). For other studies in which quantity competition among firms is modeled using EAs, we refer to Arifovic (1994), Price (1997), (Dawid & Kopel, 1998), Franke (1998), Vriend (2000), Arifovic and Maschek (2006), Wheeler et al. (2006), Casari (2008), and Maschek (2010).

The Cournot market that we consider has the following characteristics: The number of firms equals four, firms produce perfect substitutes, the demand function is linear, firms have identical cost functions, and marginal cost is constant. The inverse demand function is given by

$$p = \max\left(256 - \sum_{i=1}^{4} q_i, 0\right), \tag{5.1}$$

where p denotes the market price and q_i denotes firm i's production level. Firm i's total cost equals $c_i = 56q_i$. Hence, it follows that firm i's profit is given by

$$\pi_i = pq_i - c_i = q_i \max\left(200 - \sum_{i'=1}^{4} q_{i'}', -56\right). \tag{5.2}$$

A Nash (or Cournot) equilibrium is obtained if each firm chooses a production level that maximizes its profit given the production levels of its competitors. This means that in a Nash equilibrium $\partial \pi_i / \partial q_i = 0$ for $i = 1, \ldots, 4$. It is easy to see that the Cournot market that we consider has a Nash equilibrium in which each firm produces a quantity of 40. Each firm makes a profit of 1600 in the Nash equilibrium. In addition to a Nash equilibrium, the Cournot market that we consider also has a competitive (or Walrasian) equilibrium. This equilibrium is obtained if firms are not aware of their influence on the market price and therefore behave as price takers. In the competitive equilibrium, the four firms jointly produce a quantity of 200 and each firm makes a profit of 0.

5.3   Evolutionary algorithms

As discussed by Vriend (2000), there are two quite different ways in which EAs can be employed to model the behavior of economic agents. In the individual learning approach, each agent learns exclusively from its own experience (e.g., Arifovic, 1994; Price, 1997; Arifovic & Maschek, 2006; Casari, 2008). This is modeled through the use of a separate EA for each agent. In the social learning approach, each agent learns not

only from its own experience but also from the experience of other agents (e.g., Arifovic, 1994; Dawid & Kopel, 1998; Franke, 1998; Alkemade et al., 2006, 2007, 2009). This is modeled through the use of a single EA for all agents together. The social learning approach seems to be more popular than the individual learning approach (Arifovic & Maschek, 2006). In this chapter, we focus on the social learning approach.

An important observation about the social learning approach is made by Alkemade et al. (2006, 2007, 2009). They note that the social learning approach can be implemented in two quite different ways. On the one hand, one can employ an EA with a population size that equals the number of interacting agents (e.g., Arifovic, 1994; Dawid & Kopel, 1998; Franke, 1998; Vriend, 2000). This results in a one-to-one relationship between strategies and agents. In the case of the Cournot oligopoly market discussed in the previous section, the EA would have a population size of four (since there are four firms in the market). On the other hand, one can employ an EA with a population size that exceeds the number of interacting agents (e.g., Axelrod, 1987; Andreoni & Miller, 1995; Dawid, 1996, Sections 4.5 and 5.3; Miller, 1996). Strategies are then evaluated using some matching mechanism. In the case of the Cournot oligopoly market discussed in the previous section, the EA would have a population size greater than four. Alkemade et al. show that the two ways in which the social learning approach can be implemented can lead to very different results.[1] In this chapter, we only employ EAs with a population size that equals the number of interacting agents. We take this approach because it is very suitable for demonstrating how the use of a binary encoding of strategies can lead to artifacts.

We consider six different EAs in this chapter. We refer to these EAs as EA1 to EA6. Each of the EAs provides a slightly different model of the behavior of firms in the Cournot market discussed in the previous section. The six EAs all have the same general form. This general form is shown in Figure 5.1. Each EA works on a population of four strategies. A strategy corresponds with the production level of one of the four firms in the Cournot market. The EAs all impose the constraint that the production level of a firm must lie between 0 and 127, and they all randomly generate an initial population by drawing four strategies from a uniform distribution over all possible strategies. In each

---

[1] In the evolutionary game theory literature, similar observations have been made by various researchers, for example by Hansen and Samuelson (1988) and Rhode and Stegeman (1996) in an economic context and by Schaffer (1988) in a biological context.

---

1 Randomly generate an initial population of strategies

2 **repeat**

3   Calculate each strategy's profit

4   Calculate each strategy's fitness

5   Apply the selection operator

6   Apply the crossover and mutation operators

7 **until** a specific number of iterations have been performed

---

Figure 5.1: General form of the six EAs considered in this chapter.

iteration of an EA, the profit resulting from each of the four strategies in the current population is calculated using (5.2). Based on the profits of the four strategies, the fitness values of the strategies are calculated according to

$$f_i = \max\left(\frac{\pi_i - \mu}{\sigma} + 2, 0\right),$$ (5.3)

where μ and σ denote, respectively, the mean and the standard deviation of the profits of the strategies, that is,

$$\mu = \frac{\sum_{i=1}^{4} \pi_i}{4},$$ (5.4)

$$\sigma = \sqrt{\frac{\sum_{i=1}^{4}(\pi_i - \mu)^2}{4}}.$$ (5.5)

The above transformation from profits to fitness values is sometimes referred to as sigma scaling (e.g., Mitchell, 1996) or sigma truncation (e.g., Goldberg, 1989). The transformation is used by, for example, Axelrod (1987), Andreoni and Miller (1995),Miller (1996), and Franke (1998). The six EAs that we study all use roulette wheel selection (e.g., Goldberg, 1989; Mitchell, 1996), also known as fitness-proportionate selection. This means that in each iteration of an EA the selection operator generates a new population of strategies by randomly drawing four strategies from the old population. Strategies are drawn independently and with replacement. The probability that a strategy is drawn is proportional to the fitness of the strategy given by (5.3).

The six EAs that we consider differ from each other on the following four dimensions:

Table 5.1: Overview of the differences between the six EAs considered in this chapter.

	EA1	EA2	EA3	EA4	EA5	EA6
Strategy	integer	integer	integer	integer	real number	integer
Binary encoding	yes	yes	no	no	no	yes (Gray)
Crossover	yes	no	no	no	no	yes
Mutation	bit flip	bit flip	unif. dist.	± 1	norm. dist.	bit flip

(1) The type of strategy that is used, that is, integer production levels or real-valued production levels.

(2) Whether a binary encoding of strategies is used or not.

(3) Whether a crossover operator is used or not.

(4) The type of mutation operator that is used.

The differences between the EAs are summarized in Table 5.1. We now discuss the specific characteristics of each of the EAs.

5.3.1 EA1

EA1 is a standard genetic algorithm (e.g., Goldberg, 1989; Mitchell, 1996). It uses integer production levels and a binary encoding of strategies. A strategy is represented by a bit string of length seven. The production level corresponding to a bit string $(b_1, \ldots, b_7)$, where $b_j \in \{0, 1\}$ denotes the value of the jth bit in the string, is given by

$$q = \sum_{j=1}^{7} 2^{7-j} b_j. \tag{5.6}$$

EA1 uses a single-point crossover operator (e.g., Mitchell, 1996) with a crossover rate of 1. The mutation operator used by EA1 randomly determines for each bit in a bit string whether to invert the bit or not. The mutation rate equals 0.001, which means that each bit has a probability of 0.001 of being inverted. We note that EA1 is very similar to the EA employed by Alkemade et al. (2006, 2007, 2009).

5.3.2   EA2

EA2 is identical to EA1 except that it does not use a crossover operator.

5.3.3   EA3

Like EA2, EA3 uses integer production levels and does not use a crossover operator. Unlike EA2, EA3 does not use a binary encoding of strategies. EA3 also uses a different mutation operator than EA2. The probability that a strategy is being mutated equals 0.01. If a strategy is being mutated, it is replaced by a random new strategy that is drawn from a uniform distribution over all possible strategies. In the economic literature, EAs similar to EA3 are employed by Unver (2001), Dawid and Dermietzel (2006), and Haruvy et al. (2006).

5.3.4   EA4

EA4 is identical to EA3 except that it uses a different mutation operator. The probability that a strategy is being mutated equals 0.01. If a strategy is being mutated, the corresponding production level is either increased by one or decreased by one (both with a probability of 0.5). The increase or decrease does not take place if the resulting new production level would be below 0 or above 127.

5.3.5   EA5

Like EA3 and EA4, EA5 does not use a binary encoding of strategies and also does not use a crossover operator. Unlike EA3 and EA4, EA5 uses real-valued production levels. EA5 also uses a different mutation operator than EA3 and EA4. The probability that a strategy is being mutated equals 0.01. If a strategy is being mutated, the corresponding production level is updated according to

$$q_{\text{new}} = \min\left(\max\left(q_{\text{old}} + N\left(0, s^2\right), 0\right), 127\right), \tag{5.7}$$

where q_{old} denotes the production level before mutation, q_{new} denotes the production level after mutation, and $N(0, s^2)$ denotes a normally distributed random variable with mean 0 and standard deviation s. EA5 uses a value of 1 for the parameter s. We note that

Table 5.2: Illustration of the Gray coding used in EA6.

Production level	Bit string
0	0000000
1	0000001
2	0000011
3	0000010
4	0000110
…	…
126	1000001
127	1000000

EA5 is somewhat similar to what is referred to as an evolution strategy in the computer science literature (e.g., Beyer, 2001; Beyer & Schwefel, 2002). In the economic literature, EAs similar to EA5 are employed by Sellgren (2001), Gerding, Van Bragt, and La Poutré (2003), Lux and Schornstein (2005), and Clemens and Riechmann (2006).

5.3.6   EA6

EA6 is identical to EA1 except that it does not use an ordinary binary encoding of strategies. Instead, it uses a so-called Gray coding of strategies. Like in EA1, strategies are represented by bit strings of length seven. However, the transformation from bit strings to production levels is different from the transformation used in EA1, that is, it is different from (5.6). In EA6, the transformation from bit strings to production levels is performed in such a way that bit strings corresponding to consecutive production levels always differ by only one bit. This is referred to as a Gray coding of strategies. The Gray coding used in EA6 is illustrated in Table 5.2.

In the literature, Gray codings of strategies are used only rarely. Usually, an ordinary binary encoding of strategies is used, like in EA1 and EA2. Examples of the use of Gray codings of strategies are provided by Arifovic (1996) and Maschek (2010). Arifovic states that "the Gray coding ensures that, if a small number of bits within a binary string change, this will correspond to a small change in a decoded integer or real number" (p. 525). However, this is not correct. Even if a Gray coding of strategies is used, a change of a small number of bits in a bit string may still correspond to a large change in the decoded value. This can be seen in Table 5.2. The bit strings 0000000 and

1000000 differ by only one bit, but they correspond to two very different production levels, namely 0 and 127, respectively.

5.4   Economic interpretation of evolutionary algorithms

In many papers in which EAs are applied as an economic modeling tool, relatively little attention is paid to the economic interpretation of EAs.[2] In the present study, the economic interpretation of EAs is a central issue and hence requires serious attention. In this section, we therefore summarize the various ways in which EAs are interpreted in the literature.

An EA works on a population of strategies. What exactly does a population of strategies represent? As discussed in the previous section, we need to make a distinction between the individual learning approach and the social learning approach (Vriend, 2000).[3] In the individual learning approach, a separate EA is used for each agent. Hence, each agent has its own population of strategies. Arifovic (1994) interprets a population of strategies in the individual learning approach as "an agent's mutually competing ideas about what his behavior in a given environment should be" (p. 15). According to Dawid (1996), the interpretation of the individual learning approach has several weaknesses. For example, the individual learning approach assumes that an agent is able to determine the performance of a strategy without actually executing the strategy. This may be a strong assumption in many contexts. Price (1997), however, argues that in certain cases the assumption may be justified, in particular in the case of firms that perform scenario analysis.[4] In the social learning approach, a single EA is used for all agents together. In this approach, the population of strategies can be interpreted in two ways (Alkemade et al., 2006, 2007, 2009). In one interpretation, the population size equals the number of interacting agents and strategies and agents are related in a one-to-one manner. Each strategy then simply represents the strategy of one

---

[2]A notable exception is a paper by Chattoe (1998), in which the economic interpretation of EAs is critically discussed.

[3]Arifovic (1994) refers to these approaches as the multiple-population design and the single-population design. Chattoe (1998) refers to the approaches as the mental interpretation and the population interpretation of an EA.

[4]An alternative assumption is that agents try out the various strategies they have in mind and that they update their strategies only after they have obtained a sufficient amount of information on each strategy's performance. An assumption like this is made by Vriend (2000).

particular agent. This is the interpretation that we follow in this chapter. In the other interpretation, the population size exceeds the number of interacting agents. The population of strategies can then be seen as a pool of strategies that are commonly known to all agents. When interacting with each other, agents randomly choose a strategy from the strategy pool.

In most EAs that are employed for economic modeling, a binary encoding of strategies is used. In general, it is unclear how the use of such an encoding can be given a sensible economic interpretation. Most researchers ignore this issue. An exception is Brenner (2006), who points out that the use of a binary encoding of strategies may lead to difficulties with the interpretation of the crossover operator. A somewhat similar comment is made by Dawid (1996).

We now discuss the economic interpretation of the selection, crossover, and mutation operators of an EA. Our focus is on the social learning approach.

Following Chattoe (1998), we distinguish between two interpretations of the selection operator. The first interpretation can be used only if there is a one-to-one relationship between strategies and agents. According to this interpretation, the selection operator models the removal of unsuccessful agents from the economic environment. An example is the removal of unprofitable firms from the market due to bankruptcy. The second interpretation, which is used by most researchers, states that the selection operator models the imitation of successful strategies. According to this interpretation, agents have information on the past performance of strategies and tend to imitate those strategies that were most successful in the past. The details of this interpretation depend on the type of selection operator that is used. For example, roulette wheel or fitness-proportionate selection, which is the most commonly used selection operator, assumes that an agent has information on the past performance of all strategies. On the other hand, tournament selection, which is used in some papers (e.g., Bullard & Duffy, 1998; Van Bragt et al., 2001; Dawid & Dermietzel, 2006), assumes that an agent has information on the past performance of only a limited number of strategies. Another thing to realize is that most selection operators assume that agents update their strategies simultaneously rather than one by one. The assumption of simultaneous updating of strategies may not always be realistic. A comparison of simultaneous and non-simultaneous updating of strategies is performed by Dawid and Dermietzel (2006). They find that the

two strategy updating regimes may lead to significantly different results. We note that the issue of the appropriate strategy updating regime has also received considerable attention in the biological literature (e.g., Huberman & Glance, 1993).

The economic interpretation of the crossover and mutation operators is quite straight-forward. The crossover operator is typically interpreted as the exchange of ideas or in-formation. Hence, the crossover operator models communication between agents (or industrial espionage, as suggested by Dawid & Kopel, 1998). The mutation operator is usually interpreted as the effect of innovation. Innovation may be due to deliberate experimentation or unintended errors.

5.5   Simulation results

In this section, we report the results of the computer simulations that we have per-formed.[5] Each of the six EAs discussed in Section 5.3 was run 100 times, each time using different random numbers. Each run lasted $10,000$ iterations. The results re-ported below are fairly robust to changes in the values of the various EA parameters. By changing parameter values, somewhat different results may be obtained, but the anal-ysis will remain essentially unchanged. The results reported below are also robust to changes in the transformation from profits to fitness values. We further experimented with simulation runs that lasted one million instead of $10,000$ iterations, but this also did not affect the analysis in any fundamental way.

The results of six selected runs of EA1 are shown in Figure 5.2. Each graph in the figure corresponds with one run of EA1. The graphs display how the average production level of the four firms in the market evolves over time. In the graphs in panels (a), (b), and (c), firms' average production level stabilizes fairly quickly, respectively at a quantity of 47, 50, and 64. These results are quite typical. In almost all 100 runs of EA1, we observe that firms' average production level stabilizes at one of these three quantities. However, firms' average production level does not always remain stabilized at the same quantity during an entire run. In a few runs of EA1, we find that firms'

---

[5]The software used to obtain the results is available online at www.ludowaltman.nl/binaryencoding/. The software runs in MATLAB.

average production level switches from one stable quantity to another. Panels (d), (e), and (f) of Figure 5.2 provide examples of such runs.

In panel (a) of Figure 5.3, the results of all 100 runs of EA1 are averaged. As can be seen, on average firms' production level stabilizes at a quantity of about 52. Based on the 100 runs of EA1, the distribution of firms' average production level at the end of a run can be determined. This distribution is displayed in panel (b) of Figure 5.3. It turns out that in somewhat more than half of the runs firms' average production level after $10,000$ iterations equals the stable quantity of 50. In most other runs, firms' average production level after $10,000$ iterations equals either the stable quantity of 47 or the stable quantity of 64. There are a few runs in which firms' average production level after $10,000$ iterations does not equal one of the three stable quantities. In most of these runs, this is probably due to small disturbances caused by the mutation operator.

The results that we have obtained using EA1 are very similar to the results reported by Alkemade et al. (2006, 2007, 2009).[6] Like Alkemade et al., we find that in different EA runs firms' average production level stabilizes at different quantities. This phenomenon is referred to as premature convergence by Alkemade et al. We further find that on average firms' production level stabilizes at a quantity above 50. This means that on average firms produce a larger quantity than in the competitive equilibrium of the Cournot oligopoly market (see Section 5.2). This finding is also in agreement with the results reported by Alkemade et al.

We now turn to EA2. The results obtained using EA2 are shown in panels (c) and (d) of Figure 5.3. It is clear that the results of EA2 are quite similar to the results of EA1. Like EA1, EA2 leads to premature convergence. Firms' average production level again stabilizes at a quantity of 47, 50, or 64. The similarity between the results of EA1 and EA2 is not surprising. The only difference between the two EAs is that EA1 uses a crossover operator while EA2 does not use such an operator. In Chapter 4 of this thesis, we show mathematically that, if the mutation rate is small and some technical assumptions are satisfied, the effect of the use of a crossover operator on the results produced by an EA tends to be negligible in the long run. The results shown in Figure 5.3 are in line with this theoretical finding.

