Bounded Rationality and Heterogeneity in Economic Dynamic Models

ISBN 978 90 5170 936 0

Cover design: Crasborn Graphic Designers bno, Valkenburg a.d. Geul

This book is no. **396** of the Tinbergen Institute Research Series, established through cooperation between Thela Thesis and the Tinbergen Institute. A list of books which already appeared in the series can be found in the back.

Bounded Rationality and Heterogeneity in Economic Dynamic Models

ACADEMISCH PROEFSCHRIFT

ter verkrijging van de graad van doctor
aan de Universiteit van Amsterdam
op gezag van de Rector Magnificus
prof. mr. P.F. van der Heijden
ten overstaan van een door het college voor promoties ingestelde
commissie, in het openbaar te verdedigen in de Aula der Universiteit
op vrijdag 2 februari 2007, te 12.00 uur

door

Pietro Dino Enrico Dindo

geboren te Verona, Italië

Promotiecommissie:

Promotor: Prof. dr. C.H. Hommes

Co-promotor: Dr. C.G.H. Diks

Overige leden: Prof. dr. G. Bottazzi
Prof. dr. W.A. Brock
Prof. dr. ir. G. van der Laan
Dr. R. Ramer
Dr. J. Tuinstra

Faculteit der Economie en Bedrijfskunde

A papà, mamma e Francesca

Acknowledgments

Four years ago, when I joined CeNDEF at the Department of Quantitative Economics of the University of Amsterdam, as a Ph.D. student, the economist inside me was at its infancy. Four years later my knowledge and skills are still far from being those of a mature scholar, but now, at least, I have developed a feeling for the direction I should move in. Moreover, I have succeeded in producing this thesis. I want to take this chance to thank all those people that helped me "growing up" and achieving this result.

First of all, I would like to thank my promotor, Cars Hommes, and my co-promotor, Cees Diks. I still remember very clearly, to the detail of the shoes I was wearing, the quite informal interview I had with Cees in May 2002. We had a walk in the see-overlooking park of the campus of the International Center for Theoretical Physics in Miramare, Trieste. Apparently my shoes were appropriate enough and I came to Amsterdam few months later. I thank Cees for having trusted in me from that very first moment and for having guided my walk in the profession of science. I also remember the first meeting I had with Cars. His warm welcome made it clear to me that, although I was hired by Cees, I would have had also Cars' support for the coming years. Ex-post, I not only had his assistance but also learned a lot from interacting with him on the ground of science. It is the dream of every Ph.D. student to have two supervisors who allow the perfect mixture between freedom and protection that I enjoyed with Cees and Cars. They made my way through economics and thesis-writing a most comfortable one.

I would like to express my gratitude also to the members of the committee. Beside reading the thesis manuscript and commenting on it, they all interacted with me at various levels of my research and education. Giulio Bottazzi has shown interest in my research already at its early stages. He gave me the opportunity to present my work at the Sant' Anna School for Advanced Studies in Pisa. I am glad I will join him there soon. I greatly enjoyed the interaction with Buz Brock, who involved me in the project from which Chapter 2 originated. I thank him for the enthusiasm he showed toward willing-to-be scholars. Gerard van der Laan thought me, both at the Tinbergen Institute and at the Nake Network, most of the economic theory I know. I greatly appreciated his expertise and critical eye. I enjoyed the challenge posed by the remarks of Roald Ramer, a passionate and knowledgeable discussant. I felt very pleased the few times I managed, partly, to convince him. Jan Tuinstra gave me the opportunity to start doing research in economics. I thank him for his patience in replying to my early naive remarks. His effort proved to be very important for my future research. Chapter 3 originated from that work.

Writing and printing a Ph.D. thesis also has some costs and requires administrative assistance. With this respect I thank the Tinbergen Institute, The University of Amsterdam and the Netherlands Organization for Scientific Research (NWO) for providing their generous financial support. I also thank the secretarial office of the Department of Quantitative Economics, Andries, Kees and Ida, Bas Bouten of Resam, and the administrative staff of the Tinbergen Institute for their attentive assistance.

Boundedly rational Ph.D. students need the support of colleagues and friends to be able to end their projects. CeNDEF and the Department of Quantitative Economics provided a very comfortable working environment. My thanks go to Valentyn, Florian, Dave, Maurice K., Peter H., Marius, Maurice B., Jurek, Cesar, Yebin, Kees Jan, Peter B., Jan K., Noud, Herman, André, René and to all the other colleagues. Misha and Sander, my paranimfen, were kind enough to pay attention to my abundant, often confused, theories. Our interaction boosted both my personal and professional development. Outside the university many special people made my period as a Ph.D. students very enjoyable and cosy. Massimo, Holly and Rose, were my family at this side of the Alps. Friends, many of them being Ph.D. student, have made real that sharing of joy and pains that makes friendship so special. Special thanks to Srdjan, Marcos, Alessandro, Carlo, Silvia, Misha R., Sandra, José, Haris, Dessi, Ana, Suncica, Mattia, Astrid, Paul, Sophocles, Federico, Sebastiano and Roberta. During this period some physical exercise has been crucial for keeping my mental health. I thank my rowing team, and especially Leif-Erik for his first invitation to join in, for the many pleasant rowing sessions up and down the Amstel.

The proximity of my family and friends in Italy has played a most important role too. My parents and my sister, to whom this thesis is dedicated, have always shown appreciation and encouragement for my choice to come to Amsterdam to pursue my Ph.D. study. My Italian friends gave me the feeling that we were not forgetting of each other, both by physically reminding me of their presence in their trips to Amsterdam and by arranging last minutes meetings during my too short visits in Italy. Keeping a relationship at a thousand kilometers of distance is not easy and I am very glad that we managed to do that.

Last but not least I would like to thank Kateryna for her patience, care and love. I am a happy man beside her.

Amsterdam, December 2006

Contents

Chapter 1

General introduction and thesis outline

Economic phenomena are characterized by a strong mutual dependence between the choices of economic actors, e.g. individual consumers or firms, and the economic environment in which these choices are made and evaluated. The origin of this dependence lies in the fact that the economic environment is formed by the aggregation of all individual choices, while at the same time individual choices are affected by the overall economic environment.

In order to clarify the nature of this dependence, consider the following examples which will be recurrent throughout this thesis. In a market of a perishable consumption good, firms have to decide today how much to produce to supply to the market tomorrow. On the one hand, firms base their production decision upon tomorrow's expected profits. On the other hand, tomorrow's profits will depend on the total amount firms produce. In fact, given consumer demand, profits are a function of market clearing prices, which depend on total supply. Financial markets are another well-known example of this mutual dependence between choices and environment. The demand of investors for an asset is driven by expected future returns, while at the same time asset returns are determined, through realized asset prices, by investors' demand. Another example, with the advantage of appealing to everyone's experience, is the *El Farol* problem proposed by Arthur (1994). *El Farol* is a bar where a live music show takes place on a fixed day, say Thursday, every week. There are 100 people who would like to go to the bar on Thursday evening, but the *El Farol* bar is not big enough for all of them. In fact, the bar is too crowded when there are more than 60 people. Each Thursday evening everyone decides whether to go to the bar or not. Each decision maker is directly influencing the number of agents present at the bar, and at the same time, since that number affects his or her utility, expectations about the number of agents present at the bar influence each individual decision. Notice that in all these examples the mutual dependence between individual decisions and economic environment is mediated through a feedback relation between expectations and realizations of some variable, e.g. prices, returns, or attendance level. For this reason we can call these systems *expectational feedback systems*.

Traditional analysis in economics, i.e. neoclassical economics, has circumvented the complications arising through the type of interaction of many individuals outlined above, by assuming that all agents are *rational*. Modeling human behavior as rational implies that, in making the decision that maximizes their objectives, agents take into account the choices of others, assuming that others are doing the same. When expectations are involved, agents are actually supposed to solve the expectational feedback system and find a "fixed point" or equilibrium outcome where expectations and realizations of the same variable coincide. Assuming rational thinking leads directly to an equilibrium where choices need not be corrected, unless unanticipated changes of the exogenous parameters characterizing the environment or the decision makers, take place. Thus, if agents were rational, observed changes of economic variables, which are a function of agents' choices, should come from a response to unexpected changes in some exogenous characteristics, or fundamentals, of the economy. Empirical analysis, both of real world data as well as of laboratory experiments with human subjects, is at odds with this statement and shows that economic variables fluctuate, even when changes in the fundamentals of the economy do not occur.

In this thesis we investigate the possibility that economic fluctuations can be explained through the *interaction of boundedly rational agents*, that is, agents are not assumed to be rational and are not necessarily able to solve the mutual dependence implied by the expectational feedback. Boundedly rational agents use simple behavioral rules and adapt their behavior over time, switching from time to time to better performing rules. Since it is not a priori clear which of these simpler rules should be used, we explicitly assume that agents are *heterogeneous* and employ different rules to address the same decision problem. We let a selection mechanism, such as "survival of the fittest", discipline the class of behavioral rules.

In departing from the traditional rational approach we have two main goals. First, we want to appraise if an argument used in favor of rationality, namely that rationality is the outcome of the repeated interaction of heterogeneous boundedly rational agents, is justified. This part of the analysis has thus a theoretical motivation. Second, having shown in what respect our results differ from the rational benchmark, we want to characterize whether our interacting agents framework can reproduce empirically observed phenomena in the specific economic settings we consider.

In the context of this general background, this thesis is built around three different models discussed in three independent chapters. Every model refers to one of the three economic frameworks briefly sketched at the beginning of this introduction, for which the expectational feedback is a crucial characteristic. Chapter 2 uses a classical cobweb model, i.e. a market for a perishable consumption good, to analyze the trade-off between benefits and costs of rules with different degrees of sophistication. In fact, if sophisticated predictions, as rational expectations, are costly, agents could decide to adopt a cheaper, but simpler predictor, as long as its predictions are not too far from the realized values. We investigate the impact on commodity prices of the interaction of agents having this extra degree of freedom in choosing between cheap free riding and costly sophisticated prediction. In Chapter 3, we characterize the competition of a large group of boundedly rational agents using different strategies to repeatedly exploit the same scarce resource, when no market institutions are present to coordinate agents' actions. This chapter has been inspired directly by Arthur's *El Farol* bar problem. We compare our results with those of equilibrium rational solutions, agent-based computational simulations, and laboratory experiment with human subjects. In Chapter 4, we investigate whether, in financial markets, the interaction of boundedly rational agents, triggered by informational differences, can help explaining time series properties of empirical financial data.

In the remaining part of this introductory chapter, Section 1.1 offers a more detailed review of the theoretical background of this thesis, i.e. bounded rationality and interaction of heterogeneous agents. Section 1.2 presents each of the three remaining chapters. In that section we also offer a brief introduction to nonlinear dynamics, the main technical tool that will be used throughout this thesis. Section 1.3 concludes this introductory chapter discussing an interesting common feature of our models, namely the presence of *negative* expectational feedback. This feature proves to be helpful to interpret chapters' results.

1.1 Theoretical background

Economic theory has always faced the issue of how to model human behavior, in particular how to model individual decision making in an interactive framework. Assuming that an economic actor, e.g. an individual consumer or a firm, wants to make the decision in his best interest and that he is capable of making this judgment, the two approaches that are currently in use rely on different levels of the decision maker's *rationality*.

The approach of neoclassical economics, which we will refer to as *rational*, can be summarized using two requirements. First, the decision maker is assumed to be able to choose the alternative that maximizes utility or profit, given his beliefs about the economic environment and the actions of the other actors in the economy. Second, each decision maker is assumed to be able to predict exogenous as well as endogenous variables, so that his original beliefs are self-fulfilling. The first requirement is mostly related to the rationality of preferences and choices. The second requirement is mostly related to the rationality of expectations, that is, of agents' prediction of future variables, including other agents' actions. An implication of these two requirements is that a rational decision maker knows as much as the modeler regarding the economic framework of interest. Moreover, the rational agent is also assumed to be able to "solve" this model, that is, to make decisions such that all predictions and beliefs are consistent with the outcome of all agents' choices. Stated differently a rational agent is supposed to solve for the equilibrium of the expectational feedback system. Thus, the primary effort of a rational agent consists of searching for an equilibrium[1]. At the equilibrium rational agents do not need to revise their decisions unless an exogenous change in some variables of the model, such as an unanticipated change in the agents' preferences or in the structure of the economy, comes about.

The *bounded rationality* approach (see e.g. Conlisk, 1996 for a survey) considers the requirements that rationality poses on peoples' characteristics, both in terms of knowledge of the economic environment and of their computational capabilities, unrealistic because they are too demanding. Different notions of bounded rationality have been formulated in the literature. For example, in Simon (1957) and in Rubinstein (1998) the emphasis is on limitations of human knowledge and computational abilities in decision making, whereas Sargent (1993) questions the second aspect of rationality, that is the capability of individuals to form rational expectations. Generally speaking, a boundedly rational agent is modeled as able to choose what he perceives as the best for himself, but he does not know the exact structure of the environment. Put differently, we can think of a boundedly rational decision maker as one choosing only from a set of alternatives which is bounded by his individual perception. In particular, when predictions of future variables or choices of other economic agents are involved, bounded rationality implies that ex-ante predictions and ex-post outcomes need not coincide perfectly. Therefore a boundedly rational agent is not assumed to be able to "solve" the model. That is, he is not able to choose equilibrium outcomes where all beliefs are self-fulfilling so that his decisions are, both ex-ante and ex-post, optimal. On the contrary, a boundedly rational agent uses simple rules of thumb and keeps revising his choices as he learns about the economic environment in which he is operating, through feedback about his past decisions. In particular, expectational feedbacks relates agents' expectations about relevant variables,

[1]When more equilibria are presents, it remains to be seen if agents are able to coordinate on the same equilibrium. Coordination games are the simplest example of this problem. We do not raise this important issue in this thesis as in all the model we investigate only one equilibrium typically stands out. See also Section 1.3 of this introdusction

e.g. prices, to their realizations. As a result of the feedback structure and learning, models of bounded rationality are structurally dynamic and may or may not settle down to an equilibrium where agents have learned how to coordinate their actions.

The decision whether to model agents as rational or boundedly rational is part of the assumptions of an economic model. The scholars opting for rationality argue that it is a helpful assumption to describe the equilibrium outcome of people's economic interaction. This defence of rationality is known as the *"as if"* argument and goes back to Friedman and Savage (1948), Friedman (1953), and Alchian (1950). In particular, Friedman (1953) argues that a model should not be judged in terms of the realism of its assumptions but in terms of the realism of its predictions, and that modeling agents as rational is justified because the repeated interaction of heterogeneous boundedly rational agents leads to the same outcome *as if* agents are rational. The general underlying idea is that agents who are not rational would learn to be rational over time since incentives to behave rationally, such as higher profits or utility, are constantly at work. In fact, agents who adopt non-rational rules would be out-performed by agents using rational rules, since rational rules come from optimization and lead to higher economic performance. In summary, assuming rationality is often based on the presumption that this approach offers the equilibrium outcome of repeated interaction, the details of which are not worth being modeled.

Part of the recent interest in models of bounded rationality has been motivated by the attempt to put some structure on the *"as if"* argument. Does the interaction of boundedly rational agents leads to the same outcome *as if* agents are rational? Convergence to rational behavior has been the object of investigation of many theoretical papers in the last two decades. In macroeconomics, Sargent (1993) and Evans and Honkapohja (2001) address the possibility of agents learning to form rational expectations. The main message is that "... some rational expectations equilibria are learnable while others are not. Furthermore, convergence will in general depend on all economic parameters of the system, including policy parameters" (Bullard 2006, p. 205). Learnability of equilibria thus needs to be investigated case by case. In microeconomics, more specifically in game theory, a related issue is the possibility of learning Nash equilibria, i.e. the equilibria played by rational agents, as investigated in Fudenberg and Levine (1998). Evolutionary game theory explores the *"as if"* argument when the process of converging is regulated by evolutionary forces driven by "survival of the fittest", rather than by adaptive learning, see e.g. Weibull (1995). Whereas, according to the *"as if"* argument agents who are closer to rationality should make larger profits and thus overcome other types of behavior, cases exist, as summarized e.g. in Weibull (1994), where convergence is not attained and fluctuations around the equilibrium never vanish. In this thesis we will also encounter some examples of this behavior.

Together from appraising whether modeling the interaction of different groups of boundedly rational agents supports the *"as if"* argument, modeling agents as boundedly rational can be an important step in narrowing the differences between the predictions of economic theory and empirical data. Much evidence has been collected in the last thirty years against the practice of modeling human behavior as rational. For example, Conlisk (1996) classifies this evidence as either direct, through rationality tests on individuals, or indirect, when models assuming rational agents are at odds with empirical data. Direct evidence against rationality consists, for example, in showing that, when faced with decisions involving uncertainty which have an objectively correct answer, agents show psychological biases and failure in rationalizing the problem (see e.g. Tversky and Kahneman, 1974). As already argued by Simon (1957), agents rather use simple rules of thumb or heuristics than engage in difficult, more rational, computations. Moreover, when deciding whether to use more sophisticated rules, agents take into account deliberation costs, such as information gathering costs or information processing costs as those costs associated with the extra effort they are putting into the decision making. Indirect evidence against rationality has been collected from empirical testing of economic models built under the assumption that agents are rational. Consider the following examples that are related to the economic frameworks investigated in the rest of this thesis. In experiments of market entry games, which are similar to the *El Farol* game discussed previously, at the aggregate level agents seem to be able to coordinate on a Nash equilibrium of the game, but, at the individual level, use different simple rules which do not coincide with rational behavior. Furthermore, their interaction generates excess variability of the entry percentages with respect to the game theoretical predictions, (see e.g. Sundali, Rapoport, and Seale, 1995 or Ochs, 1990). In financial markets, prices seem to be much more volatile than justified by the movement of the underlying fundamentals (see e.g. Shiller, 1989) and returns are correlated (see e.g. Fama and French, 1988b). Moreover, expectational surveys, as Frankel and Froot (1987) or Chow (1989), argue that agents do not use rational expectations. However in general, it is still under debate whether macroeconomic fluctuations of unemployment, business cycles and growth rates are partly driven by expectations as argued in Grandmont (1985), see also Grandmont (1998) and Hommes (2004).

In deviating from rationality, and modeling agents as boundedly rational it is often assumed that agents are *heterogeneous*. Kirman (1992, 2006) summarizes some of the reasons why the assumption that agents are homogeneous, or that their heterogeneity is not relevant as their choices can be summarized by the choice of a so-called representative agent, should be discarded. One commonly referred to reason is the "no trade" argument, which states that homogeneous agents, as agents with homogeneous expectations, would have no reasons to trade among themselves. Another reason, which is particularly relevant to our theoretical background, involves an evolutionary explanation. That is, evolution and adaptation necessarily requires some variety of behavior, if one wants the selection mechanism to indicate which is the "surviving" trait. Denying heterogeneity in models with boundedly rational agents gives rise to limitations as it is not clear a priori which kind of boundedly rational behavior each agents should be endowed with. For these reasons, modeling agents as heterogeneous is becoming more and more popular in economics, as shown e.g. by the extensive surveys on analytical and computational models with heterogeneous agents in Hommes (2006) and LeBaron (2006).

In this thesis we assume that agents are heterogeneous, in the sense that they choose different simple decision rules to address the same decision problem. Generally speaking, rules can differ in terms of sophistication, where the most sophisticated rule corresponds to rationality. We also assume that the higher the sophistication of a rule, the higher the deliberation cost an agent pays in order to use it. Rules can also differ in terms of the information they use, where information can also be costly. Instead of considering fixed fractions of agents employing each rule, we let them evolve over time as a function of their "fitness". Thus, we employ an evolutionary approach where a "survival of the fittest" mechanism is at work. A rule that has performed better according to some measure, to be defined case by case, is used by a higher fraction of agents. In our models, we use two different updating mechanism for fractions. In Chapter 2 we use a discrete choice mechanism along the lines of Manski and McFadden (1981) and Anderson, de Palma, and Thisse (1993) (see also Brock and Hommes, 1997, for an early application). This updating mechanism assumes that in choosing between different strategies, agents have an idiosyncratic component that, together with the fitness measure, determines how individual choices are distributed among the different alternatives. In Chapters 3 and 4, we use the replicator dynamics of Taylor and Jonker (1978). The replicator dynamics is related to biological reproduction and the number of agents using a certain rule evolves both as a function of the current number of agents using that rule, and of the fitness of each rule. The replicator dynamics can also be motivated in the context of boundedly rational agents learning and imitating strategies in a strategic environment (see e.g. Weibull, 1995, Chapter 5 or Binmore and Samuelson, 1997).

1.2 Thesis Outline

This thesis is built around three main economic frameworks, which are developed in separate chapters. Each chapter is self-contained, with its own introduction, conclusion, notes, and appendices as needed. For this reason each chapter can be read independently from the others. A common bibliography is collected at the end of the thesis. A working paper has been extracted from each chapter: Brock, Dindo, and Hommes (2006) is based on Chapter 2, Dindo and Tuinstra (2006) on Chapter 3 and Diks and Dindo (2006) on Chapter 4. This section briefly discusses each chapter and the mathematical tools used for their analysis.

1.2.1 Deliberation costs in a cobweb model

In Chapter 2, we use the classic example of the cobweb model (see e.g. Ezekiel, 1938) to investigate the impact of endogenizing agents' choices between a costly sophisticated prediction rule, such as rational expectations, and a cheap prediction rule, such as naive expectations.

The cobweb model describes the production decisions of a producer of perishable consumption goods which take one period to be produced, such as crop or cattle. Producers want to maximize their next period's profit, which depends on the next period's market price of the good. Assume that the production technology is convex so that optimal output is an increasing function of the agent's prediction of market price. Also assume, as usual, that the aggregate demand function is downward sloping. In this case market equilibrium implies that a high (low) supply leads to a low (high) market price. Summarizing, high (low) expected prices result in a high (low) supply which clears the demand at a low (high) realized price in this system. The characteristic of this system is that the ex-post realized price is "opposite" to the ex-ante expected price.

This example offers a typical case of an economic system with a mutual dependence between individual choices and aggregate outcome. The dependence is due to the fact that producers' expectations of prices affect realized prices. When the ex-ante expected price is equal to the average ex-post realized price, that is, when the expectational feedback is at a fixed point, we have an expectational equilibrium price.

Historically, the cobweb model is an important example, because Muth (1961) uses this framework to introduce the concept of rational expectations. Muth argues that rational agents, in order to effectively optimize realized profits, should use a rational expectations predictor. This is the same as the expectational equilibrium or the "prediction of the relevant economic theory" (Muth, 1961, p. 316). Muth shows that in a cobweb model, the rational expectations predictor is the one with the highest realized profits. In other words, rational expectations are optimal expectations. Agents using other predictors would perform worse than those using rational expectations, and would, sooner or later, be wiped out of the market. However, Muth assumes that agents can choose among all possible predictors, rational or not, at no cost. Muth ignores the presence of deliberation costs, or information gathering costs, associated with more sophisticated prediction rules.

In Chapter 2, we concentrate on the cobweb model taking into account deliberation costs. Early contributions along the same line are Conlisk (1980), Sethi and Franke (1995), and Brock and Hommes (1997). In particular, in their paper Brock and Hommes introduce the concept of *adaptive rational equilibrium dynamics (ARED)*, where agents choose between a costly rational expectations forecast and a cheap naive forecast. The fractions of agents using each of the two strategies evolve over time and are endogenously coupled to the market equilibrium price dynamics. Brock and Hommes show that when the selection pressure to switch to the more profitable strategy is high, instability and complicated chaotic price fluctuations arise. Brock and Hommes call this phenomenon a rational route to randomness.

In their setting, Brock and Hommes assume that agents are *backward looking* in the sense that strategy selection is based on *experience* measured by relative past realized profits. Implicitly this means that agents, even those employing rational expectations concerning prices, use naive expectations regarding the amount of profit earned by each of the strategies. In fact, in deciding which predictor is best at maximizing expected profit, they use today's profit as a forecast of expected profit.

In Chapter 2, we model the ARED with *forward looking* agents, that is, where strategy selection is based upon *expected* profits rather than *realized* profits. As agents' objective is to maximize expected profits, assuming that agents choose a strategy based upon expected profitability seems a natural extension of the original model by Brock and Hommes. Our aim is to investigate whether forward looking behavior dampens, fosters or eliminates price fluctuations compared to backward looking behavior. In particular we analyze whether forward looking behavior has an impact on the rational route to randomness found by Brock and Hommes.

A second contribution of Chapter 2 consists in establishing an equivalence between a heterogeneous agents model with switching between two different strategies, and a representative agent framework, where the representative agent optimally chooses from a continuum of alternative predictors. As usual, predictors differ for their cost and degree of sophistication. This analysis aims at finding a correspondence between the mechanism responsible for the updating of predictor choices in a heterogeneous agents framework, and the cost function associated with a continuum of predictors in a representative agent framework. Notice that if such a correspondence exists, price fluctuations driven by strategy switching of heterogeneous firms may as well being explained by a representative firm switching between a continuum of predictors with different characteristics.

1.2.2 Competition and coordination in participation games

Chapter 3 is devoted to an analysis of repeated interaction of a large number of boundedly rational agents that are competing for the same scarce resource, when no coordinating market institution is at work. We formalize this general interaction structure, which has been inspired by the *El Farol* game, as a participation game. We concentrate on participation games for which the payoff for participating decreases as the number of participating agents increases, so that there is negative expectational feedback. Well-known examples of participation games with negative feedback analyzed in the literature are market entry games, where firms have to decide whether to enter a market and compete, or stay out of it. Another example is given by route choice games, where a group of commuters repeatedly choose, between two routes, the fastest way from their homes to their offices.

Experimental research in this area (see e.g. Sundali, Rapoport, and Seale, 1995) has aimed at appraising to which extent the aggregate participation rate emerging from the competition of many agents can be described using the traditional tools of game theory. The evidence is mixed. At the aggregate level the time average participation rate is consistent with the symmetric Nash equilibrium. However, at the individual level agents do not learn to play that Nash equilibrium, but use different deterministic rules. In particular individuals seem to employ simple rules, such as always participate, always abstain, or participate conditionally on the outcome of previous rounds. As a result the aggregate participation rate is much more volatile than would be in the case where all agents play Nash. The computational model of Arthur (1994), where 100 heterogeneous agents are choosing among different decision rules to decide whether to participate or not, gives the same results. In fact, Arthur observes convergence of the first moment of the participation rate to the symmetric Nash equilibrium, but he obtains a higher second moment, and thus a more volatile participation rate series.

Our aim is to obtain a simple analytic model that can replicate the main experimental and computational findings in the area of participation games with negative feedback. At this purpose we use the same model building guidelines as Brock and Hommes (1997) and Droste, Hommes, and Tuinstra (2002). We set up an analytic model with heterogeneous boundedly rational agents choosing between simple rules. Fractions of agents using each rule are endogenous and evolve according to the past performance of each rule as described by the replicator dynamics. We concentrate on the evolutionary dynamics produced by the competition between different deterministic rules that condition the participation decision on the outcome of the previous rounds. We characterize the resulting participation rate dynamics as the number of players increases. The interaction of a large number of players leads to complicated participation rate patterns and our aim is to check whether the average participation rate along these patterns is consistent with the symmetric Nash equilibrium and the existing experimental and computational evidence. We also investigate how the aggregation of agents' interaction is affected by the presence of agents choosing rules that try to exploit the linear autocorrelation structure of the past participation rates. Agents who try to exploit past inefficiencies of the aggregate to improve their performances, act as some kind of arbitrageurs and we investigate how their behavior affects the properties of the system.

1.2.3 Informational differences in an asset market

Asset markets, involving an extremely large number of investors of different characteristics, are a plausible context for modeling the interaction of heterogeneous boundedly rational agents. The failure of the traditional representative rational agent framework to replicate properties of asset returns, such as excess volatility, clustered volatility, correlations of returns, persistent deviations from fundamental values (see e.g. Shiller, 1989 and Fama and French, 1988a, 1988b), explains why most of the research in the area of bounded rationality and heterogeneity has been pursued in the context of financial markets as surveyed in Hommes (2006), LeBaron (2006) and Kirman (2006).

In Chapter 4 we study a market for a financial asset populated by boundedly rational agents and we concentrate on the role of informational differences. The starting point is an asset pricing model in which agents can choose among two different degrees of information on fundamentals. At the same time agents are also learning the growth rate of the dividend generating process. An inherent feature of our model is that it contains two important benchmarks as special cases. When both informational differences and learning are both discarded, our results coincide with those of the classical Gordon model (see.g. Gordon, 1962). When only informational differences are discarded our model coincides with the one of Barsky and De Long (1993).

After developing and analyzing the full model, we investigate the extent to which our model is able to explain empirical properties of asset prices. In particular we aim at offering theoretical support to the empirical evidence that the log price of a financial asset is the sum of a persistent component and a nonlinear temporary component, which switches between two different regimes. The empirical evidence for this so-called nonlinear mean reversion is documented e.g. by Gallagher and Taylor (2001) and Manzan (2003).

Chapter 4 is also closely related to the work on informational efficiency by Grossman and Stiglitz (1980). They investigate whether the price is informationally efficient in a repeated market for a single period living asset, in which agents can decide between two different degrees of information about the value of the asset return at the end of the period. They assume that both informed and uninformed agents are rational. In a financial market where agents face informational differences, the use of rational expectations poses puzzling consequences. In fact, when the information costs are positive, if agents had rational expectations the price would fully reveal the available information about future dividends and nobody would pay for information anymore. This implies that the fraction of informed agents would go to zero. However, in the limit the price would not contain information about the dividend anymore, so that it would pay to buy information again. The absence of a rational expectation equilibrium has been referred to as the Grossman-Stiglitz paradox. In a framework with rational agents one needs two sources of noise to solve the paradox. One source is the presence of noise traders who provide liquidity to the market, the other is a noisy dividend signal for the informed agents. In this case, the model of Grossman and Stiglitz leads to a *static equilibrium degree of disequilibrium*, where agents' fractions and price distribution are constant over time and a function of the exogenous noise parameters.

In Chapter 4 we analyze the case where each agent can decide whether or not to be informed about next period's dividend, but we relax the assumption of rationality. We also endogenize the dynamics of the fraction of agents choosing to buy costly information or to extract information about future dividends from the price. We investigate whether the interaction of boundedly rational agents can offer a different solution to the Grossman-Stiglitz paradox. We argue that the interaction of boundedly rational agents, triggered by informational differences, can act as a source of *endogenous noise* to the price dynamics and no other source of noise needs to be added to the system to obtain a well-defined price and well-defined fractions. Our dynamic approach aims at offering a *dynamic equilibrium degree of disequilibrium*, in contrast with the static case of Grossman and Stiglitz. Our dynamic case is strictly connected to the endogenous noise created by agents' switching between being informed and free riding on public information.

1.2.4 Computational tools for nonlinear dynamics

Whereas in economic models with rational agents there is an emphasis on equilibria, in models with bounded rationality and heterogeneity there is an emphasis on dynamics. Since most of these dynamical systems are nonlinear, the theory of nonlinear dynamical systems is an important tool of analysis. For this reason nonlinear dynamics has become a widely used instrument in recent years. Day (1994), Gandolfo (1997), and Medio and Lines (2001) are, among others, introductory textbooks with a particular emphasis on nonlinear economic dynamics. Mathematical textbooks are e.g. Guckenheimer and Holmes (1991) and Wiggins (1990). For the convenience of the reader, we here include a brief discussion of some tools from nonlinear dynamics which will be used in our study.

In this thesis, we model agents' interaction and decision making taking place at discrete times separated by a conventional time unit, called one period, and we focus on discrete time dynamical systems. Typically, one of our model equations is the expectational feedback map related to the equilibrium pricing condition, and the others are the updating rules for the fractions of agents using different decision strategies. Whereas the equilibrium pricing equation is linear in the fractions of agents, the fact that these fractions are endogenously determined yields strong nonlinearities.

Occasionally, we investigate the effect of random exogenous shocks upon the dynamics. In this case, we will refer to the original, noise-free, system as the deterministic skeleton. Whereas the impact of noise in linear systems is well understood, results concerning the effect of noise on the local and global dynamics of nonlinear systems is limited, and therefore we have to rely on simulations.

Once a system of difference equations is derived from agents' decision making and from fractions' evolution, we typically proceed as follows. First, we search for the steady state(s) of the system. The steady state(s) of our systems typically corresponds to rational behavior. After a steady state is detected, we use local stability analysis to specify for which parameters values, the interaction and adaptation of boundedly rational agents converges to it. If a steady state is unstable, we continue with the analysis of the global dynamics and investigate the occurrence of periodic or complicated chaotic fluctuations. Whereas the dynamics in linear systems can only converge to a steady state or diverge to infinity (except hairline cases), the dynamics of nonlinear systems is richer. In our examples, the dynamics can be characterized either by convergence to a stable cycle or by irregular fluctuations and is always bounded. In particular the development of chaos theory has pointed out that deterministic dynamical systems can generate erratic time series with ongoing fluctuations whose patterns resemble those of random time series. We will encounter many of these time series in the following chapters.

In this thesis, when possible, the global dynamics is characterized analytically. However, often this is difficult or impossible, and we have to use computational tools. A useful numerical tool for detecting changes in the long run dynamics as one parameter of the model changes, is the bifurcation diagram. In a bifurcation diagram a parameter is varied and, for a grid of parameter values, the system of equations is simulated and the resulting long run behavior plotted. In such diagrams one can see that for some parameter values the state variable, say the price, converges to an equilibrium value, whereas for other parameter values the state variable oscillates along a two cycle, or follows a more complicated path. An example of a bifurcation diagram is given in the left panel of Figure 1.1, which is taken from the analysis of participation games in Chapter 3. The horizontal axes indicates the parameter N, which gives the number of agents playing the game. The vertical axes represents the participation rate, i.e. the fraction of the entire population that decides to go to the bar. The bifurcation diagrams allows to compare the dynamics of the system for different values of the parameter N. One can notice that when N is small, e.g. $N = 50$, the long run behavior of the system convergences to a participation rate of 0.5. As the number of agents N increases, the long run behavior

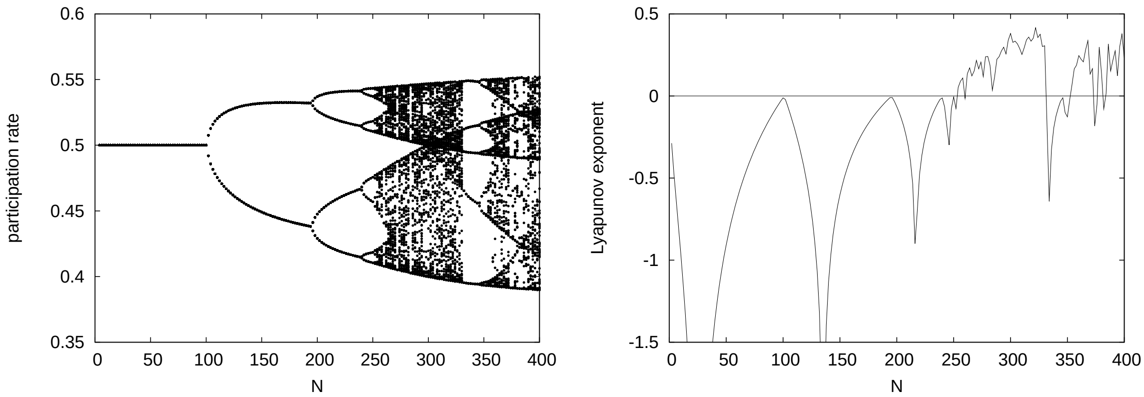

Figure 1.1: *Simulation of the dynamical systems analyzed in Chapter 3. Left panel: the bifurcation diagram with respect to the number of participants N gives the long run behavior of the participation rate as N changes. Right panel: Lyapunov exponents for different values of N. Parameters values are as in Figure 3.6. For every value of N,* 100 *iterations are used after a transient period of* 100.

of the interaction is more complicated. For example, when $N = 150$, the participation rate is oscillating between two different values. As N increases further other bifurcations occur and the participation rates follows cycles of period four, eight ... For future reference, this particular bifurcation structure with period of the cycle doubling step by step is called period-doubling bifurcation route to chaos. The bifurcations continue until we reach a point, $N \approx 250$, where many different values between 0.4 and 0.55 are possible.

A useful tool for analyzing whether the black region of the bifurcation diagram corresponds to a cycle of very large period or to more complex chaotic behavior is the largest Lyapunov exponent plot. Lyapunov exponents are used to characterize sensitive dependence on initial conditions, that is, whether a small change of the initial condition can lead to a large change of the realized state variable. When the system has a positive Lyapunov exponent there is sensitive dependence on initial conditions and the dynamics is chaotic. The right panel of Figure 1.1 gives an example of a Lyapunov exponent plot for the parameter corresponding to the bifurcation diagram. The horizontal axes indicates the parameter N. The vertical axes represents the Lyapunov exponent. When the Lyapunov exponent is negative, e.g. for $N = 50$ or $N = 150$, the system converges to a regular attractor, and, in fact, the bifurcation digram shows long run convergence to a stable cycle. When the Lyapunov exponent is positive, e.g. $N \approx 300$, the corresponding long run behavior is instead chaotic as suggested by the bifurcation diagram at the same value of N. These numerical tools will be used throughout the thesis to investigate global dynamics.

1.3 Negative expectational feedback

An interesting feature of the different economic frameworks analyzed in Chapters $2-4$ is that they all share the same type of feedback between agents' expectation of a variable, e.g. price, and its realization. In general, one can classify this expectational feedback structure as negative when positive (negative) deviations of the expectations from the expectational equilibrium price result in negative (positive) deviations of the realized price from the expectational equilibrium price. Conversely, the expectational feedback is positive when positive (negative) deviations of the expected price from the expectational feedback equilibrium result in positive (negative) deviations of the realized price from the expectational feedback equilibrium. When the expectational feedback map is differentiable, negative (positive) feedback corresponds to a negative (positive) first order derivative of the map around the expectational equilibrium. It turns out that all our examples show a *negative expectational feedback* structure.

The cobweb model of Chapter 2 is perhaps the best known example of a system with negative expectational feedback. If a producer expects a high (low) price, his optimal decision is to produce a high (low) quantity that will clear the market at a low (high) price, "opposite" to the producer's prediction. Participation games investigated in Chapter 3 also have this characteristic. Consider the *El Farol* bar problem as a concrete example. If most (few) agents believe that many agents will go to the bar, few (most) will show up, contradicting the majority belief. Finally, financial markets where agents have informational differences, as studied in Chapter 4, are also systems with negative expectational feedback. In fact when many agents believe that the information is valuable and buy it, an investor is better off not buying information as the information will be revealed by prices. On the other hand if nobody buys information, believing that it is better to extract the information from prices, an investor is better off buying information as there are so few informed agents that prices do not accurately reveal information. Our three models share the same expectational feedback structure because they all stem from essentially the same economic framework: repeated competition for a limited resource. This limited resource is the demand for crop or cattle in the cobweb model, the number of seats in the *El Farol* game, and the information about profitability of the listed firms in a financial market.

The opposite type of expectational feedback, positive feedback, is characterized by consistency of actions and beliefs. Coordination games are one example of positive feedback since when most agents think other player are going to play a certain action they play the same action thus creating consistency between beliefs and realizations. In general, probably both types of feedback play a role. For example, real financial markets have a negative feedback component due to informational differences, and a positive feedback component due to the fact that agents' demand is increasing in the expected (or anticipated) prices. Fashion cycles are also systems with mixed feedbacks, with agents copying each other at certain stages but also moving to different products at other moments.

A pair of concepts related to the expectational feedback are "strategic substitutability" and "strategic complementarity". These concepts were first developed in studies of firm interactions in Bulow, Geanakoplos, and Klemperer (1985), but later extended to the interaction of economic agents with bounded rationality (see e.g. Haltiwanger and Waldman, 1985). Strategies are substitutes if agents have an incentive to do the opposite of what most other players are doing as happens to be the case in systems with negative feedback. Strategies are complementary if agents have an incentive to imitate each other as in systems with positive feedback.

Recent experimental studies as reported by Camerer and Fehr (2006) and Heemeijer, Hommes, Sonnemans, and Tuinstra (2006) argue that convergence to a rational expectations equilibrium is more likely in economic systems with negative feedback and less likely in economic systems with positive feedback. Theoretical results of this thesis confirm, for systems with negative feedback, that overall convergence to the rational equilibrium is on average correct. Nevertheless, we also show that boundedly rational agents' interaction and adaptation trigger ongoing fluctuations around such an equilibrium. This is consistent with other experiments of systems with negative feedbacks as reported in e.g. Hommes, Sonnemans, Tuinstra, and Van de Velden (2007) for a cobweb setting, in e.g. Sundali, Rapoport, and Seale (1995) for market entry games and in e.g. Selten, Chmura, Pitz, Kube, and Schreckenberg (2006) for route choice games. This is also consistent with excess volatility in financial markets summarized in Shiller (1989). In general, these endogenous fluctuations can be characterized as irregular cycles along which rules perform better than others in different periods of time, but no rule is "dominating" the scene for every period. In this respect our results also support the observation that within systems with negative feedback incentives work in the direction of heterogeneity, that is, agents are better off if they do not imitate each other. In fact, such incentives explain why the persistence of heterogeneity, and consequently of endogenous fluctuations, is a robust characteristic of our models.

Chapter 2

Deliberation costs in a cobweb model

2.1 Introduction

The purpose of this chapter is to study a "toy model" of a compromise economy where it is costly to possess rational expectations but less costly or free to possess "simpler" expectations. Our key objective is to study prices and expectations dynamics in such an economy and to address whether switching between two predictors may or may not lead to coordination and convergence to self-fulfilling expectations. This debate is important, because instability of adaptive learning, in this case switching between predictors, may explain why markets exhibit excess volatility, i.e. are more volatile than justified by underlying economic fundamentals.

If it were very costly to purchase rational expectations our model economy would act like a temporary general equilibrium economy. Grandmont (1982) surveys the field of temporary general equilibrium theory. In this theory, at each date t, agents form expectations for relevant quantities at date $t+1$, and given these expectations, agents optimize to produce demand functions at date t. Markets then clear at date t, producing equilibrium prices and quantities at date t. At date $t+1$ the process is repeated before going into date $t+2$. A sequence of prices, quantities, and expectations, is produced. In principle, there could be feedback from equilibrium prices and quantities into expectations but the expectations need not be rational. At each date t, the actions of the agents are coordinated only by the price system. Unlike rational expectations the plans for the future which are made by the agents are not coordinated. As Grandmont (1982, p. 887) puts it, *"The aim of temporary equilibrium theory is to study the interaction through markets of different individuals in a given period, and to analyze the behavior over time of the sequence of these equilibria."*.

Instead, if rational expectations were cheap our economy would act more like a rational expectations economy. While it would be an excellent research project to study such an economy with endogenous rational choice of the "level of rationality" in a general equilibrium system at the level of generality of Grandmont (1982) and Radner (1982), it turns out to be challenging enough to study the impact of this extra level of dynamics of information choice in a simple cobweb economy as we shall do here.

Brock and Hommes (1997), BH henceforth, introduce the concept of Adaptive Rational Equilibrium Dynamics (ARED), which is an endogenous coupling between the selection of expectations rules and market equilibrium dynamics. In the ARED, the consistency requirement that has to be imposed upon learning rules is evolutionary selection of strategies, that is, agents switch to rules that have performed well in the recent past BH (1998) apply this evolutionary switching mechanism to an asset pricing model with heterogeneous beliefs; see Hommes (2006) for an extensive survey of heterogeneous agent models in economics and finance.

Another type of complicated "learning equilibrium" in which forecasting mistakes are "self-fulfilling" is the *consistent expectations equilibrium (CEE)*, introduced by Hommes (1998). In a CEE, agents use a simple linear forecasting rule in an unknown non-linear economy. A CEE arises when the sample average and sample autocorrelations of the nonlinear implied law of motion coincide with the corresponding linear belief. One possibility are chaotic CEE with chaotic price fluctuations with sample average and sample autocorrelations exactly corresponding to a stochastic AR(1) process. See also Bullard (1994), Schönhofer (1999) and Tuinstra (2003) for similar complicated learning equilibria.

As a simple illustration of the ARED concept, BH (1997) consider a cobweb model where agents can choose between two predictors: either a cheap naive predictor, equal to last observed price, or an expensive rational expectations (perfect foresight) predictor. In order to choose between the two predictors, agents compare their performance as measured by relative past *realized* profit. One of the main results in BH is that, when the selection pressure to switch between a costly sophisticated rational and cheap simple naive strategy is high, the price dynamics becomes *locally unstable* around the steady state given by the rational expectation equilibrium. As prices diverge, errors from the simple, naive strategy increase and it becomes worthwhile to switch to the costly rational strategy, pushing prices back towards the steady state. This interaction between a "close to the steady state destabilizing force" and a "far from the steady state stabilizing force" leads to complicated, chaotic price fluctuations. These complicated "learning equilibria" are driven by past realized net profits of the expectation strategies. More precisely, in the original BH (1997) model strategy selection is given by a discrete choice model with a performance or fitness measure based upon past realized profits from earlier choices of predictors. Strategy selection is thus based on "experience" or "regret", and agents tend to switch to strategies that have performed well in the (recent) past. In the original BH

(1997) model both the naive and the rational agents are *backward looking* with respect to the choice of their prediction strategy but *forward looking* in their production decision because they want to maximize expected profits. In fact, these agents implicitly use past profits as a proxy for expected future profits. One might argue that sophisticated agents should realize that other agents are making their choices in this backward looking manner and will learn to use a more sophisticated predictor selection strategy attempting to exploit the backward looking behavior. This reasoning raises the obvious question: *Will the BH (1997) instability results vanish under a concept of sophisticated forward looking predictor strategy selection*?

In this chapter, we investigate ARED with agents who are *forward looking* concerning their strategy selection. For this purpose we reconsider the same cobweb model as BH (1997), in the case where agents evaluate the different predictors based upon their *expected* profit instead of their realized profits. As one might expect in hindsight, forward looking behavior dampens some of the instabilities uncovered by the original BH (1997) work. The tendency to overshoot when "selection pressure" is high is dampened. Indeed in one case the erratic dynamics is dampened down to a stable 2-cycle. In other cases however, small amplitude chaotic price fluctuations persist.

Our setup is also related to the concept of quantal response equilibrium (QRE) introduced by McKelvey and Palfrey (1995, 1998). They also use a discrete choice model for strategy selection in a game theoretic setting, with expected payoff as the performance measure. A similar approach is used by Camerer, Ho, and Chong (2002), who study repeated games with various levels of rationality. An important difference between these game theoretic settings and BH (1997) is that in the latter case strategy selection is coupled to the dynamics of an endogenous variable, say the market price, whose realization affects the performance of all strategies.

A second contribution of this chapter is to formulate a *representative agent* version of the model, where the *deliberation cost* of more sophisticated strategies are endogenized. This approach is inspired by Simon (1955, 1957) and more recently by Evans and Ramey (1992) and especially by Dudek (2004). We formulate hybrid models based on Brock and Hommes (1997) and Dudek (2004), where a representative agent chooses optimally among predictors of different quality, where each predictor is purchased at a cost which is increasing and convex in its quality. We establish a close link between the representative agent optimizing between the benefits of sophisticated prediction rules and deliberation costs and a heterogeneous agent framework with switching of strategies. See Kirman (1992) for a critique upon the representative agent approach in economics and Hommes (2006) and LeBaron (2006) for surveys of heterogeneous agent modeling.

The chapter is organized as follows. Section 2.2 reviews some facts from BH (1997) for comparison purposes. Section 2.3 introduces forward looking behavior with respect to strategy selection. Two different cases will be discussed, one where agents have perfect foresight regarding expected profits for strategy selection and one where agents make a boundedly rational estimate of expected profits. Section 2.4 introduces a representative agent who weighs expected profits versus deliberation costs. Section 5 concludes and all proofs are contained in the Appendix to this chapter.

2.2 The BH model

Following BH (1997), we recall some key features of the ARED in the simple economic setting of the cobweb model with selection of forecasting rules based upon past realized profits. The cobweb model describes price fluctuations in a competitive market for a non-storable good which takes one period to be produced. We call p_{t+1} the price of the good at time $t+1$ and p^e_{t+1} a producer expectation at time t of the market price at time $t+1$. The demand at time $t+1$, $D(p_{t+1})$, is a linearly decreasing function:

$$D(p_{t+1}) = A - Bp_{t+1}, \qquad A > 0, B > 0. \tag{2.1}$$

The supply at time $t+1$, S, is an increasing function of the producers' forecasts at time t of the price at time $t+1$. More specifically, for every producer, the supply curve S is derived from expected profit maximization:

$$S(p^e_{t+1}) = \text{Argmax}_x(xp^e_{t+1} - c(x)), \tag{2.2}$$

where x is the amount of goods he decides to produce at time t and $c(x)$ is the production cost function. Since the decision is taken at time t, but the price is realized only at time $t+1$, a producer have to condition their decision on their expected price p^e_{t+1}. Taking a quadric cost function

$$c(x) = \frac{x^2}{2b}, \qquad b > 0 \tag{2.3}$$

gives the linear supply curve:[1]

$$S(p^e_{t+1}) = bp^e_{t+1}. \tag{2.4}$$

We assume that agents (producers) can choose between two types of predictors, $p^{e,1}_{t+1}$ and $p^{e,2}_{t+1}$. Let n^1_t and n^2_t be the fractions of agents choosing at the end of time t (or at the beginning of time $t+1$) respectively predictor 1 and predictor 2. The market clearing

[1]The general case with non-linear demand and non-linear supply is investigated in Goeree and Hommes (2000). Since it leads to similar results we restrict our analysis to the linear case.

equation at time $t+1$ is given by:

$$D(p_{t+1}) = n_t^1 S(p_{t+1}^{e,1}) + n_t^2 S(p_{t+1}^{e,2}). \tag{2.5}$$

We now have to specify how agents choose between the two prediction rules. BH (1997) assume that agents are "comparing" *past* realized profits. At time t, the realized profit of strategy i is a function of both the realized price at time t, p_t, and the predictor $p_t^{e,i}$, $i = 1, 2$, used at time $t-1$. The price p_t determines the price at which the goods are sold, while the predictor $p_t^{e,i}$ determines the amount of good produced by agent i. As a result, the *realized profit* at time t under expectations scheme i is given by

$$\pi_t^i = \pi(p_t, p_t^{e,i}) = p_t S(p_t^{e,i}) - c(S(p_t^{e,i})) = p_t b p_t^{e,i} - \frac{(b p_t^{e,i})^2}{2b} = \frac{b}{2} p_t^{e,i}(2p_t - p_t^{e,i}). \tag{2.6}$$

We consider the realistic case where one predictor, say $p^{e,1}$, is more sophisticated or of higher "quality" than the other predictor $p^{e,2}$, which is just an easy-to-use rule of thumb. The sophisticated predictor however is more "expensive" than the simple predictor because it is more difficult to compute, deliberation cost, or requires more information gathering. We call $C \geq 0$ the net costs for obtaining the sophisticated predictor. In this chapter we will refer to C as to deliberation or information costs.

As mentioned before, the choice of the predictor rule is based on the fitness or performance measure for the two strategies. BH (1997) take last period's net realized profits as the fitness measure, i.e.[2]

$$U_t^1 = \pi_t^1 - C,$$

$$U_t^2 = \pi_t^2,$$

$$\Delta U_t = \Delta \pi_t - C = \pi_t^1 - \pi_t^2 - C,$$

where ΔU_t is the difference in fitness and $\Delta \pi_t$ is the difference in realized profits (ex-cost for the sophisticated rule). The fraction of agents choosing predictor i at time t is given by a discrete choice (logit) model:

$$n_t^i = \frac{e^{\beta U_t^i}}{Z_t}, \qquad i = 1, 2, \tag{2.7}$$

[2]More generally, BH (1997) introduce memory in the fitness measure which is a weighted average of past realized profits. In this chapter, we focus on the simplest case with fitness determined by last period realized profit.

where β is the intensity of choice parameter and $Z_t = \sum_h e^{\beta U_t^h}$ is a normalization factor such that $n_t^1 + n_t^2 = 1$. The model (2.7) is derived from a random utility framework see e.g. McFadden (1973); see also BH (1997) for more details and Anderson, de Palma, and Thisse (1993) for many economic applications. The *intensity of choice* parameter β determines how eager agents are to switch to different strategies as it is inversely related to the variance noise level in the random utility model. In the extreme case $\beta = 0$, there is no switching at all and both fractions are 0.5, whereas in the other extreme case $\beta = \infty$, all agents choose the predictor with the highest fitness. In an evolutionary framework one can refer to β as the selection pressure. As β increases the selection pressure increases, so that more and more agents use the strategy with the higher fitness.

If we introduce $m_t = n_t^1 - n_t^2$ we can rewrite the market equilibrium equation as:

$$A - Bp_{t+1} = \frac{b}{2}\left(p_{t+1}^{e,1}(1+m_t) + p_{t+1}^{e,2}(1-m_t)\right)$$

where

$$m_{t+1} = \tanh\left(\frac{\beta}{2}\left(\Delta\pi_{t+1} - C\right)\right).$$

The framework introduced so far is still general in terms of the actual predictors used. Following BH (1997) we consider the choice between costly rational expectations and freely available naive expectations:

$$p_{t+1}^{e,1} = p_{t+1},$$

$$p_{t+1}^{e,2} = p_t.$$

The ARED in the case of a rational expectation predictor "versus" a naive predictor is then given by:

$$A - Bp_{t+1} = \tfrac{b}{2}\left(p_{t+1}(1+m_t) + p_t(1-m_t)\right), \tag{2.8}$$

$$m_{t+1} = \tanh\left(\tfrac{\beta}{2}\left[\tfrac{b}{2}(p_{t+1}-p_t)^2 - C\right]\right). \tag{2.9}$$

The timing in the ARED is important. First, the new market clearing price p_{t+1} is determined using the old fractions n_t^1 and n_t^2 (or difference in fractions m_t). Second, the new realized market price p_{t+1} is used to update and determine the new fractions n_{t+1}^1 and n_{t+1}^2 (or difference in fractions m_{t+1}). Notice that in the market clearing equation (2.8) p_{t+1} is only implicitly defined, but it can be solved explicitly easily by some map $p_{t+1} = F_{\beta,1}(p_t, m_t)$. The ARED of the cobweb model with rational versus naive expectations and strategy selection based upon realized profits is thus described by a two-dimensional system of nonlinear difference equations $(p_{t+1}, m_{t+1}) = F_\beta(p_t, m_t) = (F_{\beta,1}(p_t, m_t), F_{\beta,2}(p_t, m_t))$. We discuss the dynamics of F_β as the intensity of choice, or selection pressure, β, increases.

