1 Introduction

The economy is a complex system with nonlinear interactions and feedback loops. Early traces of this view date back, for example, to Schumpeter and Hayek, and to Simon. The complexity modeling paradigm has been strongly advocated since the 1980s by economists and multidisciplinary scientists from various fields, such as physics, computer science and biology, linked to the Santa Fe Institute.[1] More recently the complexity view has also drawn the attention of policy makers, who are faced with complex phenomena, irregular fluctuations and sudden, unpredictable market transitions. For example, the chairman of the FED, Ben Bernanke, noted that the 1000-point collapse of the Dow Jones Industrial Average on the afternoon of May 6, 2010, reflected the complexity of financial-market systems:

> The brief market plunge was just a small indicator of how complex and chaotic, in the formal sense, these systems have become. Our financial system is so complicated and so interactive – so many different markets in different countries and so many sets of rules. What happened in the stock market is just a little example of how things can cascade or how technology can interact with market panic.
>
> (interview Ben Bernanke, IHT, May 17, 2010).

The recent financial-economic crisis is a dramatic example of large movements, similar to critical transitions that are so characteristic for complex evolving systems. These large changes of global financial markets can hardly be viewed as a rational response to news about economic fundamentals and cannot be explained by traditional representative rational agent macro-finance models. A more compelling and intuitive explanation is that these extreme large movements have been triggered by bad economic news, and subsequently strongly amplified by an "irrational" overreaction of a heterogeneous population of boundedly rational, interacting agents. In a well-known speech the former president of the ECB, Jean-Claude Trichet, called for a new approach for policy makers to managing crises:

> First, we have to think about how to characterise the homo economicus at the heart of any model. The atomistic, optimising agents underlying existing models do not capture behaviour during a

[1] See, e.g., the early collections of papers in the Santa Fe conference proceedings Anderson et al. (1988) and Arthur et al. (1997a).

crisis period. We need to deal better with heterogeneity across agents and the interaction among those heterogeneous agents. We need to entertain alternative motivations for economic choices. Behavioural economics draws on psychology to explain decisions made in crisis circumstances. Agent-based modelling dispenses with the optimisation assumption and allows for more complex interactions between agents. Such approaches are worthy of our attention.

Second, we may need to consider a richer characterisation of expectation formation. Rational expectations theory has brought macroeconomic analysis a long way over the past four decades. But there is a clear need to re-examine this assumption. Very encouraging work is under way on new concepts, such as learning and rational inattention.

(Speech by Jean-Claude Trichet, ECB Central Banking Conference, Frankfurt, November 18, 2010)

This book presents some simple, stylized complexity models in economics. Our main focus will be an underlying *behavioral theory of heterogeneous expectations* of boundedly rational individual agents in a complex, adaptive economic environment. We will also discuss empirical validation, both at the micro and at the macro level, of a behavioral theory of heterogeneous expectations through financial time series data and laboratory experiments with human subjects. The need for an empirically grounded behavioral theory of expectations for economic dynamics has already been stressed by Herb Simon (1984, p. 54):

A very natural next step for economics is to maintain expectations in the strategic position they have come to occupy, but to build an empirically validated theory of how attention is in fact directed within a social system, and how expectations are, in fact, formed. Taking that next step requires that empirical work in economics take a new direction, the direction of micro-level investigation proposed by Behavioralism.

1.1 Economic dynamics, nonlinearity and complexity

Economic dynamics is concerned with modeling fluctuations in economic and financial variables, such as commodity prices, output growth, unemployment, interest rates, exchange rates and stock prices. Broadly speaking, there are two contrasting views concerning the main sources of economic fluctuations. According to the first, business cycles are mainly driven by "news" about economic fundamentals, that is, by random *exogenous* shocks to preferences, endowments, technology, firms' future earnings or dividends, etc. These random shocks typically act on an inherently stable (linearized) economic system. This view dates back to the 1930s, to Frisch, Slutsky and Tinbergen, who showed that a stable linear system subject to an irregular sequence of external, random shocks may produce fluctuations very similar to those observed in real business cycles.

The *linear, stable view* was criticized in the 1940s and 1950s, mainly because it did not offer an *economic* explanation of observed fluctuations, but rather attributed those fluctuations to external, non-economic forces. As an alternative, Goodwin, Hicks and Kaldor developed nonlinear, *endogenous* business cycle models, with the savings-investment mechanism as the main economic force generating business fluctuations. According to this *nonlinear view*, the economy may be intrinsically unstable and, even in the absence of external shocks, fluctuations in economic variables can arise. These early Keynesian nonlinear business cycle models, however, were criticized for at least three reasons. Firstly, the limit cycles generated by these models were much too regular to explain the sometimes highly irregular movements in economic and financial time series data. Secondly, the "laws of motion" were considered to be "ad hoc," since they had not been derived from micro foundations, i.e., from utility and profit maximization principles. A third important critique was that agents' behavior was considered as *irrational*, since their *expectations were systematically wrong* along the regular business cycles. Smart, rational traders would learn from experience to anticipate these cyclic movements and revise their expectations accordingly, and, so the story goes, this would cause the cycles to disappear.

These shortcomings triggered the rational expectations revolution in the 1960s and 1970s, inspired by the seminal papers of Muth (1961) and Lucas (1972a and b). New classical economists developed an alternative within the exogenous approach, the stochastic real business cycle (RBC) models, pioneered by Kydland and Prescott (1982). RBC models fit into the general equilibrium framework, characterized by utility-maximizing consumers, profit-maximizing firms, market clearing for all goods at all dates and all traders having rational expectations. More recently, New Keynesian Dynamic Stochastic General Equilibrium (DSGE) models have moved to the forefront of macroeconomic modeling and policy analysis (Clarida et al., 1999; Woodford 2003). Typically these DSGE models are log linearized and assume a representative rational agent framework. A representative, perfectly rational agent nicely fits into a linear view of a globally stable, and hence predictable, economy. By the late 1970s and early 1980s, the debate concerning the main source of business cycles seemed to have been settled in favor of the exogenous shock hypothesis, culminating in the currently dominating DSGE macro models for policy analysis.