---

[6]Due to an error in their computer simulations, the results reported by Alkemade et al. (2006, 2007) are not entirely correct. For a correction of the results, see Alkemade et al. (2009). For some additional comments on the results, see Waltman and Van Eck (2009).
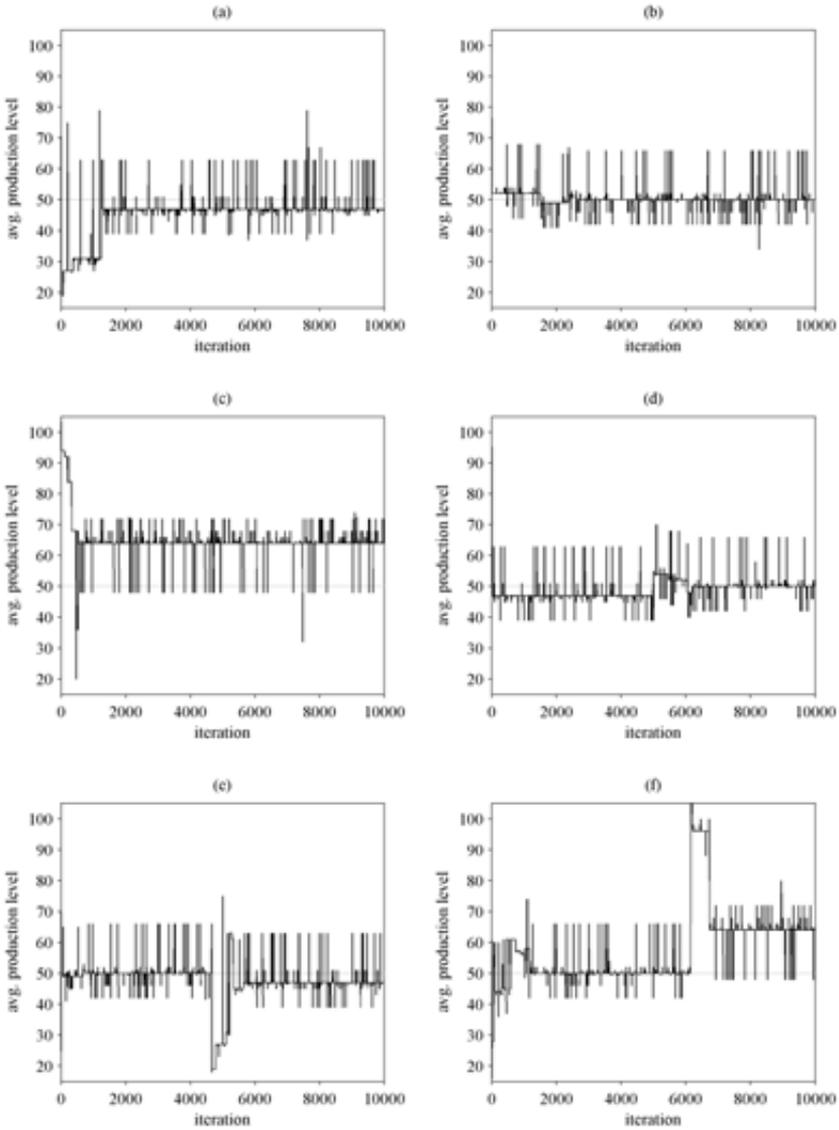
Figure 5.2: Results of six selected runs of EA1. The graphs display for each run how firms' average production level evolves over time. The horizontal line in each graph indicates the competitive equilibrium quantity of 50.
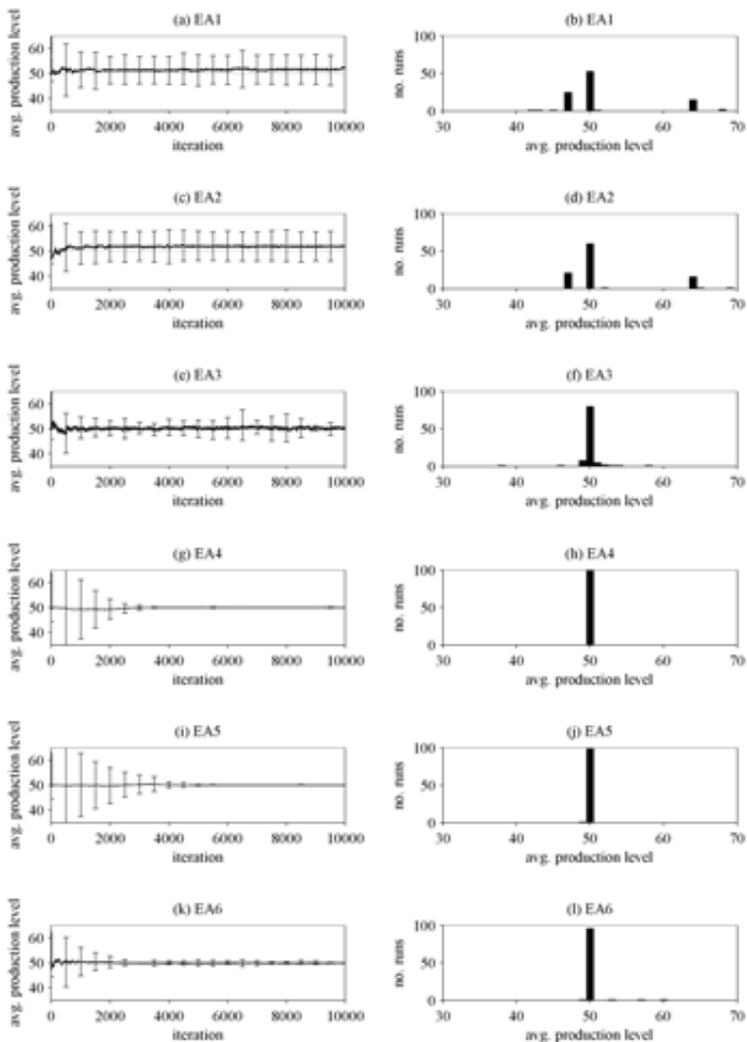
Figure 5.3: Results of EA1 to EA6. The graphs in the left panels display for each EA how firms' average production level evolves over time. The results shown in the graphs are averages over 100 EA runs. Error bars indicate standard deviations. The horizontal line in each graph indicates the competitive equilibrium quantity of 50. The histograms in the right panels display for each EA the distribution of firms' average production level at the end of a run. The distributions are based on 100 EA runs.

Finally, we consider EA3, EA4, EA5, and EA6. The results obtained using these EAs are shown in panels (e) to (l) of Figure 5.3. As can be seen, in the long run the four EAs produce quite similar results. The only noteworthy difference is that the results of EA3 are more volatile than the results of EA4, EA5, and EA6. However, this is to be expected, since EA3 uses a more disruptive mutation operator than the other three EAs (see Section 5.3). What is more interesting to look at is the difference between the results of EA3, EA4, EA5, and EA6 on the one hand and the results of EA1 and EA2 on the other hand. When looking at the aggregate results of 100 EA runs (see the left panels of Figure 5.3), it can be seen that firms' average production level stabilizes at a quantity of 50 in the case of EA3, EA4, EA5, and EA6 while it stabilizes at a quantity of about 52 in the case of EA1 and EA2. Hence, in the case of EA3, EA4, EA5, and EA6, firms on average produce the quantity associated with the competitive equilibrium of the Cournot market (see Section 5.2). In the case of EA1 and EA2, on the other hand, firms on average produce a quantity that is larger than the competitive equilibrium quantity. When looking at the results of individual EA runs (see the right panels of Figure 5.3), it turns out that firms' average production level stabilizes around a quantity of 50 in the case of EA3, EA4, EA5, and EA6 while it stabilizes at a quantity of 47, 50, or 64 in the case of EA1 and EA2. Hence, premature convergence only takes place in the case of EA1 and EA2. It does not take place in the case of the other four EAs. Based on the above observations, it is clear that, both at the aggregate level and at the level of individual runs, EA3, EA4, EA5, and EA6 produce fundamentally different results than EA1 and EA2.

The principal difference between EA1 and EA2 on the one hand and EA3, EA4, and EA5 on the other hand is that EA1 and EA2 use a binary encoding of strategies while EA3, EA4, and EA5 do not use such an encoding (see Table 5.1). EA6 also uses a binary encoding of strategies, but this is a special type of binary encoding, namely a Gray coding. Based on the results reported in this section, it seems that the use of a binary encoding of strategies can have quite significant effects. In the next section, we provide an analysis of these effects, in particular of the phenomenon of premature convergence.

5.6 Analysis of the effect of a binary encoding of strategies

For each of the six EAs considered in the previous section, it turned out that in at least a substantial number of runs firms' average production level stabilized at a quantity of 50, that is, at the quantity associated with the competitive equilibrium of the Cournot oligopoly market. There is a straightforward explanation for this finding. Suppose that in some iteration of an EA the population consists of four identical strategies, each corresponding with a production level of 50. The market price then equals firms' constant marginal cost, and each strategy therefore results in a profit of 0. Suppose now that the mutation operator changes one of the strategies in the population. We refer to this strategy as strategy A. Strategy A's production level may either increase or decrease. If strategy A's production level increases, the total quantity produced in the market will increase and, as a consequence, the market price will fall below firms' constant marginal cost. All four strategies in the population will then result in a loss, but strategy A will result in a larger loss than the other strategies. If on the other hand strategy A's production level decreases, the total quantity produced in the market will decrease and, as a consequence, the market price will rise above firms' constant marginal cost. All four strategies in the population will then result in a profit, but strategy A will result in a smaller profit than the other strategies. Hence, regardless of whether strategy A's production level increases or decreases, the fitness of strategy A will always be lower than the fitness of the other strategies in the population. As a consequence, the probability that in the next iterations of the EA strategy A remains in the population is quite low. Most likely, within one or a few iterations, the selection operator will remove strategy A from the population. The population will then return to its original state, that is, it will again consist of four identical strategies, each corresponding with a production level of 50.

The above mechanism explains why in the EAs considered in the previous section firms' average production level tends to stabilize at the competitive equilibrium quantity of 50. It should be noted that the mechanism has been discussed quite extensively in the evolutionary game theory literature. The mechanism was first discussed by Hansen and Samuelson (1988) and Schaffer (1989), and a comprehensive mathematical treatment of

the mechanism was provided by Vega-Redondo (1997). A discussion of the mechanism can also be found in a study by Vriend (2000) on EA modeling in a Cournot oligopoly environment.

As discussed in the previous section, the fundamental difference between the results obtained using EA1 and EA2 on the one hand and the results obtained using EA3, EA4, EA5, and EA6 on the other hand is that in the case of EA1 and EA2 firms' average production level can stabilize not only at a quantity of 50 but also at a quantity of 47 or 64. This phenomenon of multiple stable quantities is referred to as premature convergence by Alkemade et al. (2006, 2007, 2009). We now show that the premature convergence phenomenon is caused by the use of a binary encoding of strategies.

Suppose that in some iteration of EA1 or EA2 the population consists of four identical strategies, each corresponding with a production level of 47. Each strategy is then represented by the bit string 0101111. Suppose now that out of the 28 bits used to represent the four strategies in the population exactly one bit is inverted by the mutation operator. This means that after applying the mutation operator one of the four strategies in the population has changed while the other three strategies have not changed. We refer to the strategy that has changed as strategy A. The first column of Table 5.3 lists seven bit strings. It is clear that one of these bit strings must represent strategy A (which one depends on which bit has been inverted). For each bit string, the corresponding production level is listed in the second column of the table. The last two columns of the table list for each production level the resulting profit of strategy A as well as the resulting profit of the other three strategies in the population, that is, the strategies corresponding with a production level of 47. Profits were calculated using (5.1) and (5.2). As can be seen in the table, the profit of strategy A will always be smaller than the profit of the other three strategies, regardless of which bit has been inverted. This means that within one or a few iterations of the EA the selection operator will most likely remove strategy A from the population. The population will then return to its original state, that is, it will again consist of four identical strategies, each corresponding with a production level of 47.

In the case of EA1 and EA2, the above mechanism shows that, if the population is in a state in which each strategy corresponds with a production level of 47, the inversion of a single bit is unlikely to upset this state for more than a few iterations. Of course, things

Table 5.3: Effect of the inversion of a single bit given a population in which each strategy corresponds with a production level of 47 (represented by the bit string 0101111).

Bit string strategy A	Production level strategy A	Profit strategy A	Profit other strategies
0101110	46	598	611
0101101	45	630	658
0101011	43	688	752
0100111	39	780	940
0111111	63	-252	-188
0001111	15	660	2068
1101111	111	-5772	-2444

Table 5.4: Effect of the inversion of a single bit given a population in which each strategy corresponds with a production level of 64 (represented by the bit string 1000000).

Bit string strategy A	Production level strategy A	Profit strategy A	Profit other strategies
1000001	65	-3640	-3584
1000010	66	-3696	-3584
1000100	68	-3808	-3584
1001000	72	-4032	-3584
1010000	80	-4480	-3584
1100000	96	-5376	-3584
0000000	0	0	512

Table 5.5: Effect of the inversion of a single bit given a population in which each strategy corresponds with a production level of 48 (represented by the bit string 0110000).

Bit string strategy A	Production level strategy A	Profit strategy A	Profit other strategies
0110001	49	343	336
0110010	50	300	288
0110100	52	208	192
0111000	56	0	0
0100000	32	768	1152
0010000	16	640	1920
1110000	112	-6272	-2688

may be different when the mutation operator inverts two or more bits at the same time. However, this happens only very rarely. (Given a mutation rate of 0.001, this happens on average once in every 2692 iterations of an EA.) The above mechanism therefore explains why in the case of EA1 and EA2 firms' average production level can stabilize at a quantity of 47.

The same explanation also holds for a quantity of 64. This quantity corresponds with the bit string 1000000. The bit strings that can be obtained by inverting a single bit are listed in the first column of Table 5.4. This table has a similar structure as Table 5.3. Like in Table 5.3, the profits listed in the third column of Table 5.4 are always smaller than those listed in the fourth column. This indicates that, if the population is in a state in which each strategy corresponds with a production level of 64, the inversion of a single bit is unlikely to upset this state for more than a few iterations. Taking into account that the simultaneous inversion of two or more bits happens only very rarely, this explains why 64 is a stable quantity in the case of EA1 and EA2.

A question that remains is whether in the case of EA1 and EA2 there are other stable quantities in addition to 47, 50, and 64. To answer this question, we calculated tables similar to Tables 5.3 and 5.4 for all integer quantities between 0 and 127. It turned out that 47, 50, and 64 are the only quantities for which the inversion of a single bit always results in a smaller profit for the mutated strategy than for the three non-mutated strategies. 47, 50, and 64 are therefore the only stable quantities. All other quantities are unstable. Consider for example Table 5.5. This table was calculated for a quantity of 48. As can be seen in the table, the inversion of one of the three rightmost bits results in a larger profit for the mutated strategy (referred to as strategy A in the table) than for the three non-mutated strategies. This indicates that, given a population in which each strategy corresponds with a production level of 48, the inversion of a single bit can relatively easily trigger a transition to a completely different population. This makes 48 an unstable quantity.

We have now shown how the use of a binary encoding of strategies causes the premature convergence observed in the case of EA1 and EA2. Based on our analysis, it is clear that the phenomenon of premature convergence depends crucially on the use of a binary encoding of strategies. An obvious question then is why no premature convergence is observed in the case of EA6. Like EA1 and EA2, EA6 uses a binary encoding

of strategies. In the case of EA6, however, a special type of binary encoding is used, namely a Gray coding. Why is no premature convergence observed when a Gray coding is used? This can be explained as follows. Suppose that in some iteration of EA6 the population consists of four identical strategies. These strategies correspond with a production level of q, where q denotes an integer below 50. Suppose further that the mutation operator inverts a single bit. We refer to the strategy that has changed as strategy A. Due to the use of a Gray coding of strategies, it is always possible that the inversion of a single bit causes one of the four production levels to increase by one. (Notice that this is not the case when an ordinary binary encoding of strategies is used.) Suppose that the production level corresponding with strategy A has indeed increased by one, from q to $q + 1$. It is clear that strategy A then results in a larger profit than the other three strategies in the population. This means that, due to the effect of the selection operator, it is quite likely that strategy A will spread through the population. As a consequence, within a few iterations, all four strategies in the population may correspond with a production level of $q + 1$. This mechanism explains why any quantity below 50 is unstable in the case of EA6. A similar mechanism explains why any quantity above 50 is unstable. Hence, unlike in the case of EA1 and EA2, 50 is the only stable quantity in the case of EA6. Because of this, EA6 does not exhibit premature convergence.

5.7 Conclusions

In a paper on EA modeling, Dawid and Kopel (1998) warn that "we have to be aware of the fact that simulation results may crucially depend on implementation details which have hardly any economic meaning" (p. 311). The present study can be seen as an illustration of this important but somewhat overlooked point. In the context of quantity competition among firms, it is difficult if not impossible to give a sensible economic interpretation to the use of a binary encoding of strategies. In fact, the use of a binary encoding seems merely a relic from the genetic algorithm literature in the field of computer science. Of course, nothing would be wrong with the use of a binary encoding if its effect on the results produced by an EA were insignificant. However, our computer simulations and the subsequent analysis make clear that this need not be the case. They show that the use of a binary encoding may lead to a phenomenon known as premature

convergence. This phenomenon is an artifact that depends crucially on strategies being encoded in binary form.

Based on our findings, we conclude that in general the use of a binary encoding of strategies is undesirable. By not using a binary encoding, one avoids the risk of having to deal with all kinds of artifacts, such as the premature convergence observed by Alkemade et al. (2006, 2007, 2009).[7] For various examples of studies in which EAs are employed without using a binary encoding, we refer to Sellgren (2001), Unver (2001), Gerding et al. (2003), Lux and Schornstein (2005), Clemens and Riechmann (2006), Dawid and Dermietzel (2006), and Haruvy et al. (2006). It should be noted, however, that there are special cases in which we consider the use of a binary encoding perfectly acceptable. In a prisoner's dilemma, for example, agents have to make decisions that are intrinsically binary, namely decisions between cooperation and defection. The use of a binary encoding of strategies (e.g., Axelrod, 1987) then seems a very natural choice that is unlikely to cause any artifacts.

The more general point that we want to make is that, when one employs an EA for economic modeling, all elements of the EA should have a meaningful economic interpretation (see also Dawid & Dermietzel, 2006).[8] Many EAs employed in economic research have been adopted from the computer science literature without any substantial modification. Such EAs are likely to contain elements of which the economic interpretation is unclear. The use of a binary encoding of strategies is an example of such an element. Other EA elements of which the economic interpretation requires special attention include the population size (Alkemade et al., 2006, 2007, 2009), the selection operator (Van Bragt et al., 2001; Dawid & Dermietzel, 2006), and the strategy updating regime (Dawid & Dermietzel, 2006). As we have shown in this chapter, EA elements without a sensible economic interpretation may lead to simulation results that lack a

---

[7]Our results seem to suggest that, if one insists on the use of a binary encoding of strategies, it is advisable to use a Gray coding rather than an ordinary binary encoding. However, even though in this chapter we have not observed any artifacts of the use of a Gray coding, it seems quite well possible that such artifacts will be observed in other contexts. For example, the fact that in the case of a Gray coding the smallest and the largest decoded value differ by only one bit (see Table 5.2) seems unnatural and it may well be that this sometimes has unintended consequences.

[8]This point stands in stark contrast with one of the conclusions reached by Alkemade et al. (2006, 2007, 2009). They state that "economic model parameters and evolutionary algorithm parameters should be treated separately" (Alkemade et al., 2006, p. 367).

sound underlying economic rationale. To avoid such results, paying close attention to the economic interpretation of the various elements of an EA is absolutely essential.