A straightforward computation shows that the unique steady state of the system is:

$$E = (p^*, m^*) = \left(\frac{A}{b+B}, \tanh\left(-\frac{\beta C}{2}\right)\right).$$

Notice that the steady state value of the difference in fractions, $m^* = m^*(\beta)$, depends on β. In particular when costs for rational expectations $C > 0$, as the intensity of choice β increases, the steady state fraction of rational agents, $n^{1*} = (1 + m^*)/2$, decreases and $n^{1*} \to 0$ as $\beta \to \infty$. The economic intuition is that at the steady state p^* there is no fitness advantage in buying a rational predictor, since both predictors predict the same, correct, value. As the intensity of choice increases more and more agents choose the naive predictor. The following Theorem (Theorem 3.1 in BH (1997)) summarizes the stability of the steady state.

Theorem 2.1 *Assume that the slopes of supply and demand satisfy $b/B > 1$:*

(i) When the information costs $C = 0$, the steady state $E = (p^*, 0)$ is always globally stable;

(ii) When the information costs $C > 0$, then there exists a critical value β_1 such that for $0 \leq \beta < \beta_1$ the steady state is globally stable, while for $\beta > \beta_1$ the steady state is an unstable saddle point with eigenvalues 0 and

$$\lambda(\beta) = -\frac{b(1 - m^*(\beta))}{2B + b(1 + m^*(\beta))}.$$

At the critical value β_1 the steady state value $m^*(\beta_1) = -B/b$.

(iii) When the steady state is unstable, there exists a locally unique period 2 cycle. There exists a $\beta_2 > \beta_1$ such that the period 2 cycle is stable for $\beta_1 < \beta < \beta_2$.

The assumption $b/B > 1$ means that if all agents employ naive expectations, the market will be unstable. The case with $b/B < 1$ is straightforward since it leads to convergence to the steady state for all values in the parameter space. In the ARED with $b/B > 1$, as soon as deliberation costs are positive an increase in the intensity of choice destabilizes the system, and yields an unstable saddle point steady state and the creation of a (stable) two-cycle through a period doubling bifurcation. BH (1997) show that, as the intensity of choice further increases, the two-cycle also becomes unstable and a *rational route to* *randomness*, that is, a bifurcation route to complicated chaotic price fluctuations, occurs. The mechanism responsible for generating complicated price dynamics is the interplay between a local, close to the steady state, destabilizing force and a global, far from the steady state, stabilizing force. The interplay is fostered by the evolutionary switching mechanism. Technically, it is responsible for *homoclinic bifurcations* and the presence of strange attractors in the dynamics of prices and fractions as discussed in detail in BH (1997). Recently, the same mechanism and type of bifurcations have been shown to generate complicated price fluctuations in other frameworks, for example a Cournot duopoly model in Droste, Hommes, and Tuinstra (2002) or a financial market where informed and uniformed agents coexist as in De Fontnouvelle (2000) and in Chapter 4 of this thesis.

The economic intuition of this phenomenon can be explained in the limiting case $\beta = \infty$. In this case, in each period, *all* agents choose the predictor with the highest fitness measure, i.e. past realized profits, no matter how small or big this difference is. From (2.9), for $\beta = \infty$, the difference in fractions m_t is determined by:

$$m_t = \begin{cases} +1, & \text{if } \frac{b}{2}(p_t - p_{t-1})^2 > C, \\ \\ -1, & \text{if } \frac{b}{2}(p_t - p_{t-1})^2 \leq C, \end{cases}$$

so that p_{t+1} is given by:

$$p_{t+1} = f_\infty(p_t, p_{t-1}) = \begin{cases} p^*, & \text{if } \frac{b}{2}(p_t - p_{t-1})^2 > C, \\ \\ \frac{A}{B} - \frac{b}{B}p_t, & \text{if } \frac{b}{2}(p_t - p_{t-1})^2 \leq C. \end{cases} \tag{2.10}$$

The following Theorem (Theorem 3.2 in BH (1997)) characterizes the price dynamics in this case.

Theorem 2.2 *For $\beta = \infty$, even when the market is locally unstable (i.e. $b/B > 1$) and when deliberation or information costs $C > 0$, the system always converges to the saddle point equilibrium steady state $E = (p^*, -1)$.*

The reasoning behind the proof is quite instructive and provides a simple economic intuition. For $\beta = \infty$ in each period either all agents are rational ($m = +1$) or all agents are naive ($m = -1$). Assume e.g. that all agents are naive and let the price be close to the steady state. For $m = -1$ the price will diverge from p^* due to the local instability. As long as $m = -1$, price fluctuations become bigger (in absolute value) and errors made by the naive forecast will rapidly increase, until the point where the difference between realized profits of rational and naive expectations exceeds the costs for rational expectations. At that point, since $\beta = \infty$, all agents buy the rational predictor, so that m becomes $+1$, and in the next period the price jumps immediately onto the steady state price p^*, and remains there forever.

For future comparison with the model of Section 2.3, it is instructive to compute the difference between the maximum and the minimum observable price when $\beta = \infty$. We define this maximum difference Δ^f_∞, where the superscript $f$and the subscript ∞ refer to the map f_∞ defined in (2.10).

Lemma 2.1 *When $\beta = \infty$, the difference between the maximum and the minimum observable price is* $\Delta^f_\infty = \dfrac{2b^2}{B(B+b)}\sqrt{2C/b}$.

2.3 Strategy switching with forward looking agents

In the original BH-model agents are *backward* looking in their selection of prediction strategies since their choice is based upon *past* realized profits, but they are *forward looking* in their production decision because, given their price forecast, they maximize expected profits. In this section we discuss the model with *forward looking* agents, with both strategy selection and production decisions based upon *expected* profit. In the forward looking case, at time t the fitness measure difference used for strategy selection between the rational and naive predictors becomes:

$$\Delta U^e_{t+1} = \Delta\pi^e_{t+1} - C = \pi^{e,1}_{t+1} - \pi^{e,2}_{t+1} - C, \tag{2.11}$$

where, as before, C represents the costs for the rational expectations predictor.

We will discuss two different versions of the model with forward looking agents, depending on the way agents compute their expected profit. In Subsection 2.3.1 all agents, i.e. both agents using the rational and the naive forecast, have perfect foresight on expected profits conditional on their production decision. We will refer to this case as *perfectly forward looking* agents in strategy selection. In Subsection 2.3.2 we discuss a different version of the model where agents using the rational price forecast still have perfect foresight on expected profits, whereas the agents using the naive price forecast employ their most recent observations to make a simple estimate of their expected profit. We will refer to this case as *boundedly rational forward looking* agents.

2.3.1 Perfectly forward looking agents

Consider forward looking agents and a discrete choice random utility model, with expected profit as the fitness measure in the following way. As before, given the expected price p^e_{t+1}, the optimal supply $S(p^e_{t+1}) = bp^e_{t+1}$ is derived from expected profit maximization with quadratic production cost function. There are two price forecast strategies available, the rational forecast at cost C and the freely available naive forecast. As before, market equilibrium at time $t+1$ is given by:

$$A - Bp_{t+1} = n^1_t bp_{t+1} + n^2_t bp_t. \tag{2.12}$$

The fraction of agents who adopt strategy i, n^i_t, is determined via a discrete choice model as in (2.7), but this time the fitness of each predictor is measured in terms of the expectations at time t about profits at time $t+1$:

$$n^i_t = \frac{e^{\beta U^{e,i}_{t+1}}}{Z_t}, \tag{2.13}$$

where $Z_t = \sum_h e^{\beta U^{e,h}_{t+1}}$ is the usual normalization factor. The rational predictor is evaluated according to:

$$U^{e,1}_{t+1} = \pi^{e,1}_{t+1} - C, \tag{2.14}$$

whereas the naive predictor is evaluated according to:

$$U^{e,2}_{t+1} = \pi^{e,2}_{t+1}. \tag{2.15}$$

In both cases:

$$\pi^{e,i}_{t+1} = \pi(p_{t+1}, p^{e,i}_{t+1}) = p_{t+1}S(p^{e,i}_{t+1}) - c(S(p^{e,i}_{t+1})). \tag{2.16}$$

Notice that the expected profit, $\pi^{e,i}_{t+1}$, $i = 1, 2$, depends both on the expected price $p^{e,i}_{t+1}$, which determines the production decision at time t, and on the price that clears the market at time $t+1$, p_{t+1}. We assume that the economy works "as if" agents have *perfect foresight* on this second price. This implies that whatever their production decision is, agents have *perfect foresight on expected profits.* Stated differently, given the production decisions, fractions of the strategies are determined "as if" agents compute expected profits without errors.

At this point it is useful to discuss a potential inconsistency problem in the model. If agents have perfect foresight on the expected profit, then, at time t, wouldn't they also "know" the price forecast p_{t+1}? If they could extract the perfect foresight *price forecast* from perfect foresight on *expected profits*, there would be no incentive left to pay the information costs C for the perfect price forecast.

In order to avoid this behavioral inconsistency, we assume that there is an expert manager who has "structural knowledge" about the economy enabling him to have structural rational expectations. Each agent can choose either to run the business by himself or hire an expert manager with "structural knowledge" of the economy and rational expectations. The manager promises the agent a sure net profit, after subtracting an (unknown) cost C for his service. Agents take the decision to hire or not to hire the rational expert manager by evaluating the fitness measure difference (2.11) which is computed and announced by the expert manager and available as public information. The fraction of agents that chooses to hire the rational expert manager is determined by the difference in fitness announced by the manager, according to the random utility framework. It is important that agents do not know the division between profits and costs C in the net revenue they receive from their manager, because if they did they could derive the perfect foresight forecast from public information. Notice also that the manager is "credible" in the sense that, given the production decision of the agents, the manager delivers the net profits that he announces.

We are aware of the limitations of these assumptions and we do not claim this to be a realistic description of market behavior. Rather we view this model as an interesting theoretical benchmark with strategy switching determined by rational forward looking behavior on expected profits. In this theoretical benchmark, forward looking strategy switching is not affected by any mistakes of the agents in evaluating their expected profits. In Subsection 2.3.2 we will consider the probably more realistic case where some forward looking agents make mistakes in evaluating their expected payoff.

Under the assumption of perfect foresight on expected profits for both types, the expected profits for rational respectively naive agents are given by

$$\pi^{e,1}_{t+1} = \pi(p_{t+1}, p^{e,1}_{t+1}) = p_{t+1} S(p_{t+1}) - c(S(p_{t+1})) = \frac{b}{2} p^2_{t+1}, \tag{2.17}$$

$$\pi^{e,2}_{t+1} = \pi(p_{t+1}, p^{e,2}_{t+1}) = p_{t+1} S(p_t) - c(S(p_t)) = \frac{b}{2} p_t (2p_{t+1} - p_t), \tag{2.18}$$

where we have used $p^{e,1}_{t+1} = p_{t+1}$ and $p^{e,2}_{t+1} = p_t$. Notice that the fitness difference (2.11) becomes:

$$\Delta U^e_{t+1} = \frac{b}{2}(p_{t+1} - p_t)^2 - C. \tag{2.19}$$

The fractions of the two types are determined via a discrete choice model as before, with fitness measure difference (2.19). Working again with the difference in fractions

$m_t = n_t^1 - n_t^2$, the system is given by:

$$A - Bp_{t+1} = \frac{b}{2}\left(p_{t+1}(1+m_t) + p_t(1-m_t)\right), \tag{2.20}$$

$$m_t = \tanh\left(\frac{\beta}{2}\left[\frac{b}{2}(p_{t+1}-p_t)^2 - C\right]\right). \tag{2.21}$$

We define a *Managerial Perfect Foresight Equilibrium* to be a time path of prices p_t, produced quantities x_t^i, $i = 1, 2$, fractions of agents choosing strategy i, n_t^i, $i = 1, 2$, and point expectation of profits, $\pi_{t+1}^{e,i}$, $i = 1, 2$, that, at each date t satisfies equations (2.12–2.18) with supply equals demand and fulfilled point expectations. Co-evolution of prices and (difference in) fractions is described by (2.20–2.21). We use the extra adjective, *managerial*, because our concept of perfect foresight equilibrium requires both agents having perfect foresight (i.e. fulfilled point expectations) on the amount of net revenue (which is the net profit received by the agent after deducting the manager's operating cost, C), as well as the manager generating profits consistent with perfect foresight on the date $t+1$ price p_{t+1}.

It is useful to compare the model with perfectly forward looking agents (2.20-2.21) to the original model with backward looking agents (2.8-2.9). In the forward looking case the difference in fractions m_t is simply a one period ahead version of the backward looking case. In fact, in the forward looking case the difference in fractions of the two strategies, m_t, depends upon p_t and p_{t+1}, implying that the market clearing equation (2.20) is only implicitly defined. The following result states that (2.20-2.21) translates into a well defined, explicit *one-dimensional map* $p_{t+1} = g_\beta(p_t)$ (seeFigure 2.1).

Theorem 2.3 *Given a linear demand and a linear supply curve, for any value of the information cost $C \geq 0$, the intensity of choice $\beta \geq 0$ and the initial price p_t, (2.20-2.21) determines a well defined map $p_{t+1} = g_\beta(p_t)$. That is, there exists a unique non-negative price p_{t+1} such that, at time $t+1$, either the market is in equilibrium or the excess supply is positive at $p_{t+1} = 0$. Furthermore the map g_β is continuous.*

A straightforward calculation shows that, as before, $p^* = A/(b+B)$ is always the unique fixed point of the map g_β. The corresponding fraction difference m is again given by $m^*(\beta) = \tanh(-\beta C/2)$. The following theorem characterizes the dynamics:

Theorem 2.4 *Assume that the slopes of supply and demand satisfy $b/B > 1$.*

(i) When the information costs $C = 0$, the steady state $E = p^*$ is always globally stable.

(ii) When the information costs $C > 0$, then there exists a critical value β_1 such that for $0 \leq \beta < \beta_1$ the equilibrium is globally stable, while for $\beta > \beta_1$ the equilibrium is

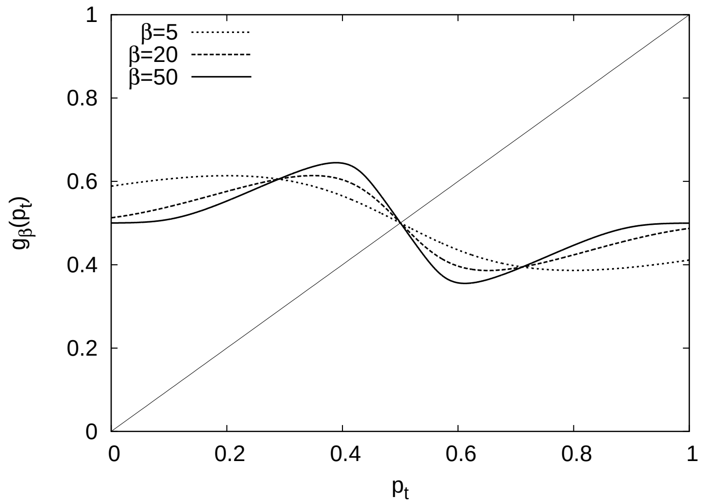

Figure 2.1: *Perfectly forward looking agents. The graph of* $p_{t+1} = g_\beta(p_t)$ *for four different values of* β*. The three points where all graphs intersect correspond to the steady state* p^* *(in the middle) and points where* $m(p_t, p_{t+1}) = 0$*, that is, where the profit difference of the two strategies is equal to the net cost* C*. Other parameters are:* $A = 1.5$, $B = 1$, $b = 2$, $C = 0.1$.

unstable with eigenvalue,

$$\lambda(\beta) = -\frac{b(1 - m^*(\beta))}{2B + b(1 + m^*(\beta))}.$$

At the critical value β_1 *the steady state value* $m^*(\beta_1) = -B/b$.

(iii) When the steady state is unstable, there exists a unique period 2 *orbit* $\{p_1, p_2\}$*. The period* 2 *orbit is globally stable for any value of* $\beta > \beta_1$.

Compare the local stability of the system with forward looking agents as specified in Theorem 2.4 with the local stability of the system with backward looking agents as specified in Theorem 2.1. Points (i) and (ii) of both Theorems imply that for small values of the switching parameter β the systems behave similarly. When the switching parameter is smaller than the primary bifurcation value β_1 the steady state is stable, while for $\beta > \beta_1$ a (stable) period 2-cycle is created. Notice that the *primary bifurcation* value β_1 is the same for both systems. For larger values of β, according to (iii), the backward looking and forward looking cases are different. In the backward looking case the 2-cycle is stable only when $\beta_1 < \beta < \beta_2$, whereas in the perfectly forward looking case the 2-cycle is stable for all values of the switching parameter $\beta > \beta_1$. The rational route to randomness (i.e. the bifurcation route to chaos) has disappeared due to the perfectly forward looking strategy switching behavior of the agents. Technically, no homoclinic bifurcation is observed and the interplay between local instability and global stability is different.

The reason that complicated dynamics does not arise is that perfectly forward looking agents do not make mistakes in the estimation of their expected profits. In contrast, backward looking agents use past realized profits to evaluate which predictor to use, and therefore may make big mistakes. In particular, these mistakes may lead to big losses for the agents who choose a naive predictor, especially when β is high. In fact in the backward looking case, when the cheap naive predictor has performed well in the last period it attracts many agents. This triggers local instability with prices deviating more and more from the steady state and bigger and bigger errors in profit estimation. When these mistakes become sufficiently large and β is high, many agents switch back to the rational predictor, thus pushing prices back very close to the steady state. But close to the steady state the cheap naive predictor works fine and is less costly than the rational predictor, so that the story repeats. When agents are perfectly forward looking concerning expected profits, big mistakes by the naive predictor are anticipated by the forward looking behavior, and therefore prices will not deviate from their steady state very far nor will they be pushed back to the steady state very closely.

An analysis of the case of an infinite switching parameter enables us to make this point precise. Remember that in the case of backward looking agents for $\beta = \infty$ the price always converges to the (locally unstable) steady state value, as stated in Theorem 2.2. As stated in the next theorem in the case of forward looking agents, in the limit case $\beta = \infty$ prices always converge to a period 2 orbit. For $\beta = \infty$ the map g_β defined by (2.20-2.21) becomes (see Figure 2.2):

$$p_{t+1} = g_\infty(p_t) = \begin{cases} p^*, & \text{if} \quad p_t \in [0, p^* - \tilde{p}] \cup [p^* + \tilde{p}, \infty), \\ p_t + \sqrt{2C/b}, & \text{if} \quad p_t \in (p^* - \tilde{p}, p^* - \delta\tilde{p}), \\ \dfrac{A}{B} - \dfrac{b}{B} p_t, & \text{if} \quad p_t \in [p^* - \delta\tilde{p}, p^* + \delta\tilde{p}], \\ p_t - \sqrt{2C/b}, & \text{if} \quad p_t \in (p^* + \delta\tilde{p}, p^* + \tilde{p}), \end{cases} \tag{2.22}$$

where $\tilde{p} = \sqrt{2C/b}$ and $\delta = B/(B+b) < 1$.

Theorem 2.5 *For $\beta = \infty$, when the market is locally unstable (i.e. $b/B > 1$) and when information costs $C > 0$ the system always converges to a period 2 orbit. A continuum of period 2 orbits exists. Furthermore, the difference between the maximum and the minimum observable price is:*

$$\Delta^g_\infty = \frac{2b}{B+b}\sqrt{\frac{2C}{b}}.$$

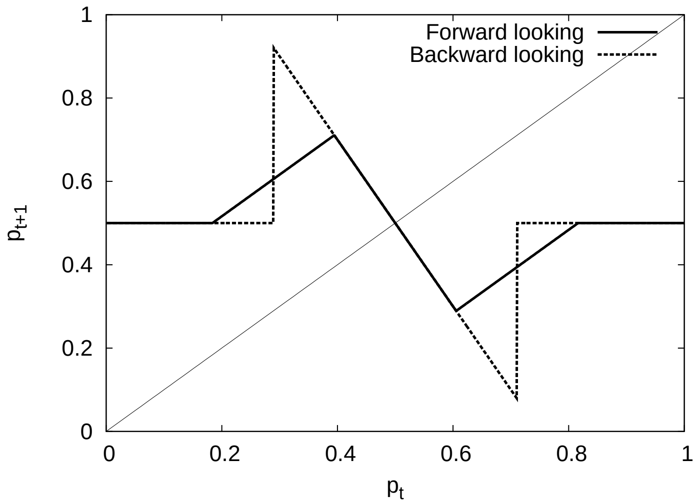

Figure 2.2: *$\beta = \infty$. Comparison of the map $p_{t+1} = g_\infty(p_t)$ in (2.22) for perfectly forward looking agents and the corresponding 1-D map in the case of backward looking agents. The latter has been obtained from the projection of the two dimensional map $p_{t+1} = f_\infty(p_t, p_{t-1})$ in (2.10) with p_{t-1} given by the inverse-image of $(p_t, -1)$ through the map $F_{\beta,1}$ (see also the Proof of Lemma 2.1 in Appendix). Other parameters: $A = 1.5$, $B = 1$, $b = 2$ $C = 0.1$.*

that is $\Delta^g_\infty = \left(\frac{B}{b}\right)\Delta^f_\infty < \Delta^f_\infty.$

The last part of Theorem 2.5 points to another difference with the original model concerning the size of fluctuations. Figure 2.2 illustrates this difference between the models with backward looking and forward looking strategy selection. The backward looking map f_∞ has a higher maximum and a lower minimum, that is $\Delta^g_\infty < \Delta^f_\infty$. A comparison of the two maps clarifies that the globally stabilizing forces are different. In the backward looking case agents make larger errors and it takes one extra time period before agents are ready to switch to the costly rational strategy. Moreover, if the errors have grown too large, all agents share the same experience and all switch to the costly rational strategy at the same time. In the forward looking case, agents anticipate large mistakes and start switching to the rational strategy already *before* the errors grow too large. Moreover, forward looking behavior prevents all agents to switch at the same point in time, but instead ensures a smooth and gradual switching to the costly rational strategy, leading to smooth dynamics and an interval of 2-cycles at some not too far away distance from the steady state with rational and naive agents co-existing (see the parts of the graph of g_∞ parallel to the diagonal). In this model, forward looking behavior based on expected profits instead of realized profits leads to smoother transitions and dampened fluctuations in the short run but also prevents the system to return (close) to the steady state price in the long run.

The presence of a regular 2-cycle for a large interval of values of the switching parameter β (e.g. Figure 2.3) raises the question whether boundedly rational agents are able to detect the regular structure from time series observations and exploit it? Stated differently, is the 2-cycle equilibrium "evolutionary stable", that is, will the cycle persist when other boundedly rational agents "invade" the system? We will not address this problem in detail here, but investigate the problem by some simple simulations, illustrated in Figure 2.4. The answer depends on whether or not additional boundedly rational rules have costs associated to them, e.g. higher costs when the rule uses more memory. In what follows, all agents, including those using the new price forecasting rules, have perfect foresight on expected profits as before. First consider the case where rational expectations at cost C versus free naive expectations leads to a stable 2-cycle. Suppose a new forecast rule $p^e_{t+1} = p_{t-1}$ (call it a period-2 rule), at costs C', $0 < C' < C$, enters the system. Along the 2-cycle this period-2 rule has in fact perfect foresight, and since it is cheaper than the rational perfect foresight rule, many agents will start using the period-2 rule. Figure 2.4 shows that the system with rational versus period-2 versus naive locks into another stable 2-cycle, with smaller amplitude because of the lower costs for the period-2 rule. The costly rational expectations rule is (almost) driven out of the market and replaced by the cheaper (but still costly) period-2 rule, but the equilibrium outcome remains a stable 2-cycle (with smaller amplitude).

Next consider the case where costly rational versus free naive expectations leads to a stable 2-cycle, and a new period-2 rule at zero costs invades the system. Figure 2.4 shows that the system with costly rational versus free naive versus free period-2 converges to a stable 3-cycle of smaller amplitude. The intuition is that along the original 2-cycle, agents switch to the period-2 rule because its forecast is the same as the rational rule but at no costs. As a consequence, the amplitude of price fluctuations decreases, which leads to smaller forecasting errors of the naive rule and thus an increase of the fraction of agents using the naive rule. With the naive and the period-2 rule both having positive fractions the system locks into a stable small amplitude 3-cycle. We can continue the story and introduce an additional type using a period-3 rule (i.e. $p_{t+1} = p_{t-2}$) entering the market. This system with 4 different rules locks into a stable 4-cycle with even smaller amplitude as shown in Figure 2.4. Similarly, adding another new type, a period-4 rule, the system locks into a stable 5-cycle with very small amplitude. Finally, adding a period-5 rule, the system stabilizes and locks into the stable steady state price. The intuition is that with more and more period-k rules at no costs, many rules gain positive weight and the system behaves as if agents use (a weighed average) of past prices as their forecast, thus stabilizing price fluctuations.

The key point of this simulation exercise is that when new boundedly rational forecasting rules "invade" the system and more memory comes at higher costs, the stable

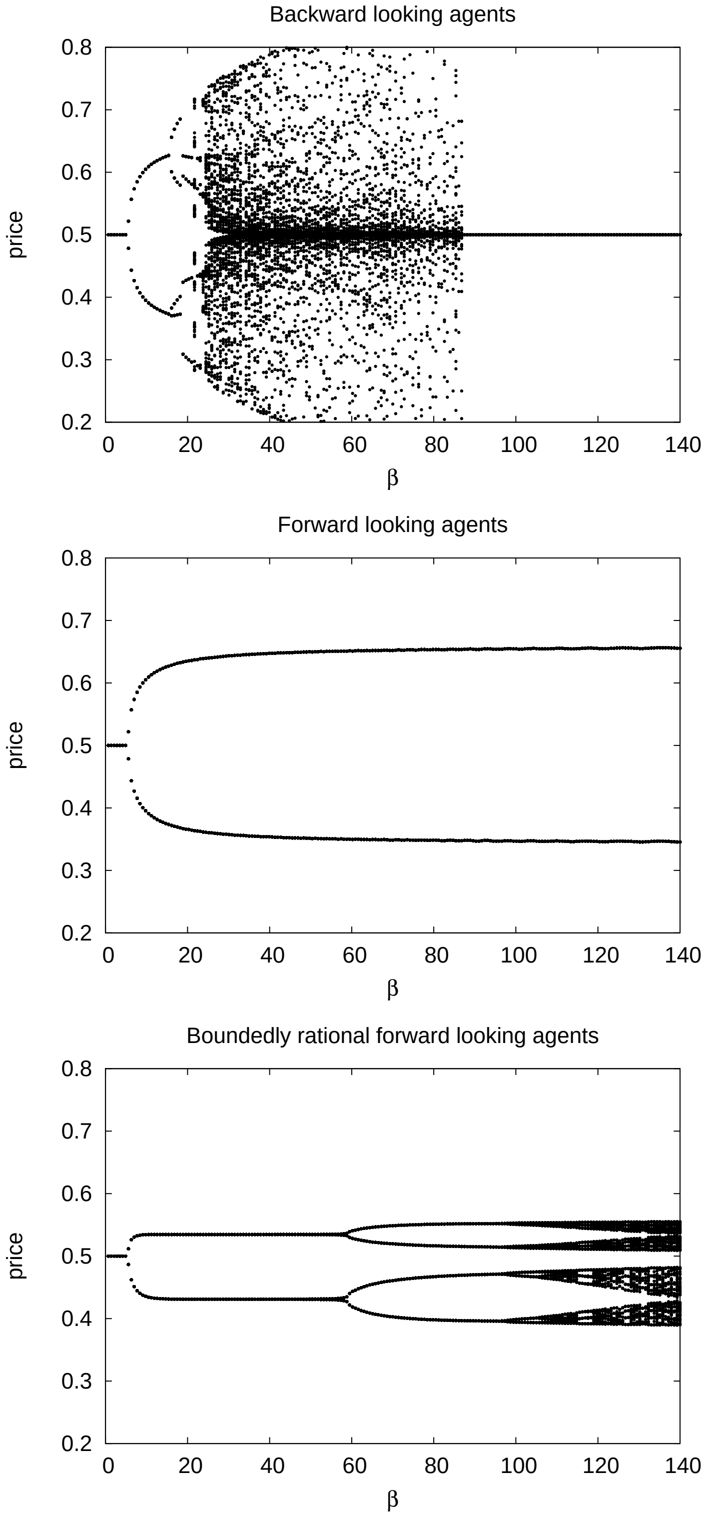

Figure 2.3: *Bifurcation diagrams with respect to β. The case of backward looking agents (top panel, Section 2.2), perfectly forward looking agents (middle panel, Subsection 2.3.1) and the case of boundedly rational forward looking agents (bottom panel, Subsection 2.3.2). The other parameters are $A = 1.5$, $B = 1.0$, $b = 2.0$ (this leads $p^* = 0.5$) and $C = 0.1$.*

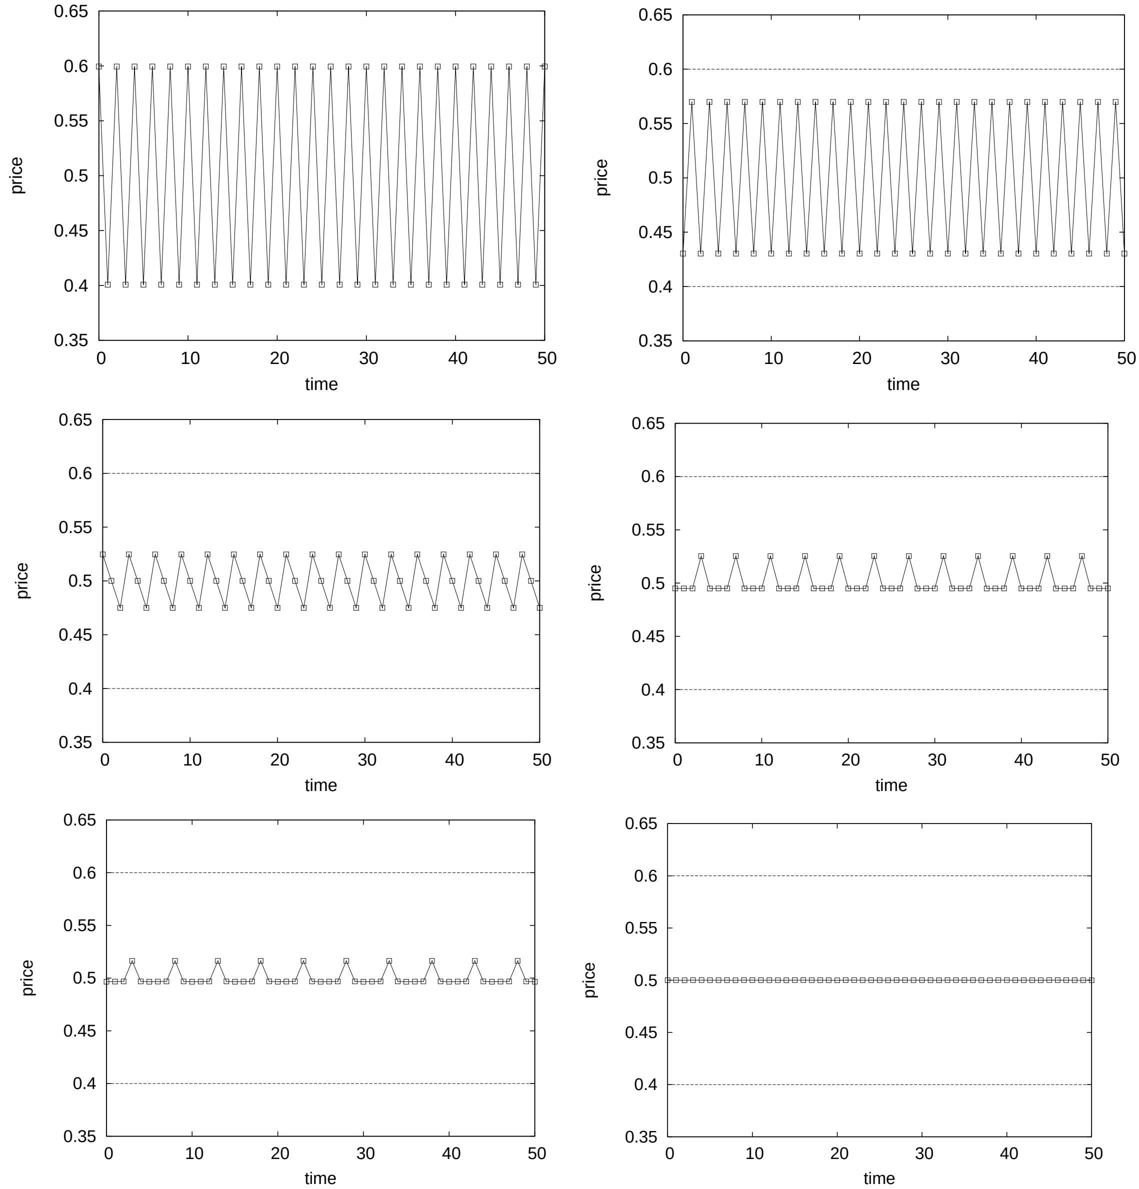

Figure 2.4: *Forward looking agents with perfect foresight. This figure illustrates the stability of the 2-cycle when naive and rational predictors are available (first row, left panel) with respect to "mutation" to other predictors (all the other panels). If a 2-period predictor emerges at a cost $C' \in (0, C)$, the 2-cycle is still stable but has a lower amplitude (first row, right panel). If a 2-period predictor emerges at no cost, the interaction of naive, rational and 2-period predictors leads to a 3-cycle (second row, left panel). In the panels that follow, at every step (from left to right and from top to bottom) a higher period predictor is introduced. This leads to creation of a 4-cycle, a 5-cycle and, finally, to the stabilization of the system. In all cases, the parameters are $A = 3$, $B = 1$, $b = 5$, $\beta = 300$, $C = 0.1$.*

2-cycle persists but has smaller amplitude. In contrast, when there are no costs associated to boundedly rational rules with more memory, price fluctuations stabilize. This situation is similar to the simulations of the El-Farol bar problem described by Arthur (1994), where "cycles are quickly arbitraged away" by cycle-detector predictors. Another example, where autocorrelation structure is arbitraged away by adding rules, is presented in Chapter 3 of this thesis. This may be a characteristic feature of systems with negative feedback from expectation to realization, as in the cobweb model.

2.3.2 Boundedly rational forward looking agents

In the previous subsection all agents were assumed to be perfectly forward looking concerning expected profits. In this subsection we back off from perfect rationality and assume that in order to choose their strategy some agents only make a boundedly rational estimate of expected profits. As in the previous subsection, we assume that each agent is faced with the choice of either hiring an expert manager to run the business or to run the business himself. In the first case, the story unfolds as in the previous subsection: by hiring an expert manager, at time t the owner is promised to receive at time $t+1$ a sure profit $U^{e,1}_{t+1}$ as computed in (2.14) and (2.17). The manager promising this last payoff uses rational expectations and keeps an (unknown) amount C of the profit of the firm for himself. As a result $U^{e,1}_{t+1} = \pi^{e,1}_{t+1} - C$. Since the agent (the owner of the firm) does not know C, he is not able to extract the rational expectations price forecast at the time he compares the two expected net profits. The computation of the fitness in the second case, when the owner decides to run the business by himself, is done in a different way than in the previous subsection. Equation (2.15) still holds but, in this case, we assume that the naive agent makes an estimate of the expected profit. We refer to this case as boundedly rational forward looking agents. We focus here on a simple case where the naive agents use their naive forecast both in the production decision and in the forecast of the expected profit. Instead of (2.18) he uses:

$$\pi^{e,2}_{t+1} = \pi(p^{e,2}_{t+1}, p^{e,2}_{t+1}) = \pi(p_t, p_t) = \frac{b}{2}(p_t)^2. \tag{2.23}$$

In Subsection 2.3.1 agents switch between prediction strategies only because one of the two is granted a higher profit. In the present case, agents are switching between the two predictors not only because naive agents make price forecasting errors, but also because they make a wrong estimate of the expected payoff they will get by using the naive predictor. As we shall see this additional error complicates the qualitative dynamics, but does not change the amplitude of the price fluctuation significantly.

Summing up, at time t, the expected fitness difference of the two predictors is:

$$\Delta U^e_{t+1} = \frac{b}{2}(p^2_{t+1} - p^2_t) - C. \tag{2.24}$$

The fraction of agents using the predictor $i = 1, 2$ at time t is as in (2.13). Using $m_t = n^1_t - n^2_t$, market equilibrium and the evolution of the fractions' difference are

$$A - Bp_{t+1} = \frac{b}{2}[p_{t+1}(1 + m_t) + p_t(1 - m_t)], \tag{2.25}$$

$$m_t = \tanh\left(\frac{\beta}{2}\left[\frac{b}{2}(p^2_{t+1} - p^2_t) - C\right]\right). \tag{2.26}$$

As in the case of perfectly forward looking agents in (2.21), the difference in fractions m_t in (2.26) depends upon p_t and p_{t+1}, so that the market clearing equation (2.25) is again implicitly defined. The following result states that (2.25-2.26) determines a well defined, explicit one-dimensional map $p_{t+1} = h_\beta(p_t)$:

Theorem 2.6 *Given linear demand and linear supply, for any value of the information cost $C \geq 0$, the intensity of choice $\beta \geq 0$ and the price p_t, the system (2.25-2.26) implies a well defined map $p_{t+1} = h_\beta(p_t)$. That is, there exists a unique minimum non-negative price p_{t+1} such that, at time $t+1$, either the market is in equilibrium or the excess supply is positive at $p_{t+1} = 0$. Sufficient conditions for continuity of the map h_β are $\beta < 2/(Ap^*)$ or , when $\beta > 2/(Ap^*)$,*

$$1 + \tanh\left(-\frac{\beta}{2}C\right) < 2\frac{B}{b}\frac{2/(Ap^*)}{\beta - 2/(Ap^*)}. \tag{2.27}$$

Notice that (2.27) is always satisfied in the limit as $\beta \to \infty$. Some graphs of the map h_β and the associated dynamics are given in Figure 2.5. The fixed point of the system is the same as before: $E = (p^*, m^*(\beta)) = (A/(b + B), tanh(-\beta C/2))$. The following theorem is the analogue of Theorems 2.1 and 2.4 in the case of forward looking boundedly rational agents:

Theorem 2.7 *Assume that the slopes of supply and demand satisfy $b/B > 1$.*

(i) When the information costs $C > 0$, then there exists a critical value β_1 such that for $0 \leq \beta < \beta_1$ the equilibrium is globally stable, while for $\beta > \beta_1$ the equilibrium is unstable with eigenvalue,

$$\lambda(\beta) = -\frac{b(1 - m^*(\beta))}{2B + b(1 + m^*(\beta))}.$$

At the critical value β_1 *the steady state value* $m^*(\beta_1) = -B/b$.

(ii) When the steady state is unstable, there exists a locally unique period 2 *orbit* $\{p_1, p_2\}$.

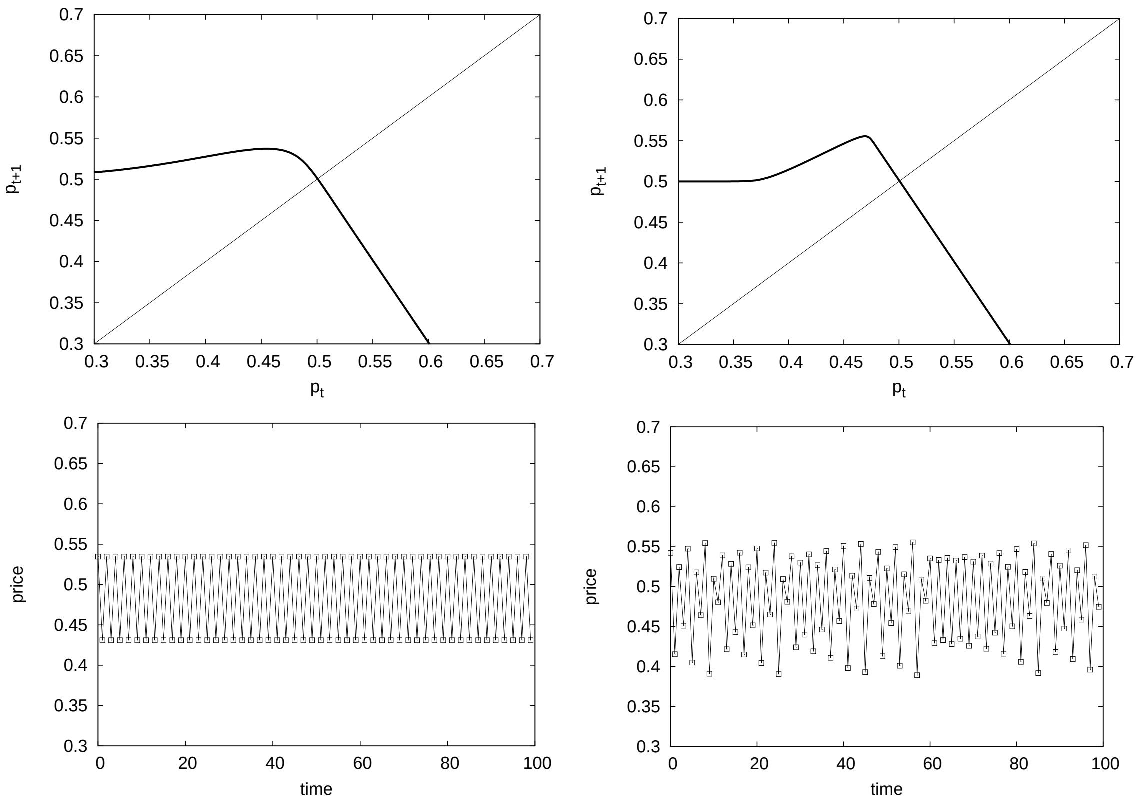

Figure 2.5: *Boundedly rational forward looking agents. On the left: map and price dynamics for* $A = 1.5$, $B = 1.0$, $b = 2.0$, $C = 0.1$ *and* $\beta = 20$. *The price dynamics converges to a 2-cycle. On the right: map and price dynamics for* $A = 1.5$, $B = 1.0$, $b = 2.0$, $C = 0.1$ *and* $\beta = 140$. *The price dynamics is chaotic.*

Notice the similarity between Theorems 2.1, 2.4 and 2.7. In all cases, the *primary bifurcation* leading to local instability is the same. After local instability sets in, the global dynamics becomes quite different however. Figure 2.3 shows the dynamics of all three cases, the original backward looking case, the perfectly forward looking case and the boundedly rational forward looking case. Both forward looking cases clearly have price fluctuations with smaller amplitude. In contrast to the perfectly forward looking case, with boundedly rational forward looking agents simulations show that a *secondary bifurcation* and a *rational route to randomness* (i.e. a bifurcation route to chaos) occur, but for higher values of β than in the original BH-model and with chaotic fluctuations with smaller amplitude. The investigation of the case $\beta = \infty$ helps us to understand the origins of this difference, the intuition behind it and its economic consequences.

For $\beta = \infty$, the difference in fractions (2.26) becomes:

$$m_{t+1} = \begin{cases} +1, & \text{if } \frac{b}{2}(p_{t+1}^2 - p_t^2) > C, \\ \\ -1, & \text{if, } \frac{b}{2}(p_{t+1}^2 - p_t^2) < C, \end{cases}$$

and, using (2.25), the price at time $t+1$ becomes

$$p_{t+1} = \begin{cases} p^*, & \text{if } p_{t+1} > \sqrt{p_t^2 + 2C/b}, \\ \\ \frac{A}{B} - \frac{b}{B}p_t, & \text{if } p_{t+1} < \sqrt{p_t^2 + 2C/b}, \\ \\ \sqrt{p_t^2 + 2C/b}, & \text{otherwise.} \end{cases}$$

If we define $\bar{p}$ such that $\sqrt{\bar{p}^2 + 2C/b} = p^*$ and $\hat{p}$ such that $A/B - (b/B)\hat{p} = \sqrt{\hat{p}^2 + 2C/b}$, the map h_∞ becomes:[3]

$$p_{t+1} = h_\infty(p_t) = \begin{cases} p^*, & \text{if } p_t < \bar{p}, \\ \\ \sqrt{p_t^2 + 2C/b}, & \text{if } p \in [\bar{p}, \hat{p}], \\ \\ \frac{A}{B} - \frac{b}{B}p_t, & \text{if } p_t > \hat{p}. \end{cases} \tag{2.28}$$

Two graphs of the map h_∞ are given in Figure 2.10 in the Appendix illustrating the proofs. The following theorem states that, if the market is sufficiently unstable, in the case with boundedly rational agents complicated dynamical behavior arises.

Theorem 2.8 *For $\beta = +\infty$, when the market is locally unstable (i.e. $b/B > 1$) and when information costs $C > 0$, there exists a value $M > 1$ (depending on p^*, C and b) such that:*

(i) when $b/B \in (1, M]$, for an interval of initial conditions, the price dynamics is bounded away from p^*.

[3]Notice that $\bar{p}$ is only defined if C is not too large. If $\bar{p}$ is well defined, so is $\hat{p}$ and $\hat{p} \in (\bar{p}, p^*)$.

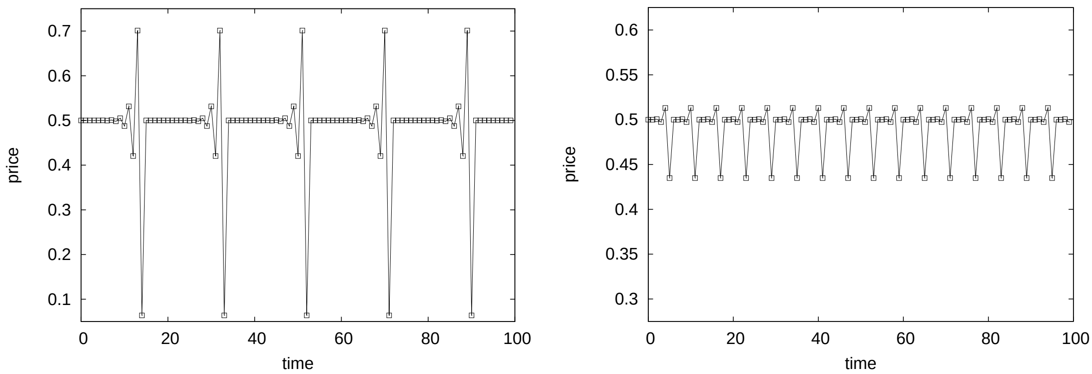

Figure 2.6: *Comparison of the effect of an homoclinic point in the case of backward looking agents (left panel) with the case of boundedly rational forward looking agents (right panel). Notice that even if the price fluctuations are qualitatively similar, they are larger and less frequent when agents are backward looking. In both cases the parameters are:* $A = 1.5$, $B = 0.85$, $b = 2.15$ *(this leads to* $p^* = 0.5$*),* $\beta = 200$ *and* $C = 0.1$.

(ii) when $b/B > M$*, there exist infinitely many homoclinic points* p*, that is, points such that* $\lim_{n\to+\infty} h_\infty^n(p) = p^*$ *and* $\lim_{n\to-\infty} h_\infty^n(p) = p^*$.

Recall from Subsection 2.3.1 that in the case with perfectly forward looking agents local instability always leads to a stable 2-cycle. In the case with boundedly rational forward looking agents more complicated dynamics arises. Mathematically this difference is explained by the fact that in the perfectly forward looking case the map g_∞ has two critical points (i.e. points where the map has a local maximum or a local minimum; see Figure 2.2), whereas in the boundedly rational forward looking case the map h_∞ has only one critical point (a local maximum) and the map is linearly decreasing for all $p > p^*$. In particular, if the local instability is strong enough (i.e. b/B is large enough), the steady state p^* of the map h_∞ has homoclinic points (i.e. points whose time path converges to p^* both forward and backward in time). As is well known, existence of homoclinic points implies complicated, chaotic dynamical behavior as illustrated in the time series of Figure 2.5 and 2.6.

From an economic viewpoint the key difference is that boundedly rational agents make errors in their estimation of expected profits. In particular, when the price $p_t > p^*$ is large, boundedly rational forward looking agents expect a high profit and therefore stick to the simple, naive strategy. This profit expectation turns out to be wrong, and the error may become so large that (almost) all agents switch to the rational strategy, pushing prices back (close) to the steady state. This interaction between local instability and global stability generates chaotic price fluctuations when the intensity of choice is high, as in case of backward looking agents. Nevertheless the size of such fluctuations remains limited as in the case of perfectly forward looking agents. Forward looking behavior thus dampens the amplitude of price fluctuations, while bounded rationality may lead to chaotic (small amplitude) fluctuations.

2.4 An optimizing representative agent

Until now we have focussed on a heterogeneous agent framework, where agents switch between a costly sophisticated and a cheap simple rule. The aim of this section is to reconcile the heterogeneous agent framework with that of a representative agent optimally choosing from a continuum of expectation rules of different quality. The representative agent outweighs the benefits of a better prediction rule against the costs of information gathering, or deliberation costs, in the spirit of Simon (1955, 1957). Our approach has been inspired by a recent chapter of Dudek (2004); see also Evans and Ramey (1992).

Consider a cobweb model with a representative supplier. As before the production decision at time t depends upon the prediction p^e_{t+1} for the price at time $t+1$. The representative supplier can optimally choose among a continuum of forecasting rules, each with different quality and costs. At time t, choosing a predictor with quality $q \in [0,1]$ corresponds to buy, at a cost $C(q)$, a signal ν_t such that

$$\nu_t = \begin{cases} p_{t+1}, & \text{with probability} \quad q, \\ \\ p_t, & \text{with probability} \quad 1-q. \end{cases} \tag{2.29}$$

The *information gathering cost function* $C(q)$ is assumed to be increasing and convex. In order to optimally choose the quality of the signal, the representative agent computes his expected net profit and maximizes it with respect to q. In general, a better signal gives a higher gross expected profit at a higher cost. At period t, the trade off between expected profits and information gathering costs sets the optimal predictor quality, q^*_t, which determines the optimal supply x_t, which, given the demand, determines the realized price and profit at time $t+1$, and so on and so forth. The purpose of this section is to investigate the equilibrium price dynamics generated by this mechanism and to compare it with the heterogeneous agent models of the previous section.

Similarly to Subsection 2.3.1, we focus on the case of a perfectly forward looking representative agent, that is, the representative agent behaves "as if" he has perfect foresight on expected profit and no systematic errors are made in estimating expected profits. Expected profits of a representative agent choosing the perfect foresight forecast with probability q and the naive forecast with probability $(1-q)$ is then given by:

$$\pi^e_{t+1}(q) = q\pi(p_{t+1}, p_{t+1}) + (1-q)\pi(p_{t+1}, p_t). \tag{2.30}$$

In this case the fitness measure for using quality q is given by:

$$U^e_{t+1}(q) = \pi^e_{t+1} - C(q). \tag{2.31}$$

Similarly to the story underlying the managerial perfect foresight equilibrium of Subsection 2.3.1, we assume that there is a continuum of managers who offer to sell a predictor of quality $q \in [0,1]$. The managers have perfect foresight on expected profits. The representative agent can not compute the expected profit by himself, but compares the fitnesses announced by the managers without knowing the value of $C(q)$ and thus being unable to derive p_{t+1} from public information.

To optimally choose the quality of the signal, the representative agent compares the fitness associated with each predictor $q \in [0,1]$. At period t, by choosing the predictor, that is, the level of q that grants him the higher fitness, the agent behaves "as if" he is solving:

$$\text{Argmax}_q\{U^e_{t+1}(q)\} = \text{Argmax}_q\{q\pi(p_{t+1}, p_{t+1}) + (1-q)\pi(p_{t+1}, p_t) - C(q)\}. \tag{2.32}$$

If this maximization problem has an interior solution, it is given by the solution of the first order condition:

$$\pi(p_{t+1}, p_{t+1}) - \pi(p_{t+1}, p_t) = \frac{b}{2}(p_{t+1} - p_t)^2 = C'(q). \tag{2.33}$$

The specific solution depends on the functional form of $C(q)$. In general, given a solution of the maximization problem at time t, which we shall call q^*_t, we can derive the implicit equation that sets the price at time $t+1$. To do that notice that, by the Law of Large Numbers, on average the representative agent produces:

$$\bar{S}(\nu_t) = S(\bar{\nu_t}) = bq^*_t p_{t+1} + b(1-q^*_t)p_t.$$

At time $t+1$ market clearing (implicitly) defines the price p_{t+1} according to:

$$A - Bp_{t+1} = bq^*_t p_{t+1} + b(1-q^*_t)p_t. \tag{2.34}$$

The price is only implicitly defined because the optimal level of q is, in general, a function of p_t and p_{t+1}, that is, $q^*_t = q^*_t(p_t, p_{t+1})$. In what follows we investigate the equilibrium price dynamics for different functional forms of the information gathering cost function $C(q)$.

2.4.1 Linear cost function

Consider first the case of a linear information gathering cost function $C(q)$. The following theorem shows that in the linear case, the price dynamics described by (2.34) exactly corresponds to the price dynamics of the heterogeneous agent model with perfectly forward looking agents in Subsection 2.3.1 when $\beta = \infty$.