1.1.1 The discovery of chaos

In mathematics and physics the view on modeling dynamic phenomena changed dramatically in the 1960s and 1970s due to the discovery of *deterministic chaos*. One of its pioneers, the MIT meteorologist Edward Lorenz (1963), discovered by computer simulations that a simple nonlinear system of three differential equations can generate highly irregular and seemingly unpredictable time series patterns.[2] Moreover, his stylized model of weather prediction was characterized by *sensitive dependence on initial conditions* (the "butterfly effect"): a small perturbation of the initial state leads to a completely different time path prediction in the medium or long run. In the 1970s, Ruelle and Takens (1971) presented a mathematical proof that a simple nonlinear system of three or four differential equations, without any external random disturbances, can indeed exhibit complicated, irregular long run dynamical behavior. They introduced

[2] See, e.g., Gleick (1987) for a stimulating historical overview of "chaos theory." It is interesting to note that one of the traditional Keynesian business cycle models from the 1950s, Hicks' classical nonlinear trade cycle model with ceilings and floors, can in fact generate irregular, chaotic time series. In particular, figures 9 and 10 in Hicks (1950, pp. 76–79), computed by hand at the time, are similar to the computer simulated chaotic series in Hommes (1995), so that in some sense Hicks was close to discovering chaos in his trade cycle model.

the notion of a *strange attractor* to describe irregular long run behavior in a nonlinear deterministic dynamical system. The discovery of deterministic chaos and strange attractors shattered the Laplacian deterministic view of perfect predictability and made scientists realize that, because initial states can only be measured with finite precision, long run prediction may be fundamentally impossible, even when the laws of motion are perfectly known.

In the 1970s, there was yet another important mathematical article with the illuminating title "Period three implies chaos" (Li and Yorke, 1975), which played a stimulating role and was particularly important for applications. Li and Yorke showed that for a large class of simple nonlinear difference equations in one single state variable, a simple sufficient "period three" condition already implies complicated, chaotic dynamical behavior. The best-known example concerns logistic population growth in biology, as described by May (1976). These and other simple mathematical examples together with the rapidly increasing availability of computers for numerical simulations led to an explosion of interest in nonlinear dynamics in mathematics, physics and other applied sciences.

The "chaos revolution" in the 1970s had its roots, however, much earlier, at the end of the nineteenth century in the famous French mathematician Henri Poincaré. In 1887 king Oskar II of Sweden promised a prize to the best essay concerning the question "Is our solar system stable?" In his prize-winning essay, Poincaré (1890) showed that the motion in a simple three-body system, a system of sun, earth and moon, need not be periodic, but may become highly irregular and unpredictable. In modern terminology he showed that chaotic motion is possible in a three-body system. Poincaré introduced the notion of a so-called *homoclinic point*, an intersection point between the stable and the unstable manifolds of an equilibrium steady state. His notion of homoclinic orbits turned out to be a key feature of complicated motion and strange attractors and may be seen as an early signature of chaos.

1.1.2 Economic applications of chaos

In the 1980s, inspired by "chaos theory" and within the tradition of endogenous business cycle modeling, economic theorists started looking for nonlinear, deterministic models generating erratic time series similar to the patterns observed in real business cycles. This search led to new, simple nonlinear business cycle models, within the Arrow–Debreu general equilibrium paradigm of optimizing behavior, perfectly competitive markets and rational expectations, generating chaotic business fluctuations (e.g., Benhabib and Day, 1982 and Grandmont, 1985; see, e.g., Lorenz, 1993 for an overview of nonlinear business cycle models and chaos). These model examples show that irregular, chaotic fluctuations can arise under the New Classical Economics paradigm in a perfectly rational representative agent framework. It turned out to be more difficult, however, to calibrate or estimate such chaotic business cycle models to real economic data.

Simultaneously, the search for nonlinearity and chaos in economics was undertaken from an empirical perspective. In physics and mathematics nonlinear methods to distinguish between truly random and deterministic chaotic time series had been developed. For example, correlation dimension tests and Lyapunov exponent tests had been developed by Takens (1981) and Grassberger and Procaccia (1983). When the correlation dimension of a time series is low, this suggests evidence for low-dimensional chaos. In economics, for example, Brock and Sayers (1988) found a correlation dimension of about 3 for macroeconomic data (postwar quarterly US unemployment rates), and Scheinkman and LeBaron (1989) a correlation dimension of about 6 for stock market data (weekly stock returns). A problem for applying these empirical methods, particularly relevant for economic data, is that they require very long time series and that they are extremely sensitive to noise. Furthermore, it turned out that time series generated by fitted stochastic alternative models, such as linear, near unit root autoregressive models for macro data or GARCH-models for stock returns, also generate low correlation dimensions of comparable size. Hence, from these empirical findings, one *cannot* conclude that there is evidence for low-dimensional, purely deterministic chaos in economic and financial data. Brock, Dechert, Scheinkman and LeBaron (1996) have developed a general test (the BDS test), based upon the notion of correlation dimension, to test for *nonlinearity* in a given time series; see Brock et al., (1991) for the basic theory, references and applications. The BDS test has become widely used, in economics but also in physics, and has high power against many nonlinear alternatives. From an empirical viewpoint, evidence for low-dimensional, purely deterministic chaos in economic and financial data is weak, but there is strong evidence for nonlinear dependence. At the same time, it seems fair to add that, because of the sensitivity to noise of these methods, the hypothesis of chaos buffeted with (small) dynamic noise has *not* been rejected either.[3] Nor has higher-dimensional chaos been rejected by these time series methods.

Empirical difficulties, both in calibrating new classical nonlinear endogenous business cycle models to economic data and in finding evidence for low-dimensional chaos in economic and financial time series, thus prevented a full embracement and appreciation of nonlinear dynamics in economics in the 1980s and early 1990s.

1.1.3 Expectations

The most important difference between economics and the natural sciences is perhaps the fact that decisions of economic agents today depend upon their *expectations* or *beliefs* about the future. To illustrate this difference, weather forecasts for tomorrow will not affect today's weather, but investors' predictions about future stock prices may affect financial market movements today. A classic example is the Dutch "tulip mania" in the seventeenth century, as described in Kindleberger (1996). The dreams and hopes of Dutch investors for excessive high returns on their investments in tulip bulbs may have exaggerated the explosion of the price of tulip bulbs by a factor of more than 20 at the beginning of 1636, and its crash back to its original level by the end of that year. Another more recent example is the "dot-com bubble," the rapid run up of stock

[3] See Hommes and Manzan (2006) for a brief recent discussion.

prices in financial markets worldwide in the late 1990s, and the subsequent crash. This rise in stock prices was triggered by good news about economic fundamentals, a new communication technology, the internet. An overoptimistic estimate of future growth of ICT industries seems to have contributed to and strongly reinforced the excessively rapid growth of stock prices in 1995–2000, leading to extreme overvaluation of stock markets worldwide, and their subsequent fall in 2000–2003. A more recent example is the 2008–2012 financial-economic crisis. It is hard to believe that the decline of worldwide financial markets in 2008 of more than 50% was completely driven by changes in economic fundamentals. Rather it seems that the large decline was strongly amplified by pessimistic expectations and market psychology. A similar observation applies to the 2011–2012 EU debt crisis. While the budget deficits of EU countries are partly caused by economic fundamentals, the sharp rise in the spread of, e.g., Italian and German bonds in 2011 seems to have been exaggerated by investors' pessimistic expectations. The predictions, expectations or beliefs of consumers, firms and investors about the future state of the economy are part of the "law of motion." The economy is a highly nonlinear *expectations feedback* system, and therefore a *theory of expectations* is a crucial part of any dynamic economic model or theory.