Chapter 6

Summary

In this thesis, various computational and game-theoretic approaches to economic modeling have been studied. Unlike traditional approaches to economic modeling, the approaches studied in this thesis do not rely on the assumption that economic agents behave in a fully rational way. Instead, economic agents are assumed to be boundedly rational. Agents for instance do not take into consideration the influence of their behavior on the behavior of others. Abandoning the assumption of full rationality has a number of consequences for the way in which economic reality is being modeled. Traditionally, economic models are mostly of a static nature and have a strong focus on deriving equilibria. Also, models are usually analyzed mathematically. In models of boundedly rational behavior, dynamic elements play a much more prominent role and there is less emphasis on equilibrium behavior. Also, to analyze models of boundedly rational behavior, researchers not only use mathematical techniques but they also rely heavily on computer simulations.

Within the field of economics, there are a number of subfields that are concerned with modeling boundedly rational behavior. The work that has been presented in this thesis relates to three of these subfields. These three subfields were referred to as evolutionary game theory, economic learning theory, and agent-based computational economics. The first two subfields are still relatively closely related to traditional approaches to economic modeling, and in particular to the traditional game-theoretic approach. The third subfield, agent-based computational economics, differs much more strongly from traditional approaches to economic modeling. This subfield has been

strongly influenced by computer science research and is dominated by studies based on computer simulations.

The core of this thesis consisted of four chapters (Chapters 2 to 5). In each of these chapters, an independent piece of research was presented and a separate contribution to the literature was made. In Chapters 2 and 3, the focus was on a substantive economic issue, namely the issue of the emergence of cooperative behavior. In Chapters 4 and 5, the focus was on methodological issues, namely issues concerning the application of evolutionary algorithms for economic modeling purposes. Each of the chapters will now be briefly summarized.

In Chapter 2, a simple model of the learning behavior of boundedly rational agents was studied. The model was adopted from the computer science literature, where it is referred to as Q-learning (Watkins, 1989; Watkins and Dayan, 1992). The Q-learning model, which belongs to the broader family of reinforcement learning models, makes only very limited assumptions about the information available to agents and the cognitive abilities of agents. In Chapter 2, the Q-learning model was studied in the context of a Cournot oligopoly market. The agents were firms that have to decide on their production level. Firms can increase their profits if they cooperate with each other by jointly decreasing the quantity they produce. The analysis of the Q-learning model was performed partly mathematically and partly using computer simulations. The main contribution of Chapter 2 consists of showing that the Q-learning model is able to explain the emergence of cooperative behavior. Many similar learning models studied in the economic literature are unable to explain this phenomenon.

In Chapter 3, a relatively recent explanation for the emergence of cooperative behavior was studied. According to this explanation, cooperative behavior can be a consequence of evolutionary dynamics combined with local interaction among spatially distributed agents. The explanation was first proposed by the biologists Nowak and May (1992) and was introduced in the economic literature by Bergstrom and Stark (1993) and Eshel, Samuelson and Shaked (1998). In the economic literature, the explanation has been studied mainly for agents that are located in a one-dimensional world and that can choose from only two actions (i.e., cooperation and defection). The advantage of this highly stylized setting is that it can be analyzed mathematically. The disadvantage

is that it is unclear whether results derived in this setting are also valid in other more complex settings.

The contribution of Chapter 3 consists of studying some of these more complex settings, in particular settings in which agents are located in a two-dimensional world and settings in which agents can choose from more than two actions. More specifically, in the models considered in Chapter 3, agents were firms that compete with each other based on price. Firms can cooperate with each other by jointly increasing their price. Because the models studied in Chapter 3 seemed to be mathematically intractable, most of the analysis was performed using computer simulations. The main finding of Chapter 3 is that the emergence of cooperative behavior depends strongly on the amount of information available to agents. Agents tend to behave most cooperatively if they have only a very limited amount of information about their neighbors. Contrary to earlier research reported in the literature, it was found that in some cases agents behave even less cooperatively than they would do according to the Nash equilibrium prediction.

Chapter 4 was concerned with the application of genetic algorithms (i.e., a specific type of evolutionary algorithms) for modeling purposes in economic research. The application of genetic algorithms for economic modeling purposes is quite popular in agent-based computational economics research. Researchers almost always use computer simulations to analyze genetic algorithm models. In Chapter 4, it was shown how genetic algorithm models can be analyzed mathematically rather than using computer simulations. The proposed approach for mathematically analyzing genetic algorithm models relies on a mathematical technique that is frequently used in the evolutionary game theory literature.

The main contribution of Chapter 4 consists of exploring the consequences of the similarities between evolutionary models studied in the area of agent-based computational economics and evolutionary models studied in the area of evolutionary game theory. As discussed in Chapter 1, these two research areas are quite separated from each other. The similarities between the models studied in the two areas suggest that the separation of the areas is somewhat artificial and may impede the further development of evolutionary modeling approaches. In Chapter 4, this point was illustrated by showing how agent-based computational economics research can benefit from a mathematical technique that is well-known in the evolutionary game theory literature. Chapter 4 also

indicated an advantage of the agent-based computational economics approach, and in particular of the use of computer simulations. In the evolutionary game theory literature, one almost always focuses on the limit case in which the probability of experimentation (or mutation) approaches zero. Focusing on this limit case is convenient from a mathematical point of view, but it may lead to less realistic modeling. In Chapter 4, it was shown how computer simulations can be used to test the sensitivity of one's modeling results to the assumption of an almost zero experimentation probability. It was found that in some cases results are quite sensitive to this assumption.

In Chapter 5, the focus was on a problematic aspect of genetic algorithm models. These models typically rely on a binary encoding of strategies. The use of such an encoding originates from the computer science literature, in which genetic algorithms were first introduced. In Chapter 5, it was argued that a binary encoding of strategies usually does not have a meaningful economic interpretation. It was also shown that the use of a binary encoding may lead to artifacts in the results of an analysis. Hence, the contribution of Chapter 5 consists of pointing out that in general it is not appropriate to use genetic algorithm models with a binary encoding of strategies for economic modeling purposes. The more general contribution of Chapter 5 consists of illustrating how the fairly high complexity of models in agent-based computational economics research combined with a somewhat ad hoc justification may lead to results that have no economic significance and that are merely artifacts of the model. The analysis presented in Chapter 5 can be seen as an argument for simplifying the evolutionary models studied in the agent-based computational economics literature. Simplifying these models would also illuminate the similarities with models studied in the evolutionary game theory literature. Hence, it may help to reduce the somewhat artificial separation between the research areas of agent-based computational economics and evolutionary game theory.

Nederlandse Samenvatting (Summary in Dutch)

Dit proefschrift gaat over computationele en speltheoretische methodes voor economisch modelleren. In tegenstelling tot de traditionele manier van economisch modelleren, zijn de in dit proefschrift onderzochte methodes niet gebaseerd op de aanname dat economische agenten volledig rationeel handelen. In plaats daarvan wordt verondersteld dat aan het handelen van economische agenten een beperkte rationaliteit ('bounded rationality') ten grondslag ligt. Agenten denken bijvoorbeeld slechts een klein aantal stappen vooruit en ze overzien bijvoorbeeld niet hoe hun eigen gedrag dat van anderen beïnvloedt. Het afstand doen van de aanname van volledige rationaliteit heeft verschillende gevolgen voor de manier waarop de economische werkelijkheid wordt gemodelleerd. Traditioneel zijn economische modellen veelal statisch van aard en sterk gericht op evenwichtssituaties. Modellen worden doorgaans wiskundig geanalyseerd. Wanneer het uitgangspunt van de beperkte rationaliteit wordt gekozen, gaan dynamische aspecten een veel belangrijker rol spelen en wordt er minder nadruk gelegd op evenwichtssituaties. Voor het analyseren van modellen die uitgaan van beperkte rationaliteit wordt naast de wiskundige benadering ook veel gebruik gemaakt van computersimulaties.

Binnen de economische wetenschap zijn er verschillende onderzoeksgebieden die zich bezighouden met modelleren onder de aanname van beperkte rationaliteit. Het werk dat in dit proefschrift wordt gepresenteerd is aan drie onderzoeksgebieden gerelateerd. In het proefschrift worden deze gebieden aangeduid als de evolutionaire speltheorie, de economische leertheorie en de agent-gebaseerde computationele economie. De eerste twee gebieden staan relatief dicht bij de traditionele manier van economisch modelleren, en in het bijzonder bij de traditionele speltheoretische benadering. Het derde

gebied, de agent-gebaseerde computationele economie, staat veel verder van de traditionele manier van economisch modelleren af. Dit gebied is sterk beïnvloed door de informatica en wordt gedomineerd door onderzoek dat gebruik maakt van computersimulaties.

Dit proefschrift bestaat uit zes hoofdstukken. Het eerste hoofdstuk biedt een algemene inleiding en het laatste hoofdstuk sluit het proefschrift af. In de vier tussenliggende hoofdstukken worden vier op zichzelf staande studies gepresenteerd. Wat deze studies gemeenschappelijk hebben, is dat ze zich alle vier bezighouden met het bestuderen van modellen die gebaseerd zijn op de aanname van beperkte rationaliteit. De studies in hoofdstuk 2 en 3 gaan over het modelleren van het ontstaan van samenwerkingsgedrag. Hoofdstuk 4 en 5 hebben een meer methodologisch karakter. De studies die in deze twee hoofdstukken worden gepresenteerd gaan over het gebruik van zogeheten evolutionaire algoritmes voor economisch modelleren. Hieronder worden hoofdstuk 2 tot en met 5 kort samengevat.

In hoofdstuk 2 wordt een eenvoudig model van het leergedrag van beperkt rationele agenten bestudeerd. Het model is overgenomen uit de informatica literatuur, waar het wordt aangeduid met de term Q-leren ('Q-learning'). Het Q-leermodel maakt slechts zeer beperkte aannames over de informatie waar agenten over beschikken en over de cognitieve vaardigheden van agenten. In hoofdstuk 2 wordt het Q-leermodel bestudeerd in de context van een Cournot oligopoliemarkt. De agenten zijn bedrijven die moeten beslissen hoeveel ze gaan produceren. Bedrijven kunnen hun winst verhogen als ze onderling samenwerken door de hoeveelheid die ze gezamenlijk produceren te verlagen. De analyse van het Q-leermodel wordt deels wiskundig en deels met computersimulaties uitgevoerd. Het belangrijkste resultaat van hoofdstuk 2 is dat het Q-leermodel in staat is om het ontstaan van samenwerkingsgedrag te verklaren. Veel vergelijkbare leermodellen die in de economische literatuur worden bestudeerd kunnen dit verschijnsel niet verklaren.

In hoofdstuk 3 wordt een relatief recente verklaring voor het ontstaan van samenwerkingsgedrag onderzocht. Volgens deze verklaring kan samenwerkingsgedrag het resultaat zijn van evolutionaire mechanismes in combinatie met lokale interactie tussen ruimtelijk georganiseerde agenten. Deze verklaring heeft begin jaren 90 van de vorige eeuw in de biologie zijn intrede gedaan en is vervolgens ook in de economische litera-

tuur geı"ntroduceerd. In de economische literatuur is de verklaring vooral onderzocht voor agenten die zich in een eendimensionale wereld bevinden en die uit slechts twee acties kunnen kiezen (wel of niet samenwerken). Het voordeel van deze sterk gestileerde opzet is dat een grondige wiskundige analyse mogelijk is. Het nadeel is dat het niet duidelijk is of resultaten die in deze opzet zijn afgeleid ook in andere complexere situaties van toepassing zijn.

De bijdrage van hoofdstuk 3 bestaat uit het onderzoeken van enkele van deze complexere situaties, in het bijzonder situaties waarin agenten zich in een tweedimensionale wereld bevinden en situaties waarin agenten uit meer dan twee acties kunnen kiezen. Meer specifiek houdt hoofdstuk 3 zich bezig met modellen waarin bedrijven op basis van prijs met elkaar concurreren. Omdat de modellen die in hoofdstuk 3 worden onderzocht wiskundig niet volledig lijken te kunnen worden geanalyseerd, wordt de analyse voor het grootste deel met computersimulaties uitgevoerd. Het voornaamste resultaat van hoofdstuk 3 is dat het ontstaan van samenwerkingsgedrag sterk afhankelijk is van de hoeveelheid informatie die agenten tot hun beschikking hebben. Agenten gedragen zich doorgaans het meest coöperatief indien ze slechts een beperkte hoeveelheid informatie hebben over hun buren. In tegenstelling tot eerder onderzoek blijkt verder dat agenten zich in bepaalde gevallen minder coöperatief gedragen dan in het Nash evenwicht.

Hoofdstuk 4 gaat over het toepassen van zogeheten genetische algoritmes (een speciaal soort evolutionaire algoritmes) voor economisch modelleren. Het gebruik van genetische algoritmes voor economisch modelleren is behoorlijk populair in het onderzoeksgebied van de agent-gebaseerde computationele economie. Onderzoekers gebruiken vrijwel altijd computersimulaties om modellen gebaseerd op genetische algoritmes te analyseren. In hoofdstuk 4 wordt uiteengezet hoe dit soort modellen wiskundig kunnen worden geanalyseerd in plaats van met computersimulaties. De voorgestelde methode voor het wiskundig analyseren van modellen gebaseerd op genetische algoritmes maakt gebruik van een wiskundige techniek die ook veelvuldig wordt gebruikt in de literatuur op het gebied van de evolutionaire speltheorie.

De belangrijkste bijdrage van hoofdstuk 4 bestaat uit het verkennen van de overeenkomsten tussen evolutionaire modellen die in het gebied van de agent-gebaseerde computationele economie worden onderzocht en evolutionaire modellen die worden on-

derzocht in het gebied van de evolutionaire speltheorie. Deze twee onderzoeksgebieden opereren in hoge mate onafhankelijk van elkaar. De overeenkomsten tussen de modellen die in de twee gebieden worden onderzocht suggereren dat de scheiding tussen de gebieden enigszins kunstmatig is en wellicht een belemmering zou kunnen vormen voor de verdere ontwikkeling van het evolutionaire modelleerparadigma. In hoofdstuk 4 wordt dit punt onderbouwd door te laten zien hoe de agent-gebaseerde computationele economie kan profiteren van een wiskundige techniek die veelvuldig wordt gebruikt in de evolutionaire speltheorie. Hoofdstuk 4 laat ook een voordeel zien van de benadering van de agent-gebaseerde computationele economie, en in het bijzonder van het gebruik van computersimulaties. In de evolutionaire speltheorie richt men zich vrijwel altijd op het limietgeval waarin de kans op experimenteren (of muteren) naar nul nadert. Hoewel het vanuit wiskundig opzicht inderdaad handig is om je op dit limietgeval te concentreren, kan dit wel tot minder realistische modellen leiden. In hoofdstuk 4 wordt getoond hoe computersimulaties kunnen worden gebruikt om te testen in welke mate modelleerresultaten afhankelijk zijn van de aanname van een experimenteerkans van vrijwel nul. In bepaalde gevallen blijkt de afhankelijkheid van deze aanname behoorlijk groot te zijn.

Hoofdstuk 5 houdt zich bezig met een problematisch aspect van modellen die gebaseerd zijn op genetische algoritmes. Deze modellen maken gewoonlijk gebruik van een zogeheten binaire codering van strategieën. Het gebruik van een dergelijk codering stamt uit de informatica literatuur, waar genetische algoritmes hun oorsprong vinden. In hoofdstuk 5 wordt beargumenteerd dat een binaire codering van strategieën doorgaans geen betekenisvolle economische interpretatie heeft. Bovendien wordt getoond hoe het gebruik van een binaire codering tot artificiële resultaten kan leiden. Hoofdstuk 5 laat dus zien dat het over het algemeen niet gewenst is om genetische algoritmes met een binaire codering van strategieën te gebruiken voor economisch modelleren. In meer algemene zin illustreert hoofdstuk 5 hoe de behoorlijk hoge complexiteit van modellen in de agent-gebaseerde computationele economie, in combinatie met een enigszins ad hoc onderbouwing, kan leiden tot resultaten die geen economische betekenis hebben en die slechts artefacten van het gebruikte model zijn. De analyse die in hoofdstuk 5 wordt gepresenteerd kan worden gezien als een pleidooi voor het vereenvoudigen van de evolutionaire modellen die in de agent-gebaseerde computationele economie wor-

den gebruikt. Het vereenvoudigen van deze modellen zou ook meer inzicht geven in de gelijkenissen met modellen die in de evolutionaire speltheorie worden gebruikt. Op die manier kan het bijdragen aan het verkleinen van de enigszins kunstmatige scheiding tussen de onderzoeksgebieden van de agent-gebaseerde computationele economie en de evolutionaire speltheorie.

Bibliography

Alkemade, F., La Poutré, J. A., & Amman, H. M. (2006). Robust evolutionary algorithm design for socio-economic simulation. *Computational Economics*, *28*(4), 355-370.

Alkemade, F., La Poutré, J. A., & Amman, H. M. (2007). On social learning and robust evolutionary algorithm design in the Cournot oligopoly game. *Computational Intelligence*, *23*(2), 162-175.

Alkemade, F., La Poutré, J. A., & Amman, H. M. (2009). Robust evolutionary algorithm design for socio-economic simulation: A correction. *Computational Economics*, *33*(1), 99-101.

Alkemade, F., Van Bragt, D. D. B., & La Poutré, J. A. (2005). Stabilization of tag-mediated interaction by sexual reproduction in an evolutionary agent system. *Information Sciences*, *170*(1), 101-119.

Alós-Ferrer, C. (2004). Cournot versus Walras in dynamic oligopolies with memory. *International Journal of Industrial Organization*, *22*, 193-217.

Andreoni, J., & Miller, J. H. (1995). Auctions with artificial adaptive agents. *Games and Economic Behavior*, *10*(1), 39-64.

Arifovic, J. (1994). Genetic algorithm learning and the cobweb model. *Journal of Economic Dynamics and Control*, *18*, 3-28.

Arifovic, J. (1996). The behavior of the exchange rate in the genetic algorithm and experimental economies. *Journal of Political Economy*, *104*(3), 510-541.

Arifovic, J., & Maschek, M. K. (2006). Revisiting individual evolutionary learning in the cobweb model: An illustration of the virtual spite-effect. *Computational Economics*, *28*(4), 333-354.

Arthur, W. B. (1991). Designing economic agents that act like human agents: A behav-

ioral approach to bounded rationality. *American Economic Review (Papers and Proceedings)*, *81*(2), 353-359.

Arthur, W. B. (1994). Inductive reasoning and bounded rationality. *American Economic Review (Papers and Proceedings)*, *84*(2), 406-411.

Ashlock, D., & Kim, E.-Y. (2008). Fingerprinting: Visualization and automatic analysis of prisoner's dilemma strategies. *IEEE Transactions on Evolutionary Computation*, *12*(5), 647-659.

Ashlock, D., Kim, E.-Y., & Leahy, N. (2006). Understanding representational sensitivity in the iterated prisoner's dilemma with fingerprints. *IEEE Transactions on Systems, Man, and Cybernetics C*, *36*(4), 464-475.