Theorem 2.9 *When the marginal information gathering costs of the representative agent is constant, i.e. , $C'(q) = C$, the system (2.33-2.34) is equivalent to the heterogeneous agent system (2.20-2.21) with $\beta = \infty$. Consequently, when the market is locally unstable the system (2.33-2.34) always converges to a period 2 orbit.*

This is a first interesting correspondence between the models: the representative agent model with a linear cost function for information gathering corresponds exactly to the heterogeneous agent model with intensity of choice $\beta = \infty$. This result may be explained by observing two key features: (1) in the limit $\beta \to \infty$ all agents choose the best predictor, so that the heterogeneous agent model reduces to a single agent model (possibly switching between strategies over time); (2) in the case of a linear information gathering function in each time period typically it is optimal for the representative agent to use an extreme signal, that is, to use either $q^* = 0$ (naive) or $q^* = 1$ (rational expectations).

2.4.2 Nonlinear cost function

What is the relation between a heterogeneous agent and a representative agent framework in the case of a general, nonlinear information gathering function $C(q)$? To answer this question, it is useful to consider an explicit example. Take as cost function:

$$C(q; \alpha) = Cq^{\alpha}, \qquad \alpha > 1,$$

where $C(0) = 0$, $C(1) = C$, $C'(0) = 0$ and $C'(1) = \alpha C$. In this case the maximization problem (2.32) may have an interior solution. In order to find it, we solve the first order condition (2.33) and get:

$$q_t^* = \left[\frac{b/2(p_{t+1} - p_t)^2}{\alpha C} \right]^{\frac{1}{\alpha - 1}}.$$

Notice that as long as the cost function is convex ($\alpha > 1$), the optimal level of q is an increasing function of p_{t+1}. This turns out to be important for the uniqueness of the market equilibrium price. The market equilibrium equation (2.34) becomes:

$$A - Bp_{t+1} = \left[\frac{b/2(p_{t+1} - p_t)^2}{\alpha C}\right]^{\frac{1}{\alpha-1}} b(p_{t+1} - p_t) + bp_t. \tag{2.35}$$

Theorem 2.10 *Let $\alpha > 1$. For any fixed value of $p_t \geq 0$, the system (2.35) implies a well-defined map $p_{t+1} = r(p_t)$. That is, there exists a unique minimum non-negative price p_{t+1} such that, at time $t+1$, either the market is in equilibrium or the excess supply is positive at $p_{t+1} = 0$.*

A straightforward calculation shows that for any $\alpha \geq 1$ and $C \geq 0$ the fixed point of the system is $p^* = A/(b+B)$, as before. Figure 2.7 shows a graph of the map in the case of a quadratic cost function ($\alpha = 2$). The following theorem characterizes the price dynamics.

Theorem 2.11 *Let $C(q) = Cq^2$. When $b/B < 1$, the price dynamics converges to the steady state equilibrium p^*. When $b/B > 1$, the steady state is locally unstable and a unique globally stable 2-cycle exists.*

It turns out to be possible to generalize Theorem 2.11 for any increasing and convex cost function $C(q)$. In fact, the condition for the existence of an interior solution of (2.32) defines four parallel lines with slope one and intercepts $\sqrt{2C'(1)/b}$, $\sqrt{2C'(0)/b}$, $-\sqrt{2C'(0)/b}$, $-\sqrt{2C'(1)/b}$; see the Proof of Theorem 2.12 for details. We can use these lines to define the price dynamics in different regions, for any given $C(q)$. In some regions delimited by these lines the representative agent is choosing either $q = 0$ or $q = 1$ (see Figure 2.7). For all other values of q the dynamics is implicitly determined by the solution of the corresponding market equilibrium equation. In terms of the deviation $x_t = p_t - p^*$ from the steady state, the map becomes:

$$x_{t+1} = r(x_t) = \begin{cases} 0 & \text{if} \quad x_t \leq -\sqrt{2C'(1)/b}, \\ \text{implicit} & \text{if} \quad x_t \in \left(-\sqrt{2C'(1)/b}, -\sqrt{2C'(0)/b}\dfrac{B}{B+b}\right), \\ -\dfrac{b}{B}x_t & \text{if} \quad x_t \in \left[-\sqrt{2C'(0)/b}\dfrac{B}{B+b}, +\sqrt{2C'(0)/b}\dfrac{B}{B+b}\right], \\ \text{implicit} & \text{if} \quad x_t \in \left(+\sqrt{2C'(0)/b}\dfrac{B}{B+b}, +\sqrt{2C'(1)/b}\right), \\ 0 & \text{if} \quad x_t \geq \sqrt{2C'(1)/b}. \end{cases} \tag{2.36}$$

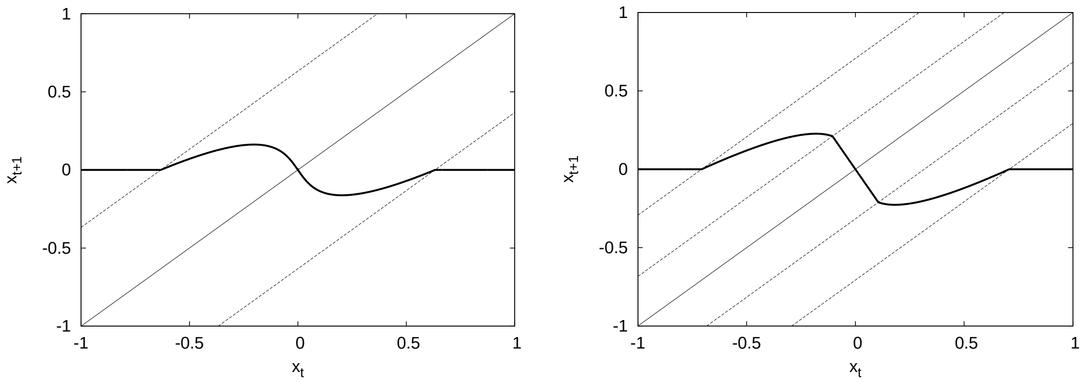

Figure 2.7: *Graphs of the map $x_{t+1} = r(x_t)$ in (2.36) for different cost functions. Left panel: $C(q) = 0.2q^2$. Right panel: $C(q) = 0.2q^2 + 0.1q$. In both case, due to the symmetry of the map around $x = 0$, the dynamics converges to a 2-cycle. The lines parallel to the diagonal delimit zones where an implicit solution exists, as specified in (2.36). Other parameters are $B = 1$ and $b = 2.0$.*

Figure 2.7 shows the graph of r in two specific cases. The next Theorem shows that only this partial knowledge of the map r is sufficient to characterize the price dynamic:

Theorem 2.12 *Let $C(q)$ be any increasing and convex information gathering cost function. When $b/B < 1$ the price dynamics converges to the stable steady state p^*; when $b/B > 1$ the steady state is locally unstable and a unique globally stable 2-cycle exists.*

Notice that according to Theorem 2.12, in the case of a perfectly forward looking representative agent concerning expected profit, in the long run the incentives to buy a perfect foresight predictor are never strong enough, no matter how strong the local instability of the market is. Stated differently, for any functional form of the cost function, either (when the market is stable) the agent always chooses a cheap naive predictor ($q = 0$) or (when the market is unstable) he switches between a naive predictor and a better, but non-perfect, predictor with $q < 1$. Before concluding, we present another interesting relationship between the representative agent and the heterogeneous agent framework of Subsection 2.3.1.

Theorem 2.13 *There exists a non-decreasing and convex information gathering cost function $C(q)$ such that the price dynamics driven by the interaction of a group of heterogeneous agents choosing between a freely available naive forecast ($q = 0$) and a costly perfect foresight forecast ($q = 1$) at constant cost C as given by (2.20-2.21) is the same as if a representative agent is operating in the market and optimally chooses a signal $q \in [0, 1]$ with cost $C(q)$ as in(2.33-2.34). For every fixed β and C the cost function $C(q)$*

is given by:

$$C(q) = \begin{cases} \displaystyle\int_{t_0}^{q} \Big(\frac{2\tanh^{-1}(2t-1)}{\beta} + C\Big)dt, & \text{if} \quad q \in (t_0, 1], \\ 0, & \text{if} \quad q \in [0, t_0]. \end{cases}$$

where $t_0 = (1 + m^*)/2$.

It can be shown that this exact correspondence between a heterogeneous agent and a representative agent framework is not only valid in the perfectly forward looking case studied in this section, but holds more generally. For example also in the case of boundedly rational forward looking agents (making mistakes in the estimation of expected profits) or in the original framework of BH (1997) (with backward looking agents and strategy switching based on realized profits). Hence, the same trade off between local instability and global stability when a group of heterogeneous agent is operating in the market exists when a representative agent optimally chooses the quality of his price predictor. This kind of theorems show that, for an outside observer, it would be impossible to distinguish between a heterogeneous agent and the corresponding representative agent economy (cf. e.g. the discussion in (Kirman, 1992)).

2.5 Conclusion

The Adaptive Rational Equilibrium dynamic can produce complicated equilibrium price dynamics, due to the interaction between a locally destabilizing force when agents use simple, cheap strategies and a far from the steady state globally stabilizing force when errors become so large that most agents switch to the costly rational forecast. In the original framework, strategy selection is driven by experience or regret, as it based upon a measure of past realized profits. In this chapter we have investigated how forward looking behavior in strategy selection may affect the co-evolving equilibrium dynamics. With forward looking behavior, the same local instability due to costly information gathering and free riding remains, but the amplitude of price fluctuations is dampened. How exactly forward looking behavior affects the globally stabilizing force depends upon how sophisticated agents are in computing expected profits. When agents are able to make a perfect forecast for expected profits, prices lock into a stable 2-cycle and the errors of the cheap naive strategy remain small enough for the population of agents never to switch completely to the costly rational forecast. If however agents can only make a boundedly rational forecast of expected profits, errors may grow big enough for (almost) all agents to switch to the costly rational strategy. In that case, the globally stabilizing force becomes strong enough to push prices back close to the (locally unstable) steady state, thus causing irregular switching between price fluctuations of low and moderate amplitude. Boundedly rational forward looking agents thus dampen the amplitude of the price oscillations but a rational route to randomness, as in the backward looking case, remains.

We also have established an equivalence relation between heterogeneous agent models with evolutionary switching of strategies and a representative agent who optimally chooses between the benefits of a high quality forecasts and the associated information gathering costs. To an outside observer it is impossible to distinguish between the heterogeneous agent model and the corresponding optimal representative agent framework.

We emphasize that we have shown these results to hold in a supply-driven commodity market, with negative expectational feedback (i.e. a high price forecast leads to high production and thus a low realized market price). An interesting question for future work is whether similar results hold for speculative asset markets. In fact, for asset markets we conjecture that forward looking behavior may actually destabilize the amplitude of price fluctuations, because of the positive expectational feedback (i.e. high expectations of future asset prices lead to increased asset demand and thus higher realized market prices). We leave this conjecture for future research.

Appendix

2.A Proofs

Proofs of Section 2.2

Proof of Lemma 2.1 The map f_∞ defined in (2.10) is two-dimensional. Define f_∞^1 as a one-dimensional projection of f_∞ according to

$$f_\infty^1(p_t) = f_\infty(p_t, (f_{\infty,1}^-)^{-1}(p_t))$$

where $f_{\infty,1}^-(\cdot) = F_{\infty,1}(\cdot, -1)$ and $F_{\infty,1}$ is the first component of the map F_β, for $\beta = \infty$, defined in Section 2.2 below equations (2.8-2.9). By applying this definition, one can easily show that

$$f_\infty^1(p_t) = \begin{cases} \frac{A}{B} - \frac{b}{B}p_t, & \text{if} \quad p_t \in [p^* - \varepsilon, p^* + \varepsilon] \\ p^*, & \text{otherwise} \end{cases}$$

where $\varepsilon = \frac{b}{b+B}\sqrt{2C/b}$. The graph of f_∞^1 is given in Figure 2.2. The lemma follows by noticing that the difference between the maximum and the minimum of the map f_∞^1 is $2\frac{b^2}{B(B+b)}\sqrt{2C/b}$. □

Proofs of Section 2.3

Proof of Theorem 2.3 From (2.20) the excess supply function $ES(p_t, p_{t+1})$ is given by:

$$ES(p_t, p_{t+1}) = \frac{b}{2}\left(p_{t+1}(1 + m(p_t, p_{t+1})) + p_t(1 - m(p_t, p_{t+1}))\right) - A + Bp_{t+1}, \qquad (2.37)$$

where

$$m(p_t, p_{t+1}) = m_t = \tanh\left(\frac{\beta}{2}\left[\frac{b}{2}(p_{t+1} - p_t)^2 - C\right]\right). \qquad (2.38)$$

Consider for a moment the function $ES(p_t, y)$. We define a function $y = \tilde{g}(p_t)$ such that $ES(p_t, y) = 0$. A straightforward computation shows that:

$$\frac{\partial ES(p_t, y)}{\partial y} = \frac{b}{2}\left(1 + m_t + \beta\frac{b}{2}(y - p_t)^2(1 - m_t)(1 + m_t)\right) + B > 0.$$

This implies that given any value of p_t we can always find a unique y such that $ES(p_t, y) = 0$. Consequently the function $y = \tilde{g}(p_t)$ is well defined and, by the implicit function theorem, the function $\tilde{g}$ is C^1. Define the map g_β as follows

$$g_\beta(p_t) = \begin{cases} \tilde{g}(p_t) & \text{if } \tilde{g}(p_t) \geq 0 \\ 0 & \text{if } \tilde{g}(p_t) < 0. \end{cases}$$

This means that when $g_\beta(p_t) = p_{t+1} > 0$, we have $ES(p_t, p_{t+1}) = 0$, whereas when $g_\beta(p_t) = p_{t+1} = 0$, there is excess supply even when the market price is $p_{t+1} = 0$. In fact $\partial ES(p_t, y)/\partial y > 0$ implies that when $\tilde{g}(p_t) < 0$, $ES(p_t, 0) > ES(p_t, \tilde{g}(p_t)) = 0$. Finally the map g is continuous but in general not differentiable at the lowest price p_t for which $g_\beta(p_t) = 0$. □

Proof of Theorem 2.4 In terms of $x_t = p_t - p^*$ the system (2.20-2.21) becomes:

$$\begin{cases} -Bx_{t+1} = \frac{b}{2}\left(x_{t+1}(1+m_t) + x_t(1-m_t)\right) \\ \\ m_t = \tanh\left(\frac{\beta}{2}\left[\frac{b}{2}(x_{t+1} - x_t)^2 - C\right]\right) \end{cases}$$

with fixed point $x^* = 0$ and corresponding $m^* = \tanh(-\beta C/2)$. Similarly as in the proof of Theorem 2.3 it follows that there exists a well defined map $x_{t+1} = \bar{g}_\beta(x_t)$, obtained from the map $p_{t+1} = g_\beta(p_t)$ by choosing p^* as the origin. To keep the notation simple we drop the bar, and write g instead of $\bar{g}$ in what follows. The following properties of the map g will be useful to prove the Theorem:

(a) g is odd, that is $g(-x) = -g(x)$, and $g(x) > 0$ when $x < 0$.

(b) When $\beta > \beta_1$ there exists a unique period 2 orbit $\{a, -a\}$, satisfying $g(g(a)) = a$ and $g(a) = -a$.

(c) $g'(0) < g'(x) < 1$, for all x.

(d) The map g has two critical points, c and $-c$. Furthermore $g'(x) < 0$ iff $x \in (-c, c)$, so that g has a local minimum at $x = c$ and a local maximum at $x = -c$.

We postpone the proof of properties $(a) - (d)$, and first use them to prove $(i) - (iii)$ in Theorem 2.4. Notice that, (a) and (d) imply that $g(c)$ is in fact a global minimum and $g(-c)$ a global maximum. Using the implicit function theorem, we get:

$$g'(x) = -\frac{\frac{\partial ES(x,y)}{\partial x}}{\frac{\partial ES(x,y)}{\partial y}}\Big|_{y=g(x)} = \frac{-b(1-\alpha)(1-m(x,y))}{b(1+m(x,y)+\alpha(1-m(x,y)))+2B}\Big|_{y=g(x)}, \tag{2.39}$$

where

$$\alpha = \frac{b}{2}\beta(y-x)^2(1+m(x,y))$$

and $ES(x, y)$ is the excess supply function defined in (2.37) and $m(x, y)$ is the difference of agents' fractions in (2.38).

Proof of *(i)*. The global stability of $x^* = 0$ when $C = 0$ follows from the fact that in this case g is a contraction. Indeed using (2.39) one gets $g'(0) = -b/(b + 2B) > -1$, so that property (c) implies $g'(x) \in (-1, 1)$ for all x.

Proof of *(ii)*. The global stability of the fixed point when $\beta < \beta_1$ follows again from the fact that g is a contraction. In fact

$$g'(0) = \lambda(\beta) = -\frac{b(1 - m(0,0))}{b(1 + m(0,0)) + 2B} > -1 \quad \text{iff} \quad \beta < \beta_1,$$

so that (c) implies $g'(x) \in (-1, 1)$. Local instability follows from the fact that $g'(0) < -1$ when $\beta > \beta_1$. The value of β_1 is determined by the condition $g'(0) = -1$ which gives $m(0,0) = m^* = -B/b$.

Proof of *(iii)*. Uniqueness of the period 2 cycle follows from (b). In order to show global stability of the 2-cycle we have to characterize the shape of the map g^2. By (d) the map g has only two critical points, $-c < 0$ (local maximum) and $c > 0$ (local minimum). From now on we concentrate on g^2 for $x > 0$, the results for $x < 0$ follow by symmetry. We look for the critical points of g^2, that is for the points where $(g^2)'(x) = g'(g(x))g'(x) = 0$. The positive critical points of g^2 are the positive critical point c of g, and points $d > 0$ such that $g(d) = -c$. We distinguish two cases.

CASE 1.: $g(c) > -c$ and $g(-c) < c$ (see Figure 2.8). In this case, because c is larger than the global maximum $g(-c)$, there is no $d > 0$ such that $g(-d) = c$. Hence c and $-c$ are the unique critical points of g^2. We claim that $(g^2)'(x) > 0$ when $x \in (-c, c)$. This follows because $x \in (-c, c)$ implies both $g'(x) < 0$ and $g(x) \in (-c, c)$ so that $(g^2)'(x) = g'(g(x))g'(x) > 0$. This, together with $g(c) > -c$, implies that at the unique point $a > 0$ for which $g(a) = -a$, we have $a < c$ and $g'(a) > 0$. When $\beta > \beta_1$, $g'(0) < -1$, so that $(g^2)'(0) > 1$. Hence, $x = a$ is the unique intersection point of $g^2(x)$ with $y = x$ and $(g^2)'(a) < 1$. $\{a, -a\}$ is therefore a locally stable 2-cycle. Since $g^2(x) > 0$ for all $x > 0$ all points (except the unstable steady state) converge and the 2-cycle is globally stable.

CASE 2.: $g(c) < -c$ and $g(-c) > c$ (see Figure 2.9). In this case there exist two other positive critical points of g^2, d_1 and d_2, $d_1 < c < d_2$, for which $g(x) = -c$. Moreover, property (b) implies that a is the unique positive intersection point of $g(x)$ with the line $y = -x$, and since $g(c) < -c$ we have $d_1 < c < a < d_2$. By symmetry also $-d_1$ and $-d_2$ are critical points of g^2 and $-d_2 < -a < -c < d_1$. Clearly $g^2(d_1)$ and $g^2(d_2)$ are local

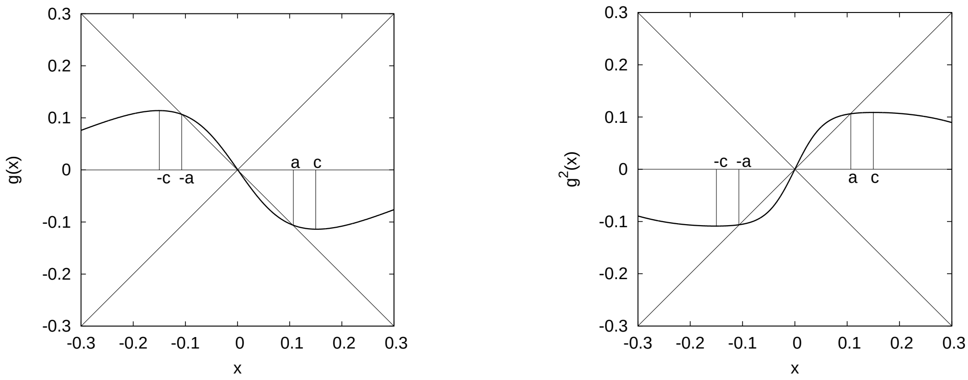

Figure 2.8: *Graph of the map g_∞ (left panel) and its second iterate g^2 (right panel) in the case when $c > a$. $A = 1.5$, $B = 1$, $b = 2$, $C = 0.1$ and $\beta = 20$.*

maxima and $g^2(c)$ is a local minimum, and using (c) and (d) we get $0 \leq (g^2)'(x) < 1$, for all $x \in [c, d_2]$. This implies that $a \in (c, d_2)$ is a locally stable fixed point of g^2. From the global shape of the graph of g^2 it follows easily that the 2-cycle is globally stable.

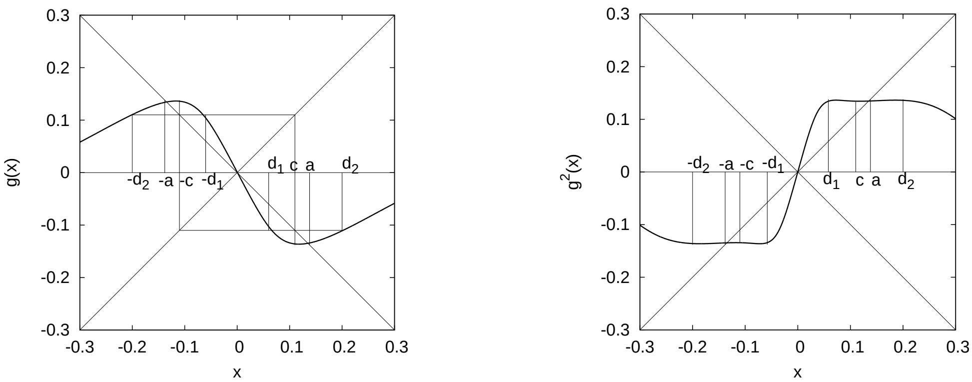

Figure 2.9: *Graph of the map g_∞ (left panel) and its second iterate g^2 (right panel) in the case when $c < a$. $A = 1.5$, $B = 1$, $b = 2$, $C = 0.1$ and $\beta = 50$.*

To conclude the proof we have to show that (a), (b), (c), (d) hold.

(a). In deviations x from the steady state the excess supply function is

$$ES(x_t, x_{t+1}) = \frac{b}{2}\left(x_{t+1}(1 + m(x_t, x_{t+1})) + x_t(1 - m(x_t, x_{t+1}))\right) + Bp_{t+1}. \quad (2.40)$$

If $ES(x, g(x)) = 0$ then also $ES(-x, -g(x)) = 0$, implying $g(-x) = -g(x)$. Moreover, when $x > 0$ ($x < 0$) the only possibility to have $ES(x, g(x)) = 0$ is $g(x) < 0$ ($g(x) < 0$).

(b). We are looking for a point a that $g(g(a)) = a$. If such a point exists, $ES(a, g(a)) = 0$ implies $ES(g(a), g(g(a))) = ES(g(a), a) = 0$ and vice versa. Hence the existence and uniqueness of a period-2 orbit $\{a, g(a)\}$ is equivalent to the existence and uniqueness of a such that $ES(a, g(a)) = ES(g(a), a) = 0$. Consider the change of variable $z = a + g(a)$, $w = g(a) - a$. We can rewrite $ES(a, g(a)) = ES((z-w)/2, (z+w)/2) = \widetilde{ES}(w, z)$. In terms of the new variables and the function $\widetilde{ES}$, the existence and uniqueness of a is equivalent to the existence and uniqueness of two points $\bar{z}$ and $\bar{w}$ such that both $\widetilde{ES}(\bar{w}, \bar{z}) = \widetilde{ES}(-\bar{w}, \bar{z}) = 0$. Writing down these two conditions as a function of the new variables one obtains:

$$\begin{cases} \bar{z} = -\frac{B+bm(\bar{w})}{B+b}\bar{w} \\ \\ \bar{z} = \frac{B+bm(\bar{w})}{B+b}\bar{w} \end{cases} \tag{2.41}$$

where, with abuse of notation, $m(w)$ is the difference of fraction in (2.38) defined as a function of $g(a) - a = w$. Notice that if a solution of (2.41) exists, then $\bar{z} = 0$ and $bm(\bar{w}) = -B$. Obviously $\bar{z} = 0$ implies $g(a) = -a$. The other equation $bm(\bar{w}) = -B$ is only possible when $\beta > \beta_1$. In fact, when $\beta < \beta_1$, $m(0) > -B/b$ so that $m(w) > m(0)$ implies $m(w) > -B/b$. When $\beta > \beta_1$, $bm(w) = -B$ has two symmetric solutions. Consequently, a unique symmetric 2-cycle $\{a, -a\}$ exists for $\beta > \beta_1$.

(c). The expression of $g'(x)$ in (2.39) can be rewritten as:

$$g'(x) = \frac{-b(1-\alpha)}{b\left(\frac{(1+m(x,y))}{(1-m(x,y))} + \alpha\right) + \frac{2B}{(1-m(x,y))}}, \tag{2.42}$$

where

$$\alpha = \frac{b}{2}\beta(y-x)^2(1+m(x,y))$$

Since the difference of fractions m is an increasing function of $(y-x)^2$, α is also an increasing function of $(y-x)^2$. Moreover $\alpha \geq 0$ and $\alpha = 0$ when $(y-x)^2 = 0$. These properties together with $m \in (-1, 1)$ imply that the denominator of $g'(x)$ is a positive increasing function of $(y-x)^2$ and that the numerator is an increasing function of $(y-x)^2$ which is negative when $(y-x)^2 = 0$. These facts imply that the minimum of $g'(x)$ is achieved when $y - x = 0$ that is when $y = x$ which implies $x = 0$. Moreover when $g'(x) > 0$, that is when $\alpha > 1$, one can easily show that the numerator is always smaller then the denominator so that $g'(x) < 1$ for all x.

(d). A critical point c satisfies $g'(c) = 0$. From (2.42), and from the facts that α is increasing in $(y-x)^2$ and $\alpha = 0$ for $(y-x)^2 = 0$, it follows that there exists a unique $(y-x)^2 = k$ such that $g' = 0$. In order to show that g has a unique critical point $c > 0$, we thus have to show that c is the unique solution of $(g(c) - c)^2 = k$. This translates into showing that $g(c) - c = -\sqrt{k}$ has a unique positive solution[4]. We claim here the more

[4]Since by (a) $g(x) < 0$ when $x > 0$, in this case the equation $g(c) - c = +\sqrt{k}$ has no positive solution.

general statement that the line $y = x - h$ and the curve $y = g(x)$ intersect only once when $x > 0$, $\forall h > 0$. Change variables from (x, y) to $(x, h = x - y)$, and consider ES as a function of the new variables. We have $ES(x, y) = ES(x, x - h) = \widehat{ES}(x, h)$. One can easily show that holds:

$$\frac{\partial \widehat{ES}(x,h)}{\partial x} = \frac{\partial (b/2((1+m(h))(x-h) + (1-m(h))x) + B(x-h))}{\partial x} = B + b > 0.$$

As a consequence we can apply the implicit function theorem and find a function t such that $\widehat{ES}(x = t(h), h) = 0$. This means that $\forall h > 0$ it exists a unique $x = t(h)$ and, as a consequence, a unique $y = x - h$ where $\widehat{ES}(x, h) = \widehat{ES}(x, x - y) = ES(x, y) = 0$. From the last expression it follows that $y = g(x)$, so that, for every given h, there exists a unique point $(x, g(x))$ with $g(x) = x - h$. From this we conlcude that $g(x) - x = -\sqrt{k}$ has a unique solution c. Property (a) implies $c > 0$. Since c is the unique positive critical point, by symmetry also $-c$ is a critical point. Being c and $-c$ the only critical points and $g'(0) < 0$ it must be that $g'(x) < 0$ iff $x \in (-c, c)$. This implies that $g(c)$ is a local minimum and $g(-c)$ a local maximum. □

Proof of Theorem 2.5 The existence of a 2-cycle can be proven along the same lines as in Theorem 2.4. Consider the explicit definition of g_∞ in (2.22). Let c and $-c$ be the critical points and $d_2 > c$ and $-d_2 < -c$ points such that $g(d_2) = -c$ and $g(-d_2) = c$ as in the proof of Theorem 2.4. A straightforward computation shows that $g^2(x) = x$ for all $x \in [-d_2, c] \cup [c, d_2]$. This implies that g has a continuum of 2-cycles. The computation of Δ^g_∞ follows directly from the definition of the map g_∞ in (2.22). □

Proof of Theorem 2.6 From (2.25) excess supply is given by

$$ES(p_t, p_{t+1}) = \frac{b}{2}[p_{t+1}(1 + m_t) + p_t(1 - m_t)] - A + Bp_{t+1} = 0,$$

where as in (2.26)

$$m_t = \tanh\left(\frac{\beta}{2}\left[\frac{b}{2}(p_{t+1}^2 - p_t^2) - C\right]\right).$$

We are looking for a function h (we drop the subscript β here) such that $ES(p_t, h(p_t)) = 0$. First, we show that there always exists a unique minimum non-negative price p_{t+1} such that, at time $t + 1$, either the market is in equilibrium or the excess supply is positive at $p_{t+1} = 0$. Second, we show that under certain parameters restrictions h is continuous.

Existence. Write $y = p_{t+1}$ and $x = p_t$, and let $p^* = A/(B + b)$ as usual and notice that

$ES(p^*, p^*) = 0$. Consider $y > x$, it can be verified that in this case:

$$\frac{\partial ES(x,y)}{\partial y} = \frac{b}{2}\left(1 + m_t + \frac{\beta}{2} by(y-x)(1-m_t)(1+m_t)\right) + B > 0. \tag{2.43}$$

Consequently we can apply the implicit function theorem and there exists a C^1 map h such that $ES(x, h(x)) = 0$. The case $y < x$ is more difficult, since $\partial ES(x,y)/\partial y$ may be zero. We can rewrite the condition $ES(x,y) = 0$ as:

$$\frac{2(A - (B+b)y)}{b(y-x)} = m(y,x) - 1 \tag{2.44}$$

where

$$m(y,x) = \tanh\left(\frac{\beta}{2}\frac{b}{2}(y^2 - x^2) - \frac{\beta}{2}C\right).$$

For $y < x$, $-2 < m - 1 < -1$ together with (2.44) imply

$$\frac{A}{B} - \frac{b}{B}x < y < \frac{2A}{2B+b} - \frac{b}{2B+b}x \tag{2.45}$$

and therefore also $x > p^*$. Eq. (2.44) and $-2 < m - 1 < -1$ also imply that when $y = 0$ there is a value $\tilde{x} \in (A/b, 2A/b)$ such that $ES(\tilde{x}, 0) = 0$. Furthermore since

$$\frac{\partial ES(x,y)}{\partial x}|_{y=0} = \frac{b}{2}\left(1 - m_t - \beta bx(y-x)(1-m_t)(1+m_t)\right)|_{y=0} > 0 \tag{2.46}$$

such a value is unique and $ES(x, 0) > 0$ for $x > \tilde{x}$, so that we can define $h(x) = 0$ for $x \in [\tilde{x}, \infty)$. Notice that for all $y \leq x$, $\partial ES(x,y)/\partial x > 0$. By the implicit function theorem there exists a C^1 function $x = s(y)$ such that $ES(s(y), y) = 0$. Clearly $s(p^*) = p^*$ and $s(0) = \tilde{x}$. Furthermore by(2.45), $s(y)$ must always be between the lines $y = A/b - b/Bx$ and $y = 2A/(2B+b) - b/(2B+b)x$. The function h we are looking for is not well defined yet since many different y values may be mapped to the same x through the map s. However, when this is the case case, we can always choose the minimum of these y values. Consequently there always exists a unique minimum non negative price $y = p_{t+1}$ such that, at time $t+1$, either the market is in equilibrium or the excess supply is positive at $p_{t+1} = 0$ so that h is well defined. Notice that $y = p_{t+1}$ is the lowest non negative price for which $ES(p_t, y) \geq 0$.

Continuity. The map h defined above may be discontinuous. A sufficient condition for continuity can be obtained by restricting parameters values such that $\partial ES(x,y)/\partial y > 0$, also for $y < x$. In order to obtain such restriction we use (2.43) and (2.44) to evaluate

$\partial ES(x,y)/\partial y > 0$ in those points (x,y) for which $ES(x,y) = 0$:

$$\frac{\partial ES(x,y)}{\partial y}|_{ES=0} = B + (1 + m(x,y))\frac{b}{2}(1 - \beta y(A - (B+b)y)) > 0. \tag{2.47}$$

Since (2.45) implies that $0 < y < p^*$, we have $1 - \beta y(A - (B+b)y)$ has a minimum value $1 - (\beta A p^*/4)$. As a consequence condition (2.47) is satisfied when

$$B + (1 + m(x,y))\frac{b}{2}(1 - \beta\frac{Ap^*}{4}) > 0. \tag{2.48}$$

This is clearly always the case when $\beta < 4/(Ap^*)$. Otherwise, when $\beta \geq 4/(Ap^*)$, since $0 \leq y < x$, condition (2.48) is satisfied when

$$B + (1 + m(0,0))\frac{b}{2}(1 - \beta\frac{Ap^*}{4}) > 0. \tag{2.49}$$

which can be rewritten as condition (2.27) in Theorem 2.6. □

Proof of Theorem 2.7 From Theorem 2.6 one can derive that the map $h_\beta(p)$ is always well defined and differentiable in a neighborhood of the point (p^*, p^*). Furthermore one can use expression (2.43) and (2.46) to compute:

$$(h_\beta)'(p^*) \quad = \quad \lambda(\beta) = -\frac{(\partial ES(x,y)/\partial x)|_{(p^*,p^*)}}{(\partial ES(x,y)/\partial y)|_{(p^*,p^*)}} = -\frac{b(1 - m^*(\beta))}{2B + b(1 + m^*(\beta))}$$

The value of β_1 is found by imposing $\lambda(\beta) = -1$. Local stability follows from $\lambda(\beta) \in (-1, 0)$ when $\beta < \beta_1$. When $\beta > \beta_1$, $\lambda(\beta) < -1$ and graphical analysis of h_β^2 shows that h_β^2 has (at least) two other intersection points with the diagonal than p^*. □

Proof of Theorem 2.8 Consider the map h_∞ given in (2.28). Call $\bar{p} = a$ and $\hat{p} = c$ that is a solves $\sqrt{a^2 + 2C/b} = p^*$ and c solves $A/B - (b/B)c = \sqrt{c^2 + 2C/b}$. One can show easily that these points always exist, as long as $2C/b < (p^*)^2$ which we assume here, are unique, and $c > a$. In terms of the points a and c, the map h_∞ can be rewritten as:

$$h_\infty(p_t) = \begin{cases} p^*, & 0 \leq p_t < a \\ \sqrt{p_t^2 + 2C/b}, & p \in [a, c] \\ p^* - (b/B)(p_t - p^*), & p_t > c. \end{cases}$$

The map h_∞ has a global maximum at the critical point c. Let d be the point such that $h_\infty(d) = a$. Notice that when $p > d$, $h^2_\infty(p) = p^*$. Two examples of the map are given in Figure 2.10.

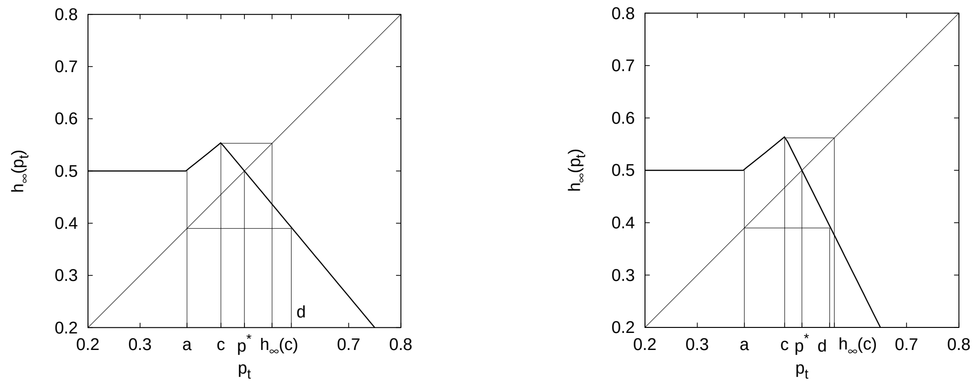

Figure 2.10: *Graphs of the map h_∞. In the left panel ($b/B = 1.2$) $h_\infty(c) < d$ implies that the price fluctuations remain bounded away from p^* in the long run. In the right panel, ($b/B = 2$), $h_\infty(c) > d$ implies that the map has (infinitely many) homoclinic points and a typical price time series converges to the locally unstable steady state p^*. Other parameters: $A = 1.5$, $b = 2$ and $C = 0.1$.*

Since $b/B > 1$, the steady state p^* is locally unstable. In what follows we show that for b/B sufficiently large, there are (infinitely many) homoclinic points, whose orbits converge to p^* both forward and backward in time. We consider two cases:

CASE 1.: $h_\infty(c) < d$ (as in the left panel of Figure 2.10). When $h_\infty(c) < d$ we have $h^2_\infty(c) < h^4_\infty(c) < p^* < h^3_\infty(c) < h_\infty(c)$. Let $I_1 = [h^2_\infty(c), h^4_\infty(c)]$ and $I_2 = [h^3_\infty(c), h_\infty(c)]$. In this case it is easy to show all initial states p_0 are "repelled" from p^* and mapped in $I_1 \cup I_2$ after some iterations. The long run dynamics is therefore contained in the set $I_1 \cup I_2$. and therefore bounded away from p^*.

CASE 2.:, $h_\infty(c) \geq d$ (as in the right panel of Figure 2.10). In this case the critical point c is a homoclinic point. It converges to p^* forward in time and also backwards in time (take successive inverse images x ¿ c, with $h^k_\infty(x) = c$, $k = 1, 2, ...$). In fact, there is an interval I of homoclinic points containing c.

To conclude the proof we have to show that there exists an $M > 1$ such that when $b/B > M$ then case 2 obtains, i.e. $h_\infty(c) \geq d$. We derive the existence of such a value M from the dependence of the point d and $h_\infty(c)$ upon b/B. Define $B/b = N$, $N \in (0, 1)$. Depending upon N, p^*, C and b, the point d and the image $h_\infty(c)$ are defined as:

$$
\begin{aligned}
d(N) &= p^* + N\left(p^* - \sqrt{(p^*)^2 - 2C/b}\right) \\
h_\infty(c; N) &= p^* + \tfrac{p^*}{1-N}\left(\sqrt{1 + \tfrac{1-N}{1+N}(2C)/b} - 1\right)
\end{aligned}
$$

One can show that $d(0) < h_\infty(c; 0)$, $d(1) > h_\infty(c; 1)$ and $d'(N) > 0$. Furthermore $(p^*)^2 >$

$2C/b$, which is assumed here to guarantee the existence of a and c, is a sufficient condition for $\partial h_\infty(c;N)/\partial N < 0$. These facts imply that there exists a unique N such that $d(N) = h_\infty(c;N)$. It follows that there exists a unique $M = 1/N$ such that: (i) when $b/B \in (1, M]$, case 1. applies, and (ii) when $b/B > M$, case 2. applies. □

Proofs of Section 2.4

Proof of Theorem 2.9 When $C'(q) = C$ the first order condition (2.33) for the optimal choice of q becomes:

$$\Delta\pi_{t+1} = \frac{b}{2}(p_{t+1} - p_t)^2 = C$$

Consequently in general there is no interior solution, but $q_t^* = 1$ if $\Delta\pi_{t+1} > C$ and $q_t^* = 0$ if $\Delta\pi_{t+1} < C$. This implies that the equilibrium price dynamics (2.34) is governed either by naïve expectations or by rational expectation, which gives exactly the same dynamical system as in (2.20-2.21) with $\beta = +\infty$, or equivalently as in (2.22). □

Proof of Theorem 2.10 Consider the market equilibrium equation (2.35):

$$A - Bp_{t+1} = \left[\frac{b/2(p_{t+1} - p_t)^2}{\alpha C}\right]^{\frac{1}{\alpha-1}} b(p_{t+1} - p_t) + bp_t.$$

Notice that, given the values of the parameters A, B, b, α, C, for any fixed value of $p_t \geq 0$ the demand (l.h.s) is a decreasing function of p_{t+1}, while the supply (r.h.s) is an increasing function of p_{t+1} provided that $\alpha > 1$. This implies that there exists a unique point $x \in \mathcal{R}$ where demand and supply are equal. Notice that this point is negative for those values of p_t for which there is excess supply at $p_{t+1} = 0$, that is when

$$(1 - \left[\frac{bp_t^2}{2\alpha C}\right]^{\frac{1}{\alpha-1}})bp_t \geq A$$

In this case we set $p_{t+1} = 0$, and otherwise $p_{t+1} = x$. □

Proof of Theorem 2.11 This is just a special case of Theorem 2.12. □

Proof of Theorem 2.12 Market clearing in (2.34) implicitly defines a map $p_{t+1} = r(p_t)$ by $ES(p_t, r(p_t)) = 0$, that is,

$$ES(p_t, r(p_t)) = q_t^* b(r(p_t) - p_t) + bp_t - Bp_{t+1} + A = 0,$$

where q_t^* solves the maximization problem (2.32), whose F.O.C is

$$\Delta\pi_{t+1} = \frac{b}{2}(p_{t+1} - p_t)^2 = C'(q). \tag{2.50}$$

The F.O.C. (2.50) gives only the interior optimal solution. From the convexity of $C(q)$ it follows that $C'(0)$ is a global minimum and $C'(1)$ is a global maximum of $C'(q)$. Hence, $q_t^* = 0$ if $\Delta\pi_{t+1} \leq C'(0)$ and $q_t^* = 1$ if $\Delta\pi \geq C'(1)$. These two conditions on $\Delta\pi_{t+1}$ define four parallel lines in the plane (p_t, p_{t+1}) with slope 1 and intercepts $\sqrt{2C'(1)/b}$, $\sqrt{2C'(0)/b}$, $-\sqrt{2C'(0)/b}$, $-\sqrt{2C'(1)/b}$. The interior solution (2.50) determines the map r only in the region between the relevant lines in the map (2.36). The implicitly defined part of r in (2.36) is the function $r(x_t)$ that solves

$$ES(x_t, r(x_t)) = q_t^* b(r(x_t) - x_t) + bx_t - Br(x_t) = 0, \tag{2.51}$$

with

$$q_t^* = q(x_t, x_{t+1}) = (C')^{-1}(\frac{b}{2}(x_{t+1} - x_t)^2).$$

Notice that, since $C'(q)$ is increasing, $(C')^{-1}$ is always a well defined function and is itself increasing. We claim that the map r has exactly the same properties (a-d) as the map g in the proof of Theorem 2.4. From property (c) it follows that when $b/B < 1$, $-1 < g'(x) < 1$, so that the map is a contraction and all orbits converge to the steady state p^*. From properties $(a - d)$ and the proof of Theorem 2.4 it follows that when $b/B > 1$ the steady state is locally unstable and a unique globally stable 2-cycle exists.

We conclude the proof by showing that properties (a-d) hold. in this case.

(a). From (2.51) it follows immediately that if $ES(x, r(x)) = 0$, then $ES(-x, -r(x)) = 0$. Hence $r(-x) = -r(x)$.

(b) We show that $y = r(x)$ has only one positive intersection a with the line $y = -x$. This follows from a change in variables $z = x + r(x)$ and $w = r(x) - x$ as in the proof of Theorem 2.4. In this case the corresponding system (2.41) has solution only when $b/B > 1$.

(c) Obviously when the map in (2.36) is explicitly defined, $r'(x) = 0$ or $r'(x) = -b/B$. In order to compute $r'(x)$ when it is implicitly defined one can use the implicit function theorem and obtain:

$$r'(x_t) = -\frac{\partial ES/\partial x_t}{\partial ES/\partial x_{t+1}} = \frac{q'(x_t, x_{t+1})(x_t - x_{t+1})^2 + q(x_t, x_{t+1}) - 1}{q'(x_t, x_{t+1})(x_t - x_{t+1})^2 + q(x_t, x_{t+1}) + \frac{B}{b}} \tag{2.52}$$

Since both q and q' are always positive, from the expression above it follows that $r'(x_t) \in [-b/B, 1)$. Notice that q' is positive iff $C'(q)$ is convex.

(d). Consider the proof of property (d) in Theorem 2.4. One can prove here that the same statement holds by replacing (2.51) with (2.37) and (2.52) with (2.42). □

Proof of Theorem 2.13 Comparing (2.34) and (2.12), we have to show that there exists a non decreasing, convex, information gathering cost function $C(q)$, such that the optimal solution of the representative agent problem (2.32) is given by:

$$q_t^* = n_{1,t} = \frac{1 + \tanh\{\frac{\beta}{2}[\frac{b}{2}(p_{t+1} - p_t)^2 - C]\}}{2}. \tag{2.53}$$

From the first order condition in (2.33) for an arbitrary cost function we obtain:

$$q_t^* = (C')^{-1}\left(\frac{b}{2}(p_{t+1} - p_t)^2\right). \tag{2.54}$$

Combining (2.53) and (2.54) it follows that:

$$(\frac{b}{2}(p_{t+1} - p_t)^2) = C'\left(\frac{1 + \tanh\left(\beta[\frac{b}{2}(p_{t+1} - p_t)^2 - C]\right)}{2}\right). \tag{2.55}$$

In terms of the variable $z = b/2(p_{t+1} - p_t)^2 \geq 0$, (2.55) becomes:

$$z = C'\Big(\frac{1 + \tanh\{\frac{\beta}{2}[z - C]\}}{2}\Big).$$

The change of variable $t = (1+\tanh(\beta(z-C)/2))/2$ gives an ordinary differential equation:

$$\frac{2\tanh^{-1}(2t - 1)}{\beta} + C = C'(t), \tag{2.56}$$

whose solution is the cost function $C(q)$ we are looking for. The restriction $z \geq 0$ implies $t \geq (1 + \tanh(-\beta C/2))/2 = t_0$. In integral form, $C(q)$ is given by:

$$C(q) = \int_{t_0}^{q}\left(\frac{2\tanh^{-1}(2t - 1)}{\beta} + C\right) dt + C_0, \quad q \in (t_0, 1]$$

Notice that $t_0 = \frac{1+m^*}{2}$, $C(t_0) = C_0$ and $C'(t_0) = 0$. This guarantees that when $p_t = p_{t+1} = p^*$ the solution of (2.54) is $q^* = (1 + m^*)/2$, the minimum fraction of rational agents at the fixed point in the heterogeneous agent model. When $q < t_0$ one can define without loss of generality, $C(q) = C_0$, and take the integration constant $C_0 = 0$. As a result one gets

$$C(q) = \begin{cases} \int_{t_0}^{q} \left(\frac{2\tanh^{-1}(2t-1)}{\beta} + C \right) dt & q \in (t_0, 1] \\ 0 & q \in [0, t_0] \end{cases}$$

By construction, the function $C(q)$ above has derivative $C'(q) = 0$ when $q \leq t_0$, and $C'(q)$ as in (2.56) when $q > t_0$. Using this plus the fact that the function $(\tanh)^{-1}(2t-1)$ is positive and increasing when $t \geq t_0$, it follows that the map $C(q)$ is non-decreasing and convex. □

Chapter 3

Competition and coordination in participation games

3.1 Introduction

In economics and social sciences many strategic decisions can be modeled as participation games. Market entry, route choice, union membership, voter turnout, public good provision are only a few examples. We can classify all of them as participation games where the two actions are either to participate, action 1, or to abstain, action 0. Typically the payoff associated with participating is a function of the number of other players participating, while the payoff associated with abstaining is fixed. One key element of participation games is the shape of the payoff function associated with the action of participating. On the one hand, cases where a group of agents compete for a limited resource, as in market entry games and route choice, can be modeled with a decreasing payoff function. The participation payoff diminishes as the number of other participating agents increases. On the other hand, cases where agents profit only if some of them decide to participate, such as public good provisions, union memberships or technology adoption, can be modeled with an increasing payoff function. The general participation game can be seen as a combination of these two types of strategic interaction. Hype cycles, for example, can be modeled with an initially increasing and thereafter decreasing participation payoff function.

Strictly related to the shape of the participation payoff is the feedback between individual motives and the subsequent aggregate participation rate. This effect is particularly relevant in situations where the participation game is played repeatedly over time, as considered here. When the participation payoff is a decreasing function in the number of other participating agents there is *negative expectational feedback*. In fact if, at time t, one individual believes that the next period few agents will participate, he increases his willingness to participate. If all agents have the same belief, the $t + 1$ participation rate will be high, which is the opposite of the individual beliefs at time t. In contrast, when the participating payoff is an increasing function in the number of participating agents there is *positive expectational feedback*. In this case the belief that many agents will participate is self reinforcing, because, if many agents share this belief, the number of participating agents will be large. The result is a positive feedback between beliefs or motives of behavior and the realization of the aggregate outcome.

Related to expectational feedback are the concept of substitutability and complementarity. (see e.g. Haltiwanger and Waldman, 1985). Strategies are substitutes if agents have an incentive to do the opposite of what other player are doing as it happens in systems with negative feedback. Strategies are complementary if agents have an incentive to imitate each other as it in systems with positive feedback. Recent experimental studies as reported in Camerer and Fehr (2006) or in Heemeijer, Hommes, Sonnemans, and Tuinstra (2006) show there indeed exist substantial differences between behavior in environments characterised by positive expectations feedback and environments characterised by negative expectations feedback.

In this chapter, we focus on participation games with *negative feedback* and we leave the analysis of positive or mixed feedback for future research. Participation games with *negative feedback* are interesting because they stem from a basic economic problem, i.e. exploitation of scarce resources, in the simplest setting, i.e. repeated interaction, without market institutions acting as coordination devices. In the general participation game with negative feedback, agents (firms or drivers) repeatedly interact to use the same resource (a market or a route) without the possibility of communicating with each other or of coordinating on a certain outcome. Each agent, thus, relies on personal beliefs and/or on aggregate quantities observed in the previous rounds, such as the participation rate history, to decide whether to participate or not. An important question we try to answer when we deal with exploitation of scarce resources is under which circumstances agents are able to use the resource efficiently. On the one hand each agent would like to exploit the available resource (a market, a road ...) alone and repeatedly. Agents thus fiercely compete to use the limited resource. On the other hand, knowing that the other agents have the same objective, agents have an incentive to implicitly coordinate their actions so that the exploitation of the scarce resource is maintained at a level from which agents can still profit. Competition and coordination are thus at work at the same time. Interestingly, competition and coordination are related to two different types of efficiency. Coordination is related to allocative efficiency: The better the agents are able to coordinate their actions the more they can profit from the use of the limited resource. Competition is related to informational efficiency: If patterns of the participation rate contain regularities, there are profit opportunities to be taken as agents can use these regularities to better forecast the future participation rate and use this information to decide when to participate and when not. If competition works at best, no regularity should be observed in the participation rate time series, that is an informationally efficient outcome is generated.

The aim of this chapter is to build a theoretical model of the interaction of many agents who are repeatedly playing a participation game with negative feedback. In particular we want to use this framework to address issues of informational and allocative efficiency that we outlined above.

3.1.1 Overview of the literature

Before we present the guidelines of our model building, we review the existing experimental, computational, and theoretical investigations of participation games with negative feedback. Existing examples of participation games with negative feedback can be divided in two classes of games. In the first class, usually referred to as market entry games, the payoff for participating is a decreasing function of the number of other agents participating while the payoff for absenting is constant. Firms deciding whether to enter a new market (participating), face competition from other firms, and thus are exposed to an uncertain payoff, only if they enter the market. When they stay out of the market (absenting), their payoff is constant and not depending on the number of other firms which have decided to enter. In the other class of negative feedback games, usually referred to as route choice games, both the payoff for participating and for absenting are a decreasing function of the number of other agents taking the same decision. Commuters traveling every day from their residential area in A to their office facilities in B have to decide whether to take road 1 (participating) or road 0 (absenting). Their traveling time, which can be associated to their payoff, is clearly decreasing with the number of other commuters who take the same route. We will see that results for these two types of games are similar.