Since the introduction of rational expectations by Muth (1961) and its popularization in macroeconomics by Lucas (1972a and b) and others, the *rational expectations hypothesis* (REH) became the dominating expectations formation paradigm in economics. According to the REH all agents are rational and take as their subjective expectation of future variables the objective prediction by economic theory. In economic modeling practice, expectations are given as the mathematical conditional expectation given all available information. Rational agents do not make "systematic mistakes" and their expectations are, on average, correct. The REH provides an elegant "fixed-point" solution to an economic expectations feedback system by imposing that, on average, expectations and realizations coincide. In the absence of exogenous shocks, rational expectations implies that agents have perfect foresight and make no mistakes at all. This shortcut solution excludes all irrationality and market psychology from economic analysis, and instead postulates that expectations are in equilibrium and perfectly self-fulfilling.

The rational expectations revolution in economics took place *before* the discovery of chaos, at least before the time that the irregular behavior and complexity of nonlinear dynamics were widely known among economists. The fact that chaos can arise in simple nonlinear systems and its implications for limited predictability, however, shed important new light on the expectations hypothesis. In a simple (linear) stable economy with a unique steady state, predictability prevails and it seems natural that agents may have rational expectations, at least in the long run. A representative, perfectly rational agent model nicely fits into a linear view of a globally stable and predictable economy. But how can agents have rational expectations or perfect foresight in a *complex, nonlinear world*, when the true law of motion is unknown and prices and quantities move irregularly on a strange attractor exhibiting sensitivity to initial conditions? A boundedly rational world view with agents using simple forecasting strategies, which may not be perfect but are at least approximately right, seems more appropriate for a complex nonlinear environment. Indeed, already around 1900 Poincaré, one of the founding fathers of nonlinear dynamics, expressed his concerns about the implications of limited predictability in nonlinear systems for economics in a letter to Walras, one of the founders of mathematical economics:[4]

> You regard men as infinitely selfish and infinitely farsighted. The first hypothesis may perhaps be admitted in a first approximation, the second may call for some reservations.

1.1.4 Bounded rationality and adaptive learning

In economics in the 1950s, Herbert Simon emphasized that rationality requires extreme assumptions concerning agents' information gathering and computing abilities. Firstly, rational agents are typically assumed to have perfect information about economic fundamentals and perfect knowledge about underlying market equilibrium equations. This assumption seems unrealistically strong, especially since the "law of motion" of the economy depends on the expectations of *all other* agents. Secondly, even if such information and knowledge were available, typically in a nonlinear market equilibrium model it would be very hard, or even impossible, to derive the rational expectations forecast analytically, and it would require quite an effort to do it computationally. As an alternative, Simon strongly argued for *bounded rationality*, with limited computing capabilities and agents using simple rules of thumb instead of perfectly optimal decision rules, as a more accurate and more realistic description of human behavior. Simon's reasoning lost against the rational expectations revolution in the 1970s, but in the last two decades similar reasoning has caused an explosion of interest in bounded rationality. Modeling a world with boundedly rational agents, who adapt their behavior and learn from past experiences over time, leads to a complex and highly nonlinear dynamic system.

A common assumption underlying models of bounded rationality is that agents do *not* know the actual "law of motion" of the economy, but instead base their forecasts upon time series *observations*. They behave like economic statisticians, forming expectations based upon time series observations, using a simple statistical model for their perceived law of motion. *Adaptive learning*, sometimes also referred to as *statistical learning*, means that agents adapt their beliefs over time by updating the parameters of their perceived law of motion according to some learning scheme (e.g., recursive ordinary least squares), as additional observations become available. The adaptive learning approach has been used extensively in macroeconomics. Sargent (1993) gives an early overview of learning in macroeconomics, while Evans and Honkapohja (2001) contains a more recent extensive and detailed treatment; see also Conlisk (1996) for a stimulating discussion of bounded rationality. An important issue that has received much attention in the literature is the *stability* of rational expectations equilibria under adaptive learning. If adaptive learning enforces convergence to a rational

[4] Front quotation in Grandmont (1998) and Ingrao and Israel (1990), from letter of October 1, 1901 of Henri Poincaré to Léon Walras.

expectations equilibrium, the REH would be more plausible as a (long run) description of the economy, since the underlying informational assumptions could be considerably relaxed. However, many examples have been found where adaptive learning does *not* converge to rational expectations, but rather settles down to some kind of "learning equilibrium" displaying endogenous, sometimes even chaotic, fluctuations and excess volatility (e.g., Bullard, 1994, Grandmont, 1998, Hommes and Sorger, 1998 and Schönhofer, 1999).

1.1.5 Heterogeneity in complex adaptive systems

The representative agent model has played a dominant role in modern economics for quite some time. Most rational expectations models assume a single, *representative agent*, representing average consumer, average firm or average investment behavior. An important motivation for the rational agent model dates back to the 1950s, to Milton Friedman (1953) who argued that non-rational agents will be driven out of the market by rational agents, who will trade against them and earn higher profits. In recent years however, this view has been challenged and heterogeneous agent models are becoming increasingly popular in finance and in macroeconomics. Kirman (1992, 2010), for example, provides an illuminating critique on representative rational agent modeling.

Bounded rationality and learning in a complex environment naturally fit with *heterogeneous expectations*, with the economy viewed as a complex evolving system composed of many different, boundedly rational, interacting agents, using different decision strategies, heuristics and forecasting rules. Heterogeneous strategies compete against each other and an evolutionary selection mechanism, e.g., through genetic algorithm learning, disciplines the class of strategies being used by individual agents. In such a complex system, expectations and realizations coevolve over time. The work at the Santa Fe Institute has played a stimulating role and the collections of papers in Anderson et al. (1988) and Arthur et al. (1997a) of Santa Fe conferences provide early examples of the complexity modeling approach in economics.

The complexity view in economics is naturally linked to *agent-based computational economics (ACE)*, characterized by agent-based computer simulation models with many heterogeneous agents; see, e.g., the recent *Handbook* of Tesfatsion and Judd (2006) for surveys of the state of the art of ACE. An advantage of agent-based models is that one can use a "bottom up" approach and build "realistic" models from micro interactions to simulate and mimic macro phenomena. However, in agent-based models with many interacting agents, the "wilderness of bounded rationality" is enormous, there are infinitely many possibilities for individual decision rules and, for a given model, it is often hard to pin down what exactly causes certain stylized facts at the macro level in agent-based micro simulations.