Ashlock, D., Smucker, M. D., Stanley, E. A., & Tesfatsion, L. (1996). Preferential partner selection in an evolutionary study of prisoner's dilemma. *Biosystems*, *37*(1–2), 99-125.

Axelrod, R. (1984). *The evolution of cooperation*. Basic Books.

Axelrod, R. (1987). The evolution of strategies in the iterated prisoner's dilemma. In L. Davis (Ed.), *Genetic algorithms and simulated annealing* (p. 32-41). Morgan Kaufmann.

Axelrod, R. (1997a). Advancing the art of simulation in the social sciences. In R. Conte, R. Hegselmann, & P. Terna (Eds.), *Simulating social phenomena* (p. 21-40). Springer.

Axelrod, R. (1997b). *The complexity of cooperation: Agent-based models of competition and collaboration*. Princeton University Press.

Axelrod, R. (2006a). Advancing the art of simulation in the social sciences. In J.-P. Rennard (Ed.), *Handbook of research on nature inspired computing for economics and management.* Idea Group.

Axelrod, R. (2006b). Agent-based modeling as a bridge between disciplines. In L. Tesfatsion & K. L. Judd (Eds.), *Handbook of computational economics, Volume 2* (p. 1565-1584). Elsevier.

Axtell, R. (2000). *Why agents? On the varied motivations for agent computing in the social sciences* (Working Paper No. 17). Center on Social and Economic Dynamics, The Brookings Institution.

Axtell, R., Axelrod, R., Epstein, J. M., & Cohen, M. D. (1996). Aligning simulation

models: A case study and results. *Computational and Mathematical Organization Theory*, *1*(2), 123-141.

Barr, J., & Tassier, T. (2010). Endogenous neighborhood selection and the attainment of cooperation in a spatial prisoner's dilemma game. *Computational Economics*, *35*(3), 211-234.

Barrett, R., Berry, M., Chan, T. F., Demmel, J., Donato, J., Dongarra, J., et al. (2008). *Templates for the solution of linear systems: Building blocks for iterative methods* (Second ed.). SIAM. (Available online at www.netlib.org/linalg/html_templates/Templates.html)

Bell, A. M. (2001). Reinforcement learning rules in a repeated game. *Computational Economics*, *18*, 89-111.

Bergin, J., & Bernhardt, D. (2005). *Cooperation through imitation* (Working Paper No. 1042). Queen's Economics Department.

Bergstrom, T. C., & Stark, O. (1993). How altruism can prevail in an evolutionary environment. *American Economic Review*, *83*(2), 149-155.

Beyer, H.-G. (2001). *The theory of evolution strategies*. Springer.

Beyer, H.-G., & Schwefel, H.-P. (2002). Evolution strategies: A comprehensive introduction. *Natural Computing*, *1*(1), 3-52.

Bilancini, E., & Boncinelli, L. (2009). The co-evolution of cooperation and defection under local interaction and endogenous network formation. *Journal of Economic Behavior and Organization*, *70*(1–2), 186-195.

Blaug, M. (2002). Ugly currents in modern economics. In U. Mäki (Ed.), *Fact and fiction in economics: Models, realism and social construction* (p. 35-56). Cambridge University Press.

Brenner, T. (1999). *Modelling learning in economics*. Edward Elgar Publishing.

Brenner, T. (2006). Agent learning representation: Advice on modelling economic learning. In L. Tesfatsion & K. L. Judd (Eds.), *Handbook of computational economics, Volume 2* (p. 895-947). Elsevier.

Brock, W. A., & Hommes, C. H. (1997). A rational route to randomness. *Econometrica*, *65*, 1059-1095.

Bullard, J., & Duffy, J. (1998). A model of learning and emulation with artificial adaptive agents. *Journal of Economic Dynamics and Control*, *22*(2), 179-207.

Camerer, C. (2003). *Behavioral game theory: Experiments in strategic interaction*. Princeton University Press.

Camerer, C., & Ho, T.-H. (1999). Experience-weighted attraction learning in normal form games. *Econometrica*, *67*, 827-873.

Casari, M. (2008). Markets in equilibrium with firms out of equilibrium: A simulation study. *Journal of Economic Behavior and Organization*, *65*(2), 261-276.

Chattoe, E. (1998). Just how (un)realistic are evolutionary algorithms as representations of social processes? *Journal of Artificial Societies and Social Simulation*, *1*(3).

Chen, H.-C., & Chow, Y. (2009). Evolutionary prisoner's dilemma games with one-dimensional local interaction and imitation. *Advances in Applied Probability*, *41*(1), 154-176.

Chong, S. Y., & Yao, X. (2005). Behavioral diversity, choices and noise in the iterated prisoner's dilemma. *IEEE Transactions on Evolutionary Computation*, *9*(6), 540-551.

Chong, S. Y., & Yao, X. (2007). Multiple choices and reputation in multiagent interactions. *IEEE Transactions on Evolutionary Computation*, *11*(6), 689-711.

Clemens, C., & Riechmann, T. (2006). Evolutionary dynamics in public good games. *Computational Economics*, *28*(4), 399-420.

Colander, D., Holt, R. P. F., & Rosser, J. B. (2004a). *The changing face of economics: Conversations with cutting edge economists*. University of Michigan Press.

Colander, D., Holt, R. P. F., & Rosser, J. B. (2004b). The changing face of mainstream economics. *Review of Political Economy*, *16*(4), 485-499.

Crowley, P. H., Provencher, L., Sloane, S., Dugatkin, L. A., Spohn, B., Rogers, L., et al. (1996). Evolving cooperation: The role of individual recognition. *Biosystems*, *37*(1–2), 49-66.

Dawid, H. (1996). *Adaptive learning by genetic algorithms: Analytical results and applications to economic models* (No. 441). Springer.

Dawid, H., & Dermietzel, J. (2006). How robust is the equal split norm? Responsive strategies, selection mechanisms and the need for economic interpretation of simulation parameters. *Computational Economics*, *28*(4), 371-397.

Dawid, H., & Kopel, M. (1998). On economic applications of the genetic algorithm: A model of the cobweb type. *Journal of Evolutionary Economics*, *8*(3), 297-315.

Dawid, H., La Poutré, H., & Yao, X. (2008). Computational intelligence in economic games and policy design. *IEEE Computational Intelligence Magazine*, *3*(4), 22-26.

Dawkins, R. (1989). *The selfish gene* (Revised ed.). Oxford University Press. (First edition published in 1976)

Dixon, H. D. (2000). Keeping up with the Joneses: Competition and the evolution of collusion. *Journal of Economic Behavior and Organization*, *43*, 223-238.

Droste, E., Hommes, C., & Tuinstra, J. (2002). Endogenous fluctuations under evolutionary pressure in Cournot competition. *Games and Economic Behavior*, *40*, 232-269.

Duffy, J. (2006). Agent-based models and human subject experiments. In L. Tesfatsion & K. L. Judd (Eds.), *Handbook of computational economics, Volume 2* (p. 949-1011). Elsevier.

Epstein, J. M. (2006). *Generative social science*. Princeton University Press.

Epstein, J. M., & Axtell, R. (1996). *Growing artificial societies: Social science from the bottom up*. MIT Press.

Erev, I., & Roth, A. E. (1998). Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria. *American Economic Review*, *88*, 848-881.

Eshel, I., Herreiner, D. K., Samuelson, L., Sansone, E., & Shaked, A. (2000). Cooperation, mimesis, and local interaction. *Sociological Methods and Research*, *28*(3), 341-364.

Eshel, I., Samuelson, L., & Shaked, A. (1998). Altruists, egoists, and hooligans in a local interaction model. *American Economic Review*, *88*(1), 157-179.

Eshel, I., Sansone, E., & Shaked, A. (1999). The emergence of kinship behavior in structured populations of unrelated individuals. *International Journal of Game Theory*, *28*(4), 447-463.

Fagiolo, G., Birchenhall, C., & Windrum, P. (2007). Empirical validation in agent-based models: Introduction to the special issue. *Computational Economics*, *30*(3), 189-194.

Fehr, E., & Fischbacher, U. (2003). The nature of human altruism. *Nature*, *425*, 785-791.

Fogel, D. B. (1993). Evolving behaviors in the iterated prisoner's dilemma. *Evolutionary Computation*, *1*(1), 77-97.

Fosco, C., & Mengel, F. (2011). Cooperation through imitation and exclusion in networks. *Journal of Economic Dynamics and Control*, *35*(5), 641-658.

Foster, D., & Young, P. (1990). Stochastic evolutionary game dynamics. *Theoretical Population Biology*, *38*(2), 219-232.

Franke, R. (1998). Coevolution and stable adjustments in the cobweb model. *Journal of Evolutionary Economics*, *8*(4), 383-406.

Freidlin, M. I., & Wentzell, A. D. (1998). *Random perturbations of dynamical systems* (Second ed.). Springer.

Fudenberg, D., & Levine, D. K. (1998). *The theory of learning in games*. MIT Press.

Gen, M., & Cheng, R. (2000). *Genetic algorithms & engineering optimization*. John Wiley & Sons.

Georges, C. (2006). Learning with misspecification in an artificial currency market. *Journal of Economic Behavior and Organization*, *60*(1), 70-84.

Gerding, E., Van Bragt, D., & La Poutré, J. A. (2003). Multi-issue negotiation processes by evolutionary simulation, validation and social extensions. *Computational Economics*, *22*(1), 39-63.

Gintis, H. (2000). *Game theory evolving: A problem-centered introduction to modeling strategic interaction*. Princeton University Press.

Gintis, H. (2009). *Game theory evolving: A problem-centered introduction to modeling strategic interaction* (Second ed.). Princeton University Press.

Goldberg, D. E. (1989). *Genetic algorithms in search, optimization, and machine learning*. Addison-Wesley.

Hansen, R. G., & Samuelson, W. F. (1988). Evolution in economic games. *Journal of Economic Behavior and Organization*, *10*(3), 315-338.

Haruvy, E., Roth, A. E., & Utku Ünver, M. (2006). The dynamics of law clerk matching: An experimental and computational investigation of proposals for reform of the market. *Journal of Economic Dynamics and Control*, *30*(3), 457-486.

Hofbauer, J., & Sandholm, W. H. (2007). Evolution in games with randomly disturbed payoffs. *Journal of Economic Theory*, *132*, 47-69.

Hoffmann, R. (1999). The independent localisations of interaction and learning in the

repeated prisoner's dilemma. *Theory and Decision*, *47*(1), 57-72.

Holland, J. H., & Miller, J. H. (1991). Artificial adaptive agents in economic theory. *American Economic Review (Papers and Proceedings)*, *81*(2), 365-370.

Huberman, B. A., & Glance, N. S. (1993). Evolutionary games and computer simulations. *Proceedings of the National Academy of Sciences*, *90*(16), 7716-7718.

Huck, S., Normann, H.-T., & Oechssler, J. (2004a). Through trial and error to collusion. *International Economic Review*, *45*, 205-224.

Huck, S., Normann, H.-T., & Oechssler, J. (2004b). Two are few and four are many: Number effects in experimental oligopolies. *Journal of Economic Behavior and Organization*, *53*, 435-446.

Ifti, M., Killingback, T., & Doebeli, M. (2004). Effects of neighbourhood size and connectivity on the spatial continuous prisoner's dilemma. *Journal of Theoretical Biology*, *231*(1), 97-106.

Ishibuchi, H., & Namikawa, N. (2005). Evolution of iterated prisoner's dilemma game strategies in structured demes under random pairing in game playing. *IEEE Transactions on Evolutionary Computation*, *9*(6), 552-561.

Jun, T., & Sethi, R. (2007). Neighborhood structure and the evolution of cooperation. *Journal of Evolutionary Economics*, *17*(5), 623-646.

Jun, T., & Sethi, R. (2009). Reciprocity in evolving social networks. *Journal of Evolutionary Economics*, *19*(3), 379-396.

Kaelbling, L. P., Littman, M. L., & Moore, A. W. (1996). Reinforcement learning: A survey. *Journal of Artificial Intelligence Research*, *4*, 237-285.

Kandori, M., Mailath, G. J., & Rob, R. (1993). Learning, mutation, and long run equilibria in games. *Econometrica*, *61*, 29-56.

Kemeny, J. G., & Snell, J. L. (1960). *Finite Markov chains*. D. Van Nostrand Company.

Kimbrough, S. O., & Lu, M. (2003). A note on Q-learning in the Cournot game. In *Proceedings of the second workshop on e-business*.

Kirchkamp, O. (1999). Simultaneous evolution of learning rules and strategies. *Journal of Economic Behavior and Organization*, *40*(3), 295-312.

Kirchkamp, O. (2000). Spatial evolution of automata in the prisoners' dilemma. *Journal of Economic Behavior and Organization*, *43*(2), 239-262.

Kirman, A. P., & Vriend, N. J. (2001). Evolving market structure: An ACE model

of price dispersion and loyalty. *Journal of Economic Dynamics and Control*, *25*, 459-502.

Kuhn, T. S. (1996). *The structure of scientific revolutions* (Third ed.). University of Chicago Press.

Lehtinen, A., & Kuorikoski, J. (2007). Computing the perfect model: Why do economists shun simulation? *Philosophy of Science*, *74*(3), 304-329.

Leombruni, R., & Richiardi, M. (2005). Why are economists sceptical about agent-based simulations? *Physica A: Statistical Mechanics and Its Applications*, *355*(1), 103-109.

Littman, M. L., & Stone, P. (2001). Leading best-response strategies in repeated games. In *Proceedings of the seventeenth international joint conference on artificial intelligence.*

Lux, T., & Schornstein, S. (2005). Genetic learning as an explanation of stylized facts of foreign exchange markets. *Journal of Mathematical Economics*, *41*(1–2), 169-196.

Marks, R. E. (1992). Breeding hybrid strategies: Optimal behaviour for oligopolists. *Journal of Evolutionary Economics*, *2*(1), 17-38.

Maschek, M. K. (2010). Intelligent mutation rate control in an economic application of genetic algorithms. *Computational Economics*, *35*(1), 25-49.

Maynard Smith, J. (1982). *Evolution and the theory of games*. Cambridge University Press.

McKelvey, R. D., & Palfrey, T. R. (1995). Quantal response equilibria for normal form games. *Games and Economic Behavior*, *10*, 6-38.

Mengel, F. (2009). Conformism and cooperation in a local interaction model. *Journal of Evolutionary Economics*, *19*(3), 397-415.

Michalewicz, Z. (1996). *Genetic algorithms + data structures = evolution programs* (Third ed.). Springer.

Miller, J. H. (1986). *A genetic model of adaptive economic behavior* (Working paper). University of Michigan.

Miller, J. H. (1996). The coevolution of automata in the repeated prisoner's dilemma. *Journal of Economic Behavior and Organization*, *29*(1), 87-112.

Miller, J. H., & Page, S. E. (2007). *Complex adaptive systems: An introduction to*

computational models of social life. Princeton University Press.

Mitchell, M. (1996). *An introduction to genetic algorithms*. MIT Press.

Mittal, S., & Deb, K. (2006). Optimal strategies of the iterated prisoner's dilemma problem for multiple conflicting objectives. In *Proceedings of the 2006 ieee symposium on computational intelligence and games* (p. 197-204).

Mookherjee, D., & Sopher, B. (1997). Learning and decision costs in experimental constant sum games. *Games and Economic Behavior*, *19*, 97-132.

Mühlenbein, H. (1991). Darwin's continent cycle theory and its simulation by the prisoner's dilemma. *Complex Systems*, *5*(5), 459-478.

Nix, A. E., & Vose, M. D. (1992). Modeling genetic algorithms with Markov chains. *Annals of Mathematics and Artificial Intelligence*, *5*(1), 79-88.

Noailly, J., Van den Bergh, J. C. J. M., & Withagen, C. A. (2009). Local and global interactions in an evolutionary resource game. *Computational Economics*, *33*(2), 155-173.

Noailly, J., Withagen, C. A., & Van den Bergh, J. C. J. M. (2007). Spatial evolution of social norms in a common-pool resource game. *Environmental and Resource Economics*, *36*(1), 113-141.

Nowak, M. A. (2006). Five rules for the evolution of cooperation. *Science*, *314*, 1560-1563.

Nowak, M. A., & May, R. M. (1992). Evolutionary games and spatial chaos. *Nature*, *359*, 826-829.

Oechssler, J. (2002). Cooperation as a result of learning with aspiration levels. *Journal of Economic Behavior and Organization*, *49*, 405-409.

Ostrom, E. (2000). Collective action and the evolution of social norms. *Journal of Economic Perspectives*, *14*(3), 137-158.

Outkin, A. V. (2003). Cooperation and local interactions in the prisoners' dilemma game. *Journal of Economic Behavior and Organization*, *52*(4), 481-503.

Price, T. C. (1997). Using co-evolutionary programming to simulate strategic behaviour in markets. *Journal of Evolutionary Economics*, *7*(3), 219-254.

Rhode, P., & Stegeman, M. (1996). A comment on "learning, mutation, and long-run equilibria in games". *Econometrica*, *64*(2), 443-449.

Roth, A. E., & Erev, I. (1995). Learning in extensive form games: Experimental data and

simple dynamic models in the intermediate term. *Games and Economic Behavior*, *8*, 164-212.

Rudolph, G. (1994). Convergence analysis of canonical genetic algorithms. *IEEE Transactions on Neural Networks*, *5*(1), 96-101.

Rudolph, G. (1998). Finite Markov chain results in evolutionary computation: A tour d'horizon. *Fundamenta Informaticae*, *35*(1–4), 67-89.

Samuelson, L. (1997). *Evolutionary games and equilibrium selection*. MIT Press.

Sandholm, T. W., & Crites, R. H. (1996). Multiagent reinforcement learning in the iterated prisoner's dilemma. *Biosystems*, *37*, 147-166.

Sarin, R., & Vahid, F. (1999). Payoff assessments without probabilities: A simple dynamic model of choice. *Games and Economic Behavior*, *28*, 294-309.

Sarin, R., & Vahid, F. (2001). Predicting how people play games: A simple dynamic model of choice. *Games and Economic Behavior*, *34*, 104-122.

Schaffer, M. E. (1988). Evolutionarily stable strategies for a finite population and a variable contest size. *Journal of Theoretical Biology*, *132*(4), 469-478.

Schaffer, M. E. (1989). Are profit-maximisers the best survivors? A Darwinian model of economic natural selection. *Journal of Economic Behavior and Organization*, *12*(1), 29-45.

Schelling, T. C. (1969). Models of segregation. *American Economic Review (Papers and Proceedings)*, *59*(2), 488-493.

Schelling, T. C. (1971). Dynamic models of segregation. *Journal of Mathematical Sociology*, *1*(2), 143-186.

Schelling, T. C. (1978). *Micromotives and macrobehavior*. W. W. Norton.

Sellgren, A. (2001). The evolution of insurance markets under adverse selection. *Journal of Evolutionary Economics*, *11*(5), 501-526.

Selten, R., & Apesteguia, J. (2005). Experimentally observed imitation and cooperation in price competition on the circle. *Games and Economic Behavior*, *51*(1), 171-192.