An early experimental investigation of market entry games is Sundali, Rapoport, and Seale (1995). They consider zero entry cost and linearly decreasing payoff for entering. The number of agents involved in the experiment is $N = 20$ and the maximum profitable capacity of the market changes in the range $[1, 19]$. For each capacity they run the experiment over 10 periods for 6 different groups. Their main finding is that the average participation rate is close to the prediction of the symmetric mixed strategy Nash equilibrium of the one stage market entry game. Nevertheless, at the individual level, agents are not playing a mixed strategy Nash equilibrium, in particular they are not randomizing their action. The individual profiles of actions of each player is very diversified, with players that are, over time and across different capacities, always participating, players that are always absenting, and players that are changing their action conditionally on the outcome of the previous rounds. Furthermore Sundali, Rapoport, and Seale detect excess volatility, that is the variance of the participation rate is higher than predicted by the game theoretical solution which is given by the symmetric mixed strategy equilibrium. Other experiments investigate whether these results are robust to changes of some characteristics of the Sundali, Rapoport, and Seale experimental set-up. Rapoport, Seale, Erev, and Sundali (1998) repeat the same market entry game for three different absenting payoffs. Erev and Rapoport (1998) extend the analysis to three different information treatments: one where agents know only their payoff, one where they know the payoff associated with both actions and one, called full information treatment, where in addition agents know the number of players which have chosen each action in the previous round. In a similar analysis of different information treatments Duffy and Hopkins (2005) consider 100 iterations in order to investigate the long run properties of agents' interaction. Only in the case of full information agents are coordinating on the pure Nash equilibrium in the long run. Zwick and Rapoport (2002) consider an hyperbolic functional form for the payoff for entering. Even though each of these experiments extends the previous contributions with respect to its own line of investigation, results from all of them are consistent with the results of Sundali, Rapoport, and Seale (1995), i.e. the first moment of the participation rate distribution converges to the prediction given by the symmetric mixed strategy Nash equilibrium, but the second moment is always larger in the experimental data than in the prediction. These results are also confirmed by experiments whose underlying structure corresponds to that of route choice games. In an early paper Meyer, Huyck, Battalio, and Saving (1992) consider a number of suppliers who face the problem of allocating their common product between two different locations without knowing the decision of other suppliers. Their main focus is to investigate history's role in coordinating decentralized allocation decisions. Their experiments "provide support for the view that disequilibrium dynamics are an important source of fluctuations in market prices and quantities" and "nevertheless, the symmetric equilibrium accurately predicted the average quantity supplied, the average price, and the average market efficiency ..." (both in Meyer, Huyck, Battalio, and Saving 1992 p. 315). Similar results are shown in a forthcoming paper by Selten, Chmura, Pitz, Kube, and Schreckenberg (2006) who report laboratory experiments where a group of 18 agents have to choose period after period between a main road and a side road. When the experiment is run as long as 200 periods, fluctuations persist till the end of the session. A significant correlation between the number of road changes of each individual and the size of the fluctuations is also detected. Individuals seem to use four different strategies: to use the main road, to use the side road, to switch roads after having received a payoff lower than the median of the other agents' payoffs, or to not switch roads in that same situation. Nevertheless the aggregate road use is consistent with the symmetric Nash equilibrium of the underlying one stage game.

Computational models of participation games with negative feedback also reproduce the results reported above. An influential contribution in this area is Arthur (1994) where the *El Farol* bar problem is introduced. *El Farol* is a bar close to the Santa Fe Institute (New Mexico, USA) where a live music show takes place on Thursday evening every week. There are N persons who want to go to the bar (participating) on Thursday evening but the *El Farol* bar is not big enough for all of them. In fact the bar is too crowded when more than bN persons are there (with b between 0 and 1). In this case it would have been better to remain at home (abstaining). When less than bN persons are attending the concert the evening is very pleasant. The *El Farol* bar problem is thus a participation game with negative feedback. Interestingly, as we will show later in this chapter, the *El Farol* game can be modeled both as a market entry game or as a route choice game. Arthur introduced the *El Farol* bar problem as a simple model of the interaction of boundedly rational agents with heterogeneous deterministic rules and used computer simulations to replicate the interaction of $N = 100$ agents relying on different predictors to decide whether to go to the bar or not. Every agent has a set of predictors and keeps track of the forecasting performance of each of them. Predictors which perform better are more likely to be chosen in the future. His simulations show that the average participation rate converges to the capacity of the bar, even if fluctuations persist in the long run. He also observes that the interaction of agents using different rules leads to a participation rate dynamics with no regularities because, using his words, "Where cycles-detector predictors are present, cycles are quickly 'arbitraged' away so there are no persistent cycles" (Arthur 1994 p. 409).

In order to explain Arthur's results, as well as the experimental results discussed above, other theoretical and computational models have been advanced. A class of them uses variations of standard game theoretical tools. Goeree and Holt (2005) consider a model based on the concept of quantal response equilibrium as in McKelvey and Palfrey (1995). Their general approach deals with participation games with negative and positive feedback and for the case of market entry games they argue that a quantal response equilibrium better describes the average aggregate participation rate as the capacity of the market changes. Zambrano (2004) aims at reproducing the main results of the simulations in Arthur (1994) using the Nash equilibria of a game that underlies Arthur's framework. Zambrano shows that what he calls a prediction game is the game theoretical representation of the bar problem and that the average participation rate of Arthur's simulations coincide with the set of mixed strategy Nash equilibria of the prediction game. Both contributions are interesting as they manage to perform an equilibrium analysis of interactions of boundedly rational agents, but they do not aim at justifying the variance around this equilibrium. Both contributions ignore the interaction between the dynamics of the aggregate and the individual choices in a repeated game framework. As a result both analysis fail to account for the excess variance that is observed in the experimental data or in Arthur's simulations.

Other computational contributions try to explain experimental results employing the reinforcement learning of Erev and Roth (1998). For example Erev and Rapoport (1998) show that a simple reinforcement learning model provides a good approximation of the aggregate behavior for a wide setting of experiments with market entry games. More recently Duffy and Hopkins (2005) show that reinforcement learning predicts agents to be divided in those always participating and those always absenting, and thus pure strategy Nash equilibria, in place of mixed strategy Nash equilibria, to prevail in the long run. Nevertheless, experiments in the same paper show long run convergence to the pure strategy Nash Equilibrium only when agents are informed not only about the aggregate and their performance, as assumed in the reinforcement learning model, but also about the payoff of all the other agents. Another contribution is Franke (2002) using reinforcement learning to replicate Arthur's results on the *El Farol* bar problem. His simulations show that the long run distribution of the probability to participate can be centered around the symmetric mixed strategy Nash Equilibrium or can be binomial with two peaks at very low and very high probabilities to participate. He concludes that the long run outcomes are rather sensitive to the model specification. Evidence from reinforcement learning is thus mixed and quite sensitive to its parameters values.

3.1.2 An evolutionary approach

Summing up, previous investigations of both classes of participation games with negative feedback, market entry games and route choice games, lead to the same conclusions. Experimental findings show that at the aggregate level coordination on a symmetric mixed strategy equilibrium of the underlying one stage game occurs, as the first moment of the participation rate is well predicted by the symmetric mixed strategy equilibrium. At the individual level however, agents are not randomizing but employ simple rules instead, such as: always participate, always abstain, or participate conditionally on the outcome of previous rounds. The interaction of these different rules creates excess volatility for the participation rate as compared to the symmetric mixed strategy equilibrium, so that the second moment of the participation rate is always larger than predicted. The variance is important because it is related to allocative efficiency: the more volatile the series the more resources are wasted. Simulations of Arthur where agents are using a large number of forecasting rules to decide whether to participate, reproduce the same results, both at the aggregate and individual level, but are computationally intensive. Furthermore, they suggest that the interaction of 'deterministic' rules is able to create informationally efficient participation rate patterns with no regularities. Theoretical models account for some of these findings but they do not address questions of informational or allocative efficiency directly. Moreover both the equilibrium models in Zambrano (2004) and Goeree and Holt (2005) and reinforcement learning models in Erev and Rapoport (1998) and Franke (2002) assume that agents are deciding whether to participate or not according to some probabilities, whereas experimental evidence suggests that a majority of agents uses simple deterministic rules.

In this chapter we model agents' interaction in a participation game with negative feedback as the evolutionary competition of simple deterministic rules that prescribe when to participate given past aggregate outcomes. The competition of these different behavioral modes, or heuristics or rules is embedded in an evolutionary setting so that the rule which has performed better in the recent past attains a higher number of followers. At this stage of our analysis this competition is governed by the replicator dynamics (see e.g. Weibull, 1995). Our aim is to obtain a relatively simple analytic model which can replicate the main experimental and computational findings. A similar approach has been applied in the study of endogenous fluctuation in financial markets (Brock and Hommes, 1998) in Cournot competition (Droste, Hommes, and Tuinstra, 2002) or in a cobweb model (Brock and Hommes, 1997) but, to our knowledge, it is new in participation games. The work on the cobweb model relates particularly to our setting as it also deals with an economic system with negative feedback between individual motives and aggregate outcome. In this area recent experimental results confirm the evidence we have summarized for participation games. Hommes, Sonnemans, Tuinstra, and Van de Velden (2007) show that agents are on average behaving according to the rational expectation equilibrium (the counterpart of the Nash equilibrium) but are deviating from it at the individual level. Moreover market price series shows irregular aperiodic fluctuations and excess volatility.

An important dimension of our investigation of participation games with negative feedback is that we explicitly take into account the size of the participation game, i.e. the number of players. In particular we are interested in the change of the game dynamics as the size of the game increases. This an important generalization since a large number of agents is usually involved at the same time in a participation game. Laboratory experiments deal with up to $20 - 25$ players and natural experiments such as the route choice problem, with many more. The analysis of the limit of a large number of players is interesting because if agents were randomizing their action independently according to the same distribution the law of large numbers would apply and we would fail to observe fluctuations of the participation rate in the limit. We will see that this is not the case in our setting, which is therefore consistent with laboratory and natural experiments. Another dimension of investigation is the number of rules used. The evidence from Arthur's simulations is that when many rules are used the participation rate series is seemingly unpredictable. Given a set K of competing rules we measure the linear autocorrelation of the participation rate and investigate whether it changes as new rules, which try to exploit the detected linear autocorrelations, are introduced. This part of the analysis is also related to the notion of evolutionary stability (see e.g. Maynard-Smith, 1974) as it questions the stability of the original set K of rules against mutation to other rules. Related issues are how allocative and informationally efficiency change as the size of the participation game or the number of rules used are increasing.

In Section 3.2 we formalize the participation game with negative feedback and use the traditional tools of game theory to analyze its equilibria. In Section 3.3 we model the competition between different rules adopted by boundedly rational agents as an evolutionary process governed by the replicator dynamics. In Section 3.4 we apply this framework to the competition of two simple rules and we discuss how the participation rate dynamics changes with the size of the game. In Section 3.5 we investigate how evolutionary competition is affected by the presence of new rules which try to exploit the regularities of the participation rate time series. In Section 3.6 we apply the evolutionary framework developed so far to highlight the differences and similarities between the route choice games and the market entry games. Section 3.7 concludes. Proofs of all results can be found in the appendix.

3.2 The participation game

In this section we model the one stage participation game with negative feedback which corresponds to a market entry game. In market entry games the payoff for participating is a function of the number of other agents participating whereas the payoff for absenting is constant. Generalization to the other class of participation games with negative feedback, route choice games, is considered in Section 3.6. After we model the market entry one stage game, we use standard game theoretical analysis to characterize its pure and mixed strategy Nash equilibria and its evolutionary stable strategies. This analysis gives us the benchmark against which to compare the evolutionary model we develop in Section 3.3.

Consider N agents, each of whom chooses whether or not to participate, that is to enter a market, to go to a bar, etc. Each agent chooses an action $a \in \{0, 1\}$, where $a = 1$ stands for participating and $a = 0$ for absenting. The action space is given by $A = \{0, 1\}^N$ and an action profile is given by $\boldsymbol{a} \in A$. The game is thus a N-player game where every player can choose between 2 actions. We refer to N, which is a parameter in our model, as the *size* of the game. A strategy s_i for player i is the probability in $[0, 1]$ with which he chooses action $a = 1$. As a result, the strategy space is given by $S = [0,1]^N$ and $\boldsymbol{s} \in S$ denotes a strategy profile. Finally, by $\boldsymbol{s}_{-i} = (s_1, \ldots, s_{i-1}, s_{i+1}, \ldots, s_N)$ we denote the strategies played by all the players but player i. The payoff function of player i is determined by the strategies played by the N players and is given by

$$\pi_i : S \to I\!R_+.$$

We assume that all players want to maximize their expected payoffs. Pure strategy payoffs of the participation game are given as:

$$\begin{aligned} \pi_i (0, \boldsymbol{a}_{-i}; N_c) &= \alpha, \\ \pi_i (1, \boldsymbol{a}_{-i}; N_c) &= \begin{cases} \alpha + \beta - \gamma & \text{if} \quad \sum_{j=1, j \neq i}^{N} a_j < N_c, \\ \\ \alpha - \gamma & \text{if} \quad \sum_{j=1, j \neq i}^{N} a_j \geq N_c. \end{cases} \end{aligned} \tag{3.1}$$

The parameter N_c (the subscript c stands for capacity) defines the capacity of the market or the bar. Choosing $a_i = 0$ always gives payoff α. Choosing $a_i = 1$ gives payoff $\alpha - \gamma$ if N_c or more of the other $N - 1$ players have chosen $a = 1$, and $\alpha + \beta - \gamma$ if less than N_c of the other $N - 1$ players have chosen $a = 1$. We assume assume that $\beta > \gamma$, otherwise playing $a_i = 0$ is a dominant strategy. In this formulation α is a payoff which is always granted, γ corresponds to the cost or effort of participation and β corresponds to the (uncertain) participation premium. In the present section, the payoff α could be set to zero without changing the results. Nevertheless, since in the dynamic approach of the following sections α does play a role, we have decided to introduce it from this earlier stage.

In many experimental market entry games, the payoff for participating is linearly decreasing in the number of other entrants. The present formulation, using a step function as in 3.1 for the payoffs of the entrants, is equivalent with the payoff function for the *El Farol* bar game used in Arthur (1994) and Franke (2002). We chose this stylized version of the market entry game since it facilitates comparing games with different sizes N and capacities N_c. Note that for $N = 2$ and $N_c = 1$, this payoff structure is similar to the well-known game of Chicken or Hawk-Dove game (see e.g. Fudenberg and Tirole, 1991, pp. 18-19).

The expected payoff of playing a mixed strategy s_i is given by

$$\begin{aligned} \pi_i (s_i, \boldsymbol{s}_{-i}; N_c, N) &= (1 - s_i)\, \alpha + s_i \left(\alpha + \beta p \left(\boldsymbol{s}_{-i}; N_c, N\right) - \gamma\right) \\ &= \alpha + s_i \left(\beta p \left(\boldsymbol{s}_{-i}; N_c, N\right) - \gamma\right), \end{aligned} \tag{3.2}$$

where

$$p\left(\boldsymbol{s}_{-i}; N_c, N\right) = \Pr\left\{n\left(\boldsymbol{s}_{-i}\right) \leq N_c - 1\right\}$$

is the probability that the number of other agents participating, $n\left(\boldsymbol{s}_{-i}\right)$, is strictly smaller than N_c. The value of $p\left(\boldsymbol{s}_{-i}; N_c, N\right)$ depends upon the strategy profile $\boldsymbol{s}_{-i}$.

3.2.1 Nash equilibria

The game with payoff structure (3.1) has many pure strategy Nash equilibria (henceforth PSNE). Any pure action profile $\boldsymbol{a}$ such that exactly N_c players play $a = 1$ and the other, $N - N_c$, players play $a = 0$ corresponds to a strict PSNE. There are $\binom{N}{N_c}$ of these PSNE. Note that these PSNE lead to an asymmetric distribution of payoffs.

Now consider mixed strategy Nash equilibria (henceforth MSNE), that is, equilibria where some players randomize between the two possible actions. We show that there exists a unique symmetric mixed strategy Nash equilibrium where all players are randomizing with $s^* \in (0,1)$ so that $\boldsymbol{s}^* = (s^*, \ldots, s^*)$ is the corresponding strategy profile. If each agent is playing action $a = 1$ with probability s^*, the probability that the number of players participating is strictly smaller than N_c is given by:

$$p\left(\boldsymbol{s}^*_{-i}; N_c, N\right) = \sum_{k=0}^{N_c-1} \binom{N-1}{k} \left(s^*\right)^k \left(1 - s^*\right)^{N-1-k}. \tag{3.3}$$

Notice that $p\left(\boldsymbol{s}^*_{-i}; N_c, N\right)$ is the cumulative distribution function evaluated at $N_c - 1$ of a binomial distribution with $N-1$ degrees of freedom and probability s^*. By the definition of a Nash equilibrium, s^* is a best response for player i only when, given $\boldsymbol{s}^*_{-i}$, agent i is indifferent between the action $a = 1$ and action $a = 0$, which occurs when, using (3.2):

$$\pi_i\left(1, \boldsymbol{s}^*_{-i}; N_c, N\right) = \alpha + \left(\beta p\left(\boldsymbol{s}^*_{-i}; N_c, N\right) - \gamma\right) = \alpha = \pi_i\left(0, \boldsymbol{s}^*_{-i}; N_c, N\right).$$

Hence the value of s^* can be derived by solving the following equation:

$$p\left(\boldsymbol{s}^*_{-i}; N_c, N\right) = \frac{\gamma}{\beta}. \tag{3.4}$$

Since we are interested in the case when N becomes large, but where the relative (market, bar) capacity remains the same, we define the parameter $b \equiv \frac{N_c}{N}$. The following proposition summarizes the properties of the MSNE.

Proposition 3.1 *For any N, $N_c < N$, α, γ, $\beta > \gamma$, there exists a unique symmetric MSNE s^* of the N-person participation game whose payoff function is specified by (3.1).*

The value of s^* solves (3.4) and does not depend upon α. Moreover we have $s^* \to b$ as $N \to \infty$ for all γ and β, and $s^* = b$ for all N when $b = 1/2$ and $\gamma/\beta = 1/2$.

The exact value of s^* depends on both the threshold value b and on the ratio γ/β. Figure 3.1 shows (a numerical approximation of) s^* as a function of N for different values of $\frac{\gamma}{\beta}$ in the case of $b = 0.5$. In the Figure, as stated in Proposition 3.1 $s^* \to b$ as $N \to \infty$, for all ratios γ/β. Furthermore $s^* = b$ for all N when $b = 1/2$ and $\gamma/\beta = 1/2$. The case of $\gamma/\beta = 1/2$ is the one where the payoff for absenting is exactly in between the payoffs of a successful and a non-successful participation. Because of this symmetry and of the invariance of s^* with respect to changes in N, we take $b = 1/2$ and $\gamma/\beta = 1/2$ as the benchmark values for our analysis.

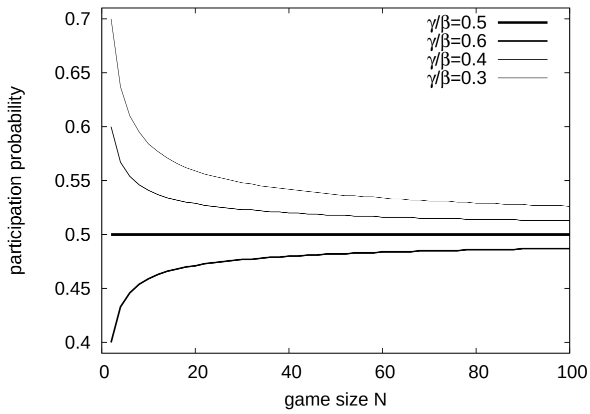

Figure 3.1: *Every curve shows an approximation of the MSNE s^*, that is a numerical solution of eq. (3.4), for N up to 100 and $b = N_c/N = 0.5$. Each curve corresponds to a different value of γ/β. The MSNE s^* is in general different from b, but converges to b as the size of the game, N, increases. From top to bottom: $\gamma/\beta = 0.3$, $\gamma/\beta = 0.4$, $\gamma/\beta = 0.5$ and $\gamma/\beta = 0.6$. In the benchmark case, that is when $\gamma/\beta = 0.5$ and $b = 0.5$, $s^* = b$ for every game size N.*

Other MSNE, where players are randomizing with different probabilities, exist. These equilibria are asymmetric with a number $M < N_c$ of agents always going and the other $N - M$ agents are randomizing with the same probability. In fact, the players randomizing are playing the symmetric MSNE of the participation game with the same payoffs, size $N - M$ and threshold $N_c - M$. These asymmetric MSNE are thus a combination of the two other types of equilibria, asymmetric PSNE and symmetric MSNE.

The PSNE has the characteristic that it distributes all available resources but implies that a symmetric group of players is able to coordinate on an asymmetric outcome, as N_c players receive $\alpha + \beta - \gamma$ and the others receive α. Asymmetric MSNE also share this characteristics. The only equilibrium which avoids this problem is the symmetric MSNE. The symmetric MSNE is in fact the only symmetric equilibrium. When players play the MSNE, the total number of participating players $\boldsymbol{N}$ (which is a random variable) has mean Ns^* and variance $Ns^*(1-s^*)$. A useful variable to consider is the (random) variable participation rate, $x_t = \frac{\boldsymbol{N}_t}{N}$, which has mean s^* and variance $\frac{s^*(1-s^*)}{N}$. Observe that as the size of the game becomes large, x_t will be equal to s^* with high probability since the variance goes to zero. The symmetric MSNE equalizes the expected payoff of participating and absenting for all players and thus leads to a lower average welfare than the pure strategy equilibrium. If a group of players were repeatedly playing the symmetric MSNE the resulting participation rate sequence, $\{x_t\}$, would be randomly distributed around s^* with variance of the order $s^*(1-s^*)/N$. Furthermore the series $\{x_t\}$ would have zero autocorrelation at all lags as agents participation decision is random and independent over time.

3.2.2 Evolutionary stability

We now discuss the game in terms of evolutionary stability. Consider an infinite population of players meeting randomly to play the N size participation game with threshold value b. Take the strategy s as the possibly mixed strategy played by all players. Suppose that a fraction ϵ of the population "mutates" his strategy, that is a fraction ϵ switches to play strategy s'. Will the fraction ϵ disappears against s^* or survive? The concept of evolutionary stability has been developed in connection with this question. In particular a symmetric strategy profile is called evolutionary stable if it cannot be invaded by any of these mutations. The original definition of evolutionary stable strategy was introduced in the context of 2-person games by Maynard-Smith (1974). The extension to N-person games is quite natural (see e.g. Palm, 1984). It turns out that the strategy s^* is the unique evolutionary stable strategy.

Proposition 3.2 *For any N, $N_c < N$, α, γ, $\beta > \gamma$, the symmetric strategy profile s^*, i.e. the unique solution of (3.4), is the unique evolutionary stable strategy of the N-person participation game whose payoff function is specified by (3.2).*

The concept of evolutionary stable strategy employed here depends on the assumption of a single infinite population whose players are randomly matched in groups of size N to play the participation game.

In Section 3.5 we will use the fact that s^* is an evolutionary stable strategy as the benchmark case against which to compare the stability of the competition of simple deterministic rules when more rules are introduced.

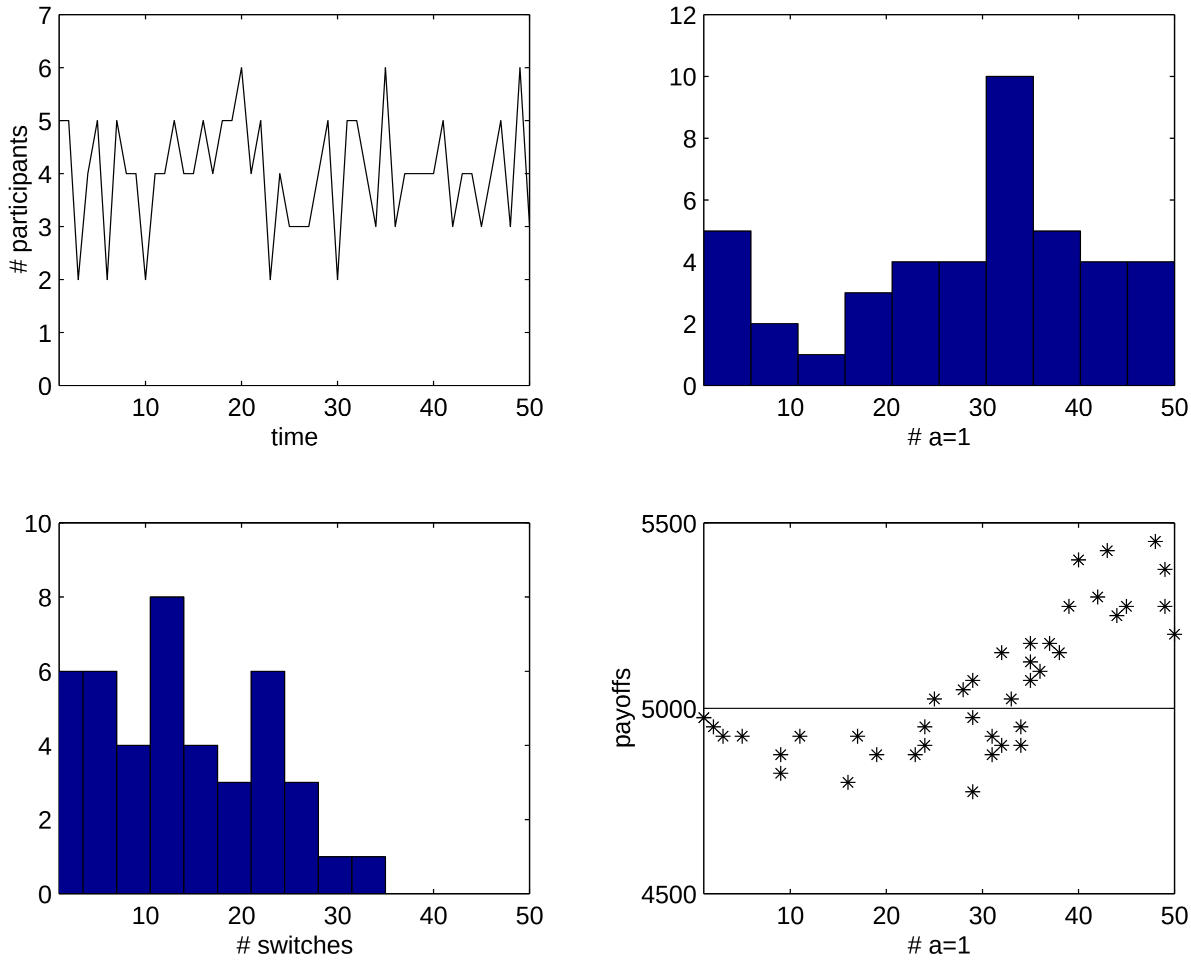

Figure 3.2: *Top left panel: Time series of number of participating players in experimental group 1. Top right panel: Histogram of the number of participations for the* 42 *players. The horizontal axes gives, out of the* 50 *periods, the total number of times an agent has chosen to participate, or* $\boldsymbol{a} = 1$*. The vertical axes gives the number of agents for each number of total participations. Bottom left panel: Histogram of the number of switches. The horizontal axes gives the total number of times an agents has chosen to switch action during the* 50 *periods. The vertical axes gives the number of agents for each number total switches. Bottom right panel: Relation between individual payoffs and total number of participations. For every agent the number of times he has chosen to participate,* $\boldsymbol{a} = 1$*, is plotted against his final payoff. When everybody plays the MSNE, the expected payoff after* 50 *rounds is* 5000 *as given by the horizontal line.*

3.2.3 Some experimental results

We will now briefly describe some experimental results on a negative feedback participation game from Heemeijer (2006). Our aim is not to provide a full-fledged analysis of the experimental data, but to illustrate and motivate the model introduced in this paper.

The experiments were conducted in October 2005 and February 2006 at the CREED laboratory of the University of Amsterdam. It involved 6 groups of $N = 7$ players, which had to make a participation decision for 50 subsequent periods. Group composition remained the same over the course of the experiment. Capacity was equal to $N_c = 4$ and payoff parameters were given by $\alpha = 100$, $\beta = 50$ and $\gamma = 25$. Moreover, a stochastic term ε_t from a symmetric triangular distribution on $[-25, 25]$ was added to the payoff for participating in every period. Subjects did know the value of α, but not those of β, γ and ε_t. Following the analysis above we have $\binom{7}{4} = 35$ PSNE. The symmetric MSNE s^* is implicitly given by $p(s^*; 4, 7) = \frac{1}{2}$. This gives $s^* \approx 0.5786$, which is slightly larger than $b = \frac{N_c}{N} = \frac{4}{7} \approx 0.5714$.

The results from this experiment are consistent with the experimental evidence discussed in the introduction. The upper left panel of Figure 3.2 shows, for the first experimental group, the dynamics of the number of participating subjects over 50 periods. Behavior in the other groups was similar. Aggregate participation decisions are unstable and keep fluctuating until the end of the experiment. Subjects did not coordinate on one of the PSNE and, at the aggregate level, the symmetric MSNE seems to provide a better description of the data. The first row of Table 3.1 shows that average participation rates in all groups are quite close to $s^* \approx 0.5786$, although there seems to be some 'underparticipation': for five of the six groups the average participation rate is somewhat lower than predicted by the symmetric MSNE.

	gr. 1	gr. 2	gr. 3	gr. 4	gr. 5	gr. 6	mean	MSNE
$\frac{1}{50}\sum_{t=1}^{50} x_t$	0.574	0.594	0.560	0.563	0.571	0.554	0.570	0.579
group switches	25	14	25	25	26	22	22.8	22.9
individual switches	15.7	8.3	15.6	13	17	14	14	23.9
% naivety	0.736	0.741	0.716	0.824	0.689	0.714	0.733	0.2438

Table 3.1: *Experimental results. The first row gives the average participation rate for each group. The second row gives the number of times, per group, that participation changed from four or less to five or more subjects or vice versa. The third row gives, per group, the number of individual switches between participating and not participating, averaged over subjects. The last row gives the percentage of individual switches that follow directly after a negative payoff experience. The last column gives the same quantities as expected from the symmetric MSNE solution.*

At the individual level, however, the symmetric MSNE is not supported by the data. The upper right panel of Figure 3.2, for example, shows a histogram of the number of times each of the 42 subjects participated. Apparently, some subjects participate almost always, for example 8 subjects have participated at least 40 times in 50 rounds. Others have almost never participated, for example 7 subjects have participated no more than 10 times. This suggests that many subjects do not randomize. Further evidence is given by the second and third row of Table 3.1, which show that, although the number of switches in aggregate participation are roughly consistent with the symmetric MSNE, individual subjects change their participation decision much less frequently than predicted by the symmetric MSNE. The histogram of the number of individual switches, depicted in the lower left panel of Figure 3.2 also provides compelling evidence that many subjects are quite reluctant to change. Instead of playing according to the symmetric MSNE, subjects seem to condition their decision on past payoffs. The last row of Table 3.1 shows that about 73% of individual switches were preceded directly by a negative payoff signal (i.e., start (stop) to participate when that gave a higher (lower) payoff in the previous period). A last observation worth noting is the evidence for a 'participation premium' in the lower right panel of Fig 3.2. Subjects that participated more often did, on average, earn higher payoffs. We will return to this issue in Section 3.6.

3.3 Evolutionary competition of deterministic rules

The experiment discussed above suggests, together with the earlier experimental evidence, that our negative feedback participating game is inherently unstable, with persistent fluctuations in the participation rate. Moreover, individual subjects do not seem to randomize their decisions and these fluctuations therefore cannot be easily attributed to mixed strategy Nash equilibria. Instead, subjects seem to base their decisions on deterministic behavioral rules. The existing behavioral models of negative feedback participation games are complex computational models, and some of them (such as reinforcement learning) do assume that agents randomize. In this section we introduce a behavioral model which is analytically tractable and consistent with the experimental findings.

3.3.1 Behavioral deterministic rules

Consider a population of $\widehat{N}$ boundedly rational agents who play the N-person participation game repeatedly, but possibly in different groups every period (random matching). For simplicity, we consider the limit $\widehat{N} \to \infty$ so that the law of large number applies and expected values of random variables can be taken as realized values. Period after period, each player behaves according to one of K different rules. Each rule specifies, given the information it uses, whether to participate or to abstain, that is $p \in \{0, 1\}$. A rule which performs better in period t is adopted by more agents the next period $t + 1$. We assume that the distribution of agents over the different K rules evolves according to the replicator dynamics mechanism. The replicator dynamics can be motivated in the context of boundedly rational agents learning and imitating which strategy to play in a strategic environment (see e.g. Weibull, 1995; Binmore and Samuelson, 1997). Furthermore, its simplicity lends it considerable appeal. The distribution of agents across rules at period $t+1$ lead to a different payoff associated to each rule and possibly to a new distribution. In what follows, after introducing what we exactly mean by rule, and how the replicator dynamics works in this context, we investigate the properties of the equilibria of such a process.

Let x_t be the participation rate, i.e. the fraction of agents having chosen to participate in period t. At every period t, a deterministic behavioral rule k specifies whether to participate or to absent for a given information set:

$$p_{k,t} = f_k\left(\mathcal{I}_{t-1}\right). \tag{3.5}$$

where the information set is given by past participation rates:

$$\mathcal{I}_{t-1} = \left\{x_{t-1}, x_{t-2}, \ldots, x_1, x_0\right\}. \tag{3.6}$$

and where

$$f_k : [0,1]^t \to \{0,1\}.$$

Some examples of simple rules are:

$$\begin{aligned}
p_{k,t} &= 1, \quad p_{k,t} = 0, \\
p_{k,t} &= \begin{cases} 1, & \text{if} \quad x_{t-1} < b, \\ 0, & \text{if} \quad x_{t-1} \geq b, \end{cases} \\
p_{k,t} &= \begin{cases} 1, & \text{if} \quad x_{t-1} \in A, \\ 0. & \text{if} \quad x_{t-1} \in A^c \equiv [0,1]/A. \end{cases}
\end{aligned}$$

The first rule indicates to always participate, whereas the second specifies to always abstain. The third and fourth rule are a function of the last participation rate, prescribing to participate if and only if the previous participation rate is smaller than the threshold b or lies in a specified set A.

The class of rules we will use most in our analysis can be written in terms of a best response to predictors of the current period fraction. That is,

$$p_{k,t} \in BR\left(g_k\left(x_{t-1}, x_t, \ldots, x_0\right)\right), \tag{3.7}$$

where $g_k(\cdot)$ is the prediction of the fraction in the current period. Notice that these predictor-based rules are a subset of a more general rule f_k.

The vector $\boldsymbol{x}_t$ represents the distribution of agents over the rules. Each component

$x_{k,t}$ is the fraction of the entire population that uses rule k in period t. That is $N_{k,t} = \widehat{N} x_{k,t}$ is the number of players using rule k at time t. Evidently we have $x_{k,t} \geq 0$ and $\sum_{k=1}^{K} x_{k,t} = 1$, that is $\boldsymbol{x}_t$ is in the $(K-1)$ −dimensional unit simplex, $\boldsymbol{x}_t \in \triangle^{K-1} = \left\{ \boldsymbol{x}_t \in I\!R_+ : \sum_{k=1}^{K} x_{kt} = 1 \right\}$. Notice that at time t aggregate participation x_t is completely characterized by $\boldsymbol{x}_t$ and $\boldsymbol{p}_t = (p_{1,t}, \ldots, p_{K,t})$. The participation rate at time t, x_t, is given by:

$$x_t = \boldsymbol{x}_t \cdot \boldsymbol{p}_t = \sum_{i=1}^{K} x_{i,t} p_{it}.$$

3.3.2 Evolutionary competition

In the previous subsection we have described a number of rules that agents use to play the game. Here we characterize how the distribution of agents across these rules evolves over time. First we define the profit of playing rule f_k. The fractions $x_{k,t}$, the strategies $p_{k,t}$ and the process of random matching of the $\widehat{N}$ players in groups of N players induce a probability distribution over payoffs $(\alpha - \gamma, \alpha, \alpha + \beta - \gamma)$, which we denote here by

$$\pi_{k,t} = \pi_k \left(\boldsymbol{x}_t, \boldsymbol{p}_t \left(\mathcal{I}_{t-1} \right); N_c, N \right), \text{ for } k = 1, \ldots, K.$$

Once we take the limit of the number of players of the population to infinity, $\widehat{N} \to \infty$, the realized payoff of the players playing each rule is equal to the ex-ante expected payoff for playing such a rule. In this case, by using equation (3.2), we can characterize the payoff associated with each rule k as a function of p_k and of the aggregate attendance x as:

$$\pi_{k,t}(x, p_{k,t}; N_c, N) = (1 - p_{k,t})\pi_0 + p_{k,t}\pi_1(x_t; N_c, N), \tag{3.8}$$

where $\pi_O = \alpha$ is the payoff for not participating (the subscript O stands for out, that is, abstain) and

$$\pi_I(x_t; N_c, N) = \alpha + \beta p(x_t; N_c, N) - \gamma \tag{3.9}$$

is the payoff for participating (the subscript I stands for in, that is, participate). The function $p(x; N_c, N)$ is the probability that the number of other agents $N-1$ participating is less than N_c when the participation rate is $x = \boldsymbol{x} \cdot \boldsymbol{p}$, that is

$$p(x; N_c, N) = \sum_{j=0}^{N_c - 1} \binom{N-1}{j} (x)^j (1-x)^{N-1-j}. \tag{3.10}$$

We have already noticed that $p(x; N_c, N)$ is the cumulative distribution function evaluated at $N_c - 1$ of a binomial distribution with parameters $N-1$ and x. Notice that since equation (3.9) establishes a one-to-one correspondence between x and π_1, the information set $\mathcal{I}_t$ defined in (3.6) is equivalent to $\{\pi_{I,t}, \pi_{I,t-1}, \ldots, \pi_{I,0}\}$.

Adding all these elements together, we can rewrite equation (3.8) as

$$\pi_{k,t}(x_t, p_{k,t}; N_c, N) = \alpha + (\beta p(x_t; N_c, N) - \gamma) p_{k,t}, \tag{3.11}$$

that is, the payoff of each rule is given by the "granted" payoff α plus p_k times the profit coming from the action of participating.

We use the payoff $\pi_{k,t}$ to characterize the evolution of the vector $\boldsymbol{x}$. Given the vector of expected or population average payoffs $\boldsymbol{\pi}_t = (\pi_{1,t}, \ldots, \pi_{k,t}, \ldots, \pi_{K,t})$ at time t, we define an evolutionary dynamics of the form

$$\boldsymbol{x}_{t+1} = \boldsymbol{H}(\boldsymbol{x}_t, \boldsymbol{\pi}_t), \tag{3.12}$$

where $\boldsymbol{H} : \triangle^{K-1} \times [\alpha - \gamma, \alpha + \beta - \gamma]^K \to \triangle^{K-1}$ is a continuous and differentiable function with $\frac{\partial H_k}{\partial \pi_k} > 0$ and $\frac{\partial H_j}{\partial \pi_k} < 0$ for $j \neq k$. That is, if the average profit of rule k increases the fraction of the population which uses rule k in the next period increases, and the fraction of other rules decreases. Note that we can extend the model into different directions by also including values of $\boldsymbol{x}_{t-\tau}$ and $\boldsymbol{\pi}_{t-\tau}$, for $\tau = 1, 2, \ldots$. An equilibrium of the evolutionary process (3.12) is defined by:

$$(\boldsymbol{x}^*, \boldsymbol{p}^*) = (x_1^*, \ldots, x_K^*, p_1^*, \ldots, p_K^*),$$

such that:

$$\boldsymbol{x}^* = \boldsymbol{H}(\boldsymbol{x}^*, \boldsymbol{\pi}^*),$$

where the equilibrium participation rate is $x^* = \boldsymbol{x}^* \cdot \boldsymbol{p}^*$. In this case $p_k^* = f_k(x^*, x^*, \ldots, x^*)$ and equilibrium profits are given by $\pi_k^* = \pi_k(\boldsymbol{x}^*, \boldsymbol{p}^*; N_c, N)$.

There exist different specifications for H_k in the literature. In this chapter we consider the replicator dynamics (see e.g. Taylor and Jonker, 1978 and Weibull, 1995) which takes the following format:

$$x_{k,t+1} = \frac{x_{k,t}\pi_{k,t}}{\sum_j x_{j,t}\pi_{j,t}}. \tag{3.13}$$

Such an updating mechanism can be interpreted as a biological reproduction model, where each period the number of agents using rule k grows proportionally to the performance of that rule, as measured by its payoff π_k. In fact equation (3.13) can be derived by assuming

$$N_{k,t+1} = N_{k,t}\pi_{k,t},$$

and using that $N_{k,t} = \widehat{N} x_{k,t}$. Notice that in this formulation, the payoffs $\pi_{k,t}$ for all k, and, as a consequence, the parameters α, β and γ, have to be normalized to meet the interpretation of a growth rate. In particular $\pi_{k,t}$ can not be negative, which in our framework translates into $\alpha > \gamma$. Notice that this way of deriving the replicator dynamics is taken directly from biological replication. Nevertheless equation (3.13) also arises in imitation processes in large populations of interacting agents, as reviewed in Chapter 4 of Weibull (1995). For this reason the replicator dynamics is a widely used updating mechanism in evolutionary economic dynamics.

It is useful to use equation (3.11) and equation (3.13) to compute the difference between today's and yesterday's fraction of agents using rule k:

$$\Delta x_{k,t+1} = \frac{x_{k,t}(\beta p(x_t; N_c, N) - \gamma)(p_{k,t} - x_t)}{\alpha + (\beta p(x_t; N_c, N) - \gamma)x_t} = \frac{x_{k,t}(\pi_{k,t} - \bar{\pi}_t)}{\bar{\pi}_t}, \tag{3.14}$$

where we have defined $\bar{\pi}_t = \sum_j x_{j,t}\pi_{j,t}$. Equation (3.14) explicitly relates the change in fractions of agents using rule k to a function of other agents payoffs. A characteristic of the replicator dynamics in discrete time is that equation (3.13) is not invariant under addition of a constant to the payoffs. In our formulation this is equivalent to say that changes in fractions of agents using one rule are not invariant under changes of α[1]. When α is small, everything else being equal, fractions are changing faster then when α is large. Moreover, when $\alpha \to \infty$ trajectories of our discrete dynamical system (3.13) approach the trajectories of the corresponding continuous time dynamical system (see e.g. Weibull, 1995, Chapter 4). For this reason the parameter α can be refered to as the inverse of the speed of adjustment of the replicator dynamics.

The condition $\Delta x_k = 0$, for all k, gives the equilibria of the replicator dynamics. They correspond both to corners of the simplex (where all agents use the same rule so that there exists one k with $x_k^* = 1$) and to interior solutions that satisfy $\pi_{j,t}^* = \pi_{k,t}^*$ for all k and j with both x_k^* and x_j^* strictly positive. The first type of fixed point of the dynamics does not correspond to a Nash equilibrium of the underlying one shot participation game. Among interior equilibria there are those where all rules prescribe the same action, that is $f_k(x^*, \ldots, x^*) = p^*$ for all k, so that $x^* = p^*$. We call these equilibria non-generic, since they arise only in the special case where the competing rules are all intersecting at one point x^* when they are evaluated at that point x^*. Other interior equilibria are those where $\beta p(x^*; N_c, N) = \gamma$. This last condition is the same as in equation (3.4) which characterizes the unique symmetric MSNE of the underlying one stage participation game. In contrast to the other, we call these equilibria generic. An example of the generic case occurs when the rule that states to always participate and the rule that states to always abstain are part of the set K of competing rules because these two rules are always behaving differently. This is the case which will be investigated in more detail in the next

[1]We remind the reader that the analysis of Section 3.2 is instead independent from changes of α.

section. The following proposition summarizes these results and characterizes the generic equilibria.

Proposition 3.3 *The system given by K rules as in (3.5), payoffs as in (3.11) and updating of fractions as in (3.13) has K fixed points for which the whole population uses the same rule. Other fixed points satisfy the property that $\pi_k^* = \pi_{k'}^*$ for all k and k' that are being played. In the generic case $x^* = s^*$, where s^* is the symmetric MSNE of Proposition 3.1.*

We are interested in the coupled dynamics of the set of rules and their evolutionary competition. Research questions are: (i) under which conditions will this dynamics lead to the symmetric Nash equilibrium? (ii) under which conditions does this dynamics lead to results which are qualitatively similar to the results obtained in the experiments run by, among others, Sundali, Rapoport, and Seale (1995), and with the computational models of bounded rationality of, e.g. Arthur (1994)? We address these questions with a concrete example in the next section.

3.4 Interaction of optimists and pessimists

In this section we study the simplest possible setup that follows from the framework we have outlined in the previous section. We consider the evolutionary competition between the following two rules:

$$p_{1,t} = 1 \text{ for all } t, \tag{3.15}$$

$$p_{0,t} = 0 \text{ for all } t. \tag{3.16}$$

The first rule specifies to participate with certainty, we call this the "optimistic" rule. The second rule specifies to abstain with certainty, which we call the "pessimistic" rule. As we discussed in the introduction the evidence from most experiments shows that a large fraction of agents is indeed using one of these rules. The simple rules (3.15) and (3.16) can also be understood as best reply rules as in (3.7). Rule (3.15) corresponds to a best reply rule (3.7) where the predictor used predicts a current period participation rate always below the threshold b; rule (3.16) corresponds instead to a best reply rule with a prediction of a participation rate always above b.

The value of the payoff associated with each rule can be derived using the general formula (3.11). As a result $\pi_{1,t}$, the payoff gained by the optimistic rule at time t, is a function of x_t and given by:

$$\pi_{1,t}\left(x_t; bN, N\right) = \alpha - \gamma + \beta p(x_t; bN, N). \tag{3.17}$$

The payoff gained by the pessimistic rule at time t, $\pi_{0,t}$, is constant and given by $\pi_O = \alpha$. Notice that π_1 depends also on the parameters b and N. In particular, for fixed values of b, the functional form of payoff of the optimistic rule with respect to the participation rate x is different for different value of the game size N. The left panel of Figure 3.3 shows examples of the payoffs of the two rules for different game sizes when $b = 1/2$.

Given the payoffs, the fraction $x_{1,t}$ of 'optimistic' players evolves according to (3.13) which gives:

$$x_{1,t+1} = \frac{x_{1,t}\pi_{1,t}}{x_{1,t}\pi_{1,t} + x_{0,t}\pi_{0,t}}.$$

Notice that since $x_t = x_{1,t}$ and $x_{0,t} = 1 - x_{1,t}$, one can rewrite the evolution of x as a first order nonlinear difference equation of the form:

$$x_{t+1} = f\left(x_t; b, N\right) = \frac{x_t(\alpha + p(x_t; bN, N)\beta - \gamma)}{x_t(p(x_t; bN, N)\beta - \gamma) + \alpha}. \tag{3.18}$$

The first order difference equation is parametrized by N, the size of the game, by the relative capacity $b = \frac{N_c}{N}$, by the payoff for absenting α, which is the inverse of the speed of adjustment of the replicator dynamics, by the cost of participating γ and by the return of a successful investment project β. The right panel of Figure 3.3 shows the function f for different values of N when the other parameters are fixed.

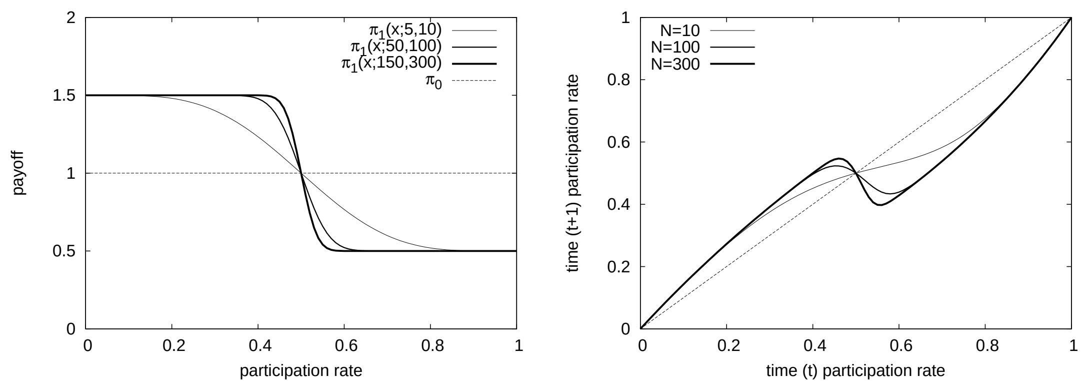

Figure 3.3: *Left panel: Payoffs of the pessimistic rule (3.16) (horizontal line) and optimistic rule (3.15) (other curves) for different values of the game size N. Right panel: Graph of the participation rate updating map (3.18) for the same different values of the game size N. In both panels the other parameters are $\alpha = 1$, $\gamma = 1/2$, $\beta = 2\gamma$ and $b = 1/2$..*

What are the steady states of the dynamics of the participation rate given by (3.18)?

First notice that, since $p_{1,t} \neq p_{0,t}$ for all t, we are in what we have called a "generic" case. By applying Proposition 3.3 we know that the steady states are given by the corner solutions $x = 0$ and $x = 1$ and by the the interior solution $x = s^*$. In fact, at s^*,

$$\pi_1(s^*; bN, N) = \alpha = \pi_0,$$

which is the requirement for the presence of an interior equilibrium. The distribution of agents which correspond to the symmetric mixed strategy Nash equilibrium we found in Section 3.3, s^*, is the value such that the two rules yield the same payoff. Notice that the interior equilibrium is *polymorphic*, that is, it is characterized by the simultaneous presence of two types of rules, a fraction s^* are optimists and the rest are pessimists.

3.4.1 Local stability

In order to characterize whether the equilibria are stable or not, it is convenient to define

$$\delta(x; b, N) = \frac{dp(x; bN, N)}{dx}.$$

Using the fact that $p(x; bN, N)$ can be written using the Beta function, as derived in the Proof of Proposition 3.4, one obtains:

$$\delta(x; b, N) = -\frac{x^{N(1-b)-1}(1-x)^{Nb-1}}{B(N(1-b), Nb)},$$

where $B(\cdot, \cdot)$ denote the Beta function which is defined as $B(a, b) = (a+b-1)!/((a-1)!(b-1)!)$. Notice that $\delta(x; b, N) < 0$ since an increase in the fraction of agents participating always decreases $\Pr\{n \leq bN\}$. We also define $\delta^* \equiv \delta(s^*; b, N)$. This partial derivative is crucial for the stability of the steady state.

The following proposition describes the stability properties of the steady states.

Proposition 3.4 *The dynamics of the participation rate given by equation (3.18) has three steady states:* 0, s^* *and* 1. *The fixed points* 0 *and* 1 *are locally unstable. The interior fixed point* s^* *is locally stable when* $\psi \equiv s^*(1 - s^*) \frac{\delta^* \beta}{\alpha} > -2$.

The stability condition in Proposition 3.4 depends, through δ^* and s^*, implicitly on b and N. We are particularly interested in characterizing this stability of s^* as a function of N. Figure 3.3 already suggests that s^* becomes unstable for N large enough. The following proposition corroborates that.

Proposition 3.5 *For any given value of* α, $\gamma < \beta$, *and* b *the dynamics of the participation rate is locally unstable around* s^* *in the limit* $N \to \infty$. *Moreover, if* $\beta = 2\gamma$ *and* $b = \frac{1}{2}$

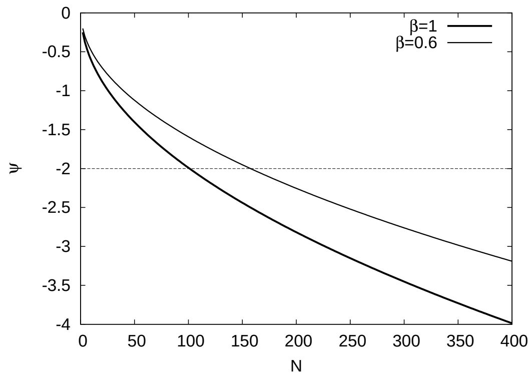

Figure 3.4: *Stability condition. Values of $\psi \equiv (1-s^*)(s^*)\delta^*\beta/\alpha$ as a function of the game size N for two different values of the participation premium β. When ψ is larger then the dotted line the steady state s^* is stable. The other parameters are $\alpha = 1$, $\gamma = \beta/2$ and $b = 1/2$.*

there exists a unique M such that equation (3.18) is locally stable around s^* if and only if $N < M$.

Figure 3.4 shows that ψ is decreasing in N and crosses -2 when N is larger than some threshold M. The intuition behind Proposition 3.5 is that as N increases the average population payoff of the optimistic rule gets closer to the step payoff of the underlying one shot game (see left panel of Figure 3.3). As a result, for any value of the payoff parameters, $\alpha, \beta > \gamma$, as N increases (3.18) becomes steeper at the steady state s^* and the system looses stability. This dependence of the dynamics upon N is due to the assumption of random matching. In fact, for $N = 2$ (with $N_c = 1$) the expected payoff function is linear in x, since $p(x; 1, 2) = 1 - x$, which is the probability of meeting a player using the pessimistic rule. However, the probability of having less than half of the players participating when each player participates (or, is an optimist) with probability larger than $\frac{1}{2}$ becomes small as N becomes large. That is, as N increases (and for given b) the function $p(x; N_c, N)$ will look more and more like a step function. We can therefore also interpret the parameter N as a measure of the shape and steepness of the payoff function at the steady state. In that case a low value of N would present a payoff function which decreases slowly as the number of participating players increases, similar to the the linear one used in the early market entry experiments (see, for example, Sundali *et al.*, 1995). A high value of N, on the other hand, would represent an expected payoff function close to the step function used in the *El Farol* bar game, with payoffs at the symmetric MSNE dropping rapidly as an extra player participates.

Summarizing, given any value of the payoffs $\alpha, \beta > \gamma$ and any value of the threshold b, there always exists a number of players N, or a size of the game, such that the system does not converge to the MSNE s^*. This is due to an overshooting effect which triggers instability only when the size of the game N is above a certain threshold.

3.4.2 Global dynamics

In this subsection we try to specify the nature of the fluctuations that characterize the unstable case. We present results primarily for the benchmark case with $b = 1/2$ and $\gamma/\beta = 1/2$.