1.1.6 Behavioral rationality and heterogeneous expectations

A good feature of the rational expectations hypothesis (REH) is that it imposes strong discipline on agents' forecasting rules and minimizes the number of free parameters in dynamic economic models. In contrast, the "wilderness of bounded rationality" in agent-based models leaves many degrees of freedom in economic modeling, and it seems far from clear which rules are the most reasonable out of an infinite class of potential behavioral rules. Stated differently in a popular phrase: "*There is only one way (or perhaps a few ways) you can be right, but there are many ways you can be wrong*." To avoid "ad hoccery," a successful bounded rationality research program needs to discipline the class of expectations and decision rules. The REH assumes *perfect consistency* between beliefs and realizations. For a successful bounded rationality research program a *reasonable* and *plausible* form of consistency between beliefs and realizations is necessary.

This book focusses on "simple" complexity models, where only a few different types of heterogeneous agents interact. Our main focus is on the role of *behavioral rationality* and *heterogeneous expectations* within stylized complexity models. Our consistency story of bounded rationality contains three important elements: (i) agents use simple decision rules, with an intuitive behavioral interpretation; (ii) agents switch between different decision rules based on evolutionary selection and learning; and (iii) the models of bounded rationality are empirically validated, at both the micro and the macro levels.

Behavioral rationality emphasizes the use of simple, intuitive decision rules – *heuristics* – with a plausible behavioral interpretation. These heuristics are not perfect and need not be optimal, but within an environment that is too complex to fully understand individual agents look for simple decision rules that perform reasonably well to a first-order approximation; for a similar approach and extensive discussions, see, e.g., the collection of papers on smart heuristics and the adaptive toolbox in Gigerenzer et al. (1999) and Gigerenzer and Selten (2001).

Two forms of learning further discipline the class of decision heuristics. First, we use the heterogeneous strategy switching framework of Brock and Hommes (1997a, 1998) of *endogenous evolutionary selection* or *reinforcement learning* among heterogeneous decision or forecasting rules. The main idea here is that agents tend to switch to rules that have performed better, according to some suitable economic performance measure such as realized profits or forecasting accuracy, in the recent past. The forecasting rules may be divided into different classes, with different degrees of rationality, ranging from simple behavioral rules such as naive or adaptive expectations, trend extrapolating rules or contrarian rules, to more sophisticated rules, such as statistical learning rules, fundamental market analysis or even rational expectations. These more sophisticated rules may be more costly – due to information-gathering costs – than alternative forecasting heuristics. The second form of learning takes place within each class of forecasting heuristics, with some parameters changing over time following some adaptive learning process. For example, within the class of trend-following heuristics, the trend coefficient or the anchor from which the trend is extrapolated may change over time and depend upon market realizations. This type of learning also has a behavioral interpretation and can be linked to the *anchor and adjustment* heuristics used in psychology (e.g., Tversky and Kahneman, 1974, Kahneman, 2003).

To discipline behavioral models and boundedly rational decision heuristics, *empirical validation both at the micro and at the macro level* is important. Laboratory experiments with human subjects, in particular experimental macroeconomics, plays a key role here, with the experimenter having full control over the type of micro interactions and the macroeconomic fundamentals. Duffy (2006, 2008a and b) provides a stimulating overview of experimental macroeconomics; the learning-to-forecast experiments surveyed in Hommes (2011) can be used to study the interactions of individual heterogeneous expectations and their aggregate effect in the laboratory.

Behavioral rationality and heterogeneous expectations naturally lead to highly *nonlinear* dynamical systems, because the fractions attached to the different rules are changing over time. Evolutionary selection of heterogeneous expectations sometimes enforces convergence to a rational expectations equilibrium. More often, however, the evolutionary system may be unstable and exhibit complicated, perpetual fluctuations, with several simple forecasting heuristics surviving evolutionary selection. In particular, we will see that when some rules act as "far from the steady state stabilizing forces" and other rules act as "close to the steady state destabilizing forces," evolutionary selection of expectations rules may lead to Poincaré's classical notion of a homoclinic orbit and may be seen as a signature of potential instability and chaos in a complex adaptive system with behaviorally rational agents.

An economy with heterogeneous, behaviorally rational agents is a highly nonlinear complex evolving system. The tools of nonlinear dynamics and complex systems are crucial to understand the behavior of markets with heterogeneous boundedly rational agents and to provide the insights to managing complex adaptive systems. This book introduces the most important analytical and computational tools in simple nonlinear complexity models and applies them to study economic dynamics with heterogeneous boundedly rational agents and learning. The remainder of this introduction gives the reader a quick overview of the contents of the book, discussing important concepts such as behavioral rationality and heterogeneous expectations in some simple examples of complex economic systems and briefly discussing their empirical validation with time series data and laboratory experiments with human subjects.

1.2 Adaptive expectations in a nonlinear economy

The simplest economic example nicely illustrating the role of expectations feedback is the "hog cycle" or *cobweb model*. Traditionally it has played a prominent role as a didactic benchmark model and has been used, for example, in the seminal article of Muth (1961) introducing rational expectations. Here we focus on the role of simple expectation rules, in particular adaptive expectations, in a *nonlinear* cobweb model.

The model is partial equilibrium and describes an independent competitive market for a non-storable consumption good, such as corn or hogs. Production takes a fixed unit of time, and suppliers therefore have to base their production decision upon their anticipation or expectation p_t^e of the market equilibrium price p_t that will prevail.

Demand, supply and market clearing are described by

$$D(p_t) = a - dp_t + \epsilon_t, \qquad a, d > 0, \tag{1.1}$$

$$S_\lambda(p_t^e) = c + \arctan(\lambda(p_t^e - \bar{p})), \qquad c, \lambda > 0, \tag{1.2}$$

$$D(p_t) = S_\lambda(p_t^e). \tag{1.3}$$

Demand D in (1.1) is a linearly decreasing function in the market price p_t, with slope $-d$, and ϵ_t is a random term representing (small) exogenous demand shocks. The supply curve S_λ in (1.2) is *nonlinear*, increasing and S-shaped, with the parameter λ tuning the nonlinearity of the supply curve and $\bar{p}$ denoting the inflection point of the nonlinear supply curve, where marginal supply assumes its maximum. It should be noted that such a nonlinear, increasing supply curve can be derived from producer's profit maximization with a convex cost function. Finally, (1.3) dictates that the price adjusts such that the market clears in each period.

To close the model we have to specify how producers form price expectations. The simplest case, studied in the 1930s, e.g., by Ezekiel (1938), assumes that producers have *naive expectations*, that is, their prediction equals the last observed price, $p_t^e = p_{t-1}$. Under naive expectations, when demand is decreasing and supply is increasing and bounded, there are only two possibilities concerning the price dynamics, depending on the ratio of marginal supply over marginal demand at the steady state price p^* (i.e., the price where demand and supply intersect):

- if $|S'(p^*)/D'(p^*)| < 1$, then the steady state is (locally) *stable* and prices converge to p^*;
- if $|S'(p^*)/D'(p^*)| > 1$, then the steady state is *unstable* and prices diverge from p^* and converge to a (stable) 2-cycle.