Silva, S. T., & Teixeira, A. A. C. (2009). On the divergence of evolutionary research paths in the past 50 years: A comprehensive bibliometric account. *Journal of Evolutionary Economics*, *19*(5), 605-642.

Stark, O., & Behrens, D. A. (2010). An evolutionary edge of knowing less (or: On the

'curse' of global information). *Journal of Evolutionary Economics*, *20*(1), 77-94.

Stewart, W. J. (1994). *Introduction to the numerical solution of Markov chains*. Princeton University Press.

Stimpson, J. L., & Goodrich, M. A. (2003). Learning to cooperate in a social dilemma: A satisficing approach to bargaining. In *Proceedings of the twentieth international conference on machine learning.*

Sugden, R. (2001). The evolutionary turn in game theory. *Journal of Economic Methodology*, *8*(1), 113-130.

Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction*. MIT Press.

Tesfatsion, L. (2003). *Agent-based computational economics* (ISU Economics Working Paper No. 1). Iowa State University.

Tesfatsion, L. (2006). Agent-based computational economics: A constructive approach to economic theory. In L. Tesfatsion & K. L. Judd (Eds.), *Handbook of computational economics, Volume 2* (p. 831-880). Elsevier.

Tesfatsion, L., & Judd, K. L. (Eds.). (2006). *Handbook of computational economics, Volume 2*. Elsevier.

Thibert-Plante, X., & Charbonneau, P. (2007). Crossover and evolutionary stability in the prisoner's dilemma. *Evolutionary Computation*, *15*(3), 321-344.

Tieman, A. F., Houba, H., & Van der Laan, G. (2000). On the level of cooperative behavior in a local-interaction model. *Journal of Economics*, *71*(1), 1-30.

Tijms, H. C. (1994). *Stochastic models: An algorithmic approach*. John Wiley & Sons.

Tijms, H. C. (2003). *A first course in stochastic models*. John Wiley & Sons.

Unver, M. U. (2001). Backward unraveling over time: The evolution of strategic behavior in the entry level British medical labor markets. *Journal of Economic Dynamics and Control*, *25*(6–7), 1039-1080.

Van Bragt, D. D. B., Van Kemenade, C., & La Poutré, J. A. (2001). The influence of evolutionary selection schemes on the iterated prisoner's dilemma. *Computational Economics*, *17*(2–3), 253-263.

Vega-Redondo, F. (1996). *Evolution, games, and economic behaviour*. Oxford University Press.

Vega-Redondo, F. (1997). The evolution of Walrasian behavior. *Econometrica*, *65*(2),

375-384.

Vose, M. D. (1999). *The simple genetic algorithm: Foundations and theory*. MIT Press.

Vriend, N. J. (2000). An illustration of the essential difference between individual and social learning, and its consequences for computational analyses. *Journal of Economic Dynamics and Control*, *24*(1), 1-19.

Waltman, L., & Kaymak, U. (2007). A theoretical analysis of cooperative behavior in multi-agent Q-learning. In *Proceedings of the 2007 ieee symposium on approximate dynamic programming and reinforcement learning* (p. 84-91).

Waltman, L., & Kaymak, U. (2008). Q-learning agents in a Cournot oligopoly model. *Journal of Economic Dynamics and Control*, *32*(10), 3275-3293.

Waltman, L., & Van Eck, N. J. (2009). Robust evolutionary algorithm design for socio-economic simulation: Some comments. *Computational Economics*, *33*(1), 103-105.

Waltman, L., Van Eck, N. J., Dekker, R., & Kaymak, U. (in press). Economic modeling using evolutionary algorithms: The effect of a binary encoding of strategies. *Journal of Evolutionary Economics*.

Watkins, C. J. C. H. (1989). *Learning from delayed rewards*. Unpublished doctoral dissertation, University of Cambridge, England.

Watkins, C. J. C. H., & Dayan, P. (1992). Q-learning. *Machine Learning*, *8*, 279-292.

Weibull, J. W. (1995). *Evolutionary game theory*. MIT Press.

Wheeler, S., Bean, N., Gaffney, J., & Taylor, P. (2006). A markov analysis of social learning and adaptation. *Journal of Evolutionary Economics*, *16*(3), 299-319.

Wilhite, A. (2006). Economic activity on fixed networks. In L. Tesfatsion & K. L. Judd (Eds.), *Handbook of computational economics, Volume 2* (p. 1014-1045). Elsevier.

Windrum, P., Fagiolo, G., & Moneta, A. (2007). Empirical validation of agent-based models: Alternatives and prospects. *Journal of Artificial Societies and Social Simulation*, *10*(2).

Xu, Y. (2006). The behavior of the exchange rate in the genetic algorithm with agents having long memory. *Journal of Evolutionary Economics*, *16*(3), 279-297.

Yao, X., & Darwen, P. (1994). An experimental study of N-person iterated prisoner's

dilemma games. *Informatica*, *18*(4), 435-450.

Young, H. P. (1993). The evolution of conventions. *Econometrica*, *61*, 57-84.

Young, H. P. (1998). *Individual strategy and social structure: An evolutionary theory of institutions*. Princeton University Press.

Young, H. P. (2004). *Strategic learning and its limits*. Oxford University Press.

Author Index

Curriculum Vitae

Ludo Waltman (1982) obtained his master's degree in Informatics & Economics with honors from Erasmus University Rotterdam in 2005. In the same year, he started his PhD research at the Econometric Institute of the Erasmus School of Economics. Ludo's PhD research is in the area of agent-based computational economics. Results of his research have been presented at various international conferences and have been published in *Computational Economics*, *Journal of Economic Dynamics and Control*, and *Journal of Evolutionary Economics*.

Since 2009, Ludo has been working as a researcher at the Centre for Science and Technology Studies of Leiden University. His research interests have shifted to the field of bibliometrics and scientometrics, where he focuses on bibliometric mapping of science and bibliometric performance indicators. His research has resulted in more than fifteen publications in the main journals of the field. Together with Nees Jan van Eck, Ludo has developed a computer program for bibliometric mapping called VOSviewer (www.vosviewer.com).

For more information, please visit www.ludowaltman.nl.

ERASMUS RESEARCH INSTITUTE OF MANAGEMENT (ERIM)

ERIM PH.D. SERIES
RESEARCH IN MANAGEMENT

ERIM Electronic Series Portal: http://hdl.handle.net/1765/1

Acciaro, M., *Bundling Strategies in Global Supply Chains*, Promotor: Prof.dr. H.E. Haralambides, EPS-2010-197-LIS, ISBN: 978-90-5892-240-3, http://hdl.handle.net/1765/19742

Agatz, N.A.H., *Demand Management in E-Fulfillment*, Promotor: Prof.dr.ir. J.A.E.E. van Nunen, EPS-2009-163-LIS, ISBN: 978-90-5892-200-7, http://hdl.handle.net/1765/15425

Alexiev, A., *Exploratory Innovation: The Role of Organizational and Top Management Team Social Capital*, Promotors: Prof.dr. F.A.J. van den Bosch & Prof.dr. H.W. Volberda, EPS-2010-208-STR, ISBN: 978-90-5892-249-6, http://hdl.handle.net/1765/20632

Althuizen, N.A.P., *Analogical Reasoning as a Decision Support Principle for Weakly Structured Marketing Problems*, Promotor: Prof.dr.ir. B. Wierenga, EPS-2006-095-MKT, ISBN: 90-5892-129-8, http://hdl.handle.net/1765/8190

Alvarez, H.L., *Distributed Collaborative Learning Communities Enabled by Information Communication Technology*, Promotor: Prof.dr. K. Kumar, EPS-2006-080-LIS, ISBN: 90-5892-112-3, http://hdl.handle.net/1765/7830

Appelman, J.H., *Governance of Global Interorganizational Tourism Networks: Changing Forms of Co-ordination between the Travel Agency and Aviation Sector*, Promotors: Prof.dr. F.M. Go & Prof.dr. B. Nooteboom, EPS-2004-036-MKT, ISBN: 90-5892-060-7, http://hdl.handle.net/1765/1199

Asperen, E. van, *Essays on Port, Container, and Bulk Chemical Logistics Optimization*, Promotor: Prof.dr.ir. R. Dekker, EPS-2009-181-LIS, ISBN: 978-90-5892-222-9, http://hdl.handle.net/1765/17626

Assem, M.J. van den, *Deal or No Deal? Decision Making under Risk in a Large-Stake TV Game Show and Related Experiments*, Promotor: Prof.dr. J. Spronk, EPS-2008-138-F&A, ISBN: 978-90-5892-173-4, http://hdl.handle.net/1765/13566

Baquero, G, *On Hedge Fund Performance, Capital Flows and Investor Psychology*, Promotor: Prof.dr.

M.J.C.M. Verbeek, EPS-2006-094-F&A, ISBN: 90-5892-131-X, http://hdl.handle.net/1765/8192

Benning, T.M., *A Consumer Perspective on Flexibility in Health Care: Priority Access Pricing and Customized Care*, Promotor: Prof.dr.ir. B.G.C. Dellaert, EPS-2011-241-MKT, ISBN: 978-90-5892-280-9, http://hdl.handle.net/1765/23670

Berens, G., *Corporate Branding: The Development of Corporate Associations and their Influence on Stakeholder Reactions*, Promotor: Prof.dr. C.B.M. van Riel, EPS-2004-039-ORG, ISBN: 90-5892-065-8, http://hdl.handle.net/1765/1273

Berghe, D.A.F. van den, *Working Across Borders: Multinational Enterprises and the Internationalization of Employment*, Promotors: Prof.dr. R.J.M. van Tulder & Prof.dr. E.J.J. Schenk, EPS-2003-029-ORG, ISBN: 90-5892-05-34, http://hdl.handle.net/1765/1041

Berghman, L.A., *Strategic Innovation Capacity: A Mixed Method Study on Deliberate Strategic Learning Mechanisms*, Promotor: Prof.dr. P. Mattyssens, EPS-2006-087-MKT, ISBN: 90-5892-120-4, http://hdl.handle.net/1765/7991

Bezemer, P.J., *Diffusion of Corporate Governance Beliefs: Board Independence and the Emergence of a Shareholder Value Orientation in the Netherlands*, Promotors: Prof.dr.ing. F.A.J. van den Bosch & Prof.dr. H.W. Volberda, EPS-2009-192-STR, ISBN: 978-90-5892-232-8, http://hdl.handle.net/1765/18458

Bijman, W.J.J., *Essays on Agricultural Co-operatives: Governance Structure in Fruit and Vegetable Chains*, Promotor: Prof.dr. G.W.J. Hendrikse, EPS-2002-015-ORG, ISBN: 90-5892-024-0, http://hdl.handle.net/1765/867

Binken, J.L.G., *System Markets: Indirect Network Effects in Action, or Inaction?*, Promotor: Prof.dr. S. Stremersch, EPS-2010-213-MKT, ISBN: 978-90-5892-260-1, http://hdl.handle.net/1765/21186

Blitz, D.C., *Benchmarking Benchmarks*, Promotors: Prof.dr. A.G.Z. Kemna & Prof.dr. W.F.C. Verschoor, EPS-2011-225-F&A, ISBN: 978-90-5892-268-7, http://hdl.handle.net/1765/22624

Bispo, A., *Labour Market Segmentation: An investigation into the Dutch hospitality industry*, Promotors: Prof.dr. G.H.M. Evers & Prof.dr. A.R. Thurik, EPS-2007-108-ORG, ISBN: 90-5892-136-9, http://hdl.handle.net/1765/10283

Blindenbach-Driessen, F., *Innovation Management in Project-Based Firms*, Promotor: Prof.dr. S.L. van de Velde, EPS-2006-082-LIS, ISBN: 90-5892-110-7, http://hdl.handle.net/1765/7828

Boer, C.A., *Distributed Simulation in Industry*, Promotors: Prof.dr. A. de Bruin & Prof.dr.ir. A. Verbraeck, EPS-2005-065-LIS, ISBN: 90-5892-093-3, http://hdl.handle.net/1765/6925

Boer, N.I., *Knowledge Sharing within Organizations: A situated and Relational Perspective*, Promotor: Prof.dr. K. Kumar, EPS-2005-060-LIS, ISBN: 90-5892-086-0, http://hdl.handle.net/1765/6770

Boer-Sorbán, K., *Agent-Based Simulation of Financial Markets: A modular, Continuous-Time Approach*, Promotor: Prof.dr. A. de Bruin, EPS-2008-119-LIS, ISBN: 90-5892-155-0, http://hdl.handle.net/1765/10870

Boon, C.T., *HRM and Fit: Survival of the Fittest!?*, Promotors: Prof.dr. J. Paauwe & Prof.dr. D.N. den Hartog, EPS-2008-129-ORG, ISBN: 978-90-5892-162-8, http://hdl.handle.net/1765/12606

Borst, W.A.M., *Understanding Crowdsourcing: Effects of Motivation and Rewards on Participation and Performance in Voluntary Online Activities*, Promotors: Prof.dr.ir. J.C.M. van den Ende & Prof.dr.ir. H.W.G.M. van Heck, EPS-2010-221-LIS, ISBN: 978-90-5892-262-5, http://hdl.handle.net/1765/ 21914

Braun, E., *City Marketing: Towards an Integrated Approach*, Promotor: Prof.dr. L. van den Berg, EPS-2008-142-MKT, ISBN: 978-90-5892-180-2, http://hdl.handle.net/1765/13694

Brito, M.P. de, *Managing Reverse Logistics or Reversing Logistics Management?*, Promotors: Prof.dr.ir. R. Dekker & Prof.dr. M. B. M. de Koster, EPS-2004-035-LIS, ISBN: 90-5892-058-5, http://hdl.handle.net/1765/1132

Brohm, R., *Polycentric Order in Organizations: A Dialogue between Michael Polanyi and IT-Consultants on Knowledge, Morality, and Organization*, Promotors: Prof.dr. G. W. J. Hendrikse & Prof.dr. H. K. Letiche, EPS-2005-063-ORG, ISBN: 90-5892-095-X, http://hdl.handle.net/1765/6911

Brumme, W.-H., *Manufacturing Capability Switching in the High-Tech Electronics Technology Life Cycle*, Promotors: Prof.dr.ir. J.A.E.E. van Nunen & Prof.dr.ir. L.N. Van Wassenhove, EPS-2008-126-LIS, ISBN: 978-90-5892-150-5, http://hdl.handle.net/1765/12103

Budiono, D.P., *The Analysis of Mutual Fund Performance: Evidence from U.S. Equity Mutual Funds*, Promotor: Prof.dr. M.J.C.M. Verbeek, EPS-2010-185-F&A, ISBN: 978-90-5892-224-3, http://hdl.handle.net/1765/18126

Burgers, J.H., *Managing Corporate Venturing: Multilevel Studies on Project Autonomy, Integration, Knowledge Relatedness, and Phases in the New Business Development Process*, Promotors: Prof.dr.ir.

F.A.J. Van den Bosch & Prof.dr. H.W. Volberda, EPS-2008-136-STR, ISBN: 978-90-5892-174-1, http://hdl.handle.net/1765/13484

Camacho, N.M., *Health and Marketing; Essays on Physician and Patient Decision-making*, Promotor: Prof.dr. S. Stremersch, EPS-2011-237-MKT, ISBN: 978-90-5892-284-7, http://hdl.handle.net/1765/23604

Campbell, R.A.J., *Rethinking Risk in International Financial Markets*, Promotor: Prof.dr. C.G. Koedijk, EPS-2001-005-F&A, ISBN: 90-5892-008-9, http://hdl.handle.net/1765/306

Carvalho de Mesquita Ferreira, L., *Attention Mosaics: Studies of Organizational Attention*, Promotors: Prof.dr. P.M.A.R. Heugens & Prof.dr. J. van Oosterhout, EPS-2010-205-ORG, ISBN: 978-90-5892-242-7, http://hdl.handle.net/1765/19882

Chen, C.-M., *Evaluation and Design of Supply Chain Operations Using DEA*, Promotor: Prof.dr. J.A.E.E. van Nunen, EPS-2009-172-LIS, ISBN: 978-90-5892-209-0, http://hdl.handle.net/1765/16181

Chen, H., *Individual Mobile Communication Services and Tariffs*, Promotor: Prof.dr. L.F.J.M. Pau, EPS-2008-123-LIS, ISBN: 90-5892-158-1, http://hdl.handle.net/1765/11141

Chen, Y., *Labour Flexibility in China's Companies: An Empirical Study*, Promotors: Prof.dr. A. Buitendam & Prof.dr. B. Krug, EPS-2001-006-ORG, ISBN: 90-5892-012-7, http://hdl.handle.net/1765/307

Damen, F.J.A., *Taking the Lead: The Role of Affect in Leadership Effectiveness*, Promotor: Prof.dr. D.L. van Knippenberg, EPS-2007-107-ORG, http://hdl.handle.net/1765/10282

Daniševská, P., *Empirical Studies on Financial Intermediation and Corporate Policies*, Promotor: Prof.dr. C.G. Koedijk, EPS-2004-044-F&A, ISBN: 90-5892-070-4, http://hdl.handle.net/1765/1518

Defilippi Angeldonis, E.F., *Access Regulation for Naturally Monopolistic Port Terminals: Lessons from Regulated Network Industries*, Promotor: Prof.dr. H.E. Haralambides, EPS-2010-204-LIS, ISBN: 978-90-5892-245-8, http://hdl.handle.net/1765/19881

Delporte-Vermeiren, D.J.E., *Improving the Flexibility and Profitability of ICT-enabled Business Networks: An Assessment Method and Tool*, Promotors: Prof. mr. dr. P.H.M. Vervest & Prof.dr.ir. H.W.G.M. van Heck, EPS-2003-020-LIS, ISBN: 90-5892-040-2, http://hdl.handle.net/1765/359

Derwall, J.M.M., *The Economic Virtues of SRI and CSR*, Promotor: Prof.dr. C.G. Koedijk,

EPS-2007-101-F&A, ISBN: 90-5892-132-8, http://hdl.handle.net/1765/8986

Desmet, P.T.M., *In Money we Trust? Trust Repair and the Psychology of Financial Compensations*, Promotors: Prof.dr. D. De Cremer & Prof.dr. E. van Dijk, EPS-2011-232-ORG, ISBN: 978-90-5892-274-8, http://hdl.handle.net/1765/23268

Diepen, M. van, *Dynamics and Competition in Charitable Giving*, Promotor: Prof.dr. Ph.H.B.F. Franses, EPS-2009-159-MKT, ISBN: 978-90-5892-188-8, http://hdl.handle.net/1765/14526

Dietvorst, R.C., *Neural Mechanisms Underlying Social Intelligence and Their Relationship with the Performance of Sales Managers*, Promotor: Prof.dr. W.J.M.I. Verbeke, EPS-2010-215-MKT, ISBN: 978-90-5892-257-1, http://hdl.handle.net/1765/21188

Dietz, H.M.S., *Managing (Sales)People towards Perfomance: HR Strategy, Leadership & Teamwork*, Promotor: Prof.dr. G.W.J. Hendrikse, EPS-2009-168-ORG, ISBN: 978-90-5892-210-6, http://hdl.handle.net/1765/16081

Dijksterhuis, M., *Organizational Dynamics of Cognition and Action in the Changing Dutch and US Banking Industries*, Promotors: Prof.dr.ir. F.A.J. van den Bosch & Prof.dr. H.W. Volberda, EPS-2003-026-STR, ISBN: 90-5892-048-8, http://hdl.handle.net/1765/1037

Eijk, A.R. van der, *Behind Networks: Knowledge Transfer, Favor Exchange and Performance*, Promotors: Prof.dr. S.L. van de Velde & Prof.dr.drs. W.A. Dolfsma, EPS-2009-161-LIS, ISBN: 978-90-5892-190-1, http://hdl.handle.net/1765/14613

Elstak, M.N., *Flipping the Identity Coin: The Comparative Effect of Perceived, Projected and Desired Organizational Identity on Organizational Identification and Desired Behavior*, Promotor: Prof.dr. C.B.M. van Riel, EPS-2008-117-ORG, ISBN: 90-5892-148-2, http://hdl.handle.net/1765/10723

Erken, H.P.G., *Productivity, R&D and Entrepreneurship*, Promotor: Prof.dr. A.R. Thurik, EPS-2008-147-ORG, ISBN: 978-90-5892-179-6, http://hdl.handle.net/1765/14004

Essen, M. van, *An Institution-Based View of Ownership*, Promotos: Prof.dr. J. van Oosterhout & Prof.dr. G.M.H. Mertens, EPS-2011-226-ORG, ISBN: 978-90-5892-269-4, http://hdl.handle.net/1765/22643

Fenema, P.C. van, *Coordination and Control of Globally Distributed Software Projects*, Promotor: Prof.dr. K. Kumar, EPS-2002-019-LIS, ISBN: 90-5892-030-5, http://hdl.handle.net/1765/360

Feng, L., *Motivation, Coordination and Cognition in Cooperatives*, Promotor: Prof.dr. G.W.J.