Figure 3.5 shows the participation rate driven by the interaction of "optimists" and "pessimists" in the case of $N = 300$ and $N_c = 150$. The time series (left panel) looks aperiodic. In fact one can show that there exists a region in the parameter space where the dynamics is chaotic. The results follows from noticing that this 1-dimensional system has a 3-cycle (see for example the left panel of Figure 3.7 for $\beta - \gamma \approx 0.6$) and by applying the Li-Yorke theorem (see Li and Yorke, 1975). Even such a simple system is able to produce complicated times series which resembles what has been obtained using the more complicated simulation models of Arthur (1994) and others. Figure 3.6 shows the effect of the change of the game size N on the dynamics of the participation rate, for every even value of N between 2 and 400 when $b = 1/2$. The bifurcation diagram (left panel) shows the long run behavior as N increases.time series of x after an initialization period. Notice that, consistently with our theoretical results, the system is locally stable if and only if the game size is lower than a certain threshold. With the parameters value used for this simulation, using Proposition 3.5, one can show that the participation rate $s^* = 1/2$ is locally stable for sizes $N \leq \widetilde{M} = 100$ (see also right panel of Figure 3.4). For games of size larger than $\widetilde{M}$, the participation rates fluctuates along a cycle of period 2. The type of cycles observed for even larger sizes are typical of a period doubling bifurcation route to chaos. The right panel of Figure 3.6 shows the Lyapunov exponent of the corresponding participation rate for each game size N. The Lyapunov exponent is used to characterize sensitive dependence on initial conditions. When the system has a positive Lyapunov exponent, e.g. when $N \approx 300$, there is sensitive dependence on initial conditions and the system is chaotic. Figure 3.7 shows the dynamics of x and the corresponding Lyapunov exponent for different values of the granted payoff α, the inverse speed of adjustment of the replicator dynamics. With small α, which correspond to an high speed of adjustment, the overshooting effect causes instability. As α increases, the speed of adjustment decreases, that is, changes in fractions per period becomes smaller and smaller, and the dynamics stabilizes to the equilibrium s^*. The right panel of Figure 3.7 shows that, also in this case, there are many values of α for which the dynamics exhibits sensitive dependence upon initial conditions. Notice that as the dynamics is independent to a rescaling of the payoffs α, β and *gamma*, an change in α by a factor ρ, from α to $\alpha\rho$, with β and γ constant, is equivalent to a change of *beta* and *gamma* to β/ρ and γ/rho respectively with α. This allows us to interpret the effect of a decrease of α, other things equal, shown in Figure 3.7, also as the effect of an increase of β and γ, other things equal.

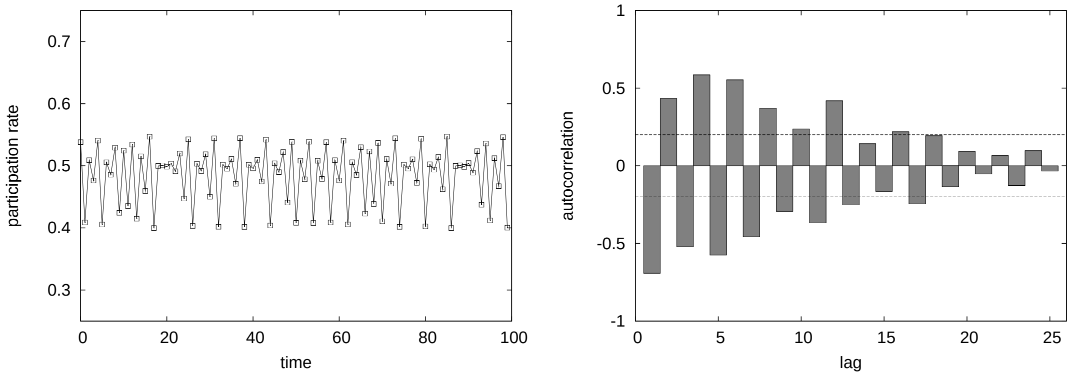

Figure 3.5: *Evolutionary competition of optimist and pessimists. Left panel:* 100 *iterations of the participation rate. Right panel: autocorrelation diagram of the series shown in the left panel. The dotted line shows the significance level of a standard autocorrelation test for* 100 *observations. Parameters values are* $N = 300$, $N_c = 150$, $\alpha = 1$, $\gamma = 0.5$ *and* $\beta = 2\gamma$.

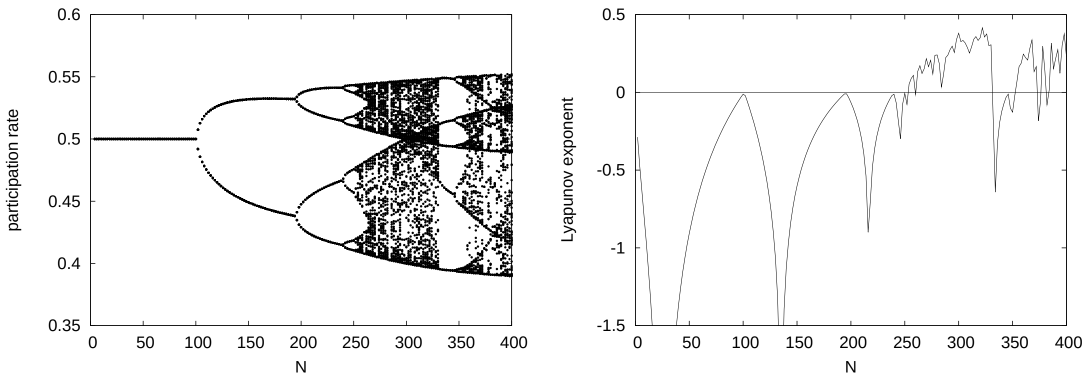

Figure 3.6: *Evolutionary competition of optimist and pessimists. Left pane: the bifurcation diagram with respect to the game size* N *gives the long run behavior of the participation rate as* N *changes. Right panel: Lyapunov exponents for different values of* N*. Parameters values are* $b = N_c/N = 0.5$, $\alpha = 1$, $\gamma = 0.5$ *and* $\beta = 2\gamma$*. For every value of* N, 100 *iterations are used after a transient period of* 100.

In this setting, the left panel of Figure 3.7 suggests that fluctuations increase and thus allocative efficiency decreases as α decreases. Nevertheless, if one fixes α, the left panel of Figure 3.6 suggests that the size of fluctuations of the participation rate is constant in N when more than, say $N = 250$ agents, are involved. This is an important difference with respect to the case when agents are playing the mixed strategy Nash equilibrium. In that case, since agents are randomizing independently, fluctuations of the participation rate approach zero as the number of agents increases. Laboratory experiments that compare fluctuations when different number of agents are playing participation games with increasing size and constant threshold level b can shed light on this issue.

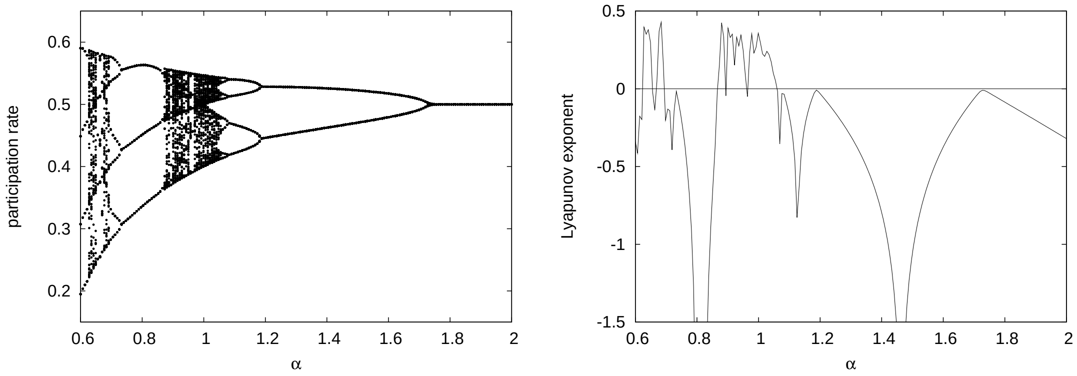

Figure 3.7: *Evolutionary competition of optimist and pessimists. Left panel: long run behavior of the participation rate for values of the granted payoff α. Right panel: Lyapunov exponents for different values of α. Parameters values are $N = 300$, $N_c = 150$, $\gamma = 0.5$ and $\beta = 2\gamma$. For every value of α 100 iterations are used after a transient period of* 100.

Another question we want to address is how the dynamics relates to the details of the competing rules. For this purpose we analyze the effect of the optimistic and pessimistic rule on the time series properties of the participation rate. The autocorrelation plot in the right panel of Figure 3.5 shows a rather strong negative first order autocorrelation and positive second order autocorrelation. The simple rules we have implemented so far are not able to profit from this structure in the data. In particular the positive second order autocorrelation is consistent with the presence of a noisy two cycle for the participation rate historical data. Agents who become aware of the two cycle structure can use this regularity to forecast when it is advantageous to participate and when not. In other words the participation rate series produced by the interaction of optimists and pessimists is not informationally efficient. In the next section, we analyze the change of the participation rate dynamics when a new rule that exploits the regularities produced by the interaction of optimists and pessimists is introduced in the evolutionary competition.

3.5 The scope for arbitrage

In the previous section we have shown that when the game size is large and/or when the net profit for a successful participation is large, the evolutionary competition of optimists and pessimists leads to ongoing fluctuations of the aggregate participation rate around the fixed point s^*. In this case the linear autocorrelation plot in Figure 3.5 shows that there is a positive autocorrelation between participation rates and (even period) lagged participation rates. One can argue that agents might try to exploit this regularity in deciding to participate. The aim of this section is to investigate how the presence of new rules trying to exploit the information contained in the data affects the participation rate dynamics. A similar analysis is performed in Hommes (1998) and in Chapter 2 for the similar negative feedback framework given by the cobweb model.

To start with, consider at time t a best response rule as in (3.7) to the prediction $x_t = x_{t-2}$. In practice, agents using this rule use:

$$p_{2,t} = BR(x_{t-2}) = \begin{cases} 1, & \text{if} \quad x_{t-2} < b, \\ 0, & \text{if} \quad x_{t-2} \geq b. \end{cases} \tag{3.19}$$

where the subscript 2 labels the rule. We call this rule "two lags best reply" since it is the best reply to the observation made two periods before. Notice that this rule is robust to a change of the information set from $\mathcal{I}_{t-1} = \{x_{t-1}, x_{t-2}\}$ to $\mathcal{I}_{t-1} = \{\pi_{I,t-1}, \pi_{I,t-2}\}$ which is the payoff for participating awarded up to two periods before. As a result, a rule which prescribe to go if $\pi_{I,t-2} > \alpha$ and to stay at home otherwise is equivalent to the two lags best reply rule. What is the result of evolutionary competition of optimists and pessimists when some agents mutate to use the two lags reply best response rule?[2]

We model evolutionary competition of these three rules along the lines of the previous section. The participation rate at time t is now given by:

$$x_t = x_{1,t} + x_{2,t} BR(x_{t-2}),$$

where x_1 is the fraction of optimists, x_2 the fraction of two lags best responders and the fraction of pessimists is $x_{0,t} = 1 - x_{1,t} - x_{2,t}$. Payoff for pessimists is always α, payoff for optimists is given by equation (3.17), whereas profits for two lags best responders are obtained using the general formula (3.8):

$$\pi_{2,t} = (1 - BR(x_{t-2}))\pi_O + BR(x_{t-2})\pi_I(x_t).$$

Once profits are realized the fractions are updated using the replicator dynamics as in (3.13) and so on. The resulting dynamical system has dimension 4, two equations govern the updating of the fraction of agents using two of the three rules and two other equations specify the lagged and twice lagged value of the participation rate. As explained at the

[2] The reader might wonder why we do not start this investigation with optimists, pessimists and one lag best responders who use $BR(x_{t-1})$. Intuitively, since the interaction of optimists and pessimists generates a participation rate with strong negative first order autocorrelation, when optimists and pessimists are present one lag best responders are very often wrong in their prediction, since x_t and x_{t-1} are negatively correlated. As a result one lag best reply are quickly wiped out. Simulations confirm this intuition.

end of Section 3.3, due to the presence of optimists and pessimists, the interior equilibrium of such a system is still non generic, and therefore, according to Proposition 3.3, given by s^*. When $x = s^*$ the strategies earn the same payoff so that there are no incentives to "imitate" other strategies. At s^* the two lags best responder is not participating because, having defined in (3.19) $p_{2,t}(b) = 0$, he behaves like a pessimist. For this reason there is a continuum of steady state fractions of agents using the three rules where $x_1^* = s^*$ and $x_0^* + x_2^* = 1 - s^*$. When is this continuum of steady states stable?

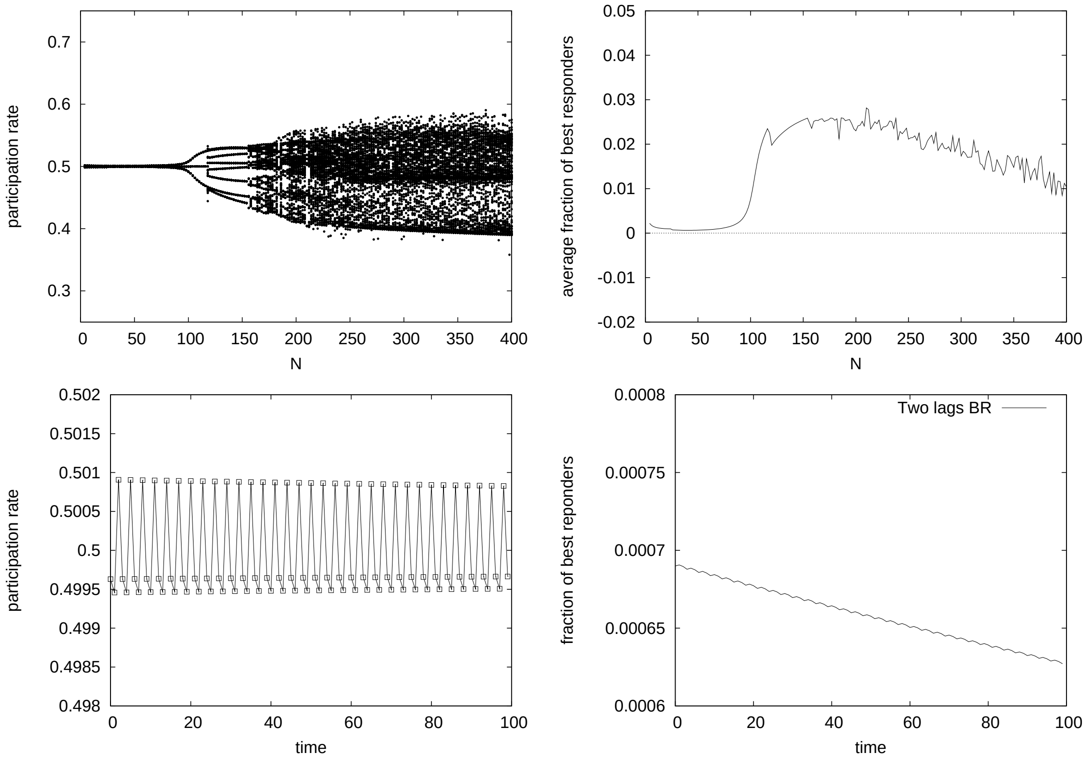

Figure 3.8: *Evolutionary competition of optimists, pessimists and two lags best responders. Top left panel: long run participation rate for different game sizes N. Top right panel: 100 periods time average of the fraction of two lags best responders for different values of N. The time average is computed along the iterations shown in the top left panel. Bottom left panel: participation rate when $N = 50$. Bottom right panel: fraction of two lags best responders when $N = 50$. In all the cases $b = 0.5$ $\beta = 2\gamma$, $\gamma = 0.5$, $\alpha = 1$.*

Given the fact that the rule $p_{2,t}$ is discontinuous and thus not differentiable, we rely on numerical simulations. The top left panel of Figure 3.8 shows the dynamics of the participation rates resulting from the interaction of optimists, pessimists and two lags best responders, for different values of the game size N. As in the case with only optimists and pessimists, the evolutionary competition produces persistent endogenous fluctuations of the participation rate when the game size, N, is large. In particular the system loses stability at the same size, $\widetilde{M} = 100$, as when only pessimists and optimists are present.

The top right panel of the same figure shows, for every size of the game, the time average of the fraction of two lagged best responders along the 100 iterations shown in the top left panel. Interestingly this average fraction is approaching zero for those sizes where the participation rate converges to s^* and stays above zero otherwise. The reason is that when the system converges to the steady state where $x = s^*$, that is when $N < 100$, it does so along a "shrinking" 3-cycle as shown in the bottom left plot of Figure 3.8. A numerical observation is that along such cycles a two lagged best responder is more often wrong than right and slowly disappears (see bottom right panel of Figure 3.8). Notice that, even if x_2 is approaching zero, it might happen that the participation rate x settles to s^* before x_2 is equal to zero. At this point x_2 stops decreasing as the system has reached a steady state. A different participation rate dynamics occurs when the game size is above $\widetilde{M}$. In this case the dynamics is often non periodic and the two lags best responders survive. The left panel of Figure 3.9 shows the resulting participation rate when $N = 300$ and the three rules are competing. Even if the time series is not periodic, the autocorrelation plot (right panel of Figure 3.9) shows that the overall linear autocorrelation structure of the participation rate resembles the one of a noisy 3-cycle: a structure that neither the original nor the new rule can profit from.

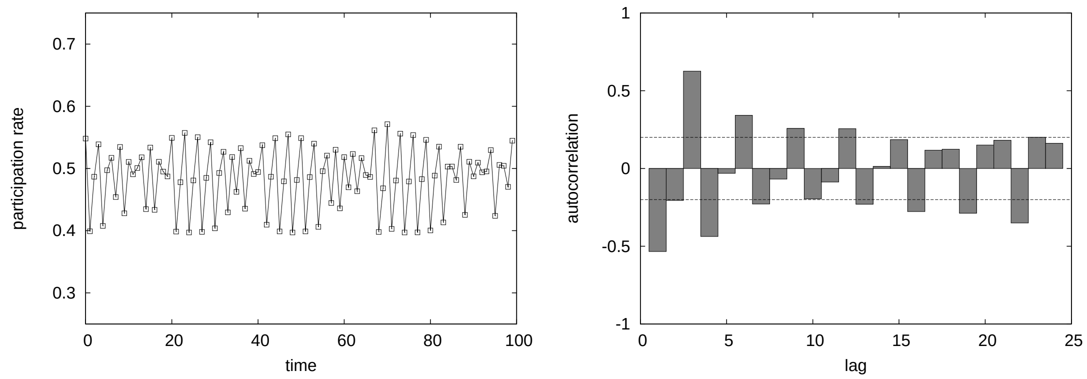

Figure 3.9: *Evolutionary competition of optimists, pessimists and two lags best responders. Left panel: participation rate for game size* $N = 300$. *Right panel: autocorrelation plot for the series shown in the left panel. The dotted line shows the significance level of a standard autocorrelation test with* 100 *observations. Parameters as in Figure 3.8.*

The natural question at this point is: what if another rule appears, a rule that best replies to the presence of a noisy 3-cycle. We define the three lags best responders as those who use:

$$p_{3,t} = BR(x_{t-3}),$$

where the subscript 3 labels the rule. The evolutionary competition of the four rules (optimists, pessimists, two and three lags best responders) can be formalized along the

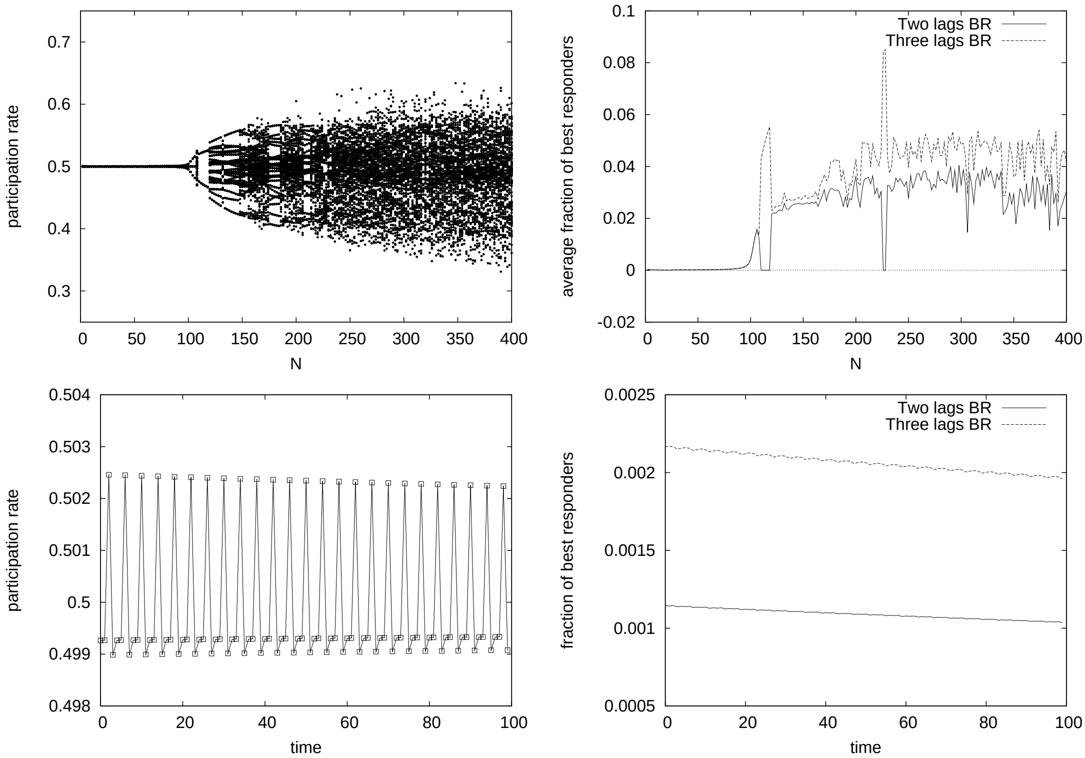

Figure 3.10: *Evolutionary competition of optimists, pessimists, two and three lags best responders. Top left panel: long run participation rate for different game sizes N. Top right panel: 100 periods time average fraction of two lags best responders for different values of game sizes N. The time average is computed along the iterations shown in the top left panel. Bottom left panel: participation rate when $N = 50$. Bottom right panel: fraction of two lags best responders when $N = 50$. In all the cases $b = 0.5$ $\beta = 2\gamma$, $\gamma = 0.5$, $\alpha = 1$.*

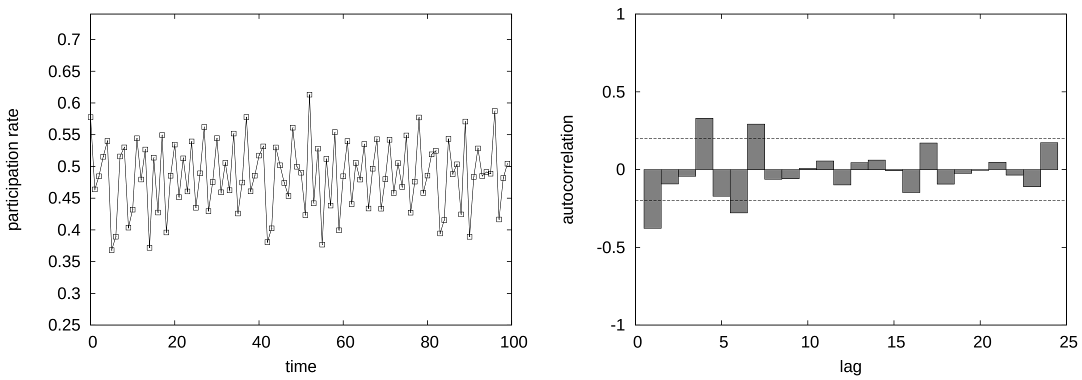

Figure 3.11: *Evolutionary competition of optimists, pessimists, two and three lags best responders. Left panel: participation rate for game size $N = 300$. Right panel: autocorrelation plot for the series shown in the left panel. The dotted line shows the significance level of a standard autocorrelation test with 100 observations. Parameters as in Figure 3.10.*

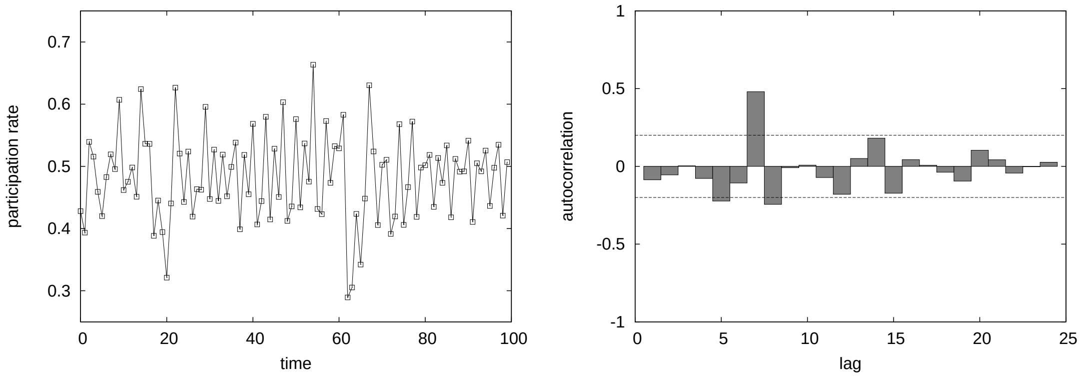

Figure 3.12: *Evolutionary competition of optimists, pessimists, two, three, four, five and six lags best responders. Left panel: participation rate for game size $N = 300$. Right panel: autocorrelation plot for the series shown in the left panel. The dotted line shows the significance level of a standard autocorrelation test with 100 observations. Parameters as in Figure 3.10.*

same lines of the previous example and results in a 6 dimensional dynamical system. The system has 2 dimensions more than the system without three lagged best responders. One additional dimension is due to the presence of a new rule, and thus of the variable x_3, the other dimension comes from the dependence of a further lag, x_{t-3}, of the aggregate participation rate. As in the competition of optimists, pessimists and two lags best responders, there exists a continuum of interior steady states where $x_1 = s^*$ and $x_0 + x_2 + x_3 = 1 - s^*$, so that $x = s^*$. Simulations of this case for different values of the game size N are shown in Figure 3.10. In the top left panel we show the participation rate resulting from 100 iterations of the four competing rules, and in the top right panel the average number of two lags best responders and of three lags best responders along these iterations. As in the competition of optimists, pessimists and two lags best responders, when $N < \widetilde{M} = 100$ the participation rate converges to s^* and the fractions of both types of best responder rules approach zero. When $N > \widetilde{M}$, instead, the dynamics is unstable and both types of best responders survive. The bottom panels of Figure 3.10 concentrate on the case of $N = 50$. The participation rate converges to $x = s^*$ along 4-cycles (left panel) and both fractions of best responders go to zero (right panel). Along a 4-cycle both best responders are more often wrong than right in their predictions and their fractions go to zero. This is not the case when the dynamics is more complicated, in this case, due to non-regular fluctuations, the best responders manage to survive. Figure 3.11 shows 100 iterations of the participation rate with the corresponding autocorrelation plot, in the case of $N = 300$. The pattern detected by linear autocorrelation resembles a four cycle, a structure from which none of the currently present strategies can profit. We are led to argue that, by adding a rule that exploits this regularity, the four lags best responders, yet another periodic structure is to be observed in the time series. Figure 3.12 shows the results of the interaction of two lags up to six lags best responders in a game of size $N = 300$. In this case a 7th order linear autocorrelation is detected. All other lags are below the significance level: they are significantly smaller and mostly within the confidence interval of the standard autocorrelation test when 100 observations are considered (dotted line).

From this exercise, we learn that in a negative feedback framework, the presence of a periodic structure in the time series can be linked to the absence of rules that exploit such a structure. The linear periodic structure which is left in the data is not exploitable by the present strategies. Another way to look at the same effect is that, using the words of Arthur, "cycles are quickly arbitraged away" (p. 409 of Arthur, 1994). Moreover, as we increase the number of competing rules the linear autocorrelation decreases at mosts lags and informationally efficiency improves. We conjecture that in systems where many different agents are competing we might even avoid to detect any significant linear structure in the data without stability emerging.

An unexpected result is that the new rules are surviving only in the unstable region. When the system is stable at the equilibrium only two rules are effectively present, optimists and pessimists. In this case the interaction of optimists and pessimists is thus "stable" against the introduction of other rules, in the sense that no other rules are managing to invade the initial population composed of optimists and pessimists. This situation corresponds to the evolutionary stability result of Proposition 3.2 of Section 3.2. The only difference is that in this context the mixed strategy has to be interpreted as a distribution of population fractions between the two actions. The interaction of optimists and pessimists is not stable against the introduction of other rules when the dynamics is unstable. In this case the new rules survive and affect the global dynamics. Notably, the introduction of new rules gives different outcomes in the stable and in the unstable region.

3.6 The participation premium

Both the market entry experiment discussed in Section 3.2 and the behavioral model from Sections 3.4 and 3.5 exhibit underparticipation, that is, the average participation rates are below the steady state value s^*, when this steady state is unstable. This underparticipation may result in a 'participation premium': the lower right panel of Figure 3.2, for example, shows that subjects in the experiment that participate more often typically earn higher average payoffs. We conjecture that this participation premium results from the payoff asymmetry in the payoffs of the two alternatives in the market entry game.

That is, players choose between a strategically uncertain alternative (participating), for which the payoff depends upon the actions of the other players and a sure alternative (not participating), where the payoff is independent of other player's actions. In this section we consider route choice games to investigate this conjecture. For the route choice game both alternatives are subject to strategic uncertainty and payoffs always depend upon the choices of the other players.

In the route choice game the payoff of player i for participating, action $a = 1$, is the same as the payoff for participating in market entry games, that is according to (3.1):

$$\pi_i\left(1, \boldsymbol{a}_{-i}; N_c\right) = \begin{cases} \alpha + \beta - \gamma, & \text{if } \sum_{j=1, j\neq i}^{N} a_j < N_c, \\ \alpha - \gamma, & \text{if } \sum_{j=1, j\neq i}^{N} a_j \geq N_c. \end{cases}$$

Now also the payoff for abstaining, action $a = 0$, is decreasing in the number of other agents choosing $a = 0$. We consider

$$\pi_i\left(0, \boldsymbol{a}_{-i}; N_c\right) = \begin{cases} \alpha - \gamma, & \text{if } \sum_{j=1, j\neq i}^{N} a_j < N_c, \\ \alpha + \beta - \gamma, & \text{if } \sum_{j=1, j\neq i}^{N} a_j \geq N_c. \end{cases}$$

Sticking to a context of route choices, action 1 can be seen as choosing route 1 whereas action 0 corresponds to choosing route 0. Route 1 is preferred when less than N_c agents are choosing it, otherwise route 0 is preferred. The expected payoff for playing a mixed strategy s_i changes from (3.2) for market entry games to:

$$\pi_i\left(s_i, \boldsymbol{s}_{-i}; N_c, N\right) = \alpha - \gamma + \beta(1 - p\left(\boldsymbol{s}_{-i}; N_c, N\right)) + s_i \beta \left(2p\left(\boldsymbol{s}_{-i}; N_c, N\right) - 1\right). \quad (3.20)$$

There are $\binom{N}{N_c}$ PSNE such that exactly N_c players participate and the other, $N - N_c$, players abstain. In contrast to market entry games, these equilibria are not strict. In fact a player playing $a = 0$ is now indifferent between $a = 0$ and $a = 1$. There are $\binom{N}{N_c+1}$ other PSNE such that $N_c + 1$ players use $a = 1$ and the rest, $N - (N_c + 1)$ players, uses $a = 0$. Also these equilibria are not strict. In fact a player playing strategy $a = 1$ is indifferent between the two pure strategies. Consequently there is an infinite number of asymmetric MSNE. Importantly, as for market entry games, there is only one symmetric MSNE. By equalizing $\pi_i(1, \boldsymbol{s}^*_{-i})$ and $\pi_i(0, \boldsymbol{s}^*_{-i})$ it can be derived that the symmetric mixed strategy equilibrium is $\boldsymbol{s}^* = (s^*, \ldots, s^*)$ with s^* satisfying:

$$p\left(s^*; N_c, N\right) = \sum_{k=0}^{N_c - 1} \binom{N-1}{k} \left(s^*\right)^k \left(1 - s^*\right)^{N-1-k} = \frac{1}{2}. \quad (3.21)$$

We now can compare (3.21) above to (3.4) which is the equivalent condition for the market entry game analyzed in Section 3.2. The symmetric MSNE does not depend upon the ratio γ/β anymore. Notice that in the route choice game, for any value of α, and for any $\gamma < \beta$, the probability s^* is such that the distribution of the participation rate of $N-1$ players playing the MSNE has $N_c - 1$ as median.

Notice that, our formulation of the route choice game is very close to the minority game, a game introduced by physicists (see e.g. Challet and Zhang, 1997 and 1998) as a setting to study coordination of "artificial" agents, using the traditional tools of statistical mechanics. In particular a minority game is a route choice game as defined in (3.20) with N odd and $N_c = (N-1)/2$. A minority game is in fact a participation game with negative feedback where only the group taking the less popular decision is awarded with a prize. See also Bottazzi and Devetag (2003) where an analysis of allocative and informational efficiency of the outcome of an experiment of a minority game with $N = 5$ players is performed.

Our route choice game can also be seen as a simplified version of the prediction game introduced an investigated by Zambrano (2004). In a prediction game every player chooses a number y in $[0, N]$. Every number is awarded with a payoff

$$\pi_y = -|y - y^*|,$$

where y^* is the total number of other players choosing a value y which is below N_c. In order to maximize his payoff a player has thus to predict how many other players are choosing a number less than N_c. The route choice game corresponds to the prediction game when each player has to choose a side, either 0 or 1, and when a player is awarded $\alpha - \beta + \gamma$ only when the total number of other players who have chosen the same side is lower than respectively $N - N_c$ or N_c. Given the correspondence of the route choice game with the prediction game in Zambrano (2004), it is not surprising that also in a prediction game at the symmetric MSNE the median of the distribution of $N-1$ players playing the MSNE is $N_c - 1$.

3.6.1 Evolutionary competition

Following the approach of Sections 3.3 and 3.4, one can model the evolutionary competition of optimists and pessimists also when the underlying structure is given by the route choice game. In this case optimists are those who are always choosing route 1, whereas pessimists are those who are always choosing route 0. The terminology is less fit than in the previous example but we hold on to it for ease of comparison. In terms of predictors, optimists are predicting that route 1 will be the least used, or that less than N_c agents predict that route 1 will be the least used, whereas pessimists are predicting that route

0 will be the least used, or that less than $N - N_c$ agents predict that route 0 will be the least used. As we have shown in Section 3.5 for the market entry game, the analysis can be extended to incorporate rules that try to exploit the linear autocorrelation structure of the participation rate time series. In Figure 3.13 we present the dynamics of the participation rate resulting from the interaction of optimists, pessimists and two lagged best responders. Also in this case two lagged best responders have an effect on the global dynamics only when the size of the game is above the local stability threshold. Notice that the condition for the local stability are different, namely the system loses stability at a lower N. The exact value where the steady state s^* loses stability can be computed as we did in Proposition 3.4 for the market entry game. Also notice that for route choice games fluctuations of the participation rate around the MSNE s^* seem to be symmetrically distributed, and underparticipation indeed disappears.

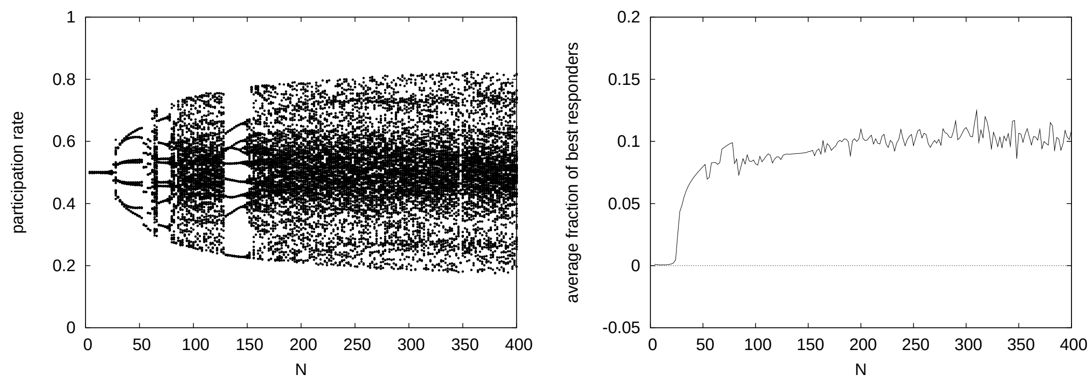

Figure 3.13: *Evolutionary competition of optimists, pessimists and two lags best responders in a route choice setting. Left panel: participation rate for different game sizes N. Right panel: 100 periods time average of the fraction of two lags best responders for different values of N. The time average is computed along the iterations shown in the left panel. Parameters are* $b = 0.5$, $\alpha = 1$, $\beta = 2\gamma$, $\gamma = 0.5$.

3.6.2 Market entry versus route choice

The following result may help us in understanding the relationship between underparticipation and the asymmetric payoff structure of the market entry game.

Proposition 3.6 *Consider an evolutionary competition between the optimistic and pessimistic rule and let $b = 1/2$ and $\beta = 2\gamma$. Denote by $z_t = x_t - s^*$ the deviations of the participation rate from the MSNE and define $\Delta z_{t+1} = z_{t+1} - z_t$. For the market entry game we can write $\Delta z_{t+1} = m(z_t)$, where $m(z) + m(-z) < 0$, for $z \notin \left\{-\frac{1}{2}, 0, \frac{1}{2}\right\}$. For the route choice game we can write $\Delta z_{t+1} = r(z_t)$, where $-r(z) = r(-z)$.*

The left panel of Figure 3.14 shows the functions $m(z)$ and $r(z)$ and illustrates that, outside the three steady states $z = -\frac{1}{2}$, $z = 0$ and $z = \frac{1}{2}$, we have $m(z) < r(z)$. This, together with Proposition 3.6 suggests that there is a tendency for z_t to be downward biased in the market entry game, since innovations in z are lower than in the route choice game (this is corroborated for example in the left panel of Figure 3.7). The origin of the 'asymmetry' of $m(z)$ lies in the denominator of (3.14), the population average payoff $\overline{\pi}_t = \sum_i x_{i,t}\pi_{i,t}$. The right panel of Figure 3.14 shows that average payoffs $\overline{\pi}$ are symmetric in x around $s^* = \frac{1}{2}$ for the route choice game, but not for the market entry game.

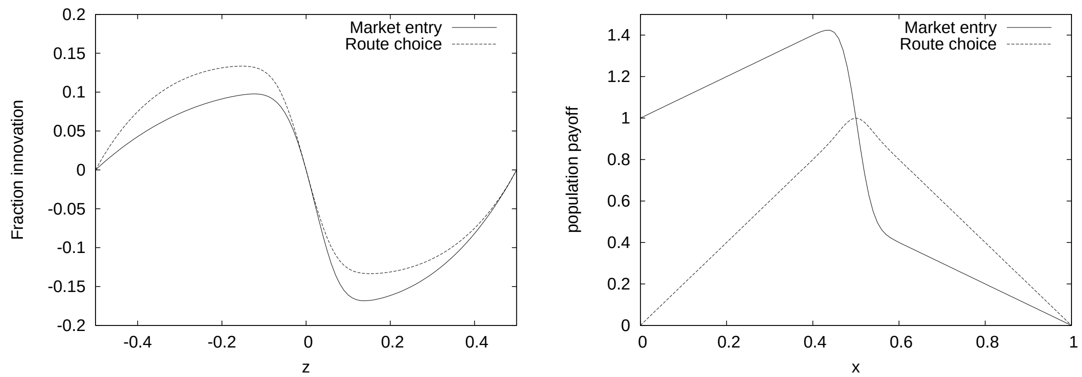

Figure 3.14: *Comparison of the evolutionary competition of optimists and pessimists for market entry games and route choice games. Left panel: Innovation function for the market entry game and for the route choice game compared. Right panel: Population average payoffs, $\overline{\pi}$, in the two cases. Parameters are $b = \frac{1}{2}$, $\alpha = 1$, $\beta = 2\gamma$ and $\gamma = \frac{1}{2}$.*

This asymmetry in $\overline{\pi}$ leads to underparticipation in the market entry game, and a participation premium for those agents using the optimistic rule, since the project is more often profitable. This is illustrated by the left panel of Figure 3.15 which shows the difference between average payoff of optimists and pessimists for the market entry and the route choice game. For the route choice game average payoffs of optimists and pessimists are, due to symmetry, always the same. For the market entry game however, just as in the experiment discussed in Section 3.2, optimists do better on average, whenever the steady state s^* is unstable.

One might expect that this payoff difference in the market entry game disappears when some "memory" is introduced. In fact when memory plays a role more agents should imitate the action of the optimists, which are performing better, and thus eliminate underparticipation. Consider the evolutionary competition of optimists versus pessimists where evolution is governed by a fitness measure that is a weighted average of past payoffs, that is, $F_t = \mu F_{t-1} + (1-\mu)\pi_{1,t-1}$. The resulting dynamical system is two dimensional,

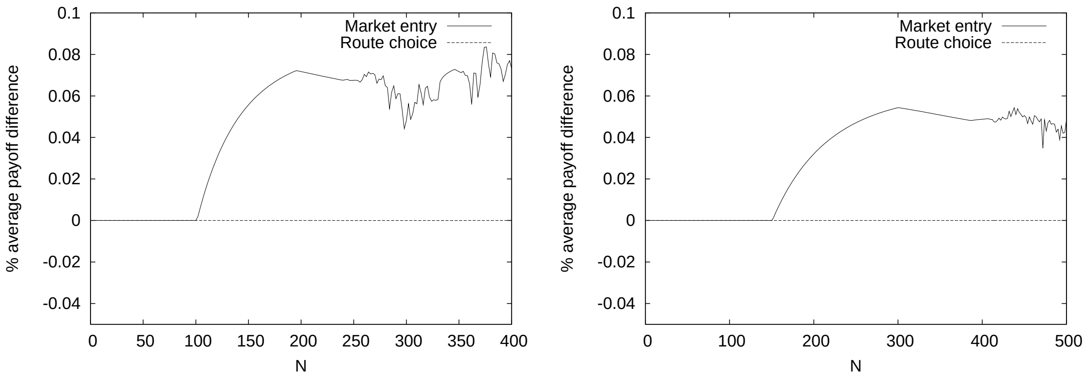

Figure 3.15: *Optimists - pessimists time average payoff differences for market entry games and route choice games. Left panel: Time average payoff difference without memory. Right panel: Time average payoff difference with memory. Parameters are* $b = \frac{1}{2}$, $\alpha = 1$, $\beta = 2\gamma$, $\gamma = \frac{1}{2}$ *and* $\mu = \frac{1}{10}$.

with variables x_t and F_t, and given by:

$$\begin{aligned} x_{t+1} &= f_1(x_t, F_t) = \frac{x_t F_t}{x_t F_t + (1 - x_t)\alpha}, \\ F_{t+1} &= f_2(x_t, F_t) = \mu F_t + (1-\mu)(\alpha - \gamma + \beta p(f_1(x_t, F_t); N_c, N).) \end{aligned} \tag{3.22}$$

The following proposition characterizes its behavior.

Proposition 3.7 *The dynamics of the participation rate and of the optimists' payoff given by the system (3.22) has three steady states:* $(0, \alpha - \gamma + \beta)$, (s^*, α) *and* $(1, \alpha - \gamma)$. *The steady states* $(0, \alpha - \gamma + \beta)$ *and* $(1, \alpha - \gamma)$ *are locally unstable. The interior steady state* (s^*, α) *is locally stable when* $\psi \equiv s^*(1 - s^*)\frac{\delta^*\beta}{\alpha} > -2\frac{1+\mu}{1-\mu}$.

The effect of memory is to stabilize the dynamics. For $\mu = 0$ we retrieve the stability condition from Proposition 3.4, but as μ increases the critical value of ψ increases. Nevertheless, when the dynamics is unstable the same difference in time averages payoffs as before emerges. The right panel of Figure 3.15 shows that when the steady state s^* is unstable over time optimists outperform pessimists. Therefore, the participation premium persists.

Summarizing, the asymmetric payoff-structure of the market entry game is indeed responsible for underparticipation and the participation premium. Note that if we assume agents are driven by payoff differences instead of absolute payoffs, the market entry game transfers naturally into a route choice problem. To see this, consider the payoff function (3.1) and define the payoff difference by $\phi_i(a_i, \boldsymbol{a}_{-i}; N_c, N) = \pi_i(a_i, \boldsymbol{a}_{-i}; N_c, N) - \pi_i(1 - a_i, \boldsymbol{a}_{-i}; N_c, N)$. This payoff difference is equal to $\beta - \gamma$ or $-\gamma$ for participating players, and equal to $\gamma - \beta$ or γ for players that do not participate, depending on the number of participating players. Clearly, $\phi_i\left(1, \boldsymbol{a}_{-i}; N_c, N\right) + \phi_i\left(0, \boldsymbol{a}_{-i}; N_c, N\right) = 0$ and choosing not to participate is subject to strategic uncertainty about payoff differences. Therefore, if players care about payoff differences between the alternatives participation rate dynamics are less stable and underparticipation is alleviated.

3.7 Conclusion

Many (economic) decision problems can be characterised as negative feedback participation games and understanding human behavior in these participation games is important. The experiment from Section 3.2 shows that, although a reasonable description at an aggregate level, the symmetric mixed strategy Nash equilibrium does not explain individual behavior very well. In particular, rather than randomizing their decisions, subjects typically use deterministic rules, possibly conditional on past outcomes. Moreover, the participation rate is inherently unstable, in all experimental groups. A series of contributions from computational economics, starting with the famous *El Farol* bar problem from Arthur (1994), also shows that complicated dynamics arise naturally in negative feedback participation games. Other complex and computationally intensive models assume that agents are randomizing their participation decision. In general, these models are difficult to study analytically and results from this literature are typically based upon numerical simulations.

In this paper we introduce an alternative type of behavioral model that is able to explain the experimental and computational results, but still is sufficiently simple to be analyzed theoretically. We consider an evolutionary competition between different deterministic behavioral rules, where players switch between these rules on the basis of past performance. For the simplest possible case, where the only available two rules are those that specify to always participate, or to never participate, respectively, the participation rate dynamics evolves according to a nonlinear one-dimensional difference equation. This difference equation can be studied analytically, and local stability of the symmetric mixed strategy Nash equilibrium turns out to depend upon the number of players. For a large number of players this simple model exhibits perpetual fluctuations in the participation rate, similar to those found in the experiments and the, much more complex, computational models. A testable prediction of our model is that these fluctuations, in contrast to the mixed strategy Nash equilibrium, are not vanishing even when a very large group of players is involved.

The erratic time series of participation rates has two other features. First, the time series exhibits certain regularities. When agents use rules that try to exploit this structure are introduced, this particular structure disappears, but fluctuations around the symmetric mixed strategy Nash equilibrium do not vanish. Instead, other (higher order) regularities are introduced. Again, adding more sophisticated behavioral rules drives out these regularities again, but does not stabilize the fluctuations, which therefore seem to be quite robust. Secondly, the time series exhibits underparticipation and a premium for participating. This is consistent with the experimental results. We establish that this is due to the asymmetry in the strategic uncertainty of the market entry game. This has interesting economic implications. In our future research we will try to use our behavioral model to explain certain economic or financial stylized facts, such as *excess volatility* and the so-called *equity premium puzzle* (Mehra and Prescott, 1985). Consider, for example, the decision to invest money in bonds, or in an index of stocks as an application of our model. The uncertainty of investing in the stock index is high and may depend on other agents choices, whereas investing in bonds is relatively safe. Our behavioral model predicts an excess return to investing in the stock index. This is consistent with the equity premium puzzle, which refers to empirical evidence that, after adjusting for risk, investing in stocks indeed is more profitable than investing in bonds.

Appendix

3.A Proofs

Proofs of Section 3.2

Proof of Proposition 3.1 The symmetric mixed strategy equilibrium s^* of the game whose payoff for strategy s is given by eq. (3.1) in Section 3.2, solves

$$\pi_i\left(1, \boldsymbol{s}^*_{-i}; N_c, N\right) = \alpha + \left(\beta p\left(s^*; N_c, N\right) - \gamma\right) = \alpha = \pi_i\left(0, \boldsymbol{s}^*_{-i}; N_c, N\right),$$

which is equivalent to:

$$p(s^*; N_c, N) = \sum_{k=0}^{N_c-1} \binom{N-1}{k} (s^*)^k (1-s^*)^{N-1-k} = \frac{\gamma}{\beta}. \tag{3.23}$$

The function $p(x; N_c, N)$ in (3.23) is the cumulative distribution function (c.d.f.) of a binomial distribution, with $N-1$ degrees of freedom and probability of participating x, evaluated at $N_c - 1$. This implies that for every $N_c < N$, it holds true that $p(0; N_c, N) = 1$ and $p(1; N_c, N) = 0$. Furthermore $dp(x; N_c, N)/dx < 0$ as if we increase the probability of participating the value of the c.d.f. at any fixed value between 0 and $N-1$ decreases. Consequently, since $\gamma/\beta < 1$, $p(x; N_c, N) = \gamma/\beta$ has a unique solution s^* for any value of $N > 1$, any $N_c \in [0, N-1]$ and any $\gamma < \beta$. Furthermore, since equation (3.23) does not depend on α, neither does its solution s^*.

The value of s^* is in general different from $b = N_c/N$. We first show that when $b = 1/2$ and $\gamma/\beta = 1/2$ then $s^* = b = 1/2$. In terms of the eq. (3.23) we have to show that $p(1/2; N/2, N) = 1/2$, for all $N > 1$. We show the equivalent statement that $p(1/2; N/2, N) + p(1/2; N/2, N) = 1$. The following chain of equalities proves this statement

$$\begin{aligned}
1 &= \left(\tfrac{1}{2} + \tfrac{1}{2}\right)^{N-1} \\
&= \textstyle\sum_{k=0}^{N-1} \binom{N-1}{k} \left(\tfrac{1}{2}\right)^k \left(\tfrac{1}{2}\right)^{N-1-k} \\
&= \textstyle\sum_{k=0}^{N/2-1} \binom{N-1}{k} \left(\tfrac{1}{2}\right)^k \left(\tfrac{1}{2}\right)^{N-1-k} + \sum_{k=N/2}^{N-1} \binom{N-1}{k} \left(\tfrac{1}{2}\right)^k \left(\tfrac{1}{2}\right)^{N-1-k} \\
&= p(1/2; N/2, N) + \textstyle\sum_{j=N-1-k=0}^{N/2-1} \binom{N-1}{j} \left(\tfrac{1}{2}\right)^{N-1-j} \left(\tfrac{1}{2}\right)^j \\
&= p(1/2; N/2, N) + p(1/2; N/2, N).
\end{aligned}$$

In the last part of this proof we show that for general values of β, of $\gamma < \beta$, of N and of $b = N_c/N \in (0,1)$ it holds true that $s^* \to b$ as $N \to \infty$. Define the random variable $\mathbf{n} = \frac{\mathbf{N-1}}{N-1}$, where $\mathbf{N-1}$ is a random variable with a Binomial distribution with probability s^* and $N-1$ degrees of freedom. Given the fact that s^* solves eq. (3.23), the γ/β percentile of the distribution of $\mathbf{n}$ is given by $bN - 1/(N-1) = b - (1-b)/(N-1)$. Also notice that the distribution of $\mathbf{n}$ has mean s^* and variance $s^*(1-s^*)/(N-1)$. Notice that when $N \to \infty$ the distribution of $\mathbf{n}$ is concentrated more and more around s^*. We reason by contradiction and show that the statement that there exists an $\varepsilon > 0$ such that when $N \to \infty$, $|b - s^*| \geq \varepsilon$, leads to a contradiction. Assume there exists an $\varepsilon > 0$ such that when $N \to \infty$ $|b - s^*| \geq \varepsilon$. This implies that either $s^* > b$ or $s^* < b$. If it were $s^* > b$, then when $N \to \infty$, $\Pr(\mathbf{n} \leq b - (1-b)/(N-1)) \to 0$. This contradicts that s^* has been chosen such that the γ/β percentile of $\mathbf{n}$, is $b - (1-b)/N$. On the other hand if it were $s^* < b$, then when $N \to \infty$, $\Pr(\mathbf{n} \leq b - (1-b)/(N-1)) \to 1$. This also contradicts that s^* has been chosen such that the γ/β percentile of $\mathbf{n}$, is $b - (1-b)/(N-1)$. We conclude that for every $\varepsilon > 0$, $|b - s^*| < \varepsilon$ as $N \to \infty$. □

Proof of Proposition 3.2 We use Proposition 4 in Palm (1984), where the concept of evolutionary stable strategy is extended to N-person games, to show that s^* is the unique evolutionary stable strategy of the game analyzed in Section 3.2. First, for any N, define the $N-1$ dimensional vector $\boldsymbol{x}_{-i}$ as $(x_1, \ldots, x_{i-1}, x_{i+1}, x_N)$. With this notation if player i plays a strategy $s_i = s$ and faces opponents playing $\boldsymbol{x}_{-i}$ then, according to eq. (3.2), his expected payoff is:

$$\pi_i(s, \boldsymbol{x}_{-i}; N_c, N) = \alpha + s\left(\beta p\left(\boldsymbol{x}_{-i}; N_c, N\right) - \gamma\right).$$

Now, since $\boldsymbol{s}^* = (s^*, \ldots, s^*)$ is a mixed strategy Nash equilibrium it holds that

$$\pi_i(s^*, \boldsymbol{s}^*_{-i}; N_c, N) = \pi_i(x, \boldsymbol{s}^*_{-i}; N_c, N) \quad \forall x \in [0,1].$$

This holds $\forall i$ since the game is symmetric. This and Proposition 4 of Palm (1984) imply that the mixed strategy s^* is an ESS if and only if:

$$\pi_i(s^*, \boldsymbol{s}^{*,x}_{-i}; N_c, N) > \pi_i(x, \boldsymbol{s}^{*,x}_{-i}; N_c, N) \quad \forall x \in [0,1] \setminus \{s^*\} \quad \text{and} \quad \forall i, \tag{3.24}$$

where the $N-1$ dimensional vector $\boldsymbol{s}^{*,x}_{-1}$ is defined as $(x, s^*, \ldots, s^*)$. Now, since:

$$\pi_i(s^*, \boldsymbol{s}^{*,x}_{-i}; N_c, N) = \alpha + s^*\left(\beta p\left(\boldsymbol{s}^{*,x}_{-i}; N_c, N\right) - \gamma\right),$$

and

$$\pi_i(x, \boldsymbol{s}_{-i}^{*,x}; N_c, N) = \alpha + x\left(\beta p\left(\boldsymbol{s}_{-i}^{*,x}; N_c, N\right) - \gamma\right),$$

inequality (3.24) becomes:

$$(s^* - x)(\beta p\left(\boldsymbol{s}_{-i}^{*,x}; N_c, N\right) - \gamma) > 0.$$

Consider player i; if one of the other $N-1$ players increases his probability x of participating while the other $(N-2$ use the same probability s^* then $p(s_{-i}^{*,x}; N-c, N)$ decreases. In differential form $dp\left(\boldsymbol{s}_{-i}^{*,x}; N_c, N\right)/dx < 0$. This, together with the fact that, by definition, s^* is the value such that $\beta p\left(\boldsymbol{s}^*{}_{-i}; N_c, N\right) - \gamma = 0$, implies that $\beta p\left(\boldsymbol{s}_{-i}^{*,x}; N_c, N\right) - \gamma > 0$ when $x > s^*$ and vice versa when $x < s^*$. As a result the inequality is always satisfied. The strategy s^* is the unique ESS because no other symmetric Nash equilibrium of the game exist. We notice that this result does not dependent by the size of the game N. □

Proofs of Section 3.3

Proof of Proposition 3.3 The equilibrium condition of eq. (3.13) is given by:

$$x_k^* = \frac{x_k^* \pi_k^*}{\sum_{j=1}^{K} x_j^* \pi_j^*}. \tag{3.25}$$

Obviously, any point with structure $x^* = (0, \ldots, 0, 1, 0, \ldots, 0)$ satisfies this equation. These are the K equilibria where all the agents use only one rule. Now suppose that at least two rules are used by a strictly positive fraction of agents. Define by $K^* \subset \{1, \ldots, K\}$ the set of indexes of these rules. Then from eq. (3.25) it follows that for any $k \in K^*$ it should hold $\pi_k^* = \sum_{j \in K^*} x_j^* \pi_j^*$, and hence π_k^* has to be the same for all $k \in K^*$. According to eq. (3.8) all the payoffs of the K^* rule are the same when:

$$(1-p_k^*)\pi_a + p_k^* \pi_p(x^*; N_c, N) = (1-p_{k'}^*)\pi_a + p_{k'}^* \pi_p(x^*; N_c, N) \quad \text{for any couple of rules} \quad k, k'$$

where $x = \sum_{j \in K^*} x_j^* p_i^*$ and $p_k^* = f_k(x^*, \ldots, x^*)$ for every k. In what we have called the non-generic case, that is when $p_k^* = p_{k'}^*$ for every k, k' in K^* the previous equation is always satisfied. In the generic case, otherwise, all the solution must have π_a equal to $\pi_p(x^*; N_c, N)$ which, is only possible when $x^* = s^*$. □

Proofs of Section 3.4

To prove Propositions 3.4 and 3.5 the following result is useful.