The unstable case is illustrated in Figure 1.1 (top panel). Due to the nonlinearity of the supply curve, in the unstable case prices converge to a stable 2-cycle, with up and down "hog cycle" price fluctuations.

In the 1960s, simple mechanical expectation rules such as naive expectations became heavily criticized as being "irrational," since these forecasts are *"systematically wrong"*. Rational farmers would discover the regular, cyclic pattern in prices, learn from their systematic mistakes, change expectations accordingly and the hog cycle would disappear, so the argument goes. Similar considerations lead Muth (1961) to the introduction of *rational expectations*, where the expected price coincides with the price predicted by economic theory. In a rational expectations equilibrium, agents use economic theory, and compute their expectations as the conditional mathematical expectation derived from the market equilibrium equations. In the cobweb model, taking conditional mathematical expectations on the left- and right-hand sides of the market equilibrium equation (1.3), one derives that the rational expectations forecast is exactly given by the steady state price p^*. In a rational expectations equilibrium, expectations are self-fulfilling and agents make no systematic forecasting errors. In a cobweb world without uncertainty (i.e., $\epsilon_t \equiv 0$), the forecast $p_t^e = p^*$ will always

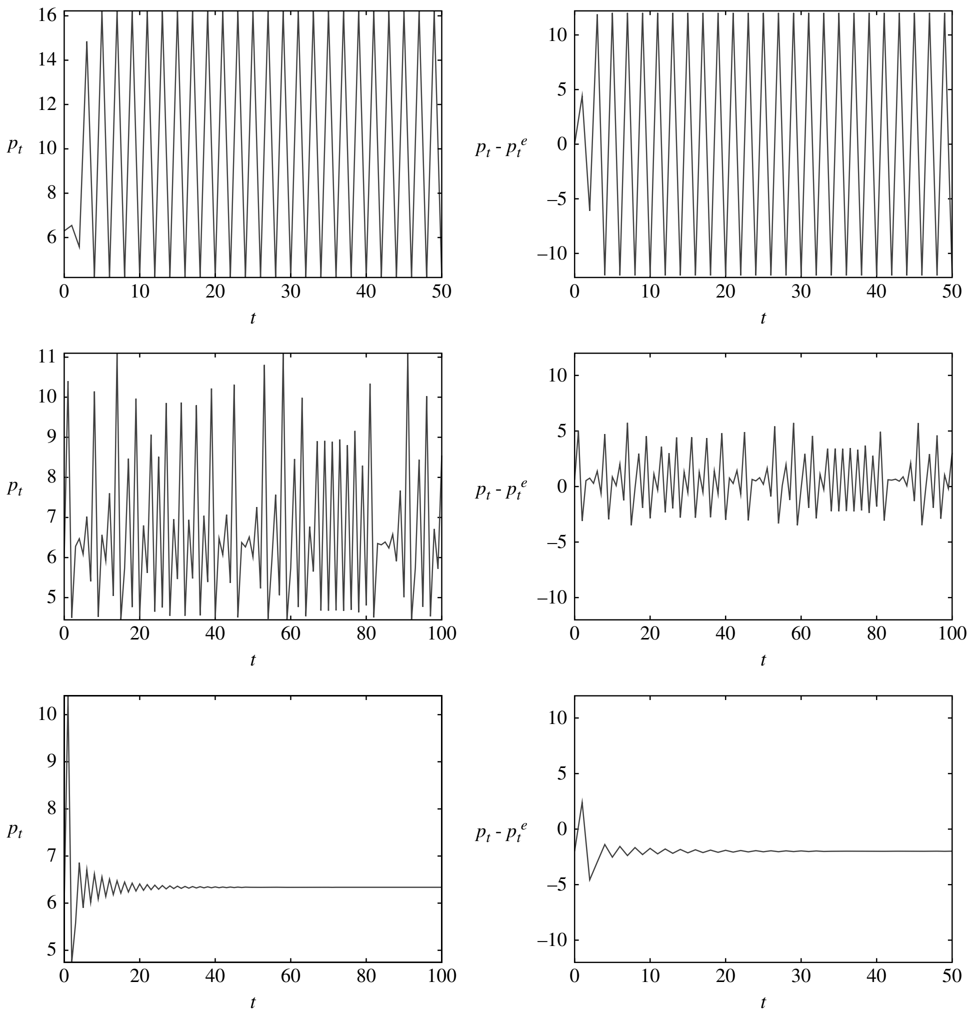

Figure 1.1. Time series of prices (left panel) and forecasting errors (right panel) in the nonlinear cobweb model with adaptive expectations for different values of expectations weight factor w: stable 2-cycle for $w = 1$ (top panel), chaotic price series for $w = 0.5$ (middle panel), and stable steady state for $w = 0.3$. Other parameter values are $\lambda = 4.8$, $c = 1.5$, $a = 4.1$, $d = 0.25$, $\bar{p} = 6$ and initial state $p_0 = 6$.

be exactly right and rational expectations coincides with perfect foresight. In a noisy cobweb world with uncertainty, the rational expectations forecast $p_t^e = p^*$ will be correct on average and agents make no systematic mistakes, since forecasting errors are proportional to the exogenous random demand shocks ϵ_t.

Now consider the case of adaptive expectations, discussed by Nerlove (1958) (but only in the case of linear demand and supply). *Adaptive expectations* are given by

$$p_t^e = (1-w)p_{t-1}^e + w p_{t-1}, \qquad 0 \le w \le 1, \tag{1.4}$$

where w is the expectations weight factor. The expected price is a weighted average of yesterday's expected and realized prices, or equivalently, the expected price is adapted by a factor w in the direction of the most recent realization. The weight factor w determines the magnitude of the "error-correction" in each period. In fact, adaptive expectations means that today's expected price is a weighted average, with geometrically declining weights, of all past prices. In the cobweb model with *linear* demand and supply curves, naive and adaptive expectations lead to the familiar "hog" cycle," characterized by up and down oscillations between a high and a low price level. In the case of *nonlinear* (but monotonic) demand and/or supply curves, things become more complicated, however. A simple computation, using (1.1–1.3) and (1.4), shows that the dynamics of expected prices becomes

$$p_t^e = f_{w,a,d,\lambda}(p_{t-1}^e) = (1-w)p_{t-1}^e + w\frac{a + \epsilon_t - c - \arctan(\lambda(p_{t-1}^e - \bar{p}))}{d}. \quad (1.5)$$

Dynamics of (expected) prices in the cobweb model with adaptive expectations is thus given by a one-dimensional (1-D) system $x_t = f_{w,a,b,\lambda}(x_{t-1})$ with four model parameters. *What can be said about the price–quantity dynamics in this nonlinear dynamic model, and how does it depend on the model parameters?*

Figure 1.1 illustrates time series of prices and corresponding forecasting errors, for different values of the expectations weight factor w. Under naive expectations ($w = 1$; top panel) prices converge to a stable 2-cycle and expectational errors are large and systematic. When agents are cautious in adapting their expectations, i.e., when the expectations weight factor is small ($w = 0.3$, bottom panel), prices converge to the RE stable steady state and forecasting errors vanish in the long run. For intermediate values of the expectations weight factor ($w = 0.5$; middle panel) prices as well as forecasting errors are chaotic. These forecasting errors are considerably smaller than under naive expectations and, because they are *chaotic*, they are much more irregular and it is more difficult for producers to learn from their errors. The degree of consistency between realizations and adaptive expectations in the chaotic case is much higher than in the 2-cycle case of naive expectations, and it may therefore be a more reasonable, boundedly rational description of market behavior.