Hendrikse, EPS-2010-220-ORG, ISBN: 90-5892-261-8, http://hdl.handle.net/1765/21680

Fleischmann, M., *Quantitative Models for Reverse Logistics*, Promotors: Prof.dr.ir. J.A.E.E. van Nunen & Prof.dr.ir. R. Dekker, EPS-2000-002-LIS, ISBN: 35-4041-711-7, http://hdl.handle.net/1765/1044

Flier, B., *Strategic Renewal of European Financial Incumbents: Coevolution of Environmental Selection, Institutional Effects, and Managerial Intentionality*, Promotors: Prof.dr.ir. F.A.J. van den Bosch & Prof.dr. H.W. Volberda, EPS-2003-033-STR, ISBN: 90-5892-055-0, http://hdl.handle.net/1765/1071

Fok, D., *Advanced Econometric Marketing Models*, Promotor: Prof.dr. Ph.H.B.F. Franses, EPS-2003-027-MKT, ISBN: 90-5892-049-6, http://hdl.handle.net/1765/1035

Ganzaroli, A., *Creating Trust between Local and Global Systems*, Promotors: Prof.dr. K. Kumar & Prof.dr. R.M. Lee, EPS-2002-018-LIS, ISBN: 90-5892-031-3, http://hdl.handle.net/1765/361

Gertsen, H.F.M., *Riding a Tiger without Being Eaten: How Companies and Analysts Tame Financial Restatements and Influence Corporate Reputation*, Promotor: Prof.dr. C.B.M. van Riel, EPS-2009-171-ORG, ISBN: 90-5892-214-4, http://hdl.handle.net/1765/16098

Gilsing, V.A., *Exploration, Exploitation and Co-evolution in Innovation Networks*, Promotors: Prof.dr. B. Nooteboom & Prof.dr. J.P.M. Groenewegen, EPS-2003-032-ORG, ISBN: 90-5892-054-2, http://hdl.handle.net/1765/1040

Gijsbers, G.W., *Agricultural Innovation in Asia: Drivers, Paradigms and Performance*, Promotor: Prof.dr. R.J.M. van Tulder, EPS-2009-156-ORG, ISBN: 978-90-5892-191-8, http://hdl.handle.net/1765/14524

Gong, Y., *Stochastic Modelling and Analysis of Warehouse Operations*, Promotors: Prof.dr. M.B.M. de Koster & Prof.dr. S.L. van de Velde, EPS-2009-180-LIS, ISBN: 978-90-5892-219-9, http://hdl.handle.net/1765/16724

Govers, R., *Virtual Tourism Destination Image: Glocal Identities Constructed, Perceived and Experienced*, Promotors: Prof.dr. F.M. Go & Prof.dr. K. Kumar, EPS-2005-069-MKT, ISBN: 90-5892-107-7, http://hdl.handle.net/1765/6981

Graaf, G. de, *Tractable Morality: Customer Discourses of Bankers, Veterinarians and Charity Workers*, Promotors: Prof.dr. F. Leijnse & Prof.dr. T. van Willigenburg, EPS-2003-031-ORG, ISBN: 90-5892-051-8, http://hdl.handle.net/1765/1038

Greeven, M.J., *Innovation in an Uncertain Institutional Environment: Private Software Entrepreneurs in Hangzhou, China*, Promotor: Prof.dr. B. Krug, EPS-2009-164-ORG, ISBN: 978-90-5892-202-1, http://hdl.handle.net/1765/15426

Groot, E.A. de, *Essays on Economic Cycles*, Promotors: Prof.dr. Ph.H.B.F. Franses & Prof.dr. H.R. Commandeur, EPS-2006-091-MKT, ISBN: 90-5892-123-9, http://hdl.handle.net/1765/8216

Guenster, N.K., *Investment Strategies Based on Social Responsibility and Bubbles*, Promotor: Prof.dr. C.G. Koedijk, EPS-2008-175-F&A, ISBN: 978-90-5892-206-9, http://hdl.handle.net/1765/1

Gutkowska, A.B., *Essays on the Dynamic Portfolio Choice*, Promotor: Prof.dr. A.C.F. Vorst, EPS-2006-085-F&A, ISBN: 90-5892-118-2, http://hdl.handle.net/1765/7994

Hagemeijer, R.E., *The Unmasking of the Other*, Promotors: Prof.dr. S.J. Magala & Prof.dr. H.K. Letiche, EPS-2005-068-ORG, ISBN: 90-5892-097-6, http://hdl.handle.net/1765/6963

Hakimi, N.A, *Leader Empowering Behaviour: The Leader's Perspective: Understanding the Motivation behind Leader Empowering Behaviour*, Promotor: Prof.dr. D.L. van Knippenberg, EPS-2010-184-ORG, http://hdl.handle.net/1765/17701

Halderen, M.D. van, *Organizational Identity Expressiveness and Perception Management: Principles for Expressing the Organizational Identity in Order to Manage the Perceptions and Behavioral Reactions of External Stakeholders*, Promotor: Prof.dr. S.B.M. van Riel, EPS-2008-122-ORG, ISBN: 90-5892-153-6, http://hdl.handle.net/1765/10872

Hartigh, E. den, *Increasing Returns and Firm Performance: An Empirical Study*, Promotor: Prof.dr. H.R. Commandeur, EPS-2005-067-STR, ISBN: 90-5892-098-4, http://hdl.handle.net/1765/6939

Hensmans, M., *A Republican Settlement Theory of the Firm: Applied to Retail Banks in England and the Netherlands (1830-2007)*, Promotors: Prof.dr. A. Jolink& Prof.dr. S.J. Magala, EPS-2010-193-ORG, ISBN 90-5892-235-9, http://hdl.handle.net/1765/19494

Hermans. J.M., *ICT in Information Services; Use and Deployment of the Dutch Securities Trade, 1860-1970*, Promotor: Prof.dr. drs. F.H.A. Janszen, EPS-2004-046-ORG, ISBN 90-5892-072-0, http://hdl.handle.net/1765/1793

Hernandez Mireles, C., *Marketing Modeling for New Products*, Promotor: Prof.dr. P.H. Franses, EPS-2010-202-MKT, ISBN 90-5892-237-3, http://hdl.handle.net/1765/19878

Hessels, S.J.A., *International Entrepreneurship: Value Creation Across National Borders*, Promotor: Prof.dr. A.R. Thurik, EPS-2008-144-ORG, ISBN: 978-90-5892-181-9, http://hdl.handle.net/1765/13942

Heugens, P.P.M.A.R., *Strategic Issues Management: Implications for Corporate Performance*, Promotors: Prof.dr.ir. F.A.J. van den Bosch & Prof.dr. C.B.M. van Riel, EPS-2001-007-STR, ISBN: 90-5892-009-9, http://hdl.handle.net/1765/358

Heuvel, W. van den, *The Economic Lot-Sizing Problem: New Results and Extensions*, Promotor: Prof.dr. A.P.L. Wagelmans, EPS-2006-093-LIS, ISBN: 90-5892-124-7, http://hdl.handle.net/1765/1805

Hoedemaekers, C.M.W., *Performance, Pinned down: A Lacanian Analysis of Subjectivity at Work*, Promotors: Prof.dr. S. Magala & Prof.dr. D.H. den Hartog, EPS-2008-121-ORG, ISBN: 90-5892-156-7, http://hdl.handle.net/1765/10871

Hooghiemstra, R., *The Construction of Reality: Cultural Differences in Self-serving Behaviour in Accounting Narratives*, Promotors: Prof.dr. L.G. van der Tas RA & Prof.dr. A.Th.H. Pruyn, EPS-2003-025-F&A, ISBN: 90-5892-047-X, http://hdl.handle.net/1765/871

Hu, Y., *Essays on the Governance of Agricultural Products: Cooperatives and Contract Farming*, Promotors: Prof.dr. G.W.J. Hendrkse & Prof.dr. B. Krug, EPS-2007-113-ORG, ISBN: 90-5892-145-1, http://hdl.handle.net/1765/10535

Huang, X., *An Analysis of Occupational Pension Provision: From Evaluation to Redesigh*, Promotors: Prof.dr. M.J.C.M. Verbeek & Prof.dr. R.J. Mahieu, EPS-2010-196-F&A, ISBN: 90-5892-239-7, http://hdl.handle.net/1765/19674

Huij, J.J., *New Insights into Mutual Funds: Performance and Family Strategies*, Promotor: Prof.dr. M.C.J.M. Verbeek, EPS-2007-099-F&A, ISBN: 90-5892-134-4, http://hdl.handle.net/1765/9398

Huurman, C.I., *Dealing with Electricity Prices*, Promotor: Prof.dr. C.D. Koedijk, EPS-2007-098-F&A, ISBN: 90-5892-130-1, http://hdl.handle.net/1765/9399

Iastrebova, K, *Manager's Information Overload: The Impact of Coping Strategies on Decision-Making Performance*, Promotor: Prof.dr. H.G. van Dissel, EPS-2006-077-LIS, ISBN: 90-5892-111-5, http://hdl.handle.net/1765/7329

Iwaarden, J.D. van, *Changing Quality Controls: The Effects of Increasing Product Variety and Shortening Product Life Cycles,* Promotors: Prof.dr. B.G. Dale & Prof.dr. A.R.T. Williams, EPS-2006-084-ORG, ISBN: 90-5892-117-4, http://hdl.handle.net/1765/7992

Jalil, M.N., *Customer Information Driven After Sales Service Management: Lessons from Spare Parts Logistics*, Promotor: Prof.dr. L.G. Kroon, EPS-2011-222-LIS, ISBN: 978-90-5892-264-9, http://hdl.handle.net/1765/22156

Jansen, J.J.P., *Ambidextrous Organizations*, Promotors: Prof.dr.ir. F.A.J. Van den Bosch & Prof.dr. H.W. Volberda, EPS-2005-055-STR, ISBN: 90-5892-081-X, http://hdl.handle.net/1765/6774

Jaspers, F.P.H., *Organizing Systemic Innovation*, Promotor: Prof.dr.ir. J.C.M. van den Ende, EPS-2009-160-ORG, ISBN: 978-90-5892-197-), http://hdl.handle.net/1765/14974

Jennen, M.G.J., *Empirical Essays on Office Market Dynamics*, Promotors: Prof.dr. C.G. Koedijk & Prof.dr. D. Brounen, EPS-2008-140-F&A, ISBN: 978-90-5892-176-5, http://hdl.handle.net/1765/13692

Jiang, T., *Capital Structure Determinants and Governance Structure Variety in Franchising*, Promotors: Prof.dr. G. Hendrikse & Prof.dr. A. de Jong, EPS-2009-158-F&A, ISBN: 978-90-5892-199-4, http://hdl.handle.net/1765/14975

Jiao, T., *Essays in Financial Accounting*, Promotor: Prof.dr. G.M.H. Mertens, EPS-2009-176-F&A, ISBN: 978-90-5892-211-3, http://hdl.handle.net/1765/16097

Jong, C. de, *Dealing with Derivatives: Studies on the Role, Informational Content and Pricing of Financial Derivatives*, Promotor: Prof.dr. C.G. Koedijk, EPS-2003-023-F&A, ISBN: 90-5892-043-7, http://hdl.handle.net/1765/1043

Kaa, G. van, *Standard Battles for Complex Systems: Empirical Research on the Home Network*, Promotors: Prof.dr.ir. J. van den Ende & Prof.dr.ir. H.W.G.M. van Heck, EPS-2009-166-ORG, ISBN: 978-90-5892-205-2, http://hdl.handle.net/1765/16011

Kagie, M., *Advances in Online Shopping Interfaces: Product Catalog Maps and Recommender Systems,* Promotor: Prof.dr. P.J.F. Groenen, EPS-2010-195-MKT, ISBN: 978-90-5892-233-5, http://hdl.handle.net/1765/19532

Kappe, E.R., *The Effectiveness of Pharmaceutical Marketing*, Promotor: Prof.dr. S. Stremersch, EPS-2011-239-MKT, ISBN: 978-90-5892-283-0, http://hdl.handle.net/1765/23610

Karreman, B., *Financial Services and Emerging Markets*, Promotors: Prof.dr. G.A. van der Knaap & Prof.dr. H.P.G. Pennings, EPS-2011-223-ORG, ISBN: 978-90-5892-266-3, http://hdl.handle.net/1765/22280

Keizer, A.B., *The Changing Logic of Japanese Employment Practices: A Firm-Level Analysis of Four Industries*, Promotors: Prof.dr. J.A. Stam & Prof.dr. J.P.M. Groenewegen, EPS-2005-057-ORG, ISBN: 90-5892-087-9, http://hdl.handle.net/1765/6667

Kijkuit, R.C., *Social Networks in the Front End: The Organizational Life of an Idea,* Promotor: Prof.dr. B. Nooteboom, EPS-2007-104-ORG, ISBN: 90-5892-137-6, http://hdl.handle.net/1765/10074

Kippers, J., *Empirical Studies on Cash Payments,* Promotor: Prof.dr. Ph.H.B.F. Franses, EPS-2004-043-F&A, ISBN: 90-5892-069-0, http://hdl.handle.net/1765/1520

Klein, M.H., *Poverty Alleviation through Sustainable Strategic Business Models: Essays on Poverty Alleviation as a Business Strategy,* Promotor: Prof.dr. H.R. Commandeur, EPS-2008-135-STR, ISBN: 978-90-5892-168-0, http://hdl.handle.net/1765/13482

Knapp, S., *The Econometrics of Maritime Safety: Recommendations to Enhance Safety at Sea*, Promotor: Prof.dr. Ph.H.B.F. Franses, EPS-2007-096-ORG, ISBN: 90-5892-127-1, http://hdl.handle.net/1765/7913

Kole, E., *On Crises, Crashes and Comovements*, Promotors: Prof.dr. C.G. Koedijk & Prof.dr. M.J.C.M. Verbeek, EPS-2006-083-F&A, ISBN: 90-5892-114-X, http://hdl.handle.net/1765/7829

Kooij-de Bode, J.M., *Distributed Information and Group Decision-Making: Effects of Diversity and Affect*, Promotor: Prof.dr. D.L. van Knippenberg, EPS-2007-115-ORG, http://hdl.handle.net/1765/10722

Koppius, O.R., *Information Architecture and Electronic Market Performance*, Promotors: Prof.dr. P.H.M. Vervest & Prof.dr.ir. H.W.G.M. van Heck, EPS-2002-013-LIS, ISBN: 90-5892-023-2, http://hdl.handle.net/1765/921

Kotlarsky, J., *Management of Globally Distributed Component-Based Software Development Projects*, Promotor: Prof.dr. K. Kumar, EPS-2005-059-LIS, ISBN: 90-5892-088-7, http://hdl.handle.net/1765/6772

Krauth, E.I., *Real-Time Planning Support: A Task-Technology Fit Perspective*, Promotors: Prof.dr. S.L. van de Velde & Prof.dr. J. van Hillegersberg, EPS-2008-155-LIS, ISBN: 978-90-5892-193-2, http://hdl.handle.net/1765/14521

Kuilman, J., *The Re-Emergence of Foreign Banks in Shanghai: An Ecological Analysis*, Promotor: Prof.dr. B. Krug, EPS-2005-066-ORG, ISBN: 90-5892-096-8, http://hdl.handle.net/1765/6926

Kwee, Z., *Investigating Three Key Principles of Sustained Strategic Renewal: A Longitudinal Study of Long-Lived Firms*, Promotors: Prof.dr.ir. F.A.J. Van den Bosch & Prof.dr. H.W. Volberda, EPS-2009-174-STR, ISBN: 90-5892-212-0, http://hdl.handle.net/1765/16207

Lam, K.Y., *Reliability and Rankings*, Promotor: Prof.dr. P.H.B.F. Franses, EPS-2011-230-MKT, ISBN: 978-90-5892-272-4, http://hdl.handle.net/1765/22977

Langen, P.W. de, *The Performance of Seaport Clusters: A Framework to Analyze Cluster Performance and an Application to the Seaport Clusters of Durban, Rotterdam and the Lower Mississippi*, Promotors: Prof.dr. B. Nooteboom & Prof. drs. H.W.H. Welters, EPS-2004-034-LIS, ISBN: 90-5892-056-9, http://hdl.handle.net/1765/1133

Langhe, B. de, *Contingencies: Learning Numerical and Emotional Associations in an Uncertain World*, Promotors: Prof.dr.ir. B. Wierenga & Prof.dr. S.M.J. van Osselaer, EPS-2011-236-MKT, ISBN: 978-90-5892-278-6, http://hdl.handle.net/1765/1

Larco Martinelli, J.A., *Incorporating Worker-Specific Factors in Operations Management Models*, Promotors: Prof.dr.ir. J. Dul & Prof.dr. M.B.M. de Koster, EPS-2010-217-LIS, ISBN: 90-5892-263-2, http://hdl.handle.net/1765/21527