Lemma 3.2 *Define,* $\delta^* = \delta(s^*; b, N,) = \frac{dp}{dx}\big|_{x=s^*}$ *then, for a fixed value of* b, δ^* *is decreasing in* N *and* $\delta^* \to -\infty$ *as* $N \to \infty$.

Proof The cumulative distribution of a binomial distribution with parameters x and $N-1$ evaluated at $bN-1$ can be written in terms of the Beta-function, $B(x,y) = \frac{(x-1)!(y-1)!}{(x+y-1)!}$, as

$$p(x; bN, N) = 1 - \frac{\int_0^x t^{Nb-1}(1-t)^{N(1-b)-1}\,dt}{B(Nb, N(1-b))}. \tag{3.26}$$

Equation (3.26) implies that we have

$$\delta(x; b, N) = \frac{\partial p(x; bN, N)}{\partial x} = -\frac{x^{Nb-1}(1-x)^{N(1-b)-1}}{B(Nb, N(1-b))}.$$

The function $\delta(x; b, N)$ has a unique maximum at $\widehat{x}_{b,N} = \frac{bN-1}{N-2}$. The associated minimum value of δ is given by get

$$\begin{aligned}
\widehat{\delta}_{b,N} = \delta(\widehat{x}_{b,N}; b, N) &= -\frac{(N-1)!}{(bN-1)!\,(N(1-b)-1)!}\left(\widehat{x}_{b,N}\right)^{bN-1}\left(1-\widehat{x}_{b,N}\right)^{N(1-b)-1} \\
&= -\frac{(N-1)\,(N-2)!}{(bN-1)!\,(N(1-b)-1)!}\left(\widehat{x}_{b,N}\right)^{bN-1}\left(1-\widehat{x}_{b,N}\right)^{N(1-b)-1}.
\end{aligned}$$

Taking the logarithm of $-\widehat{\delta}_{b,N}$ and applying the Stirling approximation formula, $\log(n!) = n\log(n) - n + \xi(n)$ where $\xi(n) \to 0$ as $n \to \infty$, we getlog

$$\begin{aligned}
\log\left(-\widehat{\delta}_{b,N}\right) &= \log(N-1) + (N-2)\log(N-2) - (N-2) \\
&\quad - (Nb-1)\log\left(\frac{Nb-1}{\widehat{x}_{b,N}}\right) \\
&\quad - (N(1-b)-1)\log\left(\frac{N(1-b)-1}{1-\widehat{x}_{b,N}}\right) + (N-2) + \xi(N) \\
&= \log(N-1) + \xi(N),
\end{aligned}$$

which goes to ∞ as $N \to \infty$ with the same speed as $\log(N)$ and therefore $\widehat{\delta}_{b,N} \to -\infty$ as $N \to \infty$. For the special case $b = \frac{1}{2}$ we have $\widehat{x}_{b,N} = \frac{1}{2}$ for all values of N and, by Proposition 3.1, $s^* = \frac{1}{2}$ for all even values of N. This implies that for $b = \frac{1}{2}$ we have $\delta^* \to -\infty$ as $N \to \infty$. For the general case with $b \neq \frac{1}{2}$, both s^* and $\hat{x}_{b,N}$ converge to b as

N goes to infinity. Moreover, $\delta(x;b,N)$ is conitnuous in x. Consequently it must be the case that also then $\delta^* \to -\infty$ as $N \to \infty$. □

Proof of Proposition 3.4 A simple computation shows that the equation $f(x;b,N) = x$, with $f(x;b,N)$ given by (3.18) has three steady states: $x=0$, $x=1$ and $x=s^*$, where s^* is the unique number solving $p(s^*;bN,N) = \gamma/\beta$. The derivative of (3.18) is given by

$$f'(x;b,N) = \frac{\alpha\left((1-x)\,x\beta\delta + (\alpha + \beta p(x;bN,N) - \gamma)\right)}{\left(x_t\left(p(x_t;bN,N)\,\beta - \gamma\right) + \alpha\right)^2}. \tag{3.27}$$

From (3.27) it follows immediately that $f'(1) = \frac{\alpha}{\alpha-\gamma} > 1$ and $f'(0) = \frac{\alpha+\beta-\gamma}{\alpha} > 1$ implying that both steady states $x=1$ and $x=0$ are locally unstable. Evaluating (3.27) at s^* gives

$$f'(s^*;b,N) = 1 + \frac{(1-s^*)\,s^*\beta\delta^*}{\alpha} = 1 + \psi,$$

where

$$\psi \equiv (1-s^*)\,s^*\delta^*\beta/\alpha. \tag{3.28}$$

Notice that the negativity of δ^* is always negative (see proof of Lemma 3.2) implies that ψ is also negative and therefore $f'(s^*;b,N) < 1$. The steady state $x = s^*$ is therefore stable if and only if $\psi > -2$. □

Proof of Proposition 3.5 First observe that for any value of α, any $\gamma > \beta$ and any $b \in (0,1)$ the derivative of $f(x;b,N)$, as given in (3.27), goes to $-\infty$ as $N \to \infty$. The latter is true since (3.27) is proportional to δ^* and from Lemma 3.2 it follows that $\delta^* \to -\infty$ as $N \to +\infty$. When $\beta = 2\gamma$ and $b = \frac{1}{2}$, we have $s^* = \frac{1}{2}$ (Proposition 3.1) for every N. As a result (3.27) can be written as

$$f'(s^*;b,N) = 1 + \frac{\delta_N^*\beta}{4\alpha},$$

where

$$\delta_N^* = -\frac{\left(\frac{1}{2}\right)^{N-2}(N-1)!}{\left(\frac{1}{2}N-1\right)!\left(\frac{1}{2}N-1\right)!}. \tag{3.29}$$

Notice that $\delta_2^* = 1$ and that for $N \geq 4$ we have $\delta_N^* = -\frac{3}{2} \times \frac{5}{4} \times \cdots \times \frac{N-1}{N-2}$. Notice that δ_N^* is monotonically decreasing in N and that $\delta_N^* \to -\infty$ as $N \to \infty$ (by Lemma 3.2). We then have that $\psi = \psi_N = -\frac{\beta}{4\alpha}\delta_N^*$. Thus there exists an integer M such that $\psi_N < -2$ when $N > M$. Since this is the local stability condition of s^*, we have proved that there exists an integer M such that s^* is locally stable if and only if $N \leq M$. □

Proofs of Section 3.6

Proof of Proposition 3.6 Recall that a function $f(x)$ is *even* when $f(-x) = f(x)$ for all x and *odd* when $f(-x) = -f(x)$ for all x. Rewriting (3.14) in terms of z gives

$$\Delta z_{t+1} = m(z_t) = \frac{n(z_t)}{d(z_t)} = \frac{\beta\left(p(z_t + \frac{1}{2}; \frac{N}{2}, N) - \frac{1}{2}\right)\left(z_t + \frac{1}{2}\right)\left(\frac{1}{2} - z_t\right)}{\left(z_t + \frac{1}{2}\right)\beta\left(p\left(z_t + \frac{1}{2}; \frac{N}{2}, N\right) - \frac{1}{2}\right) + \alpha}.$$

From (3.3) it follows that $p\left(z + \frac{1}{2}; \frac{N}{2}, N\right) + p\left(-z + \frac{1}{2}; \frac{N}{2}, N\right) = 1$ or $p\left(z + \frac{1}{2}; \frac{N}{2}, N\right) - \frac{1}{2} = -\left(p\left(-z + \frac{1}{2}; \frac{N}{2}, N\right) - \frac{1}{2}\right)$. Together with the fact that $\left(z + \frac{1}{2}\right)\left(\frac{1}{2} - z\right)$ is an even function of z this implies that $n(z)$ is odd, that is, $n(-z) = -n(z)$. Now consider $z \in \left(-\frac{1}{2}, 0\right)$. We then have

$$d(z) = \beta\left(z + \frac{1}{2}\right)\left(p\left(z + \frac{1}{2}; \frac{N}{2}, N\right) - \frac{1}{2}\right) + \alpha > 0,$$

since $p\left(z + \frac{1}{2}; \frac{N}{2}, N\right) - \frac{1}{2}$ is positive if and only if z is negative. Moreover, we then also have

$$d(-z) = -\beta\left(\frac{1}{2} - z\right)\left(p\left(z + \frac{1}{2}; \frac{N}{2}, N\right) - \frac{1}{2}\right) + \alpha > 0,$$

where the inequality follows from the fact that the $d(-z)$ is bounded from below by $-\frac{\beta}{4} + \alpha$, which is positive since $\alpha - \gamma > 0$ and $\beta = 2\gamma$. So for $z \in \left(-\frac{1}{2}, 0\right)$ we have

$$d(z) - d(-z) = \beta\left(p\left(z + \frac{1}{2}; \frac{N}{2}, N\right) - \frac{1}{2}\right) > 0.$$

Summarizing, for all $z \in \left(-\frac{1}{2}, 0\right)$ we have $d(z) > d(-z) > 0$ and $-n(-z) = n(z) > 0$ which implies $-m(-z) > m(z)$. Similarly, for $z \in \left(0, \frac{1}{2}\right)$ it can be shown that $-m(z) > m(-z)$. Combining those we find that $m(z) + m(-z) < 0$, for $z \notin \left\{-\frac{1}{2}, 0, \frac{1}{2}\right\}$.

For the route choice problem straightforward computations from the equivalent of (3.14) lead to:

$$\Delta z_{t+1} = r(z_t) = \frac{2\beta\left(p\left(z_t + \frac{1}{2}; \frac{N}{2}, N\right) - \frac{1}{2}\right)\left(z_t + \frac{1}{2}\right)\left(\frac{1}{2} - z_t\right)}{2z_t\beta\left(p\left(z_t + \frac{1}{2}; \frac{N}{2}, N\right) - \frac{1}{2}\right) + \alpha}.$$

Here the denominator is an even function and the numerator is odd, making $r(z)$ an odd function. □

Proof of Proposition 3.7 The steady states of (3.22) solve $x = \frac{xF}{xF + (1-x)\alpha}$ and $F = \alpha - \gamma + \beta p(x; N_c, N)$. The three steady states (x, F) are given by $(0, \alpha + \beta - \gamma)$, $(1, \alpha - \gamma)$

and (s^*, α), respectively. The Jacobian of (3.22), $J(x, \pi_1)$, is given by

$$\boldsymbol{J}|_{(x,F)} = \begin{pmatrix} \frac{\alpha F}{\overline{F}^2} & \frac{x(1-x)\alpha}{\overline{F}^2} \\ (1-\mu)\,\beta\delta\frac{\alpha F}{\overline{F}^2} & \mu + (1-\mu)\,\beta\delta\frac{x(1-x)\alpha}{\overline{F}^2} \end{pmatrix},$$

where $\overline{F} = xF + (1-x)\,\alpha$. It follows immediately that $J(0, \alpha - \gamma + \beta)$ has eigenvalues $\lambda_1 = \frac{F}{\alpha} = \frac{\alpha+\beta-\gamma}{\alpha} > 1$ and $\lambda_2 = \mu \in (0,1)$ and that $J(1, \alpha - \gamma)$ has eigenvalues $\lambda_1 = \frac{\alpha}{F} = \frac{\alpha}{\alpha-\gamma} > 1$ and $\lambda_2 = \mu \in (0,1)$. These boundary steady states are therefore unstable for any value of μ. The characteristic equation for the Jacobian evaluated at (s^*, α) is

$$\lambda^2 - [1 + \mu + (1-\mu)\,\psi]\,\lambda + \mu = 0, \tag{3.30}$$

where, as before, $\psi \equiv s^*(1-s^*)\frac{\delta^*\beta}{\alpha}$. When $\mu \in (0,1)$ one can show that the eigenvalues are complex as long as $-\frac{(1+\sqrt{\mu})^2}{1-\mu} < \psi < -\frac{(1-\sqrt{\mu})^2}{1-\mu}$. If this condition holds then we have$|\lambda_1| = |\lambda_2| = \mu \in (0,1)$ and the interior steady state is locally stable. If the eigenvalues are real the local stability conditions are given by $\lambda_1 < 1$ and $\lambda_2 > -1$ (where we have labeled the eigenvalues such that $\lambda_1 > \lambda_2$). From (3.30) it follows that $\psi(1-\mu) < 0$ implies that $\lambda_1 < 1$ always holds. On the other hand, $\lambda_2 > -1$ as long as $\psi > \psi^* = -2\frac{1+\mu}{1-\mu}$. Moreover, since for $\psi \leq \psi^*$ the eigenvalues are real it follows that the interior steady state (s^*, α) is locallys stable if and only if $\psi > \psi^*$. □

Chapter 4

Informational differences in an asset market

4.1 Introduction

Since the beginning of the eighties, the validity of the efficient market hypothesis has been questioned on the basis of empirical evaluation of so-called financial anomalies. Well-known examples are excess volatility, as described by Shiller (1981) and LeRoy and Porter (1981), mean reversion of asset prices, as documented by Poterba and Summers (1988) and Fama and French (1988b), and correlation between returns and lagged returns or lagged dividend yields, as shown by Shiller (1984) and Fama and French (1988a).

Stimulated by these findings, a part of the scientific community has investigated whether such anomalies can be explained by assuming that the agents operating in the market are boundedly rational. Although the exact implication of bounded rationality varies among the different models, a common characteristic is that boundedly rational agents act in an economic setting which they do not know in full detail. Furthermore, agents are often assumed to be able to optimize an objective function under certain constraints but unable to optimally anticipate the effect of their and other agents' actions. In particular, boundedly rational agents are not assumed to be able to coordinate their actions such that their beliefs are perfectly self fulfilling. In other words, expectations of boundedly rational agents need not be rational.

In order to explain fluctuations in prices that are not due to fluctuations in economic fundamentals, one class of models with boundedly rational agents concentrates on the interaction of agents choosing different expectation schemes or different investment strategies (see e.g. the survey paper of Hommes, 2006). An early example is Chiarella (1992) where a model of a stylized financial market with fundamentalists and chartists is shown to generate a number of dynamic regimes which are compatible with the empirical anomalies reported above. In a more recent paper, Brock and Hommes (1998) assume that agents do not know whether it is more profitable to predict prices by relying on fundamental information, or to extrapolate trends. In their model, agents use realized profits or a similar performance measure to decide which strategy to use. This ongoing choice for the best strategy leads to complex endogenous price fluctuations. Endogenous fluctuations can already arise in the absence of exogenous influences such as time-varying fundamentals. An advantage of this assumption is that the price dynamics can be specified in terms of deviations from a benchmark fundamental price.

A limitation of models in this class is that they typically do not take into account the direct effect of news about the economic fundamentals on agents' behavior thus excluding one of the most trivial behavioral scenarios one might deem important in asset price formation — the over- or under-reaction of agents, and hence the market, to new information. Generally speaking, market models that tend to a stable equilibrium state in the absence of news can still show fluctuations triggered by the arrival of new information. Because *a priori* we do not know if market fluctuations are necessarily self-perpetuating as in chaotic dynamics, we explicitly wish to examine the role of exogenous noise on the price dynamics, thus keeping open the possibility of scenarios where ongoing market fluctuations require repeated triggering by a sequence of exogenous shocks.

In view of this critique, there is another class of models in the literature of asset markets with boundedly rational agents, which explicitly takes into account the role of news on fundamentals in the price dynamics. Early examples are Bulkley and Tonks (1989) and Barsky and De Long (1993) who investigate the effect of agents trying to learn the growth rate of dividends from movements in the stock price. More recent examples are Timmermann (1993), Timmermann (1996) and Barucci, Monte, and Renó (2004), who assume that agents estimate parameters defining the relationship between prices and dividends. In all these cases, agents use the rational expectations relationship that would hold between endogenous variables (prices) and exogenous variables (dividends) as if the underlying parameters were known. That is, agents do not take into account that their learning effort is modifying the way dividends feed back into prices and do not take into account estimation errors. When new information about dividends becomes available, it influences returns not only directly but also indirectly as it affects the estimates of the parameters that the agents use to forecast future prices and/or dividends. The expectations in these models converge to rational expectations when the agents learn the parameters of the data generating process. A limitation of these models is that they all assume the presence of a representative agent, so that agents' interaction triggered by informational differences or by different expectations does not play a role. Moreover, due to the stochastic components associated with the incoming news about the fundamentals, results are practically always obtained by means of simulations.

The aim of this chapter is to construct a framework for examining markets with boundedly rational agents where both parameter estimation and interaction of agents play a role. When we use boundedly rational agents, we do not question that rational behavior, and especially rational expectations, can be a good approximation to the equilibrium of agents' repeated interaction. Rather, we argue that the convergence to such an equilibrium is worth investigating as it might explain part of the economic fluctuations we observe in reality. Our objective is to characterize how both parameter estimation and interaction of agents transform incoming information into realized market prices. Because it is impossible to carry out this exercise under all conceivable behavioral assumptions, we limit ourselves to a simple class of agent models, where all agents act upon the information available to them regarding fundamentals (including that revealed by prices). The agent interaction is triggered by different expectations, which can be explained by different degrees of information regarding the future value of dividends. This means that agents neither extrapolate price trends or use other chartists' rules per se, nor expect other agents to do so, so that second or higher-order expectations play no role.

We investigate the extent to which our model is able to explain empirical properties of asset prices. As it turns out, our price dynamics driven by exogenous noise leads to a simple econometric model for prices that can account for several well-documented anomalies such as autocorrelation of returns and large persistent deviations of the market price from the fundamental price in the short run but convergence to it in the long run. In fact, in line with the econometric model proposed by Summers (1986), our model leads to a (log) price which is the sum of a persistent component, proportional to the (log) dividend, and of a transitory component, proportional to the (log) dividend yield, which turns out to follow a stationary AR(1) process with a time-varying AR(1) coefficient. Our analysis shows that, whereas the transitory component follows an AR(1) process as a direct consequence of agents' learning the dividend growth rate, the fact that the AR(1) coefficient is time-varying is due to agent interaction triggered by informational differences. This offers theoretical support to the empirical evidence that the temporary component in a mean reversion model is nonlinear and switching between regimes with different rates of convergence, as documented both by Gallagher and Taylor (2001) and by Manzan (2003).

A convenient feature of our model is that it contains two important benchmarks as special cases. The first benchmark is given by the classical asset pricing model: the equilibrium price we derive coincides with the correct present value price when one discards both the role of informational differences and of agents learning the growth rate of dividends. In this case log prices follow a random walk process driven by realized dividends.

The second benchmark is given by the model developed in Barsky and De Long (1993) where log prices have a persistent component (random walk) plus a stationary component. Our equilibrium price coincides with the price derived in the model of Barsky and De Long when we discard the role of informational differences. Our model can thus be seen as an extension of the Barsky and De Long model, where agents with different information sets are active in the market.

As we consider an asset market where agents have different degrees of information, our framework is closely related to that of Grossman and Stiglitz (1980) (henceforth GS). They investigate whether the price is informationally efficient in a repeated market for a one period living asset, in which agents can decide between two different degrees of information about the value of the asset return at the end of the period. As GS, we also assume that agents operating in the asset market can decide whether or not to be informed about next period's dividend. However as Bray (1982), Hellwig (1982) and Routledge (1999), we consider a dynamic model rather than a static one. Rather than starting off by assuming that agents have rational expectations, we merely see rational expectations as a possible long run outcome of a learning process of boundedly rational agents using simpler rules. Failure of the uninformed agents to learn the relationship between prices and dividends implies deviations of the price from its fundamental value. Moreover, we assume that the fractions of informed and uninformed agents are not constant but change over time based on past performances of both strategies. The fraction of each type of agents is thus an endogenously determined variable. Another difference with the framework of GS and the papers mentioned above is that we model a market for an infinitely living asset rather than of a sequence of identical markets for a single period asset. This implies that agents need to form expectations not only on the value of the future dividend but also on the remaining value of the asset. To our knowledge, Goldbaum (2005) is the first to consider a dynamic multi-period market in a setting where agents have different degrees of information. Whereas Goldbaum assumes the dividend to be stationary in differences, in order to stay closer to real data, we assume the asset's dividend to be stationary in log-differences. Accordingly, we choose to derive our agents' demand from mean variance maximization of a constant relative risk aversion (CRRA) utility function rather than from a constant absolute risk aversion (CARA) utility function as Goldbaum (2005). Other papers that model the interaction of agents with CRRA utility function are e.g. Chiarella and He (2001) and Anufriev, Bottazzi, and Pancotto (2006).

The impact of news in a financial market, where heterogeneous boundedly rational agents are operating, has also been investigated by Hong and Stein (1999). Their model reproduces stock prices that are under-reacting in the short run, and over-reacting in the long run to the arrival of new information. In contrast to our framework, they assume that the fractions of agents of each group are constant and that the dividend generating process is stationary in levels.

Our model shows that financial markets populated by agents with different degrees of information can be seen as economic systems with negative feedback as those analyzed in the previous two chapters. This establishes a precise correspondence with the famous cobweb model, see Ezekiel (1938), and with the literature that originated from it, such as Muth (1961) and Brock and Hommes (1997). Our equation for the evolution of the dividend yield as a function of the uninformed agents' expectations has a close correspondence with the equilibrium price equation in a cobweb model in the case of linear supply and linear demand. We refer to this literature to justify the expectation formation of boundedly rational agents. In particular Brock and Hommes (1997) and Chapter 2 in this thesis show that if rational expectations come at a cost, boundedly rational agents keep switching between an expensive sophisticated and a cheap simple expectation scheme, thus generating complicated price fluctuations. Because we want to keep our model as simple as possible, we do not explicitly model agents' choice between cheap simple expectations and expensive sophisticated expectations. We concentrate on informational difference and model expectations as adaptive. An analysis where both informational differences and the role of expectation scheme choice play a role is performed by De Fontnouvelle (2000). He shows that if agents are allowed to switch among different types of expectation schemes and if rational expectations come at a cost, an asset market of the type proposed by GS leads to similar price fluctuations as Brock and Hommes (1997) found for the cobweb model. By considering informational differences as well as the choice of an expectation scheme, even in the simpler case of a one period living asset, De Fontnouvelle arrives at a rather complicated model, which is analyzed mostly by means of simulation, rather than analytically.

The chapter unfolds as follows. Section 4.2 introduces the model in terms of its three founding elements: the asset market, expectation formation and the evolution of the fractions of informed and uninformed agents. Section 4.3 analyzes the co-evolution of the market price and of the fractions of informed and uninformed agents in a world without uncertainty about future growth rates of dividends. That is, we analyze the deterministic skeleton of the system of difference equations developed in Section 4.2. Section 4.4 analyzes the full model, i.e. the evolution of the market price and of fractions of agents when uncertainty about future growth rates of dividends plays a role. There we also relate the price dynamics generated by our model to the classical asset pricing model and to the model developed in Barsky and De Long (1993). Section 4.5 concludes with a discussion of our main results. The appendix at the end of the chapter contains proofs and a micro-foundation of our model.

4.2 The model

4.2.1 The asset market

We consider a market where shares of a financial asset are traded repeatedly in discrete time periods. The market is populated by agents who believe that the discounted sum of expected future dividends constitutes a "fair" value of the asset. As in GS, every agent can decide whether or not to buy information about next period's dividend. As a result, in every period the population of agents is divided in two groups with different degrees of information concerning fundamental variables. The current setting differs from GS in that the asset represents a claim on an infinite sequence of future dividends rather than on a single dividend, that is, the asset is infinitely lived and does not perish at the end of the period. As a consequence agents, besides forming expectations on dividends, must also form expectations on future asset prices. Another important difference with respect to the GS framework is that in our model agents are boundedly rational. By this we mean that agents are unable to compute the equilibrium relationship between price and dividends that should arise in the market where informed and uninformed agents operate. The aim of this section is to characterize how, in this setting, the market price of an asset/share, p_t, and the fraction of informed agents, λ_t, co-evolve given agents' expectations and the dividend process $\{d_t\}$. In order to arrive at such a relationship we build up our model starting from stating its underlying ingredients.

We take a stochastice dividend process, $\{d_t\}$. In the benchmark case $\{d_t\}$ is given by a geometric random walk. The dividend payed at time t, d_t, is given by:

$$d_t = d_{t-1}(1 + g)\eta_t, \tag{4.1}$$

where $\{\eta_t\}$ is a sequence of independent, identically distributed (i.i.d.) random variables with mean 1 and variance σ_η^2. The constant g is referred to as the growth rate of dividends.

Agents know that the dividend is growing over time at a certain rate which they estimate using past dividend realizations. We let g^e denote their estimate (beliefs) of the growth rate of dividends. We assume that this belief is the same across agents and that agents use it for predictions "as if" it is the true value in the dividend generating process 4.1. For the moment we consider g^e as given and constant. In Subsection 4.2.2 we will discuss how agents actually estimate the growth rate g.

All agents are "fundamentalists" in the sense that they follow the present value model, i.e. the discounted sum of all future dividends is their "fair" value of the asset. The exact relationship between today's price and tomorrow's expected dividend depends on the agent's information about future dividends. The information set of agents contains past dividend and price realizations and is different across groups of agents. We denote the information set at time t for a group H as $\mathcal{F}_t^H$. The fair value, i.e. the expected discounted sum of future dividends, conditional on $\mathcal{F}_t^H$ is denoted by v_t^H:

$$v_t^H = E\left[\sum_{i=1}^{\infty} \frac{d_{t+i}}{(1+r)^i} \,\Big|\, \mathcal{F}_t^H\right]. \tag{4.2}$$

The coefficient $(1+r)$ is the discount rate or required rate of return. We assume that the discount rate is the same across agents and that agents use the same discount rate for all future periods. The latter assumption is made because we want to concentrate on price fluctuations determined by agents' interaction and learning rather than agents changing their discount rate. In general, the discount rate can be characterized as the sum of the risk free rate and the risk premium, which depends on the risk preferences of agents. In this case, to state that agents use the same constant discount rate means that agents have the same constant risk preferences. See Appendix 4.A for a derivation of the risk premium in a context where preferences of the agents are explicitly taken into account. We also assume that the discount rate is always bigger than the agents' estimate of the long run growth rate of dividends.

At any time t, each agent can decide whether to buy information about the value of d_{t+1} or not to buy it and thus remain uninformed. As a result, in every period there are two groups of agents having a different degree of information regarding the next realization of the dividend process. At time t, the informed agents, group I, are fully informed regarding d_{t+1}. This implies that their current expectation of the $t+1$ dividend is

$$d_{t,t+1}^{e,I} = d_{t+1}, \tag{4.3}$$

where the superscript e, I stands for expectations of the informed agents. We assume that they pay a fixed cost $c > 0$ per period for this information. The uninformed agents, group U, do not know d_{t+1} but can use public information, available in the form of realized dividends d_s and realized prices p_s, $s \leq t$, to form their expectations, $d_{t,t+1}^{e,U}$, about d_{t+1}. The superscript e, U stands for expectations of the uninformed agents. For example, if uninformed agents relied solely on the public belief of the dividend growth rate, they could use $d_{t,t+1}^{e,U} = (1 + g^e)d_t$. The alternative that we consider here, is that uninformed agents try to get additional information revealed by the demands of the informed agents through the current market price p_t. Uninformed agents consider the relationship between the dividend and the price to be of the form:

$$d_{t,t+1}^{e,U} = y^e p_t, \tag{4.4}$$

where y^e is the agents' belief, or estimate, of the market dividend yield, that is, the ratio between future dividend and current price. As for g^e, we start with considering y^e as given and fixed. In Subsection 4.2.2 we will discuss how agents actually form their beliefs of the market dividend yield. We let λ denote the fraction of informed agents, so that $1 - \lambda$ is the fraction of uninformed agents. A subscript t is added when we consider a time dependent λ. We use time varying fractions only from Subsection 4.2.3 where we describe how λ_t evolves endogenously.

At each time t, the ex-dividend market price of one share, p_t, is given by the following market equilibrium pricing equation:

$$p_t = \lambda v_t^I + (1 - \lambda) v_t^U, \tag{4.5}$$

where v_t^I and v_t^U denote the "fair" value of the asset conditional on the information of the informed and uninformed respectively, derived below. Under this assumption the realized price today is a weighted average, with weights equal to the fraction of each agent type, of the agent's estimate of the fair share value. Although this market equilibrium pricing equation is admittedly stylized, it can be derived by assuming that agents can choose to invest in a risky asset and in a risk free bond and use a mean variance utility to decide how much of their wealth to allocate in each investment. If one starts from such a micro-foundation of demands, equation (4.5) can be derived in a Walrasian framework where λ (1-λ) is the fraction of wealth of informed (uninformed) agents. See Appendix 4.A for a derivation of (4.5).

The next step is to derive the implication of agents' behvaior on the price dynamics in (4.5). First, we compute the fair value for the informed and for the uninformed agents.

Informed agents We have said that informed agents, like all other agents, know that the dividend is growing over time at a certain rate, which they assume to be equal to g^e and which they use for predictions "as if" it is the true value of the growth rate of the dividend process. Expectations of future share values are directly linked to expected future dividends through equation (4.2). At time t the informed agents know the value of d_{t+1} so that their information at time t is given by $\mathcal{F}_t^I = \{d_{t+1}, d_t, \ldots, p_t, p_{t-1}, \ldots\}$. Hence their expectations of future dividends are:

$$d_{t,t+j}^{e,I} = d_{t+1}(1 + g^e)^{j-1}, \quad \text{for} \quad j \geq 1. \tag{4.6}$$

Notice that agents treat their estimate g^e as if it is the true value of the growth rate of dividends, that is, they do not take into account possible estimation errors in their dividend predictions. Plugging expectations (4.6) into (4.2) we arrive at the informed agents' estimate of the value of the stock,

$$v_t^I = \frac{d_{t+1}}{(r - g^e)}. \tag{4.7}$$

The informed agents' stock valuation is thus in accordance with the Gordon model (see e.g. Gordon, 1962). Equation (4.7) implies that informed agents are behaving similarly to so-called fundamentalists in the interacting agents literature, see Hommes (2006) and LeBaron (2006) for recent surveys. In fact, the informed agents' fair value of the asset is proportional to the dividend payed at time $t+1$. For this reason, we shall refer to v_t^I, the fair price of the informed agents, as the fundamental price p_t^*, that is, we define:

$$p_t^* = \frac{d_{t+1}}{(r - g^e)}. \tag{4.8}$$

Notice that the value of the fundamental price depends on g^e, the agent's common belief of the growth rate of dividends.

Uninformed agents Boundedly rational uninformed agents try to infer the value of d_{t+1} from the market clearing price p_t. In doing so, they use their model (4.4) concerning the relationship between the current realized market price p_t and expected future dividends d_{t+1}. Combining this with agents being fundamentalists and using g^e to estimate the growth rate of dividends we obtain

$$d_{t,t+j}^{e,U} = y^e p_t (1 + g^e)^{j-1}, \quad \text{for} \quad j \geq 1, \tag{4.9}$$

which, using (4.2), and the uninformed information set $\mathcal{F}_t^U = \{d_t, d_{t-1}, \ldots, p_t, p_{t-1}, \ldots\}$, gives the uninformed agents' valuation of the value of the stock:

$$v_t^U = \frac{y^e p_t}{(r - g^e)}. \tag{4.10}$$

Notice that also for uninformed agents, there is a correspondence with the literature on interacting agents. In fact, our "fundamentalists" uninformed agents are behaving "as if" they are chartists, that is, they use current prices to estimate the value they attach to the asset. This is an important characteristic of our model and we anticipate some consequences here. Consider g^e and y^e as given. If at time t, $y^e/(r - g^e)$ is bigger than one, uninformed agents behave "as if" they are trend follower, and can drive prices well above the fundamental levels. In this case the realized dividend yield, $y_{t+1} = d_{t+1}/p_t$, will become lower than $(r - g^e)$. The converse happens when $y^e/(r - g^e) < 1$. In this case the uninformed agents behave “as if” they are contrarian. Summarizing, the uninformed agents behave “as if” they are chartists but with a different trend coefficient for different values of g^e and y^e.

Given the asset valuation for both types of agents, specified in (4.7) and (4.10), and the market equilibrium pricing equation (4.5) we get,

$$p_t = \frac{d_{t+1}}{(r-g^e)} \frac{\lambda(r-g^e)}{(r-g^e-(1-\lambda)y^e)}. \tag{4.11}$$

This equation shows that the realized price is proportional to the fundamental price p_t^* defined in (4.8) – a similar result one gets through the Gordon model – but there is an additional factor due to the presence of the uninformed agents trying to extract information from the market price. One can consider this as a generalization of the Gordon equation to a simple setting where agents have different degrees of information. Moreover, when $\lambda = 1$, i.e. all agents are informed, the realized prize p_t in (4.11) is equal to the fundamental price p_t^* in (4.8), as in the Gordon model. On the other hand when $\lambda = 0$, i.e. all agents are uninformed, either $p_t = 0$ or, if $y^e = r - g^e$, p_t is indeterminate. Thus, when $\lambda = 0$ our pricing equation (4.11) is not well defined, as the price contains no information about the dividend.

Equation (4.11) can be rewritten as a relationship between the realized dividend yield $y_{t+1} = d_{t+1}/p_t$, agents’ beliefs g^e and y^e, and the fraction of informed agents’ λ:

$$y_{t+1} = \frac{(r-g^e)}{\lambda} - \frac{(1-\lambda)}{\lambda} y^e := f(y^e, g^e, \lambda). \tag{4.12}$$

We call the map f an expectational feedback map because, given a fraction of informed agents λ and common beliefs about the growth rate of dividends g^e, it establishes a feedback between expectations of uninformed agents of the dividend yield and dividend yield realizations. Using the expectational feedback map (4.12), it can be easily derived that, for any $\lambda \in (0, 1]$, if $y^e = r - g^e$ then $y_{t+1} = r - g^e$. When this is the case, the agents’ belief regarding the dividend yield is self-fulfilling and $r - g^e$ is thus the expectational feedback equilibrium or rational expectation dividend yield. Notice that when $y^e = r - g^e$, using equation (4.11), one finds that the market price equals the fundamental price p_t^*, which we have defined in (4.8) as the fair value of the informed agents. For this reason we denote

$$y^* = r - g^e \tag{4.13}$$

as the fundamental dividend yield. The rational expectations dividend yield is thus equal to the fundamental dividend yield and, as we will specify later, it gives approximately the same price process as derived by Timmermann (1993) or by Barsky and De Long (1993) who also consider a model where agents update their estimate of the dividend growth rate g. The novelty here is that, due to the presence of informed and uninformed agents, y_{t+1} may fail to be equal to y^*. In particular the presence of uninformed agents can generate a price that differs from the fundamental price. In fact, equation (4.12) establishes a negative feedback system between the realized dividend yield y_{t+1} and its belief or estimate y^e, since $\partial f/\partial y^e = -(1-\lambda)/\lambda < 0$. This implies that positive (negative) deviations of y^e from y^*, lead to negative (positive) realized deviations of y_{t+1} from y^*. This observation links our model to the classical cobweb model (see e.g. Ezekiel, 1938, for an early treatment). In fact equation (4.12) for the price dividend ratio is the same as the equilibrium price equation in a cobweb model with linear supply and linear demand. The general asset price dynamics will be more complicated however, since, according to (4.12), asset prices will be also driven by the learning of the growth rate of the dividend process and by the evolution of agents' fractions as we shall see.

4.2.2 Expectation formation

Here we consider how agents form expectations on the growth rate of dividends and on the dividend yield. As far as the growth rate of dividends is concerned we have assumed that both informed and uninformed agents have homogeneous expectations on the dividend growth rate. We follow Barsky and De Long (1993) and assume that agents use adaptive expectations to estimate its long run value. Adaptive expectations are characterized by

$$g^e_{t,t+1} = \gamma g^e_{t-1,t} + (1-\gamma)\left(\frac{d_t}{d_{t-1}} - 1\right), \tag{4.14}$$

where $g^e_{t,t+1}$ denotes the time t belief or estimate of the long run growth rate for period $t+1$, and where we refer to γ as the memory coefficient specifying the rate at which agents discount past information. Naive expectations are obtained in the special case $\gamma = 0$, whereas $g^e_{t,t+1}$ is the mean of all past observations of y when $\gamma = (t-1)/t$. Notice that when the growth rate is time-varying, one has to update the definition of the fundamental price and of the fundamental dividend yield from expression (4.8) and (4.13) to, respectively:

$$p^*_t = \frac{d_{t+1}}{r - g^e_{t,t+1}}, \tag{4.15}$$

$$y^*_{t+1} = r - g^e_{t,t+1}. \tag{4.16}$$

The same expectation used for the estimation of the future growth rate of dividends is assumed to be used for the estimation of the value of the future dividend yield y^e. Notice that only uninformed agents need to form expectations about tomorrow's dividend yield, as informed agents already know the value of d_{t+1}. Adaptive expectations for the dividend yield are specified by:

$$y^e_{t,t+1} = \alpha y^e_{t-1,t} + (1-\alpha) y_t, \tag{4.17}$$

where, as before, $y^e_{t,t+1}$ denotes the belief or estimate of the dividend yield of time $t+1$ based on the information up to time t, and the parameter $\alpha \in [0,1]$ is, as γ, the memory parameter, which specifies the rate at which agents discount past information.

Substituting the time varying expectations of the growth rate of dividends and of the dividend yield into the expectational feedback map (4.12) one obtains:

$$y_{t+1} = f(y^e_{t,t+1}, g^e_{t,t+1}, \lambda) = \frac{(r - g^e_{t,t+1})}{\lambda} - \frac{(1-\lambda)}{\lambda} y^e_{t,t+1}. \tag{4.18}$$

This equation evaluated at time t together with (4.17) explicitly gives the dynamics for the expectations of the dividend:

$$y^e_{t,t+1} = \frac{(1-\alpha)(r - g^e_{t-1,t})}{\lambda} + \frac{\alpha + \lambda - 1}{\lambda} y^e_{t-1,t}.$$

This dynamic equation can also be written in terms of the realized dividend yield, y, using the feedback map (4.12), which gives:

$$y_{t+1} = \frac{r - g^e_{t,t+1}}{\lambda} - \frac{\alpha(r - g^e_{t-1,t})}{\lambda} + \frac{\alpha + \lambda - 1}{\lambda} y_t := h(y_t, g^e_{t,t+1}, g^e_{t-1,t} \lambda). \tag{4.19}$$

The map h establishes the dependence of the realized dividend yield y_{t+1}, on its lagged value y_t, on agents' beliefs of the growth rate of dividends at two subsequent dates, $g^e_{t,t+1}$ and $g^e_{t-1,t}$, and on the fractions of informed agents λ. We study this linear map in Section 4.3.

Notice that, given the negative expectational feedback map (4.12), adaptive expectations for the dividend yield are not rational in the sense of Muth (1961). We assume that boundedly rational agents use them for two reasons. First, empirical support in favor of the use of adaptive expectations for the dividend yield in present values mode has been given by Chow (1989). Second, Brock and Hommes (1997) and Chapter 2 of this thesis show that with endogenously determined variables as price or dividend yield, if rational expectations come at a cost, agents may switch continuously between costly rational expectations and simpler expectations at no cost. As a result, when one models this expectation choice, convergence to a rational expectations equilibrium is not ensured.

In order to keep the structural form of our model simple, we do not model agent's choice of its expectation framework and assume that boundedly rational agents use adaptive expectations. In fact adaptive expectation gives a reasonable trade-off between simplicity of use and consistency with the outcomes of the models.

4.2.3 Evolution of the fraction of informed agents

So far we have assumed that the fractions of informed and uninformed agents are fixed. In this subsection we model how these fractions change over time. The driving force behind these changes is the trade-off between the costs of information and the precision of the dividend yield estimator based on public information. Intuitively, given the costs of information, the more precise the estimate of the dividend yield, the bigger the fraction of uninformed agents. Or, given the precision of the dividend yield estimate, the higher the costs of being informed, the smaller the fraction of informed agents.

The evolution of the fraction of informed agents λ is modeled by the replicator dynamics mechanism, the same updating mechanim we have assumed in Chapter 3. We remind here that the replicator dynamics can be motivated in the context of boundedly rational agents who are learning and imitating which strategy to play in a strategic environment (see e.g. Weibull, 1995 and Binmore and Samuelson, 1997). Furthermore, in this case the replicator dynamics arises naturally in a framework where the equilibrium pricing equation (4.5) is derived from the maximization of a mean variance utility function. In fact, in this framework, as outlined in Appendix 4.B, λ is related to the fraction of the total wealth possessed by the informed agents. Although it is beyond the scope of this chapter to consider other specifications of the dynamics of λ, we believe the results to be valid more generally than just for the replicator dynamics discussed here. Since the objective of the agents is to gather information about future dividends, we assume that the success of a strategy is given by the squared forecast error of the dividend predictor. Informed agents have zero forecast error. Uninformed agents try to forecast future dividends by estimating the dividend yield, so their squared forecast error for the realization y_t is $(y^e_{t-1,t} - y_t)^2$. The costs of information are $c > 0$ per time step for the informed agents, and 0 for the uninformed. In addition, we assume that agents are granted a fixed payoff ρ. As a result, we can define the fitness π_t of the strategies at time t. The fitness of the strategy to buy information is:

$$\pi^I_t = \rho - c, \tag{4.20}$$

while the fitness of the strategy of remaining uninformed is

$$\pi_t^U = \rho - \left(\frac{y^e_{t-1,t} - y_t}{y_t^*}\right)^2. \tag{4.21}$$

Tecnically, the fixed payoff ρ assures that both (4.20) and (4.21) are always positive and can be taken as fitnesses in a replicator dynamics framework. In fact, only when fitnesses are positive, the replicator dynamics produces fractions in the interval $[0, 1]$. Notice that in the case of (4.20) a positive fitness requires $\rho > c$.

The time-varying fundamental dividend yield y_t^* defined in (4.16) is introduced in the denominator of π_t^U for normalization. Given that dividend yield y_t has y_t^* as reference value, this choice is convenient because it implies that the two fitness measures coincide when, given c, a forecasting error of $(100\sqrt{c})\%$ is made. For example, if $c = 0.1$, the two fitness measures are equal when errors of approx. 30% around y_t^* are made.

In the presence of a time varying fraction λ_t, (4.18) becomes:

$$y_{t+1} = f(y^e_{t,t+1}, g^e_{t,t+1}, \lambda_t) = \frac{r - g^e_{t,t+1}}{\lambda_t} - \frac{(1-\lambda_t)}{\lambda_t} y^e_{t,t+1}. \tag{4.22}$$

We can use this relation between y_{t+1} and $y^e_{t,t+1}$ and rewrite the fitness of the uninformed agents (4.21) as:

$$\pi_t^U = \rho - \left(\frac{y_t^* - y^e_{t-1,t}}{\lambda_{t-1} y_t^*}\right)^2. \tag{4.23}$$

Given the fitness measure of both strategies we can now specify the dynamics for λ. Under replicator dynamics the fraction λ_t of informed agents evolves according to

$$\lambda_t = (1-\delta)\frac{\lambda_{t-1}\pi_t^I}{\lambda_{t-1}\pi_t^I + (1-\lambda_{t-1})\pi_t^U} + \frac{\delta}{2}, \tag{4.24}$$

where the parameter δ is to be interpreted as a mutation or experimentation parameter (see e.g. Young and Foster, 1991; Droste, Hommes, and Tuinstra, 2002). The parameter δ is related to what we call evolutionary (or selection) pressure in the following way. When $\delta = 0$ the updating of the fractions is determined only by the selection mechanism. On the other hand when $\delta = 1$ the evolutionary pressure reaches its minimum and both fractions are $1/2$, independently on the fitness of the two strategies. Intermediate values of δ result in a convex combination of selection pressure and experimentation. Based on expressions (4.20) and (4.23), the replicator dynamics (4.24) gives,

$$\lambda_t = (1-\delta)\frac{\lambda_{t-1}(-c+\rho)}{\lambda_{t-1}(-c) - (1-\lambda_{t-1})\left(\frac{y_t^* - y^e_{t-1,t}}{\lambda_{t-1}y_t^*}\right)^2 + \rho} + \frac{\delta}{2}. \tag{4.25}$$

The parameter ρ regulates the speed of adjustment of the replicator dynamics. When ρ is large changes in λ are smaller, everything else being equal, than when ρ is small. Moreover, in the limit $\rho \to \infty$ the trajectories of 4.25 approach the trajectories of the corresponding continuos dynamical system (see e.g. Weibull, 1995, Chapter 4). For this reason we will refer to rho as to the inverse of the speed of adjustment. In Appendix 4.B we present an interpretation of this formula starting from wealth evolution of agents' maximizing a mean variance CRRA utility function.

4.2.4 Market returns

To summarize, for every shock η_t such that $d_t = d_{t-1}(1+g)\eta_t$, the full model developed so far is given by the following four equations

$$g^e_{t,t+1} = \gamma g^e_{t-1,t} + (1-\gamma)\left((1+g)\eta_t - 1\right), \tag{4.26}$$

$$y^e_{t,t+1} = \alpha y^e_{t-1,t} + (1-\alpha) y_t, \tag{4.27}$$

$$y_{t+1} = \frac{r - g^e_{t,t+1}}{\lambda_t} - \frac{(1-\lambda_t)}{\lambda_t} y^e_{t,t+1}, \tag{4.28}$$

$$\lambda_t = (1-\delta)\frac{\lambda_{t-1}(-c+\rho)}{\lambda_{t-1}(-c) - (1-\lambda_{t-1})\left(\frac{y^*_t - y^e_{t-1,t}}{\lambda_{t-1} y^*_t}\right)^2 + \rho} + \frac{\delta}{2}. \tag{4.29}$$

Equation (4.26) gives the common expectation formation regarding the growth rate of dividends g^e as defined in (4.14), taking into account of the dividend process defined in (4.1). Equation (4.27) gives the expectation formation of the dividend yield y^e by the uninformed traders as defined in (4.17). Equation (4.28) is the time varying expectational feedback map (4.22). Equation (4.29) expresses the dynamics of the fraction of informed agents λ as in (4.25). The shocks and parameters are:

- g, the (unknown) dividend growth rate;
- η_t, the shocks on the dividend growth rate;
- γ, the memory agents use to estimate the dividend growth;
- α, the memory agents use to estimate the future dividend yield;
- r, the required rate of return;

- δ, the experimentation or mutation rate;
- c, the costs of information per time step;
- ρ, the inverse speed of adjustment of the replicator dynamics.

Whenever y_t converges to its steady state value, that is, whenever $y^e_{t,t+1}$ converges to the rational expectations value $y^*_t = r - g^e_{t,t+1}$ as in (4.16), the price follows $p^*_t = d_{t+1}/r - g^e_{t,t+1}$ as in (4.15). Notice that the fundamental price depends on the changing estimates of the growth rate of dividends, and that it is the same price which has been derived by Barsky and De Long (1993). If, moreover, $g^e_{t,t+1} \to g$, i.e. beliefs about the dividend growth rate converge to its true value, the fundamental price converges to the "correct" present value price, that is $y^*_t \to r - g$ and $p^*_t \to d_{t+1}/(r - g)$. In that case the price follows a geometric random walk. If y_t fails to converge to $r - g$, deviations of the price from fundamental price can have two origins. The first is the failure of the deterministic skeleton of the system specified in (4.26–4.29) to converge to its fixed point, or equivalently adaptive expectations do not converge to rational expectations. This is related to the work of GS and to the fact that prices are not fully informative. The second possible reason is that, even if the system converges to the fixed point, it could approach an equilibrium where $g^e \neq g$. This is a situation where the fundamental price p^* is not equal to the "correct" present value price $d_{t+1}/(r - g)$. This is specifically relevant when the estimate g^e is time varying so that the system in (4.26-4.29) is stochastic. In what follows, we analyze these effects separately as well as their interplay. First, in Section 4.3, we analyze the conditions of convergence of the deterministic system dynamics of y and λ to their equilibrium values. Thereafter, in Section 4.4, we complement this analysis by investigating the effect of a time varying stochastic g^e and how the two sources interact.

4.3 Informational differences

In this section we analyze the impact of informational differences alone on the dynamics of asset prices assuming that $g^e_{t,t+1} \equiv g^e$, without learning of the divided growth rate. Technically, we analyze the system of equations (4.26–4.29) when the memory parameter $\gamma = 1$. Without loss of generality we consider only the case $g^e = g$. The generalization to $g^e \neq g$ is straightforward and implies only a shift of the level of the steady state dividend yield from $r - g$ to $r - g^e$. To simplify the notation, throughout the rest of the chapter we write $y^e_{t,t+1} \equiv y^e_{t+1}$ for the forecast of y_{t+1} made at time t.

When $g^e_{t,t+1} = g$, the system (4.26–4.29) reduces to a two dimensional (2-D) system in the variables y^e_{t+1} and λ_t:

$$y^e_{t+1} = \frac{(1-\alpha)(r-g)}{\lambda_{t-1}} + \frac{\alpha+\lambda_{t-1}-1}{\lambda_{t-1}} y^e_t, \tag{4.30}$$

$$\lambda_t = (1-\delta)\frac{\lambda_{t-1}(-c+\rho)}{\lambda_{t-1}(-c)-(1-\lambda_{t-1})\left(\frac{(r-g)-y^e_t}{\lambda_{t-1}(r-g)}\right)^2+\rho} + \frac{\delta}{2}. \tag{4.31}$$

Given the dynamics of (y^e_{t+1}, λ_t) specified by (4.30–4.31), the dynamics of the dividend yield y_{t+1} can be derived by using the expectational feedback map f defined in (4.12). Before investigating the full dynamics of (4.30–4.31) it is instructive to consider the 1-D system obtained when the fraction λ_t of informed agents is fixed to a constant value λ. Proofs of all the propositions can be found in Appendix 4.C.

4.3.1 Dividend yield dynamics

Taking $\lambda_t \equiv \lambda$, (4.30) becomes:

$$y^e_{t+1} = \frac{(1-\alpha)(r-g)}{\lambda} + \frac{\alpha+\lambda-1}{\lambda} y^e_t. \tag{4.32}$$

Given the linearity of (4.32), the analysis of the dynamics is straightforward and it is possible to compute the general solution of the difference equation. That is, given y^e_0 one can compute the value of y^e_t, for all t. The following proposition summarizes the results.

Proposition 4.8 *Given the memory parameter $\alpha \in (0,1)$, the fraction of informed agents $\lambda \in (0,1]$, and the required rate of return $r > g$, we have:*
(i) The solution of the difference equation (4.32) with initial condition y_0 is given by:

$$y^e_t = (y^e_0 - y^*)\left(\frac{\alpha+\lambda-1}{\lambda}\right)^t + y^*,$$

where

$$y^* = r - g.$$

(ii) If

$$\lambda > \bar{\lambda} \equiv \frac{1-\alpha}{2}, \tag{4.33}$$

y^e_t converges to the steady state y^*, otherwise y^e_t diverges to $\pm\infty$.