A powerful tool to investigate how the dynamical behavior of a nonlinear model depends on a single parameter is a *bifurcation diagram*. A bifurcation is a qualitative change in the dynamics as a model parameter changes. Critical transitions in complex systems arise because of some bifurcation, some qualitative change in the dynamics of the system. A bifurcation diagram shows the long run dynamical behavior as a function of a model parameter. Figure 1.2 shows a bifurcation diagram of the cobweb model with adaptive expectations with respect to the expectations weight factor w, illustrating the long run dynamics (100 iterations) after omitting a transient phase of 100 iterations.[5] For small values of w, $0 \leq w \leq 0.31$, prices converge to a stable steady

[5] Most figures in this book have been made using the E&F Chaos software for simulation of nonlinear systems, as described in Diks et al. (2008). The software is flexible and the user can, for example, easily include her own favorite nonlinear dynamic system. The software is freely downloadable at www.fee.uva.nl/cendef.

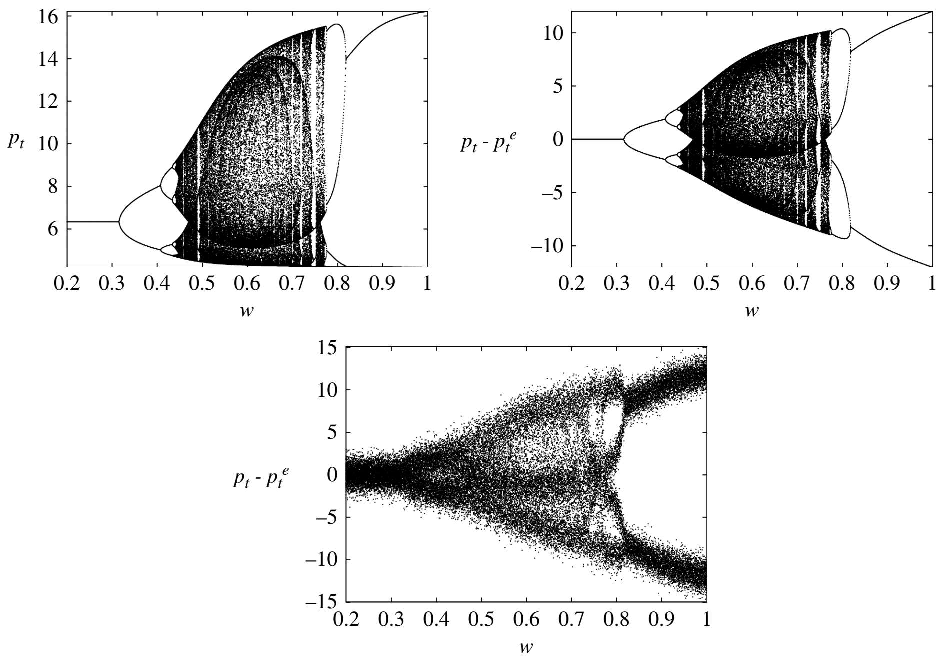

Figure 1.2. Bifurcation diagram for the expectations weight factor w, $0.2 \leq w \leq 1$, showing long run behavior of prices (top left panel) and forecasting errors (top right panel). The bottom panel shows a bifurcation diagram in the presence of small noise. Other parameter values are $\lambda = 4.8$, $c = 1.5$, $a = 4.1$, $d = 0.25$, $\bar{p} = 6$ and initial state $p_0 = 6$.

state, while for high values of w, $0.82 < w \leq 1$ (close to naive expectations) prices converge to a stable 2-cycle with large amplitude. Along the 2-cycle agents make systematic forecasting errors. For intermediate w-values however, say for $0.45 < w < 0.77$, chaotic price oscillations of moderate amplitude arise. In particular, the chaotic price fluctuations for $w = 0.5$ have been illustrated already in Figure 1.1. Figure 1.2 (bottom panel) also shows a simulation in the presence of small noise. The fine structure of the bifurcation diagram disappears, but the initial period-doubling bifurcations remain visible.

This example illustrates that a simple adaptive expectations rule in a noisy, nonlinear environment may be a reasonable forecasting strategy, which may be correct on average and which may not be easy to improve upon in a boundedly rational world.

1.3 Rational versus naive expectations

Heterogeneity of expectations among traders introduces an important *nonlinearity* into the market dynamics and is a potential source of market instability and erratic, chaotic price fluctuations. To illustrate this point by an example, we briefly discuss the cobweb model with heterogeneous expectations, rational versus naive producers, as introduced in Brock and Hommes (1997a); see Chapter 5 for a more detailed treatment. Agents can either buy a rational expectations forecast at positive information-gathering costs or freely obtain a simple, naive forecasting rule. This relates to Herbert Simon's idea to take deliberation and information-gathering costs into account in behavioral modeling. Information costs for rational expectations represent the idea that sophisticated prediction of prices, for example based upon detailed market analysis of economic fundamentals, is more costly than a simple prediction scheme, such as naive expectations or extrapolation of a price trend. The fractions of the two types change over time depending on how well both strategies performed in the recent past. Agents are boundedly rational in the sense that they tend to switch to the strategy that has performed better in the recent past.