Le Anh, T., *Intelligent Control of Vehicle-Based Internal Transport Systems*, Promotors: Prof.dr. M.B.M. de Koster & Prof.dr.ir. R. Dekker, EPS-2005-051-LIS, ISBN: 90-5892-079-8, http://hdl.handle.net/1765/6554

Le-Duc, T., *Design and Control of Efficient Order Picking Processes*, Promotor: Prof.dr. M.B.M. de Koster, EPS-2005-064-LIS, ISBN: 90-5892-094-1, http://hdl.handle.net/1765/6910

Leeuwen, E.P. van, *Recovered-Resource Dependent Industries and the Strategic Renewal of Incumbent Firm: A Multi-Level Study of Recovered Resource Dependence Management and Strategic Renewal in the European Paper and Board Industry*, Promotors: Prof.dr.ir. F.A.J. Van den Bosch & Prof.dr. H.W. Volberda, EPS-2007-109-STR, ISBN: 90-5892-140-6, http://hdl.handle.net/1765/10183

Lentink, R.M., *Algorithmic Decision Support for Shunt Planning*, Promotors: Prof.dr. L.G. Kroon & Prof.dr.ir. J.A.E.E. van Nunen, EPS-2006-073-LIS, ISBN: 90-5892-104-2, http://hdl.handle.net/1765/7328

Li, T., *Informedness and Customer-Centric Revenue Management*, Promotors: Prof.dr. P.H.M. Vervest & Prof.dr.ir. H.W.G.M. van Heck, EPS-2009-146-LIS, ISBN: 978-90-5892-195-6,

http://hdl.handle.net/1765/14525

Liang, G., *New Competition: Foreign Direct Investment and Industrial Development in China*, Promotor: Prof.dr. R.J.M. van Tulder, EPS-2004-047-ORG, ISBN: 90-5892-073-9, http://hdl.handle.net/1765/1795

Liere, D.W. van, *Network Horizon and the Dynamics of Network Positions: A Multi-Method Multi-Level Longitudinal Study of Interfirm Networks*, Promotor: Prof.dr. P.H.M. Vervest, EPS-2007-105-LIS, ISBN: 90-5892-139-0, http://hdl.handle.net/1765/10181

Loef, J., *Incongruity between Ads and Consumer Expectations of Advertising*, Promotors: Prof.dr. W.F. van Raaij & Prof.dr. G. Antonides, EPS-2002-017-MKT, ISBN: 90-5892-028-3, http://hdl.handle.net/1765/869

Londoño, M. del Pilar, *Institutional Arrangements that Affect Free Trade Agreements: Economic Rationality Versus Interest Groups*, Promotors: Prof.dr. H.E. Haralambides & Prof.dr. J.F. Francois, EPS-2006-078-LIS, ISBN: 90-5892-108-5, http://hdl.handle.net/1765/7578

Lovric, M., *Behavioral Finance and Agent-Based Artificial Markets*, Promotors: Prof.dr. J. Spronk & Prof.dr.ir. U. Kaymak, EPS-2011-229-F&A, ISBN: 978-90-5892-258-8, http://hdl.handle.net/1765/22814

Maas, A.A., van der, *Strategy Implementation in a Small Island Context: An Integrative Framework*, Promotor: Prof.dr. H.G. van Dissel, EPS-2008-127-LIS, ISBN: 978-90-5892-160-4, http://hdl.handle.net/1765/12278

Maas, K.E.G., *Corporate Social Performance: From Output Measurement to Impact Measurement*, Promotor: Prof.dr. H.R. Commandeur, EPS-2009-182-STR, ISBN: 978-90-5892-225-0, http://hdl.handle.net/1765/17627

Maeseneire, W., de, *Essays on Firm Valuation and Value Appropriation*, Promotor: Prof.dr. J.T.J. Smit, EPS-2005-053-F&A, ISBN: 90-5892-082-8, http://hdl.handle.net/1765/6768

Mandele, L.M., van der, *Leadership and the Inflection Point: A Longitudinal Perspective*, Promotors: Prof.dr. H.W. Volberda & Prof.dr. H.R. Commandeur, EPS-2004-042-STR, ISBN: 90-5892-067-4, http://hdl.handle.net/1765/1302

Markwat, T.D., *Extreme Dependence in Asset Markets Around the Globe*, Promotor: Prof.dr. D.J.C. van Dijk, EPS-2011-227-F&A, ISBN: 978-90-5892-270-0, http://hdl.handle.net/1765/22744

Meer, J.R. van der, *Operational Control of Internal Transport*, Promotors: Prof.dr. M.B.M. de Koster & Prof.dr.ir. R. Dekker, EPS-2000-001-LIS, ISBN: 90-5892-004-6, http://hdl.handle.net/1765/859

Mentink, A., *Essays on Corporate Bonds*, Promotor: Prof.dr. A.C.F. Vorst, EPS-2005-070-F&A, ISBN: 90-5892-100-X, http://hdl.handle.net/1765/7121

Meuer, J., *Configurations of Inter-Firm Relations in Management Innovation: A Study in China's Biopharmaceutical Industry*, Promotor: Prof.dr. B. Krug, EPS-2011-228-ORG, ISBN: 978-90-5892-271-1, http://hdl.handle.net/1765/22745

Meyer, R.J.H., *Mapping the Mind of the Strategist: A Quantitative Methodology for Measuring the Strategic Beliefs of Executives*, Promotor: Prof.dr. R.J.M. van Tulder, EPS-2007-106-ORG, ISBN: 978-90-5892-141-3, http://hdl.handle.net/1765/10182

Miltenburg, P.R., *Effects of Modular Sourcing on Manufacturing Flexibility in the Automotive Industry: A Study among German OEMs*, Promotors: Prof.dr. J. Paauwe & Prof.dr. H.R. Commandeur, EPS-2003-030-ORG, ISBN: 90-5892-052-6, http://hdl.handle.net/1765/1039

Moerman, G.A., *Empirical Studies on Asset Pricing and Banking in the Euro Area*, Promotor: Prof.dr. C.G. Koedijk, EPS-2005-058-F&A, ISBN: 90-5892-090-9, http://hdl.handle.net/1765/6666

Moitra, D., *Globalization of R&D: Leveraging Offshoring for Innovative Capability and Organizational Flexibility*, Promotor: Prof.dr. K. Kumar, EPS-2008-150-LIS, ISBN: 978-90-5892-184-0, http://hdl.handle.net/1765/14081

Mol, M.M., *Outsourcing, Supplier-relations and Internationalisation: Global Source Strategy as a Chinese Puzzle*, Promotor: Prof.dr. R.J.M. van Tulder, EPS-2001-010-ORG, ISBN: 90-5892-014-3, http://hdl.handle.net/1765/355

Mom, T.J.M., *Managers' Exploration and Exploitation Activities: The Influence of Organizational Factors and Knowledge Inflows*, Promotors: Prof.dr.ir. F.A.J. Van den Bosch & Prof.dr. H.W. Volberda, EPS-2006-079-STR, ISBN: 90-5892-116-6, http://hdl.handle.net/1765

Moonen, J.M., *Multi-Agent Systems for Transportation Planning and Coordination*, Promotors: Prof.dr. J. van Hillegersberg & Prof.dr. S.L. van de Velde, EPS-2009-177-LIS, ISBN: 978-90-5892-216-8, http://hdl.handle.net/1765/16208

Mulder, A., *Government Dilemmas in the Private Provision of Public Goods*, Promotor: Prof.dr. R.J.M. van Tulder, EPS-2004-045-ORG, ISBN: 90-5892-071-2, http://hdl.handle.net/1765/1790

Muller, A.R., *The Rise of Regionalism: Core Company Strategies Under The Second Wave of Integration*, Promotor: Prof.dr. R.J.M. van Tulder, EPS-2004-038-ORG, ISBN: 90-5892-062-3, http://hdl.handle.net/1765/1272

Nalbantov G.I., *Essays on Some Recent Penalization Methods with Applications in Finance and Marketing*, Promotor: Prof. dr P.J.F. Groenen, EPS-2008-132-F&A, ISBN: 978-90-5892-166-6, http://hdl.handle.net/1765/13319

Nederveen Pieterse, A., *Goal Orientation in Teams: The Role of Diversity*, Promotor: Prof.dr. D.L. van Knippenberg, EPS-2009-162-ORG, http://hdl.handle.net/1765/15240

Nguyen, T.T., *Capital Structure, Strategic Competition, and Governance*, Promotor: Prof.dr. A. de Jong, EPS-2008-148-F&A, ISBN: 90-5892-178-9, http://hdl.handle.net/1765/14005

Nielsen, L.K., *Rolling Stock Rescheduling in Passenger Railways: Applications in Short-term Planning and in Disruption Management*, Promotor: Prof.dr. L.G. Kroon, EPS-2011-224-LIS, ISBN: 978-90-5892-267-0, http://hdl.handle.net/1765/22444

Niesten, E.M.M.I., *Regulation, Governance and Adaptation: Governance Transformations in the Dutch and French Liberalizing Electricity Industries*, Promotors: Prof.dr. A. Jolink & Prof.dr. J.P.M. Groenewegen, EPS-2009-170-ORG, ISBN: 978-90-5892-208-3, http://hdl.handle.net/1765/16096

Nieuwenboer, N.A. den, *Seeing the Shadow of the Self*, Promotor: Prof.dr. S.P. Kaptein, EPS-2008-151-ORG, ISBN: 978-90-5892-182-6, http://hdl.handle.net/1765/14223

Nijdam, M.H., *Leader Firms: The Value of Companies for the Competitiveness of the Rotterdam Seaport Cluster*, Promotor: Prof.dr. R.J.M. van Tulder, EPS-2010-216-ORG, ISBN: 978-90-5892-256-4, http://hdl.handle.net/1765/21405

Ning, H., *Hierarchical Portfolio Management: Theory and Applications*, Promotor: Prof.dr. J. Spronk, EPS-2007-118-F&A, ISBN: 90-5892-152-9, http://hdl.handle.net/1765/10868

Noeverman, J., *Management Control Systems, Evaluative Style, and Behaviour: Exploring the Concept and Behavioural Consequences of Evaluative Style*, Promotors: Prof.dr. E.G.J. Vosselman & Prof.dr. A.R.T. Williams, EPS-2007-120-F&A, ISBN: 90-5892-151-2, http://hdl.handle.net/1765/10869

Noordegraaf-Eelens, L.H.J., *Contested Communication: A Critical Analysis of Central Bank Speech*, Promotor: Prof.dr. Ph.H.B.F. Franses, EPS-2010-209-MKT, ISBN: 978-90-5892-254-0,

http://hdl.handle.net/1765/21061

Nuijten, I., *Servant Leadership: Paradox or Diamond in the Rough? A Multidimensional Measure and Empirical Evidence*, Promotor: Prof.dr. D.L. van Knippenberg, EPS-2009-183-ORG, http://hdl.handle.net/1765/21405

Oosterhout, J., van, *The Quest for Legitimacy: On Authority and Responsibility in Governance*, Promotors: Prof.dr. T. van Willigenburg & Prof.mr. H.R. van Gunsteren, EPS-2002-012-ORG, ISBN: 90-5892-022-4, http://hdl.handle.net/1765/362

Oosterhout, M., van, *Business Agility and Information Technology in Service Organizations*, Promotor: Prof,dr.ir. H.W.G.M. van Heck, EPS-2010-198-LIS, ISBN: 90-5092-236-6, http://hdl.handle.net/1765/19805

Oostrum, J.M., van, *Applying Mathematical Models to Surgical Patient Planning*, Promotor: Prof.dr. A.P.M. Wagelmans, EPS-2009-179-LIS, ISBN: 978-90-5892-217-5, http://hdl.handle.net/1765/16728

Osadchiy, S.E., *The Dynamics of Formal Organization: Essays on Bureaucracy and Formal Rules*, Promotor: Prof.dr. P.P.M.A.R. Heugens, EPS-2011-231-ORG, ISBN: 978-90-5892-273-1, http://hdl.handle.net/1765/23250

Otgaar, A.H.J., *Industrial Tourism: Where the Public Meets the Private*, Promotor: Prof.dr. L. van den Berg, EPS-2010-219-ORG, ISBN: 90-5892-259-5, http://hdl.handle.net/1765/21585

Ozdemir, M.N., *Project-level Governance, Monetary Incentives and Performance in Strategic R&D Alliances*, Promoror: Prof.dr.ir. J.C.M. van den Ende, EPS-2011-235-LIS, ISBN: 978-90-5892-282-3, http://hdl.handle.net/1765/23550

Paape, L., *Corporate Governance: The Impact on the Role, Position, and Scope of Services of the Internal Audit Function*, Promotors: Prof.dr. G.J. van der Pijl & Prof.dr. H. Commandeur, EPS-2007-111-MKT, ISBN: 90-5892-143-7, http://hdl.handle.net/1765/10417

Pak, K., *Revenue Management: New Features and Models*, Promotor: Prof.dr.ir. R. Dekker, EPS-2005-061-LIS, ISBN: 90-5892-092-5, http://hdl.handle.net/1765/362/6771

Pattikawa, L.H, *Innovation in the Pharmaceutical Industry: Evidence from Drug Introduction in the U.S.*, Promotors: Prof.dr. H.R.Commandeur, EPS-2007-102-MKT, ISBN: 90-5892-135-2, http://hdl.handle.net/1765/9626

Peeters, L.W.P., *Cyclic Railway Timetable Optimization*, Promotors: Prof.dr. L.G. Kroon & Prof.dr.ir. J.A.E.E. van Nunen, EPS-2003-022-LIS, ISBN: 90-5892-042-9, http://hdl.handle.net/1765/429

Pietersz, R., *Pricing Models for Bermudan-style Interest Rate Derivatives*, Promotors: Prof.dr. A.A.J. Pelsser & Prof.dr. A.C.F. Vorst, EPS-2005-071-F&A, ISBN: 90-5892-099-2, http://hdl.handle.net/1765/7122

Pince, C., *Advances in Inventory Management: Dynamic Models*, Promotor: Prof.dr.ir. R. Dekker, EPS-2010-199-LIS, ISBN: 978-90-5892-243-4, http://hdl.handle.net/1765/19867

Poel, A.M. van der, *Empirical Essays in Corporate Finance and Financial Reporting*, Promotors: Prof.dr. A. de Jong & Prof.dr. G.M.H. Mertens, EPS-2007-133-F&A, ISBN: 978-90-5892-165-9, http://hdl.handle.net/1765/13320

Popova, V., *Knowledge Discovery and Monotonicity*, Promotor: Prof.dr. A. de Bruin, EPS-2004-037-LIS, ISBN: 90-5892-061-5, http://hdl.handle.net/1765/1201

Potthoff, D., *Railway Crew Rescheduling: Novel Approaches and Extensions*, Promotors: Prof.dr. A.P.M. Wagelmans & Prof.dr. L.G. Kroon, EPS-2010-210-LIS, ISBN: 90-5892-250-2, http://hdl.handle.net/1765/21084

Pouchkarev, I., *Performance Evaluation of Constrained Portfolios*, Promotors: Prof.dr. J. Spronk & Dr. W.G.P.M. Hallerbach, EPS-2005-052-F&A, ISBN: 90-5892-083-6, http://hdl.handle.net/1765/6731

Prins, R., *Modeling Consumer Adoption and Usage of Value-Added Mobile Services*, Promotors: Prof.dr. Ph.H.B.F. Franses & Prof.dr. P.C. Verhoef, EPS-2008-128-MKT, ISBN: 978/90-5892-161-1, http://hdl.handle.net/1765/12461

Puvanasvari Ratnasingam, P., *Interorganizational Trust in Business to Business E-Commerce*, Promotors: Prof.dr. K. Kumar & Prof.dr. H.G. van Dissel, EPS-2001-009-LIS, ISBN: 90-5892-017-8, http://hdl.handle.net/1765/356

Quak, H.J., *Sustainability of Urban Freight Transport: Retail Distribution and Local Regulation in Cities*, Promotor: Prof.dr. M.B.M. de Koster, EPS-2008-124-LIS, ISBN: 978-90-5892-154-3, http://hdl.handle.net/1765/11990

Quariguasi Frota Neto, J., *Eco-efficient Supply Chains for Electrical and Electronic Products*, Promotors: Prof.dr.ir. J.A.E.E. van Nunen & Prof.dr.ir. H.W.G.M. van Heck, EPS-2008-152-LIS, ISBN: 978-90-5892-192-5, http://hdl.handle.net/1765/14785

Radkevitch, U.L, *Online Reverse Auction for Procurement of Services*, Promotor: Prof.dr.ir. H.W.G.M. van Heck, EPS-2008-137-LIS, ISBN: 978-90-5892-171-0, http://hdl.handle.net/1765/13497

Rijsenbilt, J.A., *CEO Narcissism; Measurement and Impact*, Promotors: Prof.dr. A.G.Z. Kemna & Prof.dr. H.R. Commandeur, EPS-2011-238-STR, ISBN: 978-90-5892-281-6, http://hdl.handle.net/1765/23554

Rinsum, M. van, *Performance Measurement and Managerial Time Orientation*, Promotor: Prof.dr. F.G.H. Hartmann, EPS-2006-088-F&A, ISBN: 90-5892-121-2, http://hdl.handle.net/1765/7993

Roelofsen, E.M., *The Role of Analyst Conference Calls in Capital Markets*, Promotors: Prof.dr. G.M.H. Mertens & Prof.dr. L.G. van der Tas RA, EPS-2010-190-F&A, ISBN: 978-90-5892-228-1, http://hdl.handle.net/1765/18013

Romero Morales, D., *Optimization Problems in Supply Chain Management*, Promotors: Prof.dr.ir. J.A.E.E. van Nunen & Dr. H.E. Romeijn, EPS-2000-003-LIS, ISBN: 90-9014078-6, http://hdl.handle.net/1765/865

Roodbergen, K.J., *Layout and Routing Methods for Warehouses*, Promotors: Prof.dr. M.B.M. de Koster & Prof.dr.ir. J.A.E.E. van Nunen, EPS-2001-004-LIS, ISBN: 90-5892-005-4, http://hdl.handle.net/1765/861

Rook, L., *Imitation in Creative Task Performance*, Promotor: Prof.dr. D.L. van Knippenberg, EPS-2008-125-ORG, http://hdl.handle.net/1765/11555

Rosmalen, J. van, *Segmentation and Dimension Reduction: Exploratory and Model-Based Approaches*, Promotor: Prof.dr. P.J.F. Groenen, EPS-2009-165-MKT, ISBN: 978-90-5892-201-4, http://hdl.handle.net/1765/15536

Roza, M.W., *The Relationship between Offshoring Strategies and Firm Performance: Impact of Innovation, Absorptive Capacity and Firm Size*, Promotors: Prof.dr. H.W. Volberda & Prof.dr.ing. F.A.J. van den Bosch, EPS-2011-214-STR, ISBN: 978-90-5892-265-6, http://hdl.handle.net/1765/22155

Rus, D., *The Dark Side of Leadership: Exploring the Psychology of Leader Self-serving Behavior*, Promotor: Prof.dr. D.L. van Knippenberg, EPS-2009-178-ORG, http://hdl.handle.net/1765/16726