Notice that whenever the expected dividend yield y^e_{t+1} converges to its steady state y^*, also the realized dividend yield y_t converges to y^* through the feedback map (4.12). At the steady state y^* the price equals the fundamental price p^*_t in (4.8) and thus fully reveals the information concerning the future dividend. The shaded area in Figure 4.1 shows the stability region of (4.32) in the parameter space, (α, λ), whereas the white area shows the unstable region. The shaded area is divided in two gray scales. In the lighter region, the convergence of the expected dividend yield to its the steady state y^* is oscillatory, whereas in the darker, the convergence is monotone. Notice that the border between the stability and the instability region is characterized by oscillatory behavior of the expected dividend yield y_t^e, and thus of the realized dividend yield y_t. This implies that failure of the price to fully reveal the fundamental information should be characterized by price fluctuations with negative autocorrelation.

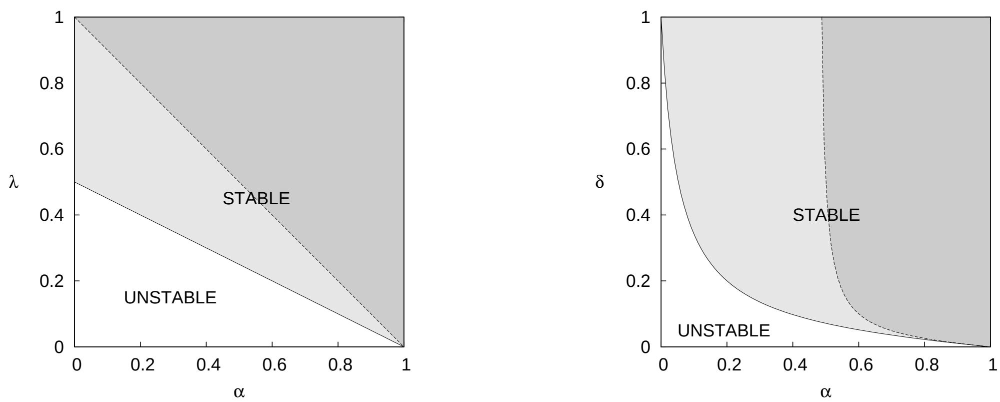

Figure 4.1: *Left panel: Stability region for the* 1*-D dynamical system in (4.32). The expected dividend yield converges to the steady state* $y^* = r - g$ *only for values of* (α, λ) *in the shaded area. In the darker region the convergence to* y^* *is monotone, whereas in the lighter region the convergence is oscillatory. Right panel: Stability region in the 2-D dynamical system in (4.30-4.31) as a function of the mutation rate* δ *and memory* α*, when* $k = \rho/c = 10$*. As in the left panel, the solid line marks the border of the stability region, while the dotted line marks the border of the region where the convergence of the expected dividend yield to* y^* *is monotone (darker region) or oscillatory (lighter region).*

4.3.2 Dividend yield and fractions dynamics

In general, the fraction of informed agents λ_t is time dependent and the dynamics of the dividend yield and of the fraction of informed agents is given by the nonlinear system (4.30–4.31). We use local stability analysis to characterize the behavior of the state variables (y_t^e, λ_t) near the steady state of (4.30–4.31).

First consider the case of zero information costs. The following proposition characterizes the steady states of the system and their local stability.

Proposition 4.9 *Consider zero information costs,* $c = 0$*. Given the memory parameter* $\alpha \in (0, 1)$*, the inverse speed of adjustment* $\rho > 0$ *and the required rate of return* $r > g$*, we*

have:
(i) When $\delta \in (0,1)$*, the unique steady state of the system (4.30–4.31) is the point* (y^*, λ^*)*, where*

$$y^* = r - g \quad \text{and} \quad \lambda^* = \frac{1}{2}.$$

This steady state is globally stable for every parameter combination.
(ii) When $\delta = 0$*, every point* (y^*, λ^*) *with* $y^* = r - g$ *and* $\lambda^* \in (0, 1]$ *is a steady state. For a given* $\alpha \in (0, 1)$ *the set of steady states such that* $\lambda^* > \bar{\lambda} = (1 - \alpha)/2$*, as in (4.33), is globally stable.*

Thus, with zero information cost, the expected dividend yield y_t^e converges to the fundamental value y^* defined in (4.13) and the fraction of informed agents is either $1/2$, when $\delta > 0$, or any value that satisfies (4.33) when $\delta = 0$. As a consequence the expectational feedback map (4.12) converges to the rational expectations equilibrium and prices are fully informative in the long run. This is not the case when one considers positive information costs:

Proposition 4.10 *Consider positive information costs,* $c > 0$*. Given the memory parameter* $\alpha \in (0, 1)$*, the experimentation level* $\delta \in (0, 1)$*, the ratio* k *between* ρ*, the inverse speed of adjustment and* c*, the cost of being informed, such that* $k = \rho/c > 1$ *and the required rate of return* $r > g$*, we have:*
(i) The unique steady state of the system (4.30–4.31) is the point (y^*, λ^*) *where*

$$y^* = r - g$$

and

$$\lambda^* = \frac{2 - \delta + 2k\delta - \sqrt{-8k\delta + (2 - \delta + 2k\delta)^2}}{4}. \tag{4.34}$$

Moreover, $\lambda^* \in (0, 1/2)$*.*
(ii) The Jacobian of (4.30–4.31) at the steady state is diagonal and given by

$$\mathbf{J}|_{(y^*,\lambda^*)} = \begin{pmatrix} \dfrac{\alpha + \lambda^* - 1}{\lambda^*} & 0 \\ 0 & (1-\delta)\dfrac{k(k-1)}{(k-\lambda^*)^2} \end{pmatrix}. \tag{4.35}$$

If

$$\delta > \bar{\delta} \equiv \frac{(1+\alpha)}{1 + \frac{2\alpha k}{(1-\alpha)}} = \frac{(1+\alpha)}{1 + \frac{2\alpha\rho}{(1-\alpha)c}}, \tag{4.36}$$

the steady state (y^*, λ^*) *is locally stable. This condition corresponds to the stability condition (4.33) of the 1-D dynamical system (4.32). That is the condition* $\delta > \bar{\delta}$*, and* $\lambda^* > \bar{\lambda}$ *are equivalent.*

The local stability condition (4.36) is represented in terms of the parameters (α, δ) in the right panel of Figure 4.1 for $k = \rho/c = 10$. In the white area the steady state (y^*, λ^*) is unstable, while in the shaded area the steady state is stable. Recall that when the expected dividend yield y_t^e converges to y^*, the realized dividend yield y_t also converges to y^*. Therefore, stability of the steady state (y^*, λ^*) implies convergence of the price to the fundamental price p_t^* and thus to a fully informative price.

Notice the steady state is well defined and locally stable, i.e. prices are fully informative, only when the experimentation rate δ is not too small, $\delta > \bar{\delta}$. Under this condition there always exists a fraction of agents that are prepared to buy fundamental information. It can also happen that the steady state is well defined but not locally stable, that is prices are not fully revealing. In fact, for any $\alpha \in (0, 1)$ there exists a sufficiently small mutation rate $0 < \delta < \bar{\delta}$ such that the prices are not fully revealing and the system is unstable. In particular, the definition of $\bar{\delta}$ in eq. (4.36) of Proposition 4.10 shows that, for a given α, the stability region of the system (4.30–4.31) shrinks, when the cost of information c increases, or when the speed of adjustment $1/\rho$ increases. In the limit case of no experimentation rate, $\delta = 0$, from (4.34) one would get $\lambda^* = 0$ so that the steady state y^* is not even defined. Summarizing, for a positive information cost, the price are not always fully informative in the long run, dependently on the size of the experimentation rate as compared to $\bar{\delta}$ in (4.36). Moreover, prices are never fully informative when the experimentation rate $\delta = 0$.

What happens to the dynamics of the expected dividend yield and of the fraction of informed agents when the steady state is unstable and prices do not converge to be fully informative? In order to answer this question we analyze the global dynamics of the system (4.30–4.31) for small experimentation levels, $\delta < \bar{\delta}$. Whereas in the 1-D system the expected and realized dividend yield diverge unboundedly and unrealistically, when the stability conditions (4.33) and (4.36) do not hold, in the 2-D system our simulations show the emergence of bounded aperiodic cycles. The top left and top right panels of Figure 4.2 show a typical evolution of the uninformed agents' expected dividend yield, y^e, and of the fraction of informed agents, λ, respectively. At time $t = 0$, the fraction of informed agents is above the dotted line, which marks the value $\bar{\lambda}$ in (4.33) and gives the stability condition for the steady state of the y_t^e dynamics. As a result, at $t = 0$ both the value of y_t^e and, through (4.12), the value of y_t, are close to their steady state value, y^*. This implies that the price is close to being fully informative, there is no advantage in buying information so that the fraction of informed agents decreases. This process continues until the fraction of informed agents is smaller than the value $\bar{\lambda}$. At this moment there are so few informed agents that the asset price starts to diverge from the fundamental. The dynamics of the expected dividend yield y_t^e is unstable and y_t^e starts to

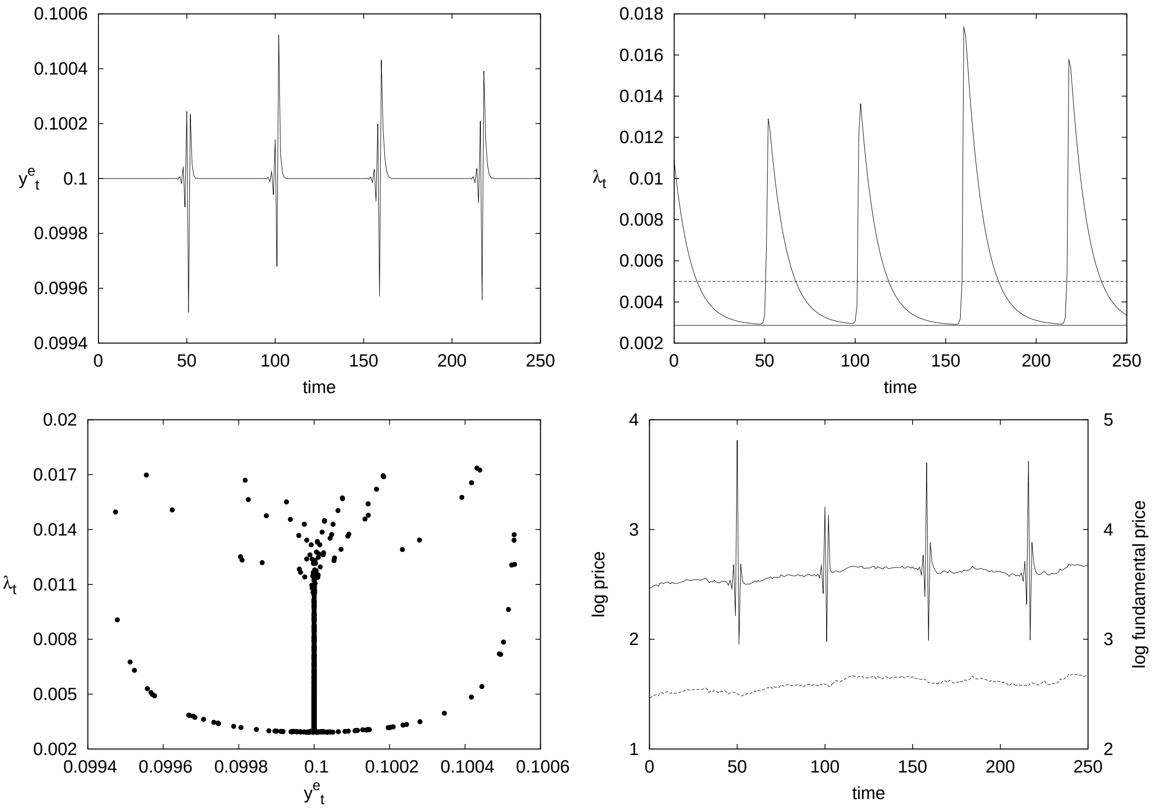

Figure 4.2: *Top panels: time series for the expected dividend yield y_t^e (left panel) and the fraction of informed agents λ_t (right panel) from (4.30-4.31). In the top right panel the dotted line corresponds to $\bar{\lambda}$ in (4.33), that is to the critical stability border, whereas the solid line corresponds to the steady state λ^* in (4.34). The bottom left panel shows the state space representation (y_t^e, λ_t). In the bottom right panel, the log price dynamics derived form the dynamics of y_t^e and λ_t . Fundamental prices are given by the dotted line and realized prices by the continuous line. Notice that the two price series levels should be read using two different scales. The left scale gives the value of log prices whereas the right scale gives the value of fundamental log prices. The dividend process is characterized by $\sigma_\eta = 0.02$. Parameter values are $\alpha = 0.99$, $c = 0.1$, $\rho = 1$ (so that $k = \rho/c = 10$), $r = 0.1$, $g = 0$ and $\delta = 0.000575$.*

diverge from the steady state y^*. The fraction of informed agents continues to decrease until the price carries so little information about p^* that informed agents are better off. Eventually, paying the cost of being informed leads to a higher fitness than using a freely available estimate with a large error. As a result, the fraction of informed agents grows sharply, see e.g. the top right plot around period $t = 50$. The fraction of informed agents reverts to a region where the price is sufficiently informative so that y_t^e returns to values close to y^*. As time continues the process repeats, with λ decreasing again, and so on and so forth. The left and right bottom panels of Figure 4.2 show, respectively, the dynamics we have just illustrated in the (y^e, λ) space and the corresponding dynamics of the log price compared to the log fundamental price.

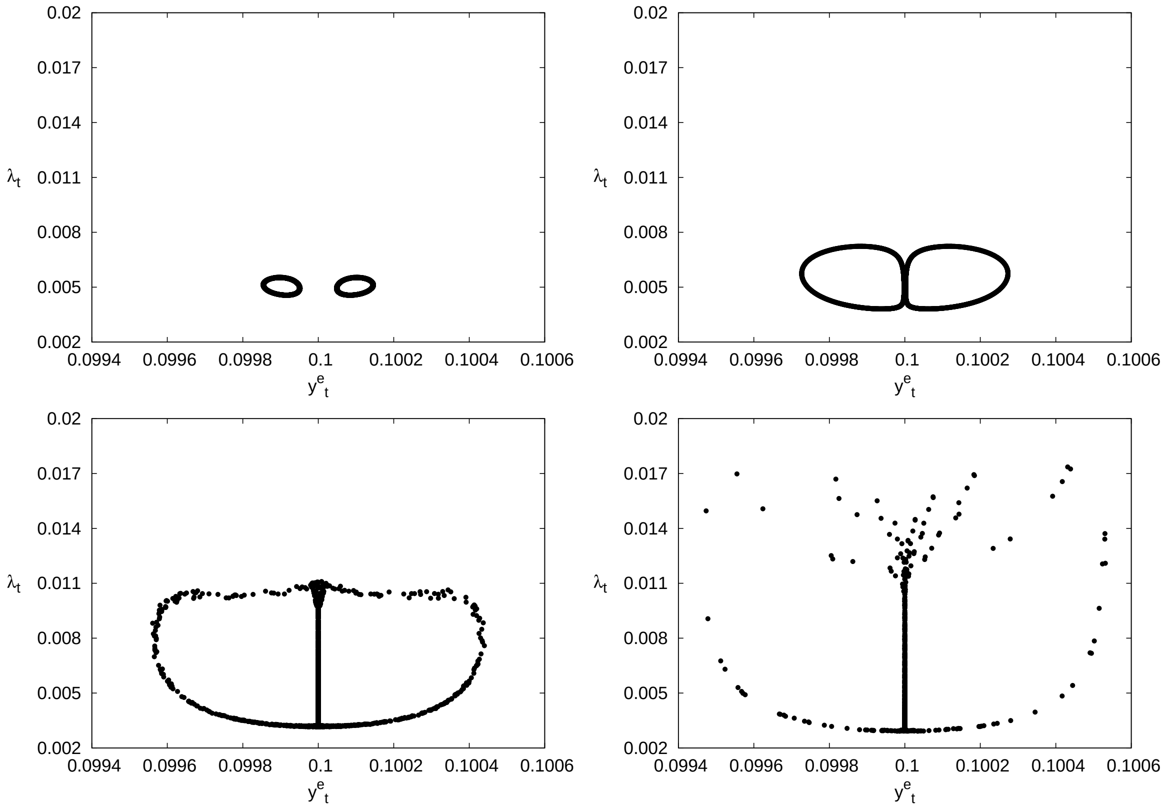

Figure 4.3: *Phase plots of the expected dividend yield* y_t^e *and the fraction of informed agents* λ_t *produced by the system of difference equations (4.30-4.31). Parameter values are* $\alpha = 0.99$, $c = 0.1$, $\rho = 1$ *(so that* $k = \rho/c = 10$*),* $r = 0.1$ *and* $g = 0$*. Top left panel:* $\delta = 0.000675$*, top right panel:* $\delta = 0.00066$*. Bottom left panel:* $\delta = 0.000625$*, bottom right panel:* $\delta = 0.000575$*.*

The phase plots shown in Figure 4.3 suggest that the fluctuations of y_t^e and λ_t, are associated with a so-called homoclinic bifurcation. Similar phenomena are encountered in other multidimensional nonlinear systems and emerge from the interplay between local instability and global stability of the dynamics. Brock and Hommes (1997) and Droste, Hommes, and Tuinstra (2002) present other economic frameworks where homoclinic bifurcation arise. They also offer detailed discussions of the mathematical aspects behind these interesting phenomena.

Before turning to an economic interpretation of the results of this section and to comparison with the results of GS, it is instructive to characterize the convergence of the fraction of informed agents in the stability region. Close to the equilibrium λ turns out to change very slowly.

Proposition 4.11 *Given a memory parameter* $\alpha \in (0, 1)$*, an experimentation rate* $\delta \in (0, 1)$ *and a inverse speed of adjustment larger than the cost of information,* $\rho > c$*, if we* *call ν_2 the eigenvalue characterizing the dynamics of λ in a neighborhood of (y^*, λ^*), we have*

$$1 > \nu_2 > (1-\delta)\left(1 - \frac{c}{\rho}\right).$$

This proposition shows that when the experimentation rate δ is small and the ratio between the inverse of the speed of adjustment and the cost of information, k, is large, the value of ν_2 is very close to one. As a result, when the system is stable, changes in the value of the fraction of informed agents λ are very slow, and hence λ is very persistent. We will find confirmation of this statement in the next section, when we appraise the effect of agent learning the growth rate of dividends on the dynamics.

4.3.3 Economic interpretation

In a financial market, when agents face informational differences, the use of homogeneous rational expectations poses puzzling consequences. This is the paradox as found by GS, they consider an asset pricing model where agents can either buy information on fundamentals or try to extract such information from the asset price. In a repeated single period model with fully rational agents, they show that an equilibrium value of the fraction of informed agents for which the price fully reveals the information about future dividends can not exist. If an equilibrium existed then nobody would pay for the information and prices could not possibly reveal any information. In order to resolve this paradox, GS need two key assumptions: (1) the supply of shocks is noisy (this is equivalent to saying that there are noise traders in the market) and (2) the informational content of the dividend signal is not perfect. Under these assumptions a rational expectations "equilibrium degree of disequilibrium", where informed and uninformed agents coexist and prices fail to be fully informative, exists.

The model analyzed in this section is inspired by that of GS, but there are three important differences. First, our agents are not rational but boundedly rational, that is, they do not use rational but adaptive expectations. Second, we consider a multi period model where future returns are determined by capital gains in addition to dividends, and agents form expectations about both future prices and dividends. Third, the fraction of informed and uninformed agents are dynamic variables in our model. The results from our model are as follow. When the experimentation rate is positive, $\delta > 0$, there are enough agents willing to buy information, even if they could extract it from prices, so that an equilibrium where prices are fully informative is well defined. This equilibrium is (locally) stable only when the experimentation rate is big enough, $\delta > \bar{\delta}$. In this case, prices converge to fully informative prices and the GS paradox is resolved by experimentation (or noise), as in GS. On the other hand, when $\delta < \bar{\delta}$ fluctuations around the fully informative equilibrium price arise. This is also the case for a zero experimentation rate ($\delta = 0$). In this case, an equilibrium is not even defined, and fluctuations of prices and agents' fractions are observed. How do these two cases, $\delta = 0$ and $\delta < \bar{\delta}$, relate to the GS paradox? The dynamics is now responsible for resolving the GS paradox, as, even if no equilibrium emerges in the long run, orbits of the system are still well defined. Along these orbits prices fluctuate between being close to fully informative and being non informative. By assuming that agents are boundedly rational and that fractions are endogenously determined, the learning process of the uninformed boundedly rational agents balances the cost of information and the informational content of the price, resolving the GS paradox. In our framework we thus obtain a "dynamic equilibrium degree of disequilibrium" due to endogenous price fluctuations produced by the interaction of boundedly rational agents.

De Fontnouvelle (2000) and Goldbaum (2005) are earlier contributions where bounded rationality and learning offer an explanation for the existence of an equilibrium degree of disequilibrium. Their framework differs from ours in many ways, most importantly in that they consider a dividend process which follows a random walk rather than a geometric random walk as we do here. Furthermore, their resulting systems of the joint evolution of the asset price and the fractions of agents is fairly complicated so that their analysis is performed only via simulations.

4.4 Informational differences and parameter learning

In the previous section we have assumed that the agents' estimate of the growth rate of dividends is constant. As a result, the equation for the dividend yield is fully deterministic. In this section we analyze the simultaneous impact of informational differences and of agents' learning of the growth rate of dividends as new information about the fundamentals becomes available. As a result we have to deal with a stochastic system. A similar analysis has been performed by Barsky and De Long (1993) and Timmermann (1993), among others, in a context where there are no informational differences among agents. In particular Barsky and De Long (1993) also assume that agents use adaptive expectations to estimate g. Recalling the results from Section 4.2, adaptive expectations are specified by (4.14) which, when the dividend follows a geometric random walk with innovations $(1 + g)\eta_t$ gives:

$$g^e_{t+1} = \gamma g^e_t + (1 - \gamma) \left((1 + g)\eta_t - 1\right).$$

This stochastic equation, together with the evolution of the dividend yield, its expectations, and the fraction of informed agents as specified in (4.26-4.29), lead to a stochastic version of the deterministic skeleton (4.30-4.31) namely

$$g^e_{t+1} = \gamma g^e_t + (1-\gamma)\left((1+g)\eta_t - 1\right), \tag{4.37}$$

$$y^e_{t+1} = \frac{(1-\alpha)(r-g^e_t)}{\lambda_{t-1}} + \frac{\alpha+\lambda_{t-1}-1}{\lambda_{t-1}} y^e_t, \tag{4.38}$$

$$\lambda_t = (1-\delta)\frac{\lambda_{t-1}(-c+\rho)}{\lambda_{t-1}(-c) - (1-\lambda_{t-1})\left(\frac{(r-g^e_t)-y^e_t}{\lambda_{t-1}(r-g)}\right)^2 + \rho} + \frac{\delta}{2}. \tag{4.39}$$

Shocks $\{\eta_t\}$ on the growth rate of dividends are the stochastic component that drives the co-evolution of agents expectations of the growth rate of dividend, of the dividend yield, and of the fraction of informed agents. Given the evolution of the expected growth rate of dividend, g^e_{t+1}, of the expected dividend yield, y^e_{t+1}, and of the fraction of informed agents, λ_t, the dividend yield itself, y_{t+1}, is set by the feedback map (4.22). Before we start with the analysis of the impact of shocks on the dynamics of (4.37-4.39), we show that our model contains two important benchmarks as special cases.

Classical Asset Pricing model The first benchmark is the classical asset pricing model, which assumes that all agents know the dividend growth rate g, and that agents use rational expectations. In this case, if only few agents are informed about d_{t+1}, the other can correctly extract this information from prices, and the market price and the market dividend yield are given by:

$$p^{**}_t = \frac{d_{t+1}}{r-g}, \quad y^{**}_{t+1} = r-g. \tag{4.40}$$

As before, the price p^{**}_t denotes the "correct" present value price as given by the Gordon model (see e.g. Gordon, 1962). Our model boils down to the classical asset pricing model when at every time t all agents are informed, $\lambda_t = 1$, and use the correct dividend growth rate, $g^e_t = g$. In fact taking $\lambda_t = 1$ and $g^e_t = g$ in equation (4.38) and using the feedback map (4.22) to transform the expected dividend yield y^e_{t+1} into realized dividend yield y_{t+1}, gives $y_{t+1} = r-g$ which corresponds to y^{**}_{t+1} defined in (4.40).

Barsky and De Long model The second benchmark is the model proposed by Barsky and De Long (1993). They consider agents without informational differences who have to form expectations about the growth rate g. In their case the price and the dividend yield are given by:

$$p^*_t = \frac{d_{t+1}}{r-g^e_{t+1}}, \quad y^*_{t+1} = r - g^e_{t+1}, \tag{4.41}$$

where g^e_{t+1} is given by adaptive expectations as in (4.37). These are respectively the fundamental price and the fundamental dividend yield defined in (4.15) and (4.16). Our model corresponds to the model of Barsky and De Long (1993) under the assumptions that all the agents that at every time t all agents are informed, $\lambda_t = 1$. In fact by substituting $\lambda_t = 1$ in (4.38) and using the feedback map (4.22) one retrieves $y_{t+1} = r - g^e_{t+1}$, with g^e_{t+1} defined in (4.37), which corresponds to y^*_{t+1} defined in (4.41).

A way of evaluating the differences between our model and the two benchmarks is to write an evolution equation for the dividend yield as a function of lagged dividend yields and shocks on the growth rate of dividends for each model. In the classical asset pricing model the dividend yield is constant and given by $y^{**} = r - g$. In the Barsky and De Long model, one can use $y^*_{t+1} = r - g^e_{t+1}$ to rewrite (4.37) as:

$$y^*_{t+1} = (1-\gamma)(1+r) + \gamma y^*_t - (1-\gamma)(1+g)\eta_t. \tag{4.42}$$

That is, the dividend yield follows an AR(1) process with shocks given by the shocks $\{\eta_t\}$ on the growth rate of dividends. The memory parameter γ is related to both the AR(1) coefficient and the variance of the innovations, $(1-\gamma)^2(1+g)^2\sigma^2_\eta$. The mean of the process is independent of the memory parameter and equal to the constant classical asset pricing dividend yield $y^{**} = r - g$.

In our model, fixing for the moment the value of λ, the map h defined in (4.19), i.e. the dynamic equation (4.38) for the expected dividend yield y^e_{t+1} rewritten in terms of the realized dividend yield y_{t+1} using the feedback map (4.22), gives:

$$y_{t+1} = \frac{(r_t - g^e_{t+1}) - \alpha(r_{t-1} - g^e_t)}{\lambda} + \frac{\alpha + \lambda - 1}{\lambda} y_t. \tag{4.43}$$

This equation expresses that, when the fraction of informed agents is fixed at λ, the dividend yield follows an AR(1) process with shocks that are correlated with the shocks of the growth rate of dividends. As for the dividend process implied by the Barsky and De Long model in (4.43), the long run mean of the dividend process implied by our model is given by the classical asset pricing dividend yield y^{**}. When $\gamma = \alpha$, that is, when agents use the same memory parameter to estimate the growth rate of dividend and the dividend yield, (4.43) has a simple appealing formulation:

$$y_{t+1} = \frac{(1-\gamma)((1+r) - (1+g)\eta_t)}{\lambda} + \frac{\gamma + \lambda - 1}{\lambda} y_t.$$

If we define

$$\tilde{\gamma} = \tilde{\gamma}(\lambda) \equiv 1 - (1-\gamma)/\lambda = \gamma - (1-\gamma)\frac{1-\lambda}{\lambda}, \tag{4.44}$$

we can rewrite (4.43) as

$$y_{t+1} = (1 - \tilde{\gamma}(\lambda))(1 + r) + \tilde{\gamma}(\lambda) y_t - (1 - \tilde{\gamma}(\lambda))(1 + g)\eta_t. \qquad (4.45)$$

The result is that, when $\gamma = \alpha$, our model specified by (4.45) and the model of Barsky and De Long (1993) specified by (4.42) differ only in the value of the memory parameter γ. Since γ is the real memory agents use to discount new information, we can refer to $\tilde{\gamma}(\lambda)$ as the effective memory. The definition (4.44) shows that the effective memory has two components, one given by the real memory and the other related to the presence of uninformed agents. This second effect becomes less important as more informed agents are present in the market. The general result is that $\tilde{\gamma}(\lambda)$ in (4.44) is an increasing function of λ with $\tilde{\gamma}(\lambda) \leq \gamma$ for all λ, and $\tilde{\gamma}(1) = \gamma$. That is, the presence of uninformed agents is equivalent to all agents being informed and using an effective memory which is lower then the real memory. The value of $\tilde{\gamma}$ determines both the AR(1) coefficients and the variance of the shocks but not the long run mean which is always $y^{**} = r - g$. In particular the lower is the effective memory, the higher is the impact of the shocks on the dynamics of the dividend yield and the faster is the reversion of the process to its mean. That is, a lower effective memory creates a bigger short run effect and a smaller long run effect. Also, since the effective memory $\tilde{\gamma}$ is a function of λ, our model allows for variation of the memory parameter as the fraction of informed agents λ varies. Changes in λ have an impact on $\tilde{\gamma}$ and thus on the variance of shocks and on the speed of convergence. In what follows we explore the importance of both the effective memory being lower than the real memory and the effective memory being time varying on the dynamics of the dividend yield implied by our model (4.45) compared to the two benchmarks in (4.40) and (4.43).

4.4.1 Nonlinear mean reversion

Our model (4.45) clearly differs, both structurally and regarding parameter values, from that of Barsky and De Long in (4.42) when the fraction of informed agents is time varying. If this is the case, our model implies an AR(1) process for the dividend yield where both the rate of convergence of the dividend yield to its mean and the variance of shocks are time varying. This consideration links our model to the econometric analysis of nonlinear mean reversion that has recently been proposed to characterize fluctuations of stock indices. By using the fact that $p_t^{**} \equiv d_{t+1}/(r - g)$, i.e. the price implied by model (4.40), and defining $x_t = y_{t+1}/(r - g)$, given the definition of the dividend yield one can write:

$$\log(p_t) = \log(d_{t+1}) - \log(y_{t+1}) = \log(p_t^{**}) - \log(x_t).$$

If x_t is close to its long run average of 1 one can rewrite the previous expression and expand the logarithm around one. Using the variable $z_t = 1 - x_t$ one gets

$$\log(p_t) \approx \log(p_t^{**}) + z_t, \tag{4.46}$$

where the dynamics of z_t can be easily derived using its definition in terms of x_t, the definition of x_t and (4.45). The resulting dynamics of the component z_t of the log price is given by:

$$z_t = \tilde{\gamma}(\lambda_t) z_{t-1} + \frac{(1 - \tilde{\gamma}(\lambda_t))(1+g)}{(r-g)}(\eta_{t-1} - 1). \tag{4.47}$$

This equation shows that we have a model whose realized log price in (4.46) is the sum of a persistent component $\log(p_t^{**})$, which follows a random walk with drift, and of a temporary component, z_t, which follows a stationary autoregressive process (4.47) with a time-varying AR(1) coefficient $\tilde{\gamma}(\lambda_t)$. Empirical investigation of the properties of stock prices are in accordance with this statement. Both Gallagher and Taylor (2001) and Manzan (2003) reject the null hypothesis that the temporary component in a mean reversion model follows a stationary process with fixed parameters. In particular Gallagher and Taylor (2001) show that quarterly data of the logarithm of the dividend yield of the index SP500 are well fitted by an ESTAR(4) (Exponentially Smooth Transition AR) ARCH(1) model whose two regimes have AR(1) coefficients equal to 0.72 and 0.20 respectively. As the model of Barsky and De Long (1993) suggests, the fact that the dividend yield follows an autoregressive process might be related to the agents' learning of the model parameters. In addition to this effect, our model suggests that changing "learning" coefficients and heteroskedasticity can be related to agent interaction. In fact, both the AR(1) coefficient and the shocks variance in (4.47) are a function of $\tilde{\gamma}(\lambda_t)$ which is a nonlinear function of the time varying fraction of informed agents λ_t.

4.4.2 Simulation study

In presenting the qualitative effect of the shocks on the growth rate of dividends on the dividend yield and fraction dynamics, we proceed by analyzing the impact of a single shock η_t, and then by analyzing the cumulative impact of a sequence of shocks $\{\eta_t\}$. We present results not only for dividend yields and prices generated by our model, but also for dividend yields and prices generated by the classical asset pricing model (4.40) and by the model of Barsky and De Long (4.41). In addition we also present results for a model similar to that of Barsky and De Long with the difference that real memory is taken as the average memory of the time varying effective memory generated by our model. We refer to this model as the "modified" Barsky and De Long model and its series of prices and dividend yields by $p^*_{\tilde{\gamma}}$ and $y^*_{\tilde{\gamma}}$. We use the "modified" Barsky and De Long model to appraise the role of time variability of effective memory. We simulate all models with dividends generated according to (4.1), that is, $d_{t+1} = d_t(1 + g)\eta_{t+1}$, where $\{\eta_t\}$ is a sequence of i.i.d. log normal shocks with mean zero and variance σ_η.

It is instructive to start the analysis by comparing the effect of a single shock on the dividend growth rate g. We perform this analysis for the four different models, namely for our model, the model of Barsky and De Long, the "modified" Barsky and De Long model, and the classical asset pricing model. Figure 4.4 shows the resulting price dynamics. In all the cases there is an initial overreaction followed by convergence to the equilibrium value, which is given by p^{**}, the price implied by the classical asset pricing model.

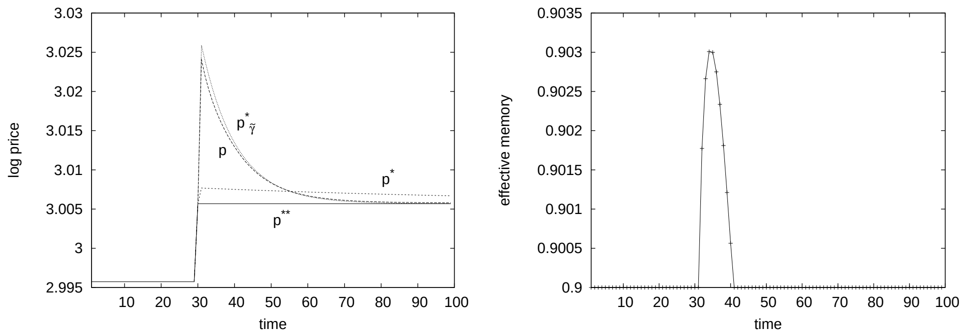

Figure 4.4: *Effect of a single shock on the growth rate g. Before and after the shock, $g = 0$. The shock is $\eta_{31} = 0.01$. Left panel: log prices as a function of time. The time series p^{**} gives the value of the log price as implied by the classical asset pricing model (4.40). The time series p^* and p give respectively the value of the log price as implied by Barsky and De Long model (4.42) and by our model (4.37-4.39) respectively. The time series $p^*_{\tilde{\gamma}}$ gives the log price implied by the " modified" Barsky and De Long model (4.42) when $\gamma = \tilde{\gamma}(\lambda_0)$, where λ_0 is the fraction of informed agents before the shock. Right panel: evolution of the effective memory $\tilde{\gamma}$. The parameter values are $\gamma = \alpha = 0.99$, $\delta = 0.02$, $c = 0.1$ and $\rho = 1.0$.*

Since, for $\lambda < 1$, the effective memory $\tilde{\gamma}$ is lower than the real memory γ, the variance of the shocks is larger in our model than in the model of Barsky and De Long, so that the overreaction is more pronounced. At the same time, when $\tilde{\gamma} < \gamma$, the value of the autoregressive coefficient is closer to zero so that convergence is faster. The overall effect is that the shock has a higher short run impact but a shorter half life for p than for p^*. The right panel of Figure 4.4 shows the response of the effective memory $\tilde{\gamma}$ to changes in λ. From the Jacobian of the 2-D system (see Proposition 2) we know that changes in y only have second order effects on λ, and as a result changes in λ are negligible in the short run. But, from Proposition 3, we also know that the eigenvalue ν_2 is close to one so that changes in λ are very persistent. Both results are confirmed by the changes in $\tilde{\gamma}$ shown in the right panel. A confirmation of the fact that one shock has no considerable consequence on changes of $\tilde{\gamma}$ comes from the time series for $p^*_{\tilde{\gamma}}$ shown in the left panel. The price $p^*_{\tilde{\gamma}}$ is the price obtained using the "modified" Barsky and De Long model, that is equation (4.42) with $\gamma = \tilde{\gamma}(\lambda_0)$, where λ_0 is the value of the fraction of informed agents before the shock. The overall comparison of the dynamics of p, p^* and $p^*_{\tilde{\gamma}}$ shows that in the single shock case the fact that the effective memory is lower than the real memory plays an important role whereas the fact that the effective memory is time varying is negligible, i.e. p is close to $p^*_{\tilde{\gamma}}$. Notice also that with informed agents in the market, the price anticipates the shock on the dividend, i.e. the price takes into account the change in the dividend before such a change is realized and much before such change has an effect on the value of the effective memory.

We now turn to investigating the effect of a sequence of shocks. Figure 4.5 shows the impact of a sequence of 500 i.i.d. shocks $\{\eta\}$. The growth rate of dividends, g, and the variance of the growth rate shocks, σ^2_η, are taken in accordance with historical quarterly data of the S&P500 index for the period 1880-2005 (source: Shiller database available from R. J. Shiller's homepage). The discount rate r is taken such that $y^{**} = 0.05$, that is the price implied by the present value model is 20 times the value of the dividend. If we think of quarters, 500 dividends correspond to 125 years. The top left panel shows the time series of the dividend yield y generated by our model whereas the right panel shows the time series of y^* generated by the model of Barsky and De Long. In both cases the horizontal line represents the long run mean $y^{**} = r - g$. The same results as for a single shock emerge: the dynamics of the dividend yield is less persistent in our model where the fraction of informed agents is time varying and smaller than in the model of Barsky and De Long (1993). Also, deviations from y^{**} are larger. The central and bottom rows offer a comparison of the systems in terms of log prices.

How important is the fact that the effective memory is time varying? The left panel of Figure 4.6 shows the changes in the effective memory for the same simulation run. These changes are due to changes in the fraction of informed agents λ via the transformation $\tilde{\gamma}(\lambda)$ defined in (4.44). As a confirmation of our previous results and of our theoretical analysis, changes in $\tilde{\gamma}$ (that is, changes in λ) are rather persistent. The right panel of Figure 4.6 shows deviations of log prices generated by our model and log prices generated by the "modified" Barsky and De Long model. We call this last series $y^*_{\tilde{\gamma}}$. Notice that deviations of up to more than ten percent arise. Our conclusion is that when subsequent shocks are present, both the fact that the effective memory is lower than the real memory and that the effective memory is time varying play an important role. Naturally, these properties are dependent on the choice of updating mechanism for λ_t and hence of the fitness measures as presented in Subsection 4.2.3. We do not claim that the mechanism we propose here to characterize the changes in the fraction of informed agents is more

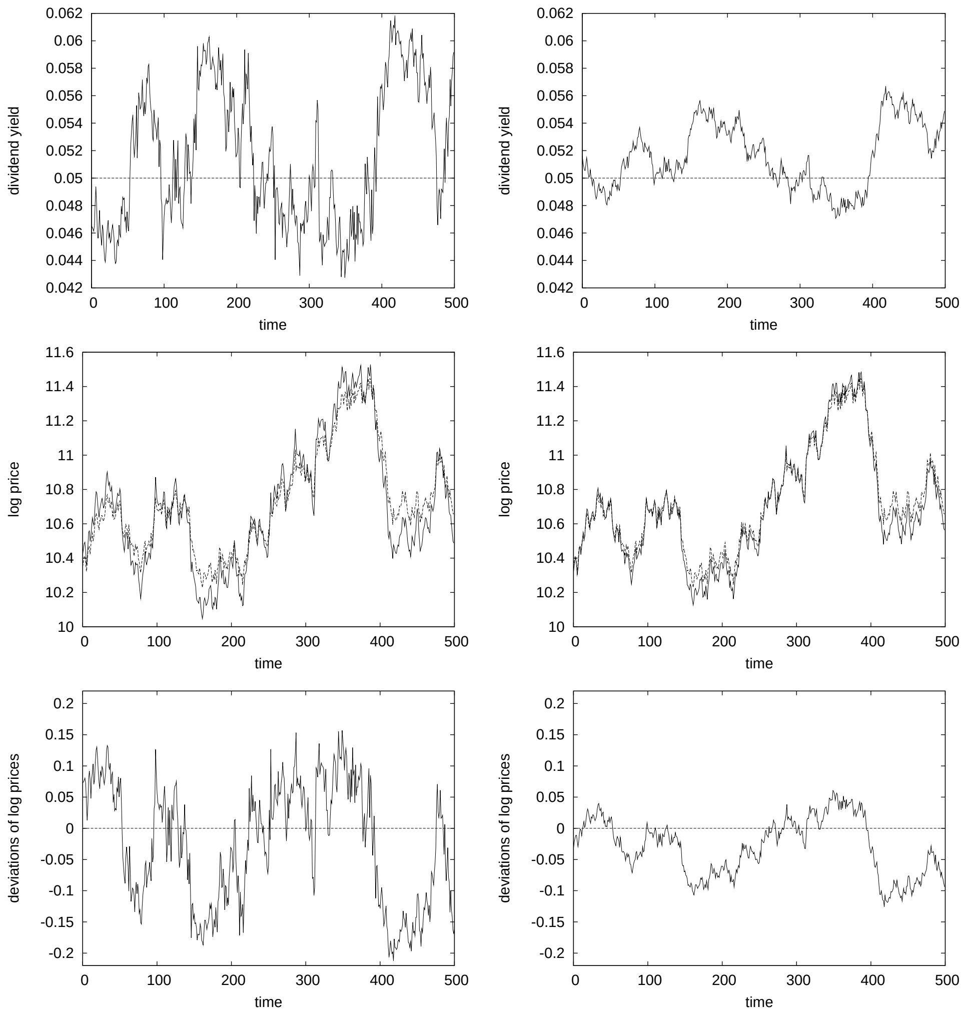

Figure 4.5: *Top left panel: time series of the dividend yield y_t generated by our model (4.37-4.39) (solid line) compared with the benchmark y^{**} (4.40) (horizontal dotted line). Top right panel: time series of the dividend yield y^* as in Barsky and De long (4.42) (solid line) compared with y^{**}. Middle left panel: logarithm of price corresponding to y, $\log(p)$ (solid line), and logarithm of the price corresponding to y^{**}, $\log(p^{**})$ (dotted line). Middle right panel: logarithm of price corresponding to y^*, $\log(p^*)$ (solid line), and logarithm of the price implied corresponding to y^{**}. The bottom panels give the deviations of the log prices series shown in the middle panels. Values of parameters are $\alpha = \gamma = 0.99$, $\rho = 1$, $c = 0.1$, $\delta = 0.02$ (these three parameters imply $\lambda^* \approx 0.09$), $\sigma_\eta = 0.04$, $\mu = 0.003$. The discount rate is $r = 0.05 + g$.*

realistic than others. We merely offer a qualitative argument to show that time varying informational differences might explain the nonlinearity in the mean reversion that has been shown to exist in the empirical literature.

Another way of comparing the various models is to check for correlation in the time series of returns produced by the evolution of y^* and y. The left panel of Figure 4.7 shows

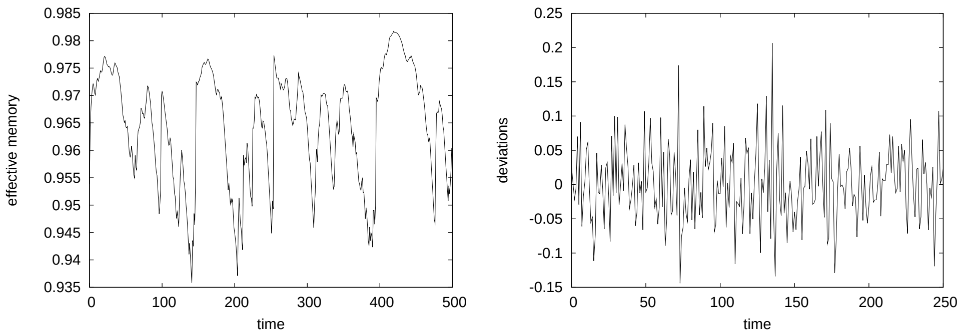

Figure 4.6: *Left panel: values of the effective memory $\tilde{\gamma}$. Right panel: deviations of the two series of log prices generated by our model y and by the "modified" Barsky and De Long model $y^*_{\tilde{\gamma}}$. Deviations are due to fluctuations of $\tilde{\gamma}$ around its mean value, $\tilde{\gamma} = 0.9644$. Other parameters as in Figure 4.5, in particular the real memory is $\gamma = 0.99$.*

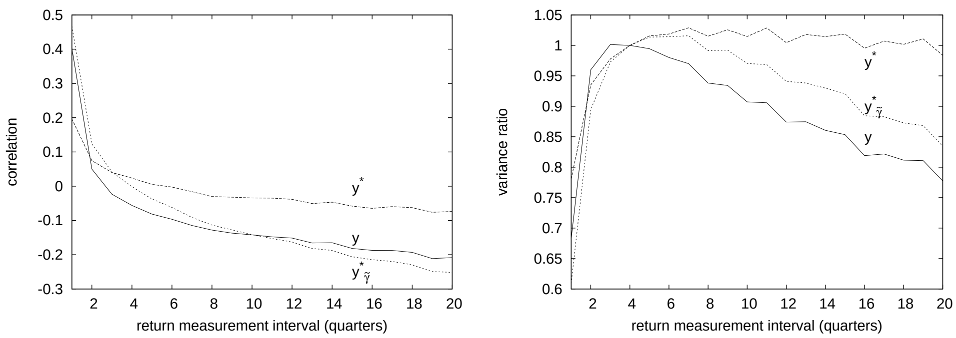

Figure 4.7: *Left panel: autocorrelation of the log return time series for a typical run of our model (4.37-4.39). Right Panel, variance test of log lagged returns. $v(\Delta t) = (\sigma^2(q^{\Delta t})/\Delta t)/(\sigma^2(q^4_y)/4)$ where $\sigma^2(q^4_y) = 0.00355$) and $q^{\Delta t}$ is the total return over a period Δt generated by the model in system (4.37-4.39). In both plots three lines refer to data generated with our model, y, to the model of Barsky and De Long in (4.42), y^*, and to the "modified" Barsky and De Long model. Parameters are the same as for Figure 4.5. Both panels refer to a simulation of* 200, 000 *periods.*

the autocorrelation of the asset log return series q_t,

$$q_t = \log(p_t + d_t) - \log(p_{t-1}), \tag{4.48}$$

for a typical run of our model (4.37-4.39). The autocorrelation of returns shows that our model and the "modified" Barsky and De Long model have higher short term autocorrelation and lower long term autocorrelation. Such results are in accordance with the results shown in the left panel of Figure 4.4: if the effective memory is lower than the real memory, shocks have a higher short run impact but a shorter half life for p then for p^*.

A test that has been used in the literature to evaluate the statistical importance of departure of the models from a random walk with drift is the variance ratio test. The variance ratio has been used by Poterba and Summers (1988) to appraise the mean reversion properties of stock prices. Under the null hypothesis that log prices follows a random walk (possibly with drift) the variance of the series of returns in (4.48) is a linear function of the return time span. Results of the variance ratio test for our model and for its restrictions are given in the right panel of Figure 4.7. The results suggest that both the fact that the effective memory is lower than the real memory and but also that the effective memory $\tilde{\gamma}(\lambda)$ is time varying affect the statistical time series properties of lagged returns. Further research will be devoted to investigating these issues in greater detail and relate them to the statistical properties of empirical market returns.

4.5 Discussion and conclusion

We have investigated the combined effect of informational differences and learning in a stylized asset market model where agents are boundedly rational. As far as the theoretical guidelines behind our model are concerned, we have shown that our model naturally and parsimoniously extends and links many other contributions in this fields. In particular, we refer to papers that concentrate on informational differences, such as Grossman and Stiglitz (1980), that analyze the impact of learning, such as Barsky and De Long (1993), and that investigate the interaction of agents who are using different predictor schemes or different strategies, such as Brock and Hommes (1998).

In Section 4.3 we have analyzed the impact of informational differences alone and we have given a dynamic solution of the GS paradox. We have shown that our model results in a “dynamic equilibrium degree of disequilibrium”, with a pattern of a-periodic oscillations where prices are not fully informative. Moreover, the informational content of the price is time varying and switches repeatedly between being nearly fully informative and hardly informative. The economic intuition behind this is quite clear: when the fraction of informed agents is sufficiently high, we are in a region where the adaptive expectations of boundedly rational uninformed agents are converging to rational expectations and the price is close to being fully informative. This pushes down the fraction of informed agents. As a result, with only few agents being informed, the uninformed agents using adaptive expectations do not converge to rational expectations anymore and the price starts to diverge from its fundamental value and, as a consequence, carries little information. This creates incentives to buy information and pushes the fraction of informed agents up again, and the story repeats. This trade-off between local instability (when too few agents are informed) and global stability (when many agents are informed) leads to complicated dynamic behavior. In our model, the “dynamic equilibrium degree of disequilibrium” is therefore a time-varying learning equilibrium where prices fluctuate between being close to fully revealing and being uninformative, and agents switch between costly information gathering and free riding.

In Section 4.4, using the insights gained from the analysis of Section 4.3, we have investigated the role of informational differences in a model where boundedly rational agents are learning the growth rate of dividends. We have compared our model to that of Barsky and De Long (1993), who have studied a similar framework without taking informational differences into account. Both models lead to mean reversion behavior of stock prices, i.e. in both cases the log price is given by the sum of a persistent component, which follows a random walk with drift, and of a temporary component, which follows a stationary AR(1) process. The innovative contribution of our model consists in showing that including the effect of informational differences has an impact on the parameters of the temporary component. In particular, both the linear coefficient and the variance of the shocks become time-varying, as they are both dependent on the time-varying fraction of informed agents. As a result, our model has the same features of recent empirical investigations of mean reversion conditions for stock prices, such as Gallagher and Taylor (2001) and Manzan (2003) where time varying parameters of the transitory component have been detected.

It is beyond the scope of this chapter to calibrate our model to reproduce the stock price evolution given the historical dividend process. Our theoretical model is based on several simplifying assumptions. Nevertheless we find it instructive to note that time varying parameters in a mean reversion model can be related to what in general may be referred to as agents' interaction, which, in our case, is triggered by informational differences. That agent interaction can be responsible for nonlinearity in the behavior of stock prices is also argued by Boswijk, Hommes, and Manzan (2007), who estimate a modified version of the model of Brock and Hommes (1998) using yearly data of the index SP500. Further efforts to characterize the effect of informational differences, for example by linking it to other observable characteristics such as the volume of transactions, might provide a basis for the design of new econometric tests for structure in financial time series.

Appendix

4.A Derivation of the equilibrium price equation

This appendix provides a micro-foundation of the equilibrium price equation (4.5). Consider a group of agents choosing at every time t whether to invest in a risk free asset, whose single period return is r_f, or to invest in a risky asset, whose single period return rate, r_{t+1}, depends on dividend payed at time $t+1$, d_{t+1}, and on the price (or remaining value) of the asset at time $t+1$, p_{t+1}:

$$r_{t+1} = \frac{p_{t+1} + d_{t+1} - p_t}{p_t}. \tag{4.49}$$

At every time t, we assume that each agent maximizes a CRRA mean variance utility function in order to decide which fraction x_t of his wealth to invest in the risky asset. The CRRA utility function to be maximized is

$$U(x_t) = E_t[x_t r_{t+1} + (1 - x_t) r_f] - \frac{\beta}{2} V_t[x_t r_{t+1}],$$

where E_t and V_t denote, respectively, the mean and the variance conditional on the information available at time t and β is the coefficient of risk aversion which we assume constant across agents. Assume also that for each agent $V_t[r_{t+1}] = \sigma^2$. The solution of the maximization of $U(x_t)$ gives:

$$x_t = \frac{E_t[r_{t+1} - r_f]}{\beta\sigma^2}, \tag{4.50}$$

as the fraction of wealth to be invested in the risky asset at time t. Consider now the case where assumptions $(i) - (iv)$ of Section 4.2 hold. Since informed and uninformed agents have different information, they have different expectations $E_t[r_{t+1}]$ and therefore different demands for risky assets. Call x_t^I (x_t^U) and w_t^I (w_T^U), respectively, the fraction of wealth to be invested at time t and the wealth at time t of the informed (uninformed) agents. Assume a net positive supply of shares s_t and call p_t the price of each share at time t. The Walrasian equilibrium equation at time t is given by:

$$s_t p_t = x_t^I w_t^I + x_t^U w_t^U. \tag{4.51}$$

Now, define λ_t as the fraction of wealth, or market power, of the informed agents at time t. This implies $\lambda_t = w_t^I / w_t$ with $w_t = w_t^I + w_t^U$. Call θ_t the average proportion, at time t, of wealth invested in the risky assets, i.e. $\theta_t = s_t p_t / w_t$. The Walrasian equilibrium equation (4.51) becomes

$$\theta_t = \lambda_t x_t^I + (1-\lambda_t) x_t^U,$$

which, given the expression for x in (4.50) and rearranging terms, becomes

$$r_f + \theta_t \beta \sigma^2 = \lambda_t E[r_{t+1}|\mathcal{F}_t^I] + (1-\lambda_t) E[r_{t+1}|\mathcal{F}_t^U], \tag{4.52}$$

where the informed and uniformed agents condition their expectations of the return of the risky asset on different information sets. Eq. (4.52) shows that θ_t is related to the risk premium required by the community of traders to hold the risky asset.