To be more concrete, suppose that in the cobweb economy there are two types of producers, with different price expectations. At the moment of their production decision, producers can either buy the rational expectations price forecast at positive information costs C, or freely obtain the naive forecast. The two forecasting rules are

$$p^e_{1,t} = p_t, \tag{1.6}$$

$$p^e_{2,t} = p_{t-1}. \tag{1.7}$$

Rational agents have perfect foresight, while naive agents use the last observation as their forecast. In a cobweb world with rational versus naive expectations, the market equilibrium price is determined by demand and aggregate supply of both groups, i.e.,

$$D(p_t) = n_{1,t} S(p_t) + n_{2,t} S(p_{t-1}), \tag{1.8}$$

where $n_{1,t}$ and $n_{2,t}$ represent the fractions of producers holding rational respectively naive expectations. Notice that rational agents have perfect knowledge about the market equilibrium equation (1.8). Hence, rational traders not only have exact knowledge about prices and their own beliefs, but in a heterogeneous world they must also have perfect knowledge about expectations or beliefs of *all other* traders. We take a linear demand curve as before and, to keep the model as simple as possible, also a linear supply curve $S(p^e) = sp^e$. Market clearing in this two-type cobweb economy then yields

$$a - dp_t = n_{1,t} s p_t + n_{2,t} s p_{t-1}. \tag{1.9}$$

The second part of the model describes how the fractions of rational and naive producers are updated over time. The basic idea is that fractions are updated according to evolutionary fitness. Producers are boundedly rational in the sense that most of them will choose the forecasting rule which has highest fitness as measured by an economic performance measure, such as realized profits. To simplify the discussion, we focus here on the case where predictor selection is based upon last period's squared forecasting error plus the costs for obtaining that forecasting rule.[6] After the market equilibrium price

[6] As we will see in Chapter 5, Section 5.2 for linear demand and supply curves the fitness measure minus squared prediction error is, up to a constant factor, identical to realized profits.

has been revealed by (1.9), the new updated fractions of rational and naive producers will be given by a discrete choice or logit model:

$$n_{1,t+1} = \frac{e^{-\beta C}}{e^{-\beta C} + e^{-\beta(p_t - p_{t-1})^2}}, \tag{1.10}$$

$$n_{2,t+1} = \frac{e^{-\beta(p_t - p_{t-1})^2}}{e^{-\beta C} + e^{-\beta(p_t - p_{t-1})^2}}. \tag{1.11}$$

Note that these fractions add up to one. The key feature of the evolutionary selection or reinforcement learning scheme (1.10–1.11) is that the rule that performs better will attract more followers. More precisely, as long as the squared forecasting error $(p_t - p_{t-1})^2$ from naive expectations is smaller than the per period costs C for rational expectations, the majority of producers will "free ride" and not bother to buy the rational expectations forecast. But as soon as squared prediction errors for naive expectations become larger than the per period information-gathering costs for rational expectations, most producers will switch prediction strategy and buy the rational expectations forecast. The parameter β is called the *intensity of choice*, and it measures how fast the mass of traders will switch to the optimal prediction strategy. In the special case $\beta = 0$, both fractions will be constant and equal, and producers never switch strategy. In the other extreme case, $\beta = +\infty$, in each period *all* producers will use the same, optimal strategy. We call this latter case the *neoclassical limit*, since it represents the highest degree of rationality with respect to strategy selection based upon past performance in a heterogeneous world.

Now suppose that the market is *unstable* under naive expectations, that is, as long as all producers are naive, prices will diverge from the steady state price p^*. This situation is quite common and arises when the sensitivity of production decisions to expected price changes is larger than the sensitivity of consumers to price changes. The evolutionary dynamics exhibits a *rational route to randomness*, that is, a bifurcation route to strange attractors occurs, when the intensity of choice to switch to optimal forecasting strategies becomes larger. Figure 1.3 illustrates chaotic time series of prices and fractions as well as a strange attractor in the phase space.

The economic intuition behind the complicated evolutionary dynamics is simple. Suppose that we start from a situation where prices are close to the steady state p^* and almost all producers are naive. With prices close to the steady state, forecasting errors of naive expectations will be small, and therefore most producers will remain naive. Prices start fluctuating and will diverge from the steady state, so that the forecasting errors from naive expectations will increase over time. At some point, these forecasting errors will become larger than the costs for rational expectations. If the intensity of choice to switch strategies is high, then most producers will switch to rational expectations, causing prices to return close to the steady state. But with prices close to the steady state, it makes no sense to buy a rational expectations forecast, and most producers will become naive again. Hence, boundedly rational switching between forecasting

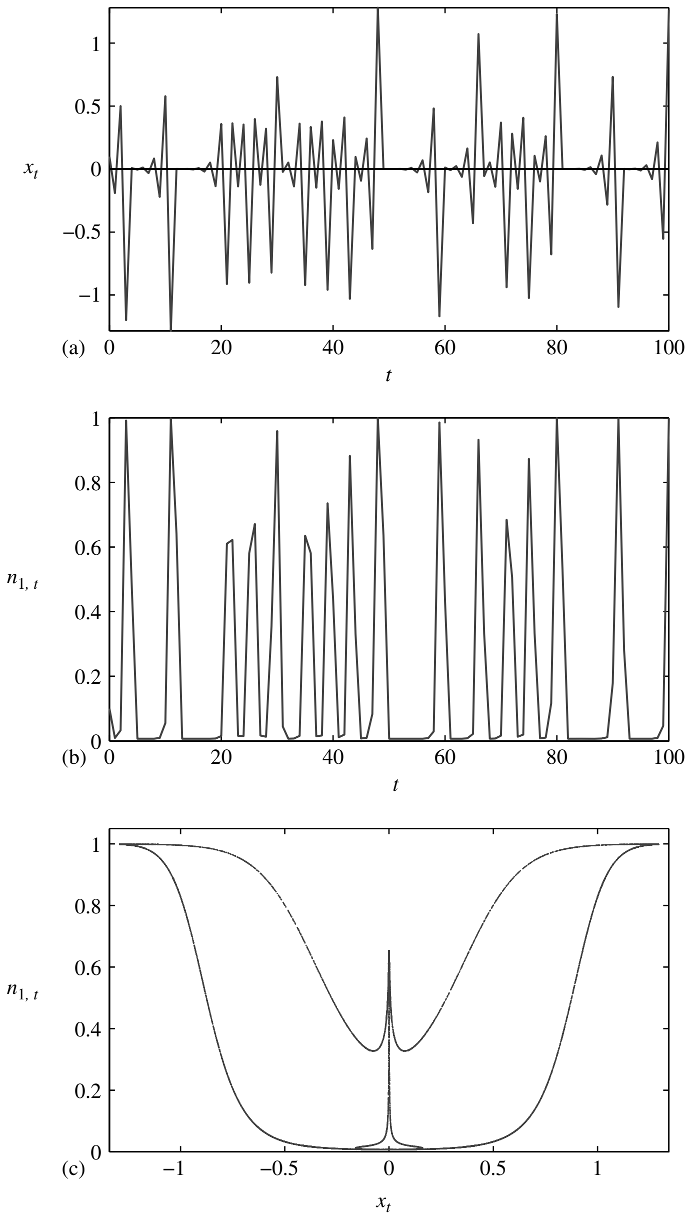

Figure 1.3. Chaotic time series of price deviations $x_t = p_t - p^*$ from steady state (top) and fractions $n_{1,t}$ of rational agents (middle) and corresponding strange attractor in the (x, n_1)-phase space (bottom). The market switches irregularly between an unstable phase of "free riding" with naive expectations dominating and a stable phase of costly rational expectations. Parameters are $\beta = 5$, $a = 0$, $d = 0.5$, $s = 1.35$ and $C = 1$.

strategies leads to an irregular switching between “cheap, destabilizing free riding” and “costly, sophisticated stabilizing prediction.”

Heterogeneous expectations in a simple linear cobweb economy lead to a natural *nonlinearity*, because the time-varying fractions of the different trader types appear as multiplicative factors in the market equilibrium equation (1.8). The economic evolutionary interaction between a “close to the steady state destabilizing force” when most agents adopt the cheap, simple strategy, and a “far from the steady state stabilizing force” when most agents switch to the costly, sophisticated strategy, is in fact closely related to Poincaré’s notion of a homoclinic orbit, which may be seen as a signature of potential instability and chaos in an evolutionary system with boundedly rational agents. Indeed, in Chapter 5, Section 5.2 we will see that for a high intensity of choice, the nonlinear evolutionary adaptive system is close to having a *homoclinic orbit*, Poincaré’s classical notion nowadays known to be a key feature of chaotic systems. The nonlinear adaptive evolutionary system describing strategy selection of a population of boundedly rational agents thus incorporates a simple economic mechanism leading to instability and chaos. In particular, a rational choice between cheap free riding and costly sophisticated prediction may lead to highly erratic equilibrium price fluctuations.

1.4 Adaptive learning

For commonly used simple expectations rules, such as naive or adaptive expectations, the parameters of the rule are fixed. *Adaptive learning*, sometimes also called *statistical learning*, refers to the more flexible situation with time-varying parameters, where agents try to learn the parameters of their forecasting rule as new observations become available over time. As a simple example, suppose agents use a linear AR(1) forecasting rule, of the form

$$p_t^e = \alpha + \rho(p_{t-1} - \alpha), \tag{1.12}$$

with two parameters α and ρ. This linear rule has a simple *behavioral* interpretation. The parameter α represents agents’ belief about the long run average of prices, while ρ represents the belief about the first-order autocorrelation coefficient, that is, the persistence (or anti-persistence) of the price series. When ρ is positive, agents believe that if the last observed price is above average, the next price will also be above average. On the other hand, when ρ is negative, agents are *contrarians*, that is, they believe that if the last observed price is above (below) average, the next price will be below (above) average. But what are the “true” or “optimal” parameters α and ρ of the linear rule in a complex market? In general, agents do *not* know the “true” parameters of their perceived law of motion, but they may try to learn the optimal parameters as additional observations become available. A more flexible forecasting rule with time-varying parameters is

$$p_t^e = \alpha_{t-1} + \rho_{t-1}(p_{t-1} - \alpha_{t-1}). \tag{1.13}$$

A simple example of an adaptive learning rule is given by

$$\alpha_t = \frac{1}{t+1}\sum_{i=0}^{t} p_i\,, \qquad t \geq 1 \tag{1.14}$$

$$\rho_t = \frac{\sum_{i=0}^{t-1}(p_i - \alpha_t)(p_{i+1} - \alpha_t)}{\sum_{i=0}^{t}(p_i - \alpha_t)^2}, \qquad t \geq 1, \tag{1.15}$$

where α_t is the *sample average* and ρ_t is the first-order *sample autocorrelation coefficient*. We refer to this adaptive learning schema as *sample autocorrelation learning* (SAC learning)[7]; see Chapter 4, Subsection 4.7.2 for a more detailed discussion. Here, we emphasize the behavioral interpretation of SAC learning. Agents try to learn the long run average α_t and the first-order autocorrelation or the "degree of persistence" of their linear forecasting rule. Hence, in a complex, nonlinear environment, agents try to match the first two moments, the long run average and the first-order autocovariance, to observed time series data.

The price dynamics in the cobweb model with linear demand and supply and SAC learning is given by

$$p_t = \frac{a - sp_t^e}{d}, \tag{1.16}$$

with the expected price p_t^e given by SAC learning (1.13), (1.14) and (1.15). The model with learning is a *nonlinear* system. When demand and supply are monotonic, however, that is, demand is decreasing and supply is increasing, the system has nice properties and always converges to the RE steady state, as illustrated in Figure 1.4.

But are all individual agents sophisticated enough to use such a statistical adaptive learning rule? Stated differently, in an unknown complex environment will individual agents coordinate on a simple adaptive learning procedure to enforce convergence of aggregate price behavior to the rational expectations benchmark?

1.4.1 Cobweb learning-to-forecast experiments

Hommes et al. (2007) ran *learning-to-forecast* laboratory experiments with human subjects to address this question; see Chapter 8 for a much more detailed discussion of learning-to-forecast experiments. Participants in the experiments were asked to predict next periods' market price of an unspecified good. The realized price p_t in the experiment was determined by the (unknown) cobweb market equilibrium equation

$$D(p_t) = \frac{1}{K}\sum_{i=1}^{K} S(p_{i,t}^e), \tag{1.17}$$

[7] SAC learning has been introduced in Hommes and Sorger (1998) and is closely related to recursive ordinary least squares (OLS)learning, which is used extensively in the literature on adaptive learning, see Evans and Honkapohja (2001).

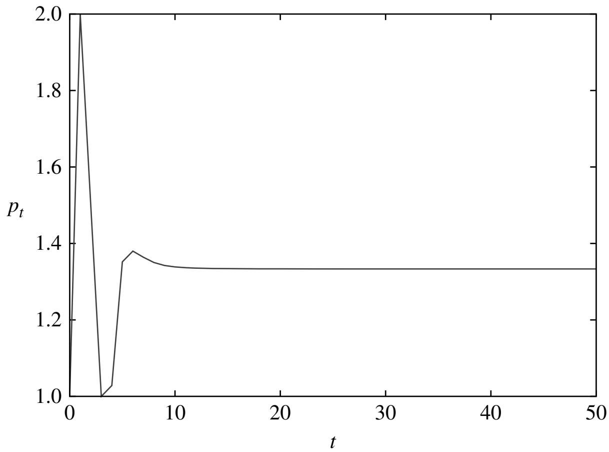

Figure 1.4. Price series under sample autocorrelation (SAC) learning converging to the unique rational expectations equilibrium. Parameters are $a = 2$, $d = 0.5$, $s = 2$ and $p_0 = 1$.

where $D(p_t)$ is the demand for the good at price p_t, K is the size of the group, $p^e_{i,t}$ is the price forecast by participant i and $S(p^e_{i,t})$ is the supply of producer i, derived from profit maximization given the forecast by participant i. Demand and supply curves D and S were fixed during the experiments