Samii, R., *Leveraging Logistics Partnerships: Lessons from Humanitarian Organizations,* Promotors: Prof.dr.ir. J.A.E.E. van Nunen & Prof.dr.ir. L.N. Van Wassenhove, EPS-2008-153-LIS, ISBN:

978-90-5892-186-4, http://hdl.handle.net/1765/14519

Schaik, D. van, *M&A in Japan: An Analysis of Merger Waves and Hostile Takeovers,* Promotors: Prof.dr. J. Spronk & Prof.dr. J.P.M. Groenewegen, EPS-2008-141-F&A, ISBN: 978-90-5892-169-7, http://hdl.handle.net/1765/13693

Schauten, M.B.J., *Valuation, Capital Structure Decisions and the Cost of Capital,* Promotors: Prof.dr. J. Spronk & Prof.dr. D. van Dijk, EPS-2008-134-F&A, ISBN: 978-90-5892-172-7, http://hdl.handle.net/1765/13480

Schellekens, G.A.C., *Language Abstraction in Word of Mouth,* Promotor: Prof.dr.ir. A. Smidts, EPS-2010-218-MKT, ISBN: 978-90-5892-252-6, http://hdl.handle.net/1765/21580

Schramade, W.L.J., *Corporate Bonds Issuers,* Promotor: Prof.dr. A. De Jong, EPS-2006-092-F&A, ISBN: 90-5892-125-5, http://hdl.handle.net/1765/8191

Schweizer, T.S., *An Individual Psychology of Novelty-Seeking, Creativity and Innovation,* Promotor: Prof.dr. R.J.M. van Tulder, EPS-2004-048-ORG, ISBN: 90-5892-077-1, http://hdl.handle.net/1765/1818

Six, F.E., *Trust and Trouble: Building Interpersonal Trust Within Organizations*, Promotors: Prof.dr. B. Nooteboom & Prof.dr. A.M. Sorge, EPS-2004-040-ORG, ISBN: 90-5892-064-X, http://hdl.handle.net/1765/1271

Slager, A.M.H., *Banking across Borders*, Promotors: Prof.dr. R.J.M. van Tulder & Prof.dr. D.M.N. van Wensveen, EPS-2004-041-ORG, ISBN: 90-5892-066–6, http://hdl.handle.net/1765/1301

Sloot, L., *Understanding Consumer Reactions to Assortment Unavailability*, Promotors: Prof.dr. H.R. Commandeur, Prof.dr. E. Peelen & Prof.dr. P.C. Verhoef, EPS-2006-074-MKT, ISBN: 90-5892-102-6, http://hdl.handle.net/1765/7438

Smit, W., *Market Information Sharing in Channel Relationships: Its Nature, Antecedents and Consequences*, Promotors: Prof.dr.ir. G.H. van Bruggen & Prof.dr.ir. B. Wierenga, EPS-2006-076-MKT, ISBN: 90-5892-106-9, http://hdl.handle.net/1765/7327

Sonnenberg, M., *The Signalling Effect of HRM on Psychological Contracts of Employees: A Multi-level Perspective*, Promotor: Prof.dr. J. Paauwe, EPS-2006-086-ORG, ISBN: 90-5892-119-0, http://hdl.handle.net/1765/7995

Sotgiu, F., *Not All Promotions are Made Equal: From the Effects of a Price War to Cross-chain*

Cannibalization, Promotors: Prof.dr. M.G. Dekimpe & Prof.dr.ir. B. Wierenga, EPS-2010-203-MKT, ISBN: 978-90-5892-238-0, http://hdl.handle.net/1765/19714

Speklé, R.F., *Beyond Generics: A closer Look at Hybrid and Hierarchical Governance*, Promotor: Prof.dr. M.A. van Hoepen RA, EPS-2001-008-F&A, ISBN: 90-5892-011-9, http://hdl.handle.net/1765/357

Srour, F.J., *Dissecting Drayage: An Examination of Structure, Information, and Control in Drayage Operations*, Promotor: Prof.dr. S.L. van de Velde, EPS-2010-186-LIS, http://hdl.handle.net/1765/18231

Stam, D.A., *Managing Dreams and Ambitions: A Psychological Analysis of Vision Communication*, Promotor: Prof.dr. D.L. van Knippenberg, EPS-2008-149-ORG, http://hdl.handle.net/1765/14080

Stienstra, M., *Strategic Renewal in Regulatory Environments: How Inter- and Intra-organisational Institutional Forces Influence European Energy Incumbent Firms,* Promotors: Prof.dr.ir. F.A.J. Van den Bosch & Prof.dr. H.W. Volberda, EPS-2008-145-STR, ISBN: 978-90-5892-184-0, http://hdl.handle.net/1765/13943

Sweldens, S.T.L.R., *Evaluative Conditioning 2.0: Direct versus Associative Transfer of Affect to Brands*, Promotor: Prof.dr. S.M.J. van Osselaer, EPS-2009-167-MKT, ISBN: 978-90-5892-204-5, http://hdl.handle.net/1765/16012

Szkudlarek, B.A., *Spinning the Web of Reentry: [Re]connecting reentry training theory and practice*, Promotor: Prof.dr. S.J. Magala, EPS-2008-143-ORG, ISBN: 978-90-5892-177-2, http://hdl.handle.net/1765/13695

Tempelaar, M.P., *Organizing for Ambidexterity: Studies on the Pursuit of Exploration and Exploitation through Differentiation, Integration, Contextual and Individual Attributes*, Promotors: Prof.dr.ing. F.A.J. van den Bosch & Prof.dr. H.W. Volberda, EPS-2010-191-STR, ISBN: 978-90-5892-231-1, http://hdl.handle.net/1765/18457

Teunter, L.H., *Analysis of Sales Promotion Effects on Household Purchase Behavior*, Promotors: Prof.dr.ir. B. Wierenga & Prof.dr. T. Kloek, EPS-2002-016-MKT, ISBN: 90-5892-029-1, http://hdl.handle.net/1765/868

Tims, B., *Empirical Studies on Exchange Rate Puzzles and Volatility*, Promotor: Prof.dr. C.G. Koedijk, EPS-2006-089-F&A, ISBN: 90-5892-113-1, http://hdl.handle.net/1765/8066

Tiwari, V., *Transition Process and Performance in IT Outsourcing: Evidence from a Field Study and*

Laboratory Experiments, Promotors: Prof.dr.ir. H.W.G.M. van Heck & Prof.dr. P.H.M. Vervest, EPS-2010-201-LIS, ISBN: 978-90-5892-241-0, http://hdl.handle.net/1765/19868

Tröster, C., *Nationality Heterogeneity and Interpersonal Relationships at Work*, Promotor: Prof.dr. D.L. van Knippenberg, EPS-2011-233-ORG, http://hdl.handle.net/1765/23298

Tuk, M.A., *Is Friendship Silent When Money Talks? How Consumers Respond to Word-of-Mouth Marketing*, Promotors: Prof.dr.ir. A. Smidts & Prof.dr. D.H.J. Wigboldus, EPS-2008-130-MKT, ISBN: 978-90-5892-164-2, http://hdl.handle.net/1765/12702

Tzioti, S., *Let Me Give You a Piece of Advice: Empirical Papers about Advice Taking in Marketing*, Promotors: Prof.dr. S.M.J. van Osselaer & Prof.dr.ir. B. Wierenga, EPS-2010-211-MKT, ISBN: 978-90-5892-251-9, hdl.handle.net/1765/21149

Vaccaro, I.G., *Management Innovation: Studies on the Role of Internal Change Agents*, Promotors: Prof.dr. F.A.J. van den Bosch & Prof.dr. H.W. Volberda, EPS-2010-212-STR, ISBN: 978-90-5892-253-3, hdl.handle.net/1765/21150

Valck, K. de, *Virtual Communities of Consumption: Networks of Consumer Knowledge and Companionship*, Promotors: Prof.dr.ir. G.H. van Bruggen & Prof.dr.ir. B. Wierenga, EPS-2005-050-stocktickerMKT, ISBN: 90-5892-078-X, http://hdl.handle.net/1765/6663

Valk, W. van der, *Buyer-Seller Interaction Patterns During Ongoing Service Exchange*, Promotors: Prof.dr. J.Y.F. Wynstra & Prof.dr.ir. B. Axelsson, EPS-2007-116-MKT, ISBN: 90-5892-146-8, http://hdl.handle.net/1765/10856

Verheijen, H.J.J., *Vendor-Buyer Coordination in Supply Chains*, Promotor: Prof.dr.ir. J.A.E.E. van Nunen, EPS-2010-194-LIS, ISBN: 90-5892-234-2, http://hdl.handle.net/1765/19594

Verheul, I., *Is There a (Fe)male Approach? Understanding Gender Differences in Entrepreneurship*, Promotor: Prof.dr. A.R. Thurik, EPS-2005-054-ORG, ISBN: 90-5892-080-1, http://hdl.handle.net/1765/2005

Verwijmeren, P., *Empirical Essays on Debt, Equity, and Convertible Securities*, Promotors: Prof.dr. A. de Jong & Prof.dr. M.J.C.M. Verbeek, EPS-2009-154-F&A, ISBN: 978-90-5892-187-1, http://hdl.handle.net/1765/14312

Vis, I.F.A., *Planning and Control Concepts for Material Handling Systems*, Promotors: Prof.dr. M.B.M. de Koster & Prof.dr.ir. R. Dekker, EPS-2002-014-LIS, ISBN: 90-5892-021-6,

http://hdl.handle.net/1765/866

Vlaar, P.W.L., *Making Sense of Formalization in Interorganizational Relationships: Beyond Coordination and Control*, Promotors: Prof.dr.ir. F.A.J. Van den Bosch & Prof.dr. H.W. Volberda, EPS-2006-075-STR, ISBN 90-5892-103-4, http://hdl.handle.net/1765/7326

Vliet, P. van, *Downside Risk and Empirical Asset Pricing*, Promotor: Prof.dr. G.T. Post, EPS-2004-049-F&A, ISBN: 90-5892-07-55, http://hdl.handle.net/1765/1819

Vlist, P. van der, *Synchronizing the Retail Supply Chain*, Promotors: Prof.dr.ir. J.A.E.E. van Nunen & Prof.dr. A.G. de Kok, EPS-2007-110-LIS, ISBN: 90-5892-142-0, http://hdl.handle.net/1765/10418

Vries-van Ketel E. de, *How Assortment Variety Affects Assortment Attractiveness:A Consumer Perspective*, Promotors: Prof.dr. G.H. van Bruggen & Prof.dr.ir. A. Smidts, EPS-2006-072-MKT, ISBN: 90-5892-101-8, http://hdl.handle.net/1765/7193

Vromans, M.J.C.M., *Reliability of Railway Systems*, Promotors: Prof.dr. L.G. Kroon, Prof.dr.ir. R. Dekker & Prof.dr.ir. J.A.E.E. van Nunen, EPS-2005-062-LIS, ISBN: 90-5892-089-5, http://hdl.handle.net/1765/6773

Vroomen, B.L.K., *The Effects of the Internet, Recommendation Quality and Decision Strategies on Consumer Choice*, Promotor: Prof.dr. Ph.H.B.F. Franses, EPS-2006-090-MKT, ISBN: 90-5892-122-0, http://hdl.handle.net/1765/8067

Waal, T. de, *Processing of Erroneous and Unsafe Data*, Promotor: Prof.dr.ir. R. Dekker, EPS-2003-024-LIS, ISBN: 90-5892-045-3, http://hdl.handle.net/1765/870

Waard, E.J. de, *Engaging Environmental Turbulence: Organizational Determinants for Repetitive Quick and Adequate Responses*, Promotors: Prof.dr. H.W. Volberda & Prof.dr. J. Soeters, EPS-2010-189-STR, ISBN: 978-90-5892-229-8, http://hdl.handle.net/1765/18012

Wall, R.S., *Netscape: Cities and Global Corporate Networks*, Promotor: Prof.dr. G.A. van der Knaap, EPS-2009-169-ORG, ISBN: 978-90-5892-207-6, http://hdl.handle.net/1765/16013

Wang, Y., *Information Content of Mutual Fund Portfolio Disclosure*, Promotor: Prof.dr. M.J.C.M. Verbeek, EPS-2011-242-F&A, ISBN: 978-90-5892-285-4, http://hdl.handle.net/1765/1

Watkins Fassler, K., *Macroeconomic Crisis and Firm Performance*, Promotors: Prof.dr. J. Spronk & Prof.dr. D.J. van Dijk, EPS-2007-103-F&A, ISBN: 90-5892-138-3, http://hdl.handle.net/1765/10065

Weerdt, N.P. van der, *Organizational Flexibility for Hypercompetitive Markets: Empirical Evidence of the Composition and Context Specificity of Dynamic Capabilities and Organization Design Parameters*, Promotor: Prof.dr. H.W. Volberda, EPS-2009-173-STR, ISBN: 978-90-5892-215-1, http://hdl.handle.net/1765/16182

Wennekers, A.R.M., *Entrepreneurship at Country Level: Economic and Non-Economic Determinants*, Promotor: Prof.dr. A.R. Thurik, EPS-2006-81-ORG, ISBN: 90-5892-115-8, http://hdl.handle.net/1765/7982

Wielemaker, M.W., *Managing Initiatives: A Synthesis of the Conditioning and Knowledge-Creating View*, Promotors: Prof.dr. H.W. Volberda & Prof.dr. C.W.F. Baden-Fuller, EPS-2003-28-STR, ISBN: 90-5892-050-X, http://hdl.handle.net/1765/1042

Wijk, R.A.J.L. van, *Organizing Knowledge in Internal Networks: A Multilevel Study*, Promotor: Prof.dr.ir. F.A.J. van den Bosch, EPS-2003-021-STR, ISBN: 90-5892-039-9, http://hdl.handle.net/1765/347

Wolters, M.J.J., *The Business of Modularity and the Modularity of Business*, Promotors: Prof. mr. dr. P.H.M. Vervest & Prof.dr.ir. H.W.G.M. van Heck, EPS-2002-011-LIS, ISBN: 90-5892-020-8, http://hdl.handle.net/1765/920

Wubben, M.J.J., *Social Functions of Emotions in Social Dilemmas*, Promotors: Prof.dr. D. De Cremer & Prof.dr. E. van Dijk, EPS-2009-187-ORG, http://hdl.handle.net/1765/18228

Xu, Y., *Empirical Essays on the Stock Returns, Risk Management, and Liquidity Creation of Banks*, Promotor: Prof.dr. M.J.C.M. Verbeek, EPS-2010-188-F&A, http://hdl.handle.net/1765/18125

Yang, J., *Towards the Restructuring and Co-ordination Mechanisms for the Architecture of Chinese Transport Logistics*, Promotor: Prof.dr. H.E. Harlambides, EPS-2009-157-LIS, ISBN: 978-90-5892-198-7, http://hdl.handle.net/1765/14527

Yu, M., *Enhancing Warehouse Perfromance by Efficient Order Picking*, Promotor: Prof.dr. M.B.M. de Koster, EPS-2008-139-LIS, ISBN: 978-90-5892-167-3, http://hdl.handle.net/1765/13691

Zhang, X., *Strategizing of Foreign Firms in China: An Institution-based Perspective*, Promotor: Prof.dr. B. Krug, EPS-2007-114-ORG, ISBN: 90-5892-147-5, http://hdl.handle.net/1765/10721

Zhang, X., *Scheduling with Time Lags*, Promotor: Prof.dr. S.L. van de Velde, EPS-2010-206-LIS, ISBN:

978-90-5892-244-1, http://hdl.handle.net/1765/19928

Zhou, H., *Knowledge, Entrepreneurship and Performance: Evidence from Country-level and Firm-level Studies*, Promotors: Prof.dr. A.R. Thurik & Prof.dr. L.M. Uhlaner, EPS-2010-207-ORG, ISBN: 90-5892-248-9, http://hdl.handle.net/1765/20634

Zhu, Z., *Essays on China's Tax System*, Promotors: Prof.dr. B. Krug & Prof.dr. G.W.J. Hendrikse, EPS-2007-112-ORG, ISBN: 90-5892-144-4, http://hdl.handle.net/1765/10502

Zwan, P.W. van der, *The Entrepreneurial Process: An International Analysis of Entry and Exit*, EPS-2011-234-ORG, ISBN: 978-90-5892-277-9, http://hdl.handle.net/1765/23422

Zwart, G.J. de, *Empirical Studies on Financial Markets: Private Equity, Corporate Bonds and Emerging Markets*, Promotors: Prof.dr. M.J.C.M. Verbeek & Prof.dr. D.J.C. van Dijk, EPS-2008-131-F&A, ISBN: 978-90-5892-163-5, http://hdl.handle.net/1765/12703

COMPUTATIONAL AND GAME-THEORETIC APPROACHES FOR MODELING BOUNDED RATIONALITY

This thesis studies various computational and game-theoretic approaches to economic modeling. Unlike traditional approaches to economic modeling, the approaches studied in this thesis do not rely on the assumption that economic agents behave in a fully rational way. Instead, economic agents are assumed to be boundedly rational. Abandoning the assumption of full rationality has a number of consequences for the way in which economic reality is being modeled. Traditionally, economic models are mostly of a static nature and have a strong focus on deriving equilibria. Also, models are usually analyzed mathematically. In models of boundedly rational behavior, dynamic elements play a much more prominent role and there is less emphasis on equilibrium behavior. Also, to analyze models of boundedly rational behavior, researchers not only use mathematical techniques but they also rely heavily on computer simulations.

This thesis presents four studies into the modeling of boundedly rational behavior of economic agents. Two studies are concerned with investigating the emergence of cooperation among boundedly rational agents. One study focuses on cooperation among firms in a Cournot oligopoly market, while the other study examines cooperation in a spatial model of price-competing firms. The other two studies in this thesis are concerned with methodological issues in the use of evolutionary algorithms for economic modeling purposes. One study shows how evolutionary algorithms can be analyzed mathematically rather than using computer simulations. The other study criticizes the use of a so-called binary encoding in evolutionary algorithms.

ERiM

The Erasmus Research Institute of Management (ERIM) is the Research School (Onderzoekschool) in the field of management of the Erasmus University Rotterdam. The founding participants of ERIM are the Rotterdam School of Management (RSM), and the Erasmus School of Economics (ESE). ERIM was founded in 1999 and is officially accredited by the Royal Netherlands Academy of Arts and Sciences (KNAW). The research undertaken by ERIM is focused on the management of the firm in its environment, its intra- and interfirm relations, and its business processes in their interdependent connections.

The objective of ERIM is to carry out first rate research in management, and to offer an advanced doctoral programme in Research in Management. Within ERIM, over three hundred senior researchers and PhD candidates are active in the different research programmes. From a variety of academic backgrounds and expertises, the ERIM community is united in striving for excellence and working at the forefront of creating new business knowledge.

Erasmus Research Institute of Management - ERiM

ERIM PhD Series
Research in Management