If we assume that both informed and uninformed agents are fundamentalists, their expectations of the future price is equal to the discounted sum of all future dividends, i.e., as specified in equation (4.2),

$$E\left[p_{t+1}|\,\mathcal{F}_t^H\right] = v_{t+1}^H = E\left[\sum_{i=1}^{\infty} \frac{d_{t+1+i}}{(1+r)^i}\,\middle|\,\mathcal{F}_t^H\right].$$

for the general information set $\mathcal{F}_t^H$. This implies that the informed agents use:

$$E[r_{t+1}|\mathcal{F}_t^I] = \frac{d_{t+1}(1+r)}{(r-g^e)p_t} - 1, \tag{4.53}$$

whereas the uninformed agents use:

$$E[r_{t+1}|\mathcal{F}_t^U] = \frac{p_t y_{t+1}^e (1+r)}{(r-g^e)p_t} - 1. \tag{4.54}$$

As a result the equilibrium equation (4.52) becomes:

$$p_t \frac{1 + r_f + \theta_t \beta \sigma^2}{1+r} = \lambda_t \frac{d_{t+1}}{(r-g^e)} + (1-\lambda_t) \frac{p_t y_{t+1}^e}{(r-g^e)}. \tag{4.55}$$

At this point, by fixing $r - r_f = \theta_t \beta \sigma^2$, that is by imposing that the asset excess return required by the agents, which is an exogenous variable of the model, is equal to the endogenous variable $\theta_t \beta \sigma^2$, we get Eq. (4.5) which solved for p_t gives Eq. (4.11).

By fixing $r - r_f = \theta_t \beta \sigma^2$ we are implicitly assuming that the endogenous expected equilibrium return of our model is given by the exogenous parameter r. In fact, when the price is informationally efficient, the resulting expected and realized dividend yield are equal to $y^* = r - g^e$ so that using Eq. (4.54), or equivalently (4.53), to compute the expected equilibrium asset return we get:

$$E[r_{t+1}|\mathcal{F}_t^U] = E[r_{t+1}|\mathcal{F}_t^I] = \frac{y^e(1+r)}{(r-g^e)} - 1 = r.$$

In order to derive the asset return, r_t, endogenously as a function of the exogenous parameters r_f, σ^2, β, s_t and of the discount rate r, one should solve the equilibrium price equation (4.55), without fixing $s_t p_t / w_t$ proportional to $r - r_f$. Levy, Levy, and Solomon (1994) are, to our knowledge, the first to perform this kind of analysis. They use computer simulations to investigate the evolution of wealth and prices in an asset market where, as in our framework, agents are using CRRA utility function and the underlying dividend process follows a geometric random walk. In a recent paper Anufriev and Dindo (2006) offer analytic support of their simulations. Other analytic models where agents are using a CRRA utility function are Chiarella and He (2001) and Anufriev, Bottazzi, and Pancotto (2006). Here, having assumed that $E_t[r_{t+1}] = r$, we are fixing the long run asset return and we concentrate on the properties of the fluctuations induced by agents interaction around this long run equilibrium level.

4.B Derivation of the dynamics of fractions

The micro-foundation of the equilibrium price equation (4.55) offers an appealing interpretation of λ_t as the fraction of wealth of the informed agents, and a natural way to endogenize its evolution. In fact the wealth fraction at time t, λ_t, is endogenously determined as a function of the fraction at time $t-1$, λ_{t-1}, the fraction of wealth invested in the risky asset by both groups of agents at time t, x_{t-1}^I and x_{t-1}^U, and of the return of the market at time t, r_t in (4.49). Using the definition of the fraction of wealth of the informed agent $\lambda_t = w_t^I/(w_t^I + w_t^U)$ and wealth evolution

$$w_t^H = w_{t-1}^H(1+r_f) + w_{t-1}^H(r_t - r_f)x_{t-1}^H, \quad \text{for} \quad H = I, U,$$

one can derive the equation that governs the evolution of the fraction of wealth of the informed agents λ_t:

$$\lambda_t = \frac{\lambda_{t-1}\pi_t^I}{\lambda_{t-1}\pi_t^I + (1-\lambda_{t-1})\pi_t^U}, \tag{4.56}$$

where

$$\pi_t^I = (1+r_f) + (r_t - r_f)x_{t-1}^I, \tag{4.57}$$

and

$$\pi_t^U = (1+r_f) + (r_t - r_f)x_{t-1}^U, \tag{4.58}$$

are the realized excess profit per unit of wealth for informed and for uninformed agents respectively. When the realized single period return of the asset r_{t+1} is higher than r_f, if the informed agents invest a higher (lower) share of their wealth compared to the share of uninformed agents, their fraction of wealth increases (decreases) compared to the fraction of wealth of the uninformed agents. Equation (4.56) corresponds to the replicator dynamics equation given in Eq. (4.24) when $\delta = 0$. To obtain exactly (4.24) one has to further assume that (4.20) and (4.21) can be used as proxies of the realized profits per unit of wealth for respectively the informed agent, as (4.57), and for the uninformed agent, as (4.58). In fact one has to assume that the dynamics of the fractions is driven by the forecasting error of the uninformed compared to the costs of information for the informed, rather than by their realized profits. Investigation of this second framework would lead to a more complicated system due to the presence of r_t, and thus of both y_t and p_t/p_{t-1}, in the expression of agents' profits. To conclude the correspondence between Eq. (4.56) and Eq. (4.24) we have to discuss the case $\delta \neq 0$. Assume that, in every period, a number of agents which holds a fraction δ of total agents' wealth exits the market and is replaced by new agents with the same amount of wealth. Also assume that these new agents split evenly between being informed and being uninformed. This would mean that at period t the total fraction of informed agents is given by:

$$\lambda_t = (1-\delta)\frac{\lambda_{t-1}((1+r_f)+\pi_t^I)}{(1+r_f)+\lambda_{t-1}\pi_t^I+(1-\lambda_{t-1})\pi_t^U}+\frac{\delta}{2}, \tag{4.59}$$

which is as (4.24) in Subsection 4.2.3 for every $\delta \in [0,1]$. Notice that irrespectively of the fitness measure, realized profits or forecasting errors, both expressions (4.24) and (4.59) for the fraction of informed agents λ_t have the same dependence on the previous fraction of informed agents λ_{t-1}.

4.C Proofs

Proofs of Section 4.3

Proof of Proposition 4.8 Given the linear difference equation in (4.32), that is

$$y_t^e = \frac{r-g}{\lambda} + \frac{\alpha+\lambda-1}{\lambda} y_{t-1}^e,$$

and the initial condition y_0, from the theory of linear systems follows that

$$y_t^e = (y_0^e - (r-g))\left(\frac{\alpha+\lambda-1}{\lambda}\right)^t + (r-g),$$

is the unique solution. The solution converges to $y^* = r - g$ as long as $\lambda > (1-\alpha)/2$ otherwise it diverges to $\pm\infty$. □

Proof of Proposition 4.9 When $c = 0$ and $\delta > 0$ solving for the fixed point of (4.30-4.31) leads to $y^* = r - g$ and to $\lambda^* = 1/2$. When $c = 0$ and also $\delta = 0$, solving for the fixed point of (4.30-4.31) leads to $y^* = r - g$ and to any $\lambda^* \in (0, 1]$. For global stability we proceed by showing that for every initial condition the dynamics converges to (y^*, λ^*). Rewrite equation (4.30) in deviations from the fixed point $y^* = r - g$ and get:

$$y^e_{t+1} - y^* = \frac{\alpha + \lambda_{t-1} - 1}{\lambda_{t-1}}(y^e_t - y^*). \tag{4.60}$$

Whenever $\lambda_{t-1} > (1-\alpha)/2 = \bar{\lambda}$, $y^e \to y^*$ as the linear coefficient of equation (4.60) is always lower then one. Since $\alpha \in (0, 1)$, the condition $\lambda > \bar{\lambda}$ is always satisfied when $\lambda_{t-1} \geq 1/2$. We show that, in the long run, it always holds $\lambda \geq 1/2$. Take first $\lambda_{t-1} \geq 1/2$, by using equation (4.31) one can derive that:

$$\lambda_t - \frac{1}{2} = (1-\delta)\frac{\lambda_{t-1}\rho - (1-\lambda_{t-1})\left(\rho - \left(\frac{(r-g)-y^e_t}{\lambda_{t-1}(r-g)}\right)^2\right)}{2\left(-(1-\lambda_{t-1})\left(\frac{(r-g)-y^e_t}{\lambda_{t-1}(r-g)}\right)^2 + \rho\right)} \geq 0, \tag{4.61}$$

if one assumes that fractions are always positive, i.e. that

$$\rho - \left(\frac{(r-g)-y^e_t}{\lambda_{t-1}(r-g)}\right)^2 > 0.$$

As a consequence when $\lambda_{t-1} \geq 1/2$ also $\lambda_t \geq 1/2$. If, on the other hand, $\lambda_{t-1} < 1/2$, one can similarly show that the condition $(\lambda_t - \lambda_{t-1}) > 0$ holds, implying that λ converges to one point of the set $[1/2, 1]$. As a consequence also when $\lambda_{t-1} < 1/2$, $\lambda \geq 1/2$ in the long run. To conclude the proof of global stability notice that when $y^e \to y^*$ also $\lambda \to \lambda^*$ via equation (4.31). □

Proof of Proposition 4.10 When $c > 0$, solving for the fixed point of (4.30-4.31) leads to $y^* = r - g$ and to λ^* solution of the following second order equation

$$c\lambda^2 + (c\delta/2 - c - \delta\rho)\lambda + \delta\rho/2 = 0,$$

which can be shown to have two real roots. Take $\alpha \in (0,1)$, $\delta \in (0,1)$ and $k = \rho/c > 1$. From:

$$2 - \delta + 2k\delta - 1 > 0,$$

it follows that

$$\frac{2 - \delta + 2k\delta + \sqrt{(2 - \delta + 2k\delta)^2 - 8k\delta}}{4} > 1.$$

That is one real root is always larger then 1 and thus not in the co-domain of our state variable λ. The other solution can be shown to be always in the interval $[0, 1/2]$. In fact

$$0 < \lambda^* = \frac{2 - \delta + 2k\delta - \sqrt{(2 - \delta + 2k\delta)^2 - 8k\delta}}{4} < \frac{1}{2},$$

reduces to

$$0 < 8k\delta \quad \text{and} \quad -4(1 - \delta) < 0.$$

Both inequalities ara always satisfied. The Jacobian follows from evaluating the derivatives of (4.30-4.31) at the fixed point (y^*, λ^*). For the stability condition notice that the matrix is diagonal and the second eigenvalue, $\nu_2 \in (0,1)$. In fact since $\delta < 1$, $k > 1 > 0.5 > \lambda^*$ one has:

$$0 < \nu_2 = (1 - \delta)\frac{k(k-1)}{(k - \lambda^*)^2} < \frac{k(k-1)}{k(k-1) + (0.5)^2} < 1. \tag{4.62}$$

The value of the first eigenvalue, ν_1, depends upon the value of λ^*. This eigenvalue is the same as the linear coefficient of equation (4.19), that is, it is $\nu_1 < 1$ given $\alpha \in (0,1)$ and $\lambda^* \in (0,1)$ and $\nu_1 > -1$ as long as

$$\lambda^* > \frac{(1 - \alpha)}{2}.$$

Given the value of λ^* one can check that the previous inequality is satisfied if and only if

$$\delta > \frac{(1 + \alpha)}{1 + \frac{2\alpha k}{(1-\alpha)}}.$$

□

Proof of Proposition 4.11 The matrix $\mathbf{J}|_{(y^*,\lambda^*)}$ in (4.35) is diagonal. As a result the dynamics of y and λ around (y^*, λ^*) can be linearized along the orthogonal basis with eigenvalues given by the diagonal entries of the matrix. Thus the eigenvalues that governs the dynamics of λ, (4.62), is given by the entry $(2,2)$ of the matrix (4.35). We recall from the previous proof that

$$\nu_2 = (1-\delta)\frac{k(k-1)}{(k-\lambda^*)^2},$$

and we have already shown that $\nu_2 < 1$. The lower bound, $\nu_2 > (1-\delta)(1-1/k)$, follows from the previous expression and $\lambda^* > 0$ for all $\delta > 0$ and for all α. □

Chapter 5

Summary

In this thesis we aim at appraising the effect of modeling human decision making as boundedly rational in a number of economic settings. In particular we concentrate on cases where a large number of individuals or firms compete to use the same limited resource, being it a market, a highway, a bar or the information of profitability of a firm. A characteristic of all these frameworks is that agents' participation produces some negative externalities on the utility of profitability of other participating agents. In these systems, if agents were rational, they would account for the effects due to their interaction and coordinate their actions to the equilibrium level where they would all gain the same profit. The nature of this equilibrium will depend upon the exogenous variables that characterize the fundamentals of the model. Convergence to the equilibrium need not *a priori* take place when agents are boundedly rational. In this case the interaction of agents may or may not settle down to the equilibrium. Modeling agents as boundedly rational could thus help explaining the observed size and variability of fluctuations of economic variables, such as prices or profits, even when no changes of the underlying fundamentals occurs.

Boundedly rational agents can be specified as using simple rules of thumb in their decision making. An important aspect is that the population of agents is heterogeneous, that is, agents can choose from different rules to solve the same problem. The set of rules is disciplined by a selection mechanism, where the best performing rule, measured according some fitness rule, attracts the most number of agents. This feature implies that our models are dynamics, with agents switching among the different rules at different periods of time.

In departing from the traditional approach where a representative rational agent is present, we have two main goals. First, we appraise when a frequently used argument in favor of rationality, namely that rationality is the outcome of the repeated interaction of heterogeneous boundedly rational agents, is justified. This part of the analysis has thus a theoretical motivation. Second, having shown in what respect our results differ from the rational benchmark, we characterize whether our interacting agents framework can reproduce empirically observed phenomena in the specific economic settings we consider.

An important role in triggering the switching of rules is played by the feedback between agents' expectations of some relevant economic variables and their realizations. This feedback system, which we call expectational feedback, translates agents' interaction into a mutual dependence between choices of economic actors and the environment against which these choices are evaluated. Having to do with this feedback, rational agents are usually assumed to have rational expectations, that is, to find actions such that expectations and realizations are consistent and the system is at equilibrium. Boundedly rational agents, instead, are not assumed to be aware of the expectational feedback so that their interaction may or may not settle down to an equilibrium. It is the dynamics that tell us when convergence is observed and when not. In this respect, as the functional dependence of fractions of agents in terms of economic variables and previous fractions is typical nonlinear, our systems are nonlinear so that a number of different types of behavior other than convergence and divergence can be observed, such as bounded erratic fluctuations. For the systems we analyze the most common outcome is given by irregular fluctuations where agents keep switching between different rules, each rule being most chosen at different periods of time and no rule dominating in the long run.

Inside this framework, in chapter 2, we concentrate on the cobweb model and we analyze the impact of increasing deliberation costs as the sophistication of a decision rule increases. In a cobweb model producers of a perishable consumption good have to decide how much to produce at every period. In order to do so they rely on expectations of next periods prices. The characteristic of this model is that high (low) expected prices trigger high (low) producers output which clears the market at a low (high) expected prices. Agents with rational expectation are able to detect this effect and set an output level such that ex-ante and ex-post prices are the same. In our framework we assume that agents who use this rational expectations rule face a high deliberation cost. An alternative is to use simpler rules such as naive expectations, that is, the price of today is taken as the predictor of the price of tomorrow. This rules are available at a lower deliberation cost. We investigate the effect on prices on agents switching between the two prediction strategies.

In previous contributions it is assumed that agents are *backward looking* in the sense that strategy selection is based on

experience measured by past realized profits. In particular, Brock and Hommes (1997) show that when the selection pressure to switch to the more profitable strategy is high, instability and complicated chaotic price fluctuations arise. Brock and Hommes call this phenomenon a rational route to randomness. In Chapter 2, we develop a cobweb model with strategy switching, where strategy selection is based upon *expected* profits rather than *realized* profits. When the deliberation cost is high enough, we find that the same local instability of the equilibrium price. Nevertheless, this time, the amplitude of price fluctuations is dampened. In particular the route to randomness result of Brock and Hommes is not always observed. Our conclusion is that key elements of the route to randomness result are unanticipated errors of agents using rational and naive strategies. We show that when these unanticipated errors are not present the rational route to randomness disappears.

In Chapter 3 we analyze the interaction of boundedly rational agents using heterogeneous rules all repeatedly competing to use the same limited resource when no market institutions are present. We formalize this general interaction structure as a participation game. We concentrate on participation games for which the payoff for participating decreases as the number of participating agents increases. Well-known examples of these type of participation games analyzed in the literature are market entry games, where firms have to decide whether to enter a market and compete, or stay out of it. Another example is given by route choice games, where a group of commuters repeatedly choose, between two routes, the fastest way from their homes to their offices. A third example is give by the *El Farol* bar problem of Arthur (1994).

Our aim is to obtain a simple analytic model that can replicate the main experimental and computational findings in the area of participation games with negative feedback. We set up an analytic model with heterogeneous boundedly rational agents choosing between simple rules. Fractions of agents using each rule are endogenous and evolve according to the past performance of each rule as described by the replicator dynamics. We concentrate on the evolutionary dynamics produced by the competition between different deterministic rules that condition the participation decision on the outcome of the previous rounds. We characterize the resulting participation rate dynamics as the number of players increases. The interaction of a large number of players leads to complicated participation rate patterns and we argue that the average participation rate along these patterns is consistent with the symmetric Nash equilibrium and the existing experimental and computational evidence. We also investigate how agents' interaction is affected by the presence of agents choosing rules that try to exploit the linear autocorrelation structure of the past participation rates. Agents who try to exploit past inefficiencies of the aggregate to improve their performances, act as some kind of arbitrageurs and we show that their behavior does not affect the stability of the system and the size of the fluctuations.

In Chapter 4 we study a market for a financial asset populated by boundedly rational agents and we concentrate on the role of informational differences. The starting point is an asset pricing model in which agents can choose among two different degrees of information on fundamentals. At the same time agents are also learning the growth rate of the dividend generating process.

After having developed and analyzed the full model, we investigate the extent to which our model is able to explain empirical properties of asset prices such as excess volatility, clustered volatility, correlations of returns, persistent deviations from fundamental values. In particular we offer theoretical support to the empirical evidence that the log price of a financial asset is the sum of a persistent component and a nonlinear temporary component, which switches between two different regimes. The empirical evidence for this so-called nonlinear mean reversion is documented in the literature, see e.g. by Gallagher and Taylor (2001) and Manzan (2003).

This chapter is closely related to the work on informational efficiency by Grossman and Stiglitz (1980) who show that when agents can choose between being informed or not, there is no rational expectation equilibrium. This result has been referred to as the Grossman-Stiglitz paradox. In order to solve it Grossman and Stiglitz assume an exogenous noise process. In this case, the model of Grossman and Stiglitz leads to a *static equilibrium degree of disequilibrium*, where agents' fractions and price distribution are constant over time and a function of the exogenous noise parameters.

In our analysis we relax the assumption of agents' rationality. We also endogenize the dynamics of the fraction of agents choosing to buy costly information or to extract information about future dividends from the price. We show that the interaction of boundedly rational agents, triggered by informational differences, can act as a source of *endogenous noise* to the price dynamics and can offer a different solution to the Grossman-Stiglitz paradox. Since our approach is dynamics, we offer a *dynamic equilibrium degree of disequilibrium*, in contrast with the static solution of Grossman and Stiglitz.

Since all the models we consider are related to the same economic framework, the repeated exploitation of a limited resource, they all share the same type of expectational feedback. In particular, we can call it a negative expectational feedback as positive (negative) deviations from the equilibrium at one time results in a negative (positive) deviations at the other time. The case where positive (negative) deviations are followed by positive (negative) relations is instead called positive expectational feedback.

Theoretical results of this thesis show, for systems with negative feedback, that overall convergence to the rational equilibrium is on average correct. Nevertheless, we also show that boundedly rational agents' interaction and adaptation trigger ongoing fluctuations around such an equilibrium. This is consistent with experiments of systems with negative feedbacks as reported for market entry games and in route choice games. This is also consistent with excess volatility in financial markets. In general, these endogenous fluctuations can be characterized as irregular cycles along which rules perform better than others in different periods of time, but no rule is "dominating" the scene for every period. In this respect our results also support the observation that within systems with negative feedback incentives work in the direction of heterogeneity, that is, agents are better off if they do not imitate each other. In fact, such incentives explain why the persistence of heterogeneity, and consequently of endogenous fluctuations, is a robust characteristic of our models.

Bibliography

Alchian, A. A. (1950). Uncertainty, evolution and economic theory. *Journal of Political Economy* **58**, 211–221.

Anderson, S., de Palma, A. and Thisse, J. (1993). *Discrete choice theory of product differentiation*. MIT Press, Cambridge, MA.

Anufriev, A., Bottazzi, G. and Pancotto, F. (2006). Equilibria, stability and asymptotic dominance in a speculative market with heterogeneous agents. *Journal of Economic Dynamics and Control* **30**, 1787–1835.

Anufriev, M. and Dindo, P. (2006). Equilibrium return and agents' survival in a multiperiod asset market: Analytic support of a simulation model. In *Advances in Artificial Economies*, C. Bruun, ed. Springer Verlag, Berlin, pp. 269–282.

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

Barsky, R. B. and De Long, J. B. (1993). Why does the stock market fluctuate? *Quarterly Journal of Economics* **58**, 291–311.

Barucci, E., Monte, R. and Renó, R. (2004). Asset price anomalies under bounded rationality. *Computational Economics* **23**, 255–269.

Binmore, K. and Samuelson, L. (1997). Muddling through: Noisy equilibrium selection. *Journal of Economic Theory* **74**, 235–265.

Boswijk, H. P., Hommes, C. H. and Manzan, S. (2007). Behavioral heterogeneity in stock prices. *Journal of Economic Dynamics and Control*, forthcoming.

Bottazzi, G. and Devetag, G. (2003). A laboratory experiment on the minority game. *Physica A* **324**, 124–132.

Bray, M. (1982). Learning, estimation, and the stability of rational expectations. *Journal of Economic Theory* **26**, 318–339.

Brock, W. A., Dindo, P. and Hommes, C. H. (2006). Adaptive rational equilibrium with forward looking agents. *International Journal of Economic Theory*, fothcoming.

Brock, W. A. and Hommes, C. H. (1997). Rational routes to randomness. *Econometrica* **65**, 1059–1095.

Brock, W. A. and Hommes, C. H. (1998). Heterogeneous beliefs and routes to chaos in a simple asset pricing model. *Journal of Economic Dynamics and Control* **22**, 451–481.

Bulkley, G. and Tonks, J. (1989). Are U.K. stock prices excessively volatile? Trading rules and variance bound tests. *Economic Journal* **99**, 1083–1098.

Bullard, J. B. (1994). Learning equilibria. *Journal of Economic Theory* **64**, 468–485.

Bullard, J. B. (2006). The learnability criterion and monetary policy. *Federal Reserve Bank of St. Luis Review* **88**, 203–217.

Bulow, J. I., Geanakoplos, J. D. and Klemperer, P. D. (1985). Multimarket oligopoly: Strategic substitutes and complements. *Journal of Political Economy* **93**, 488–511.

Camerer, C. F. and Fehr, E. (2006). When does 'Economic Man' dominate social behavior. *Science* **311**, 47–52.

Camerer, C. F., Ho, T.-H. and Chong, J. K. (2002). Sophisticated EWA learning and strategic teaching in repeated games. *Journal of Economic theory* **104**, 137–188.

Challet, D. and Zhang, Y. (1997). Emergence of cooperation and organization in an evolutionary game. *Physica A* **246**, 407–418.

Challet, D. and Zhang, Y. (1998). On the minority game: Analytical and numerical studies. *Physica A* **246**, 514–532.

Chiarella, C. (1992). The dynamics of speculative behavior. *Annals of Operations Research* **37**, 101–123.

Chiarella, C. and He, X. (2001). Asset price and wealth dynamics under heterogeneous expectations. *Quantitative Finance* **1**, 509–526.

Chow, G. C. (1989). Rational versus adaptive expectations in present value models. *Review of Economics and Statistics* **71**, 376–384.

Conlisk, J. (1980). Costly optimizers versus cheap imitators. *Journal of Economic Behavior and Organization* **1**, 275–293.

Conlisk, J. (1996). Why bounded rationality? *Journal of Economic Literature* **34**.

Day, R. H. (1994). *An introduction to dynamical systems and market mechanism. Complex Economic Dynamics, Vol I.* MIT Press, Cambridge, MA.

De Fontnouvelle, P. (2000). Information dynamics in financial markets. *Macroeconomics Dynamics* **4**, 139–169.

Diks, C. and Dindo, P. (2006). Informational differences and learning in an asset market with boundedly rational agents. CeNDEF Working paper 06-11, University of Amsterdam.

Dindo, P. and Tuinstra, J. (2006). A behavioral model for participation games with negative feedback. CeNDEF Working paper 06-10, University of Amsterdam.

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

Dudek, M. K. (2004). Expectation formation and endogenous fluctuations in aggregate demand. Mimeo, Warsaw School of Economics (SGH).

Duffy, J. and Hopkins, E. (2005). Learning, information and sorting in market entry games: Theory and evidence. *Games and Economic Behavior* **51**, 31–62.

Erev, I. and Rapoport, A. (1998). Coordination, 'magic', and reinforcement learning in a market entry game. *Games and Economic Behavior* **23**, 146–175.

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

Evans, G. W. and Honkapohja, S. (2001). *Learning and Expectations in Macroeconomics.* Princeton University Press, Princeton.

Evans, G. W. and Ramey, G. (1992). Expectation calculation and macroeconomic dynamics. *American Economic Review* **82**, 207–224.

Ezekiel, M. (1938). The cobweb theorem. *Quarterly Journal Of Economics* **52**, 255–280.

Fama, E. F. and French, K. R. (1988a). Dividend yields and expected stock returns. *Journal of Financial Economics* **22**, 3–25.

Fama, E. F. and French, K. R. (1988b). Permanent and temporary componenents of stock prices. *Journal of Political Economy* **96**, 246–273.

Franke, R. (2002). Reinforcement learning in the El farol model. *Journal of Economic Behavior & Organization* **1503**, 1–22.

Frankel, J. A. and Froot, K. A. (1987). Using survey data to test standard propositions regarding exchange rate expectations. *American Economic Review* **77**, 133–153.

Friedman, M. (1953). *Essays in Positive Economics*. University of Chicago Press, Chicago.

Friedman, M. and Savage, L. (1948). The utility analysis of choices involving risk. *Journal of Political Economy* **56**, 279–304.

Fudenberg, D. and Levine, D. (1998). *The theory of learning in games*. MIT Press, Cambridge, MA.

Fudenberg, D. and Tirole, J. (1991). *Game Theory*. MIT Press, Cambridge, MA.

Gallagher, L. A. and Taylor, M. P. (2001). Risky arbitrage, limits of arbitrage, and nonlinear adjustment in the dividend-price ratio. *Economic Inquiry* **39**, 524–536.

Gandolfo, G. (1997). *Economic Dynamics*. Springer-Verlag, Berlin, 3rd ed.

Goeree, J. and Holt, C. (2005). An explanation of anomalous behavior in models of political participation. *American Political Science Review* **99**, 201–213.

Goeree, J. and Hommes, C. H. (2000). Heterogeneous beliefs and the non-linear cobweb model. *Journal of Economic Dynamics and Control* **24**, 761–798.

Goldbaum, D. (2005). Market efficiency and learning in an endogenously unstable environment. *Journal of Economic Dynamics and Control* **29**, 953–978.

Gordon, M. J. (1962). *The Investment, Financing, and Valuation of the Corporation*. Homewood, Ill.: R.D. Irwin.

Grandmont, J. M. (1982). Temporary general equilibrium theory. In *Handbook of Mathematical Economics*, K. J. Arrow and M. D. Intriligator, eds., vol. 2, chap. 19. Elsevier, North-Holland, pp. 879–922.

Grandmont, J. M. (1985). On endogenous competitive business cycles. *Econometrica* **53**, 995–1046.

Grandmont, J. M. (1998). Expectation formation and stability in large socio-economic systems. *Econometrica* **66**, 741–781.

Grossman, S. J. and Stiglitz, J. E. (1980). On the impossibility of informationally efficient markets. *American Economic Review* **70**, 393–408.

Guckenheimer, J. and Holmes, P. (1991). *Nonlinear oscillations, dynamical systems, and bifurcations of vector fileds*. Springer-Verlag, New York.

Haltiwanger, J. and Waldman, M. (1985). Rational expectations and the limits of rationalty: An analysis of heterogeneity. *American Economic Review* **75**, 326–340.

Heemeijer, P., Hommes, C., Sonnemans, J. and Tuinstra, J. (2006). Price stability and volatility in markets positive and negative expectations feedback: An experimental investigation. CeNDEF Working paper 06-05, University of Amsterdam.

Hellwig, M. F. (1982). Rational expectations with conditioning on past prices: A mean-variance example. *Journal of Economic Theory* **26**, 279–312.

Hommes, C. (2004). Economic dynamics. In *Encyclopedia of Nonlinear Science*. Routeledge.

Hommes, C. (2006). Heterogeneous agents models in economics and finance. In *Handbook of Computational Economics*, L. Tesfatsion and K. L. Judd, eds., vol. 2, chap. 23. Elsevier, North Holland, pp. 1109–1186.

Hommes, C., Sonnemans, J., Tuinstra, J. and Van de Velden, H. (2007). Learning in cobweb experiments. *Macroeconomic Dynamics*, forthcoming.

Hommes, C. H. (1998). On the consistency of backward-looking expectations: The case of the cobweb. *Journal of Economic Behavior and Organization* **33**, 333–362.

Hong, H. and Stein, J. C. (1999). A unified theory of underreaction, momentum trading and overreaction in asset markets. *Journal of Finance* **54**, 2143–2184.

Kirman, A. P. (1992). Whom or what does the representative individual represent? *Journal of Economic Perspectives* **6**, 117–136.

Kirman, A. P. (2006). Heterogeneity in economics. *Journal of Economic Interaction and Coordination* **1**, 89–117.

LeBaron, B. (2006). Agent based model in finance. In *Handbook of Computational Economics*, L. Tesfatsion and K. L. Judd, eds., vol. 2, chap. 24. Elsevier, North Holland, pp. 1187–1234.

LeRoy, S. F. and Porter, R. D. (1981). The present-value relation: Tests based on implied variance bounds. *Econometrica* **49**, 555–574.

Levy, M., Levy, H. and Solomon, S. (1994). A microscopic model of the stock market: Cycles, booms, and crashes. *Economic Letters* **45**, 103–111.

Li, T.-Y. and Yorke, J. (1975). Period three implies choas. *The American Mathematical Monthly* **82**, 985–999.

Manski, C. and McFadden, D. (1981). *Structural Analysis of Discrete Data with Econometric Applications*. MIT Press, Cambridge, MA.

Manzan, S. (2003). Nonlinear mean reversion in stock prices. CeNDEF Working paper 03-02, University of Amsterdam.

Mas-Colell, A., Whinston, M. D. and Green, J. R. (1995). *Micoreconomic Theory*. Oxford University Press, Oxford.

Maynard-Smith, J. (1974). The theory of games and the evolution of animal conflicts. *Journal of Theoretical Biology* **47**, 209–221.

McFadden, D. (1973). Conditional logit analysis of qualitative choice behavior. In *Frontiers in Econometrics*, P. Zarembka, ed. New York: Academic Press, NY, pp. 105–142.

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

McKelvey, R. and Palfrey, T. (1998). Quantal response equilibria for extensive form games. *Experimental Economics* **1**, 9–41.

Medio, A. and Lines, M. (2001). *Non-linear Dynamics. A primer*. Cambridge University Press, Cambridge.

Mehra, R. and Prescott, E. (1985). The equity premium: A puzzle. *Journal of Monetary Economics* **15**, 145–161.

Meyer, D., Huyck, J. V., Battalio, R. and Saving, T. (1992). History's role in coordinating decentralized allocation decisions. *Journal of Political Economy* **100**, 292–316.

Muth, J. F. (1961). Rational expectations and the theory of price movements. *Econometrica* **29**, 315–335.

Ochs, J. (1990). The coordination problem in decentralized markets: An experiment. *Quarterly Journal of Economics* **105**, 545–559.

Palm, G. (1984). Evolutionary stable strategies and game dynamics for n-person games. *Journal of Mathematical Biology* **19**, 329–334.

Poterba, J. M. and Summers, L. H. (1988). Mean reversion in stock prices. *Journal of Financial Economics* **22**, 27–59.

Radner, R. (1982). Equilibrium under uncertainty. In *Handbook of Mathematical Economics*, K. J. Arrow and M. D. Intriligator, eds., vol. 2, chap. 20. Elsevier, North-Holland, pp. 923–1006.

Rapoport, A., Seale, D., Erev, I. and Sundali, J. (1998). Equilibrium play in large groups market entry games. *Management Science* **44**, 119–141.

Routledge, B. R. (1999). Adaptive learning in financial markets. *Review of Financial Studies* **12**, 1165–1202.

Rubinstein, A. (1998). *Modeling Bounded Rationality.* MIT Press, Cambridge, MA.

Sargent, T. (1993). *Bounded Rationality in Macroeconomics.* Clarendon Press, Oxford.

Schönhofer, M. (1999). Chaotic learning equilibria. *Journal of Economic Thoery* **89**, 1–20.

Selten, R., Chmura, T., Pitz, T., Kube, S. and Schreckenberg, M. (2006). Commuters route choice behavior. *Games and Economic Behavior*, forthcoming.

Sethi, R. and Franke, R. (1995). Behavioural heterogeneity under evolutionary pressure: Macroeconomic implications of costly optimization. *The Economic Journal* **105**, 583–600.

Shiller, R. (1989). *Market Volatility.* MIT Press, Cambridge, MA.

Shiller, R. J. (1981). Do stock prices move too much to be justified by subsequent changes in dividends? *American Economic Review* **71**, 421–436.

Shiller, R. J. (1984). Stock prices and social dynamics. *Brookings Papers on Economic Activity* **2**, 457–510.

Simon, H. (1957). *Models of man.* John Wiley & Sons, New York.

Simon, H. A. (1955). A behavioral model of rational choice. *Quarterly Journal of Economics* **69**, 99–118.

Summers, L. H. (1986). Does the stock market rationally reflect fundamental values. *Journal of Finance* **41**, 591–602.

Sundali, J., Rapoport, A. and Seale, D. (1995). Coordination in market entry games with symmetric players. *Organizational Behaviour and Human Decision Processes* **64**, 203–218.

Taylor, P. and Jonker, L. (1978). Evolutionary stable strategies and game dynamics. *Mathematical Biosciences* **40**, 145–156.

Timmermann, A. (1993). How learning in financial markets generates excess volatility and predictability of excess returns. *Quarterly Journal of Economics* **108**, 1135–1145.

Timmermann, A. (1996). Excess volatility and predictability of stock prices in autoregressive dividend models with learning. *Review of Economics Studies* **63**, 523–557.

Tuinstra, J. (2003). Beliefs equilibria in an overlapping generations model. *Journal of Economic Behavior and Organization* **50**, 145–164.

Tversky, A. and Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science* **185**, 1124–1131.

Weibull, J. (1995). *Evolutionary Game Theory.* MIT Press, Cambridge, MA.

Weibull, J. W. (1994). The 'as if' approach to game theory: Three positive results and four obstacles. *European Economic Review* **38**, 868–881.

Wiggins, S. (1990). *Introduction to applied non-linear dynamical systems and chaos.* Springer-Verlag, New York.

Young, H. P. and Foster, D. (1991). Cooperation in the short and in the long-run. *Games and Economic Behavior* **3**, 145–156.

Zambrano (2004). The interplay between analytics and computation in the study of congestion externalities: The case of the el farol problem. *Journal of Public Economic Theory* **6**, 375–395.

Zwick, R. and Rapoport, A. (2002). Tacit coordination in a decentralized market entry game with fixed capacity. *Experimental Economics* **5**, 253–272.

Samenvatting (Summary in Dutch)

In dit proefschrift stellen we ons tot doel om de gevolgen te evalueren van het modelleren van menselijke beslissingen in termen van begrensde rationaliteit in verschillende economische situaties. We concentreren ons in het bijzonder op gevallen waarin een groot aantal individuen of bedrijven wedijveren om dezelfde schaarse goederen, of dat nu betrekking heeft op de aandelenmarkt, een snelweg, het aantal vrije plaatsen in een café of op informatie omtrent de winstgevendheid van een bedrijf. Al deze voorbeelden delen de eigenschap dat de keuze van een agent om het schaarse goed te benutten negatieve externaliteiten heeft voor de overige aanwezige agenten. Binnen dergelijke systemen zouden rationele agenten rekening houden met de effecten van hun interactie, en hun acties zo coördineren dat er een evenwichtssituatie zou onstaan waarin alle agenten even goed af zouden zijn. De aard van dit evenwicht zal afhangen van de details van de effecten van exogene variabelen op de fundamentele variabelen van het model. Convergentie naar een evenwicht treedt niet altijd op in het geval dat agenten begrensd rationeel zijn. Het modelleren van agenten als begrensd rationeel kan dus een mogelijke verklaring bieden voor de grootte en de veranderlijkheid van waargenomen fluctuaties van economische variabelen, zoals prijzen of winsten, zelfs als de fundamentele variabelen niet veranderen.

Begrensd rationele agenten kunnen worden gespecificeerd door middel van een aantal vuistregels met behulp waarvan zij beslissingen nemen. Een belangrijk aspect is heterogeniteit van de populatie van agenten, wat inhoudt dat agenten verschillende regels kunnen kiezen om hetzelfde probleem op te lossen. De verzameling regels is onderhevig aan een selectiemechanisme, waarbij de regel die, volgens een bepaalde norm, het best heeft gepresteerd, het grootste aantal aanhangers onder agenten krijgt. Deze eigenschap zorgt ervoor dat onze modellen dynamisch gedrag vertonen waarin agenten hun regelkeuze vaak herzien.

Met het verlaten van de traditionele aanpak met een representatieve agent beogen we twee hoofddoelstellingen. Ten eerste evalueren we wanneer een veel gehoord argument voor rationaliteit, namelijk dat rationaliteit de uitkomst is van herhaalde interactie van heterogene begrensd rationele agenten, gerechtvaardigd is. Dit deel van de analyse heeft dus een theoretische motivatie. Ten tweede bekijken we, nadat we hebben laten zien in welke zin onze resultaten afwijken van de rationele standaard, of modellen met interagerende agenten in staat zijn empirisch waargenomen fenomenen te reproduceren in de specifieke economische situaties die we beschouwen.

Bij het bepalen van de regelkeuze door agenten is er een belangrijke rol weggelegd voor de terugkoppeling van verwachtingen omtrent economische variabelen van agenten, door middel van gerealiseerde waarden van deze variabelen. Deze terugkoppeling van verwachtingen vertaalt de interactie tussen agenten naar een wederzijdse afhankelijkheid tussen enerzijds de keuzen van de agenten, en anderzijds de omgeving waarin deze keuzen worden geëvalueerd. Voor rationele agenten zou dit inhouden dat hun verwachtingen en bijbehorende acties zodanig zijn dat zij consistent zijn met toekomstige realisaties, en het systeem in evenwicht is. Begrensd rationele agenten, daarentegen, hoeven zich niet bewust te zijn van de terugkoppeling van verwachtingen, en dientengevolge hoeft hun interactie niet tot evenwicht te leiden. De resulterende dynamica zal bepalen of er evenwicht onstaat of niet. Omdat de fractie van agenten die kiest voor een bepaalde strategie typischerwijs geen lineaire functie is van voorgaande fracties en waargenomen economische variabelen, verkrijgen we een niet-lineair systeem waarin een verscheidenheid aan gedrag kan optreden, waaronder chaotische fluctuaties. Voor de systemen die wij analyseren bestaat het gedrag typisch uit begrensde onregelmatige fluctuaties waarin agenten voortdurend hun regelkeuze blijven aanpassen, met als gevolg dat verschillende regels op verschillende tijden het meest gekozen worden, en het lange termijn gedrag niet door één regel wordt bepaald.

Binnen dit raamwerk concentreren we ons in hoofdstuk 2 op het spinnewebmodel en analyseren de gevolgen van toenemende afwegingskosten voor meer geavanceerde beslissingsregels. In een spinnewebmodel moeten de producenten van een niet-houdbaar goed in elke periode beslissen hoeveel ze in de komende periode willen produceren. Deze keuze hangt af van hun verwachtingen omtrent toekomstige prijzen. In dit model leiden hoge (lage) verwachte prijzen tot hoge (lage) productie en daarmee tot lage (hoge) gerealiseerde toekomstige prijzen. Rationele agenten zijn in staat om dit effect te detecteren en zullen precies zoveel produceren dat de ex-ante en ex-post prijzen met elkaar overeenkomen. Wij zullen hier aannemen dat agenten die dergelijke verwachtingen hebben daar hoge afwegingskosten voor moeten betalen. Het alternatief is om eenvoudige regels te gebruiken zoals naïeve verwachtingen waarvoor de verwachte toekomstige prijs gelijk is aan de huidige prijs. Dergelijke regels zijn te gebruiken tegen veel lagere afwegingskosten. We onderzoeken de gevolgen voor het gedrag van prijzen wanneer agenten herhaaldelijk tussen deze twee voorspellers kunnen kiezen.

In eerder werk werd vaak aangenomen dat agenten *terugkijken*, waarmee wordt bedoeld dat de gekozen strategie gebaseerd is op in het verleden gerealiseerde winstgevendheid van de strategieën. In het bijzonder hebben Brock en Hommes (1997) laten zien dat wanneer de gevoeligheid van agenten voor het maken van een onjuiste keuze toeneemt, instabiliteit en complexe chaotische prijsfluctuaties kunnen ontstaan. Brock en Hommes noemen dit fenomeen een *rational route to randomness* (rationele weg tot willekeurige fluctuaties). In hoofdstuk 2 beschouwen we een spinnewebmodel waarin de strategiekeuze gebaseerd is op *verwachte* winsten in plaats van *gerealiseerde* winsten. Als de afwegingskosten voldoende hoog zijn, vinden we een soortgelijke lokale instabiliteit van de evenwichtsprijs. Dit keer blijft de amplitude van de prijsfluctuaties echter begrensd. In het bijzonder treedt de 'route to randomness' van Brock en Hommes hier niet altijd op. Onze conclusie is dat de belangrijkste ingrediënten van de 'route to randomness' onvoorziene fouten van zowel rationele als naïeve agenten zijn. We tonen aan dat in de afwezigheid van dergelijke fouten de 'rational route to randomness' niet optreedt.

In hoofdstuk 3 analyseren we de interactie tussen heterogene, begrensd rationele, agenten die herhaaldelijk om hetzelfde schaarse middel wedijveren in afwezigheid van marktregulerende instanties. We formaliseren deze algemene interactiestructuur als een deelnamespel. We concentreren ons op deelnamespelen waarin de winst van deelname afneemt naarmate meerdere agenten besluiten om deelenemen. Bekende voorbeelden van dergelijke spelen die in de literatuur beschreven zijn, zijn markttoetredingsspelen, waarin bedrijven moeten besluiten om toe te treden tot een bepaalde markt en daarbinnen te gaan concurreren, of afzijdig te blijven van die markt. Een ander voorbeeld is dat van routekeuzespelen, waarin een forensen herhaaldelijk kiezen tussen twee mogelijk routes met als doel om zo snel mogelijk van hun huis naar hun kantoor te komen. Een derde voorbeeld bestaat uit het *El Farol* bar probleem van Arthur (1994).

Ons doel is om een eenvoudig analytisch model op te stellen dat de belangrijkste empirische en numerieke resultaten kan reproduceren op het gebied van deelnamespelen met negatieve terugkoppeling. We stellen een analytisch model op met begrensd rationele heterogene agenten die tussen eenvoudige regels kiezen. De fracties van agenten die elk van de regels gebruiken is endogeen en ontwikkeld zich op grond van prestaties uit het verleden volgens de *replicator* dynamica. We concentreren ons op de evolutionaire dynamica die voortkomt uit de competitie tussen een aantal deterministische regels die de keuze voor deelname laten afhangen van de uitkomsten in vorige ronden. We bestuderen de resulterende dynamica van de fractie deelnemende agenten als het aantal spelers toeneemt. De interactie tussen een groot aantal spelers leidt tot ingewikkelde patronen in de deelnamefractie, en we maken aannemelijk dat de gemiddelde deelnamefractie consistent is met het symmetrische Nash evenwicht en met bestaand experimenteel en numeriek bewijs. We onderzoeken tevens hoe de interactie tussen agenten wordt beïnvloed door de aanwezigheid van agenten die de lineaire autocorrelatiestructuur van deelnamefracties uit het verleden proberen uit te buiten. Agenten die proberen om gebruik te maken van dergelijke inefficiënties gedragen zich in feite als arbitrageurs, en we laten zien dat hun aanwezigheid noch de stabiliteit van het systeem noch de grootte van de fluctuaties beïnvloedt.

In hoofdstuk 4 bestuderen we een model voor een aandelenmarkt die bevolkt wordt door begrensd rationele agenten, en concentreren ons op de rol van informatieverschillen onder agenten. We gaan uit van een model voor aandelenprijzen waarin agenten kunnen kiezen tussen twee gradaties van informatie omtrent toekomstige dividenden. Tevens is de gemiddelde mate van groei van het dividend genererende proces niet bekend bij de agenten, maar dienen zij deze te achterhalen uit gerealiseerde dividenden.

Nadat we het model hebben opgezet en geanaliseerd, bekijken we in welke mate het model in staat is om empirische eigenschappen van aandelenprijzen te reproduceren, zoals 'excess volatility' (bovenmatige beweeglijkheid), volatiliteitsclustering, autocorrelaties en de persistentie van afwijkingen van fundamentele waarden. In het bijzonder leiden we een theoretisch model af dat een verklaring biedt voor het empirische bewijs dat de logaritme van aandelenprijzen bestaan uit een som van een persistente component en een niet-lineaire vergankelijke component. Het empirische bewijs voor dergelijke bewegingen rondom de fundamentele prijs is beschreven in de literatuur door, onder anderen, Callagher en Taylor (2001) en Manzan (2003).

Dit hoofdstuk is nauw verwant aan het werk op het gebied van informationele efficiëntie door Grossman en Stiglitz (1980) die aantonen dat wanneer agenten kunnen kiezen tussen al dan niet geïnformeerd worden, er geen rationele verwachtingen evenwicht bestaat. Dit resultaat staat bekend als de Grossman-Stiglitz paradox. Grossman en Stiglitz lossen dit op door een exogeen ruisproces te veronderstellen. In dat geval leidt hun model tot een *statisch evenwicht van uit-evenwicht-zijn*, waarin zowel de fractie van geïnformeerde agenten als de prijsverdeling constant zijn over de tijd, met eigenschappen die bepaald worden door die van de exogene ruisbron.

In onze analyse laten we de aanname van rationaliteit los. Verder laten we de fracties agenten die willen betalen voor informatie en die informatie uit prijzen willen achterhalen, zich endogeen ontwikkelen. We tonen aan dat de interactie van begrensd rationele agenten sich in dit geval gedraagd als een *endogene* ruis op de prijsdynamica, die zodoende een alternatieve oplossing biedt voor de Grossman-Stiglitz paradox. Omdat onze opzet dynamisch is, leidt dit tot een *dynamisch evenwicht van uit-evenwicht-zijn* in plaats van de statische oplossing van Grossman en Stiglitz.

Omdat de behandelde modellen alle gerelateerd zijn aan hetzelfde economische raamwerk, het herhaaldelijk exploiteren van een schaars goed, hebben al deze modellen eenzelfde type terugkoppeling van verwachtingen. We noemen dit een een negatieve terugkoppeling van verwachtingen omdat positieve (negatieve) afwijkingen van evenwicht op het ene moment leiden tot negatieve (positieve) afwijkingen op het volgende moment. In het geval van positieve terugkoppeling van verwachtingen, daarentegen, worden positieve (negatieve) afwijkingen gevolgd door positieve (negatieve).

De theoretische resultaten in dit proefschrift tonen aan dat voor modellen met een negatieve terugkoppeling gemiddeld gesproken convergentie naar het rationele evenwicht optreedt. Het blijkt echter ook dat de interactie en aanpassingen van begrensd rationele agenten aanleiding geven tot voortdurende fluctutaties rond dit evenwicht. Dit is consistent met experimenten met negatieve terugkoppeling zoals die beschreven zijn voor spelen voor markttoetreding en routekeuze. Dit is ook consistent met bovenmatige beweeglijkheid in financiële markten. Over het algemeen kunnen deze endogene fluctuaties worden gekarakteriseerd als onregelmatige banen waarlangs bepaalde regels soms beter presteren, en andere weer op andere tijden, maar er geen regel "dominant" is in elke periode. In die zin ondersteunen onze resultaten de waarneming dat binnen systemen met negatieve terugkoppeling prikkels in de richting van heterogeniteit werken. Dat wil zeggen dat agenten beter af zijn als ze elkaar niet imiteren. Dit is de reden dat de persistentie van de heterogeniteit en de endogene fluctuaties een robuuste karakteristieke eigenschap van onze modellen vormen.

The Tinbergen Institute is the Institute for Economic Research, which was founded in 1987 by the Faculties of Economics and Econometrics of the Erasmus Universiteit Rotterdam, Universiteit van Amsterdam and Vrije Universiteit Amsterdam. The Institute is named after the late Professor Jan Tinbergen, Dutch Nobel Prize laureate in economics in 1969. The Tinbergen Institute is located in Amsterdam and Rotterdam. The following books recently appeared in the Tinbergen Institute Research Series:

346 P.D.M.L. CARDOSO, The future of old-age pensions: Its implosion and explosion.
347 C.J.H. BOSSINK, To go or not to go…? International relocation willingness of dual-career couples.
348 R.D. VAN OEST, Essays on quantitative marketing models and Monte Carlo integration methods.
349 H.A. ROJAS-ROMAGOSA, Essays on trade and equity.
350 A.J. VAN STEL, Entrepreneurship and economic growth: Some empirical studies.
351 R. ANGLINGKUSUMO, Preparatory studies for inflation targeting in post crisis Indonesia.
352 A. GALEOTTI, On social and economic networks.
353 Y.C. CHEUNG, Essays on European bond markets.
354 A. ULE, Exclusion and cooperation in networks.
355 I.S. SCHINDELE, Three essays on venture capital contracting.
356 C.M. VAN DER HEIDE, An economic analysis of nature policy.
357 Y. HU, Essays on labour economics: Empirical studies on wage differentials across categories of working hours, employment contracts, gender and cohorts.
358 S. LONGHI, Open regional labour markets and socio-economic developments: Studies on adjustment and spatial interaction.
359 K.J. BENIERS, The quality of political decision making: Information and motivation.
360 R.J.A. LAEVEN, Essays on risk measures and stochastic dependence: With applications to insurance and finance.
361 N. VAN HOREN, Economic effects of financial integration for developing countries.
362 J.J.A. KAMPHORST, Networks and learning.
363 E. PORRAS MUSALEM, Inventory theory in practice: Joint replenishments and spare parts control.
364 M. ABREU, Spatial determinants of economic growth and technology diffusion.
365 S.M. BAJDECHI-RAITA, The risk of investment in human capital.
366 A.P.C. VAN DER PLOEG, Stochastic volatility and the pricing of financial derivatives.
367 R. VAN DER KRUK, Hedonic valuation of Dutch Wetlands.
368 P. WRASAI, Agency problems in political decision making.

369 B.K. BIERUT, Essays on the making and implementation of monetary policy decisions.
370 E. REUBEN, Fairness in the lab: The effects of norm enforcement in economic decisions.
371 G.J.M. LINDERS, Intangible barriers to trade: The impact of institutions, culture, and distance on patterns of trade.
372 A. HOPFENSITZ, The role of affect in reciprocity and risk taking: Experimental studies of economic behavior.
373 R.A. SPARROW, Health, education and economic crisis: Protecting the poor in Indonesia.
374 M.J. KOETSE, Determinants of investment behaviour: Methods and applications of meta-analysis.
375 G. MÜLLER, On the role of personality traits and social skills in adult economic attainment.
376 E.H.B. FEIJEN, The influence of powerful firms on financial markets.
377 J.W. GROSSER, Voting in the laboratory.
378 M.R.E. BRONS, Meta-analytical studies in transport economics: Methodology and applications.
379 L.F. HOOGERHEIDE, Essays on neural network sampling methods and instrumental variables.
380 M. DE GRAAF-ZIJL, Economic and social consequences of temporary employment.
381 O.A.C. VAN HEMERT, Dynamic investor decisions.
382 Z. ŜAŜOVOVÀ, Liking and disliking: The dynamic effects of social networks during a large-scale information system implementation.
383 P. RODENBURG, The construction of instruments for measuring unemployment.
384 M.J. VAN DER LEIJ, The economics of networks: Theory and empirics.
385 R. VAN DER NOLL, Essays on internet and information economics.
386 V. PANCHENKO; Nonparametric methods in economics and finance: dependence, causality and prediction.
387 C.A.S.P. SÀ, Higher education choice in The Netherlands: The economics of where to go.
388 J. DELFGAAUW, Wonderful and woeful work: Incentives, selection, turnover, and workers' motivation.
389 G. DEBREZION, Railway impacts on real estate prices.
390 A.V. HARDIYANTO, Time series studies on Indonesian rupiah/USD rate 1995–2005.
391 M.I.S.H. MUNANDAR, Essays on economic integration.
392 K.G. BERDEN, On technology, uncertainty and economic growth.
393 G. VAN DE KUILEN, The economic measurement of psychological risk attitudes.
394 E.A. MOOI, Inter-organizational cooperation, conflict, and change.
395 A. LLENA NOZAL, On the dynamics of health, work and socioeconomic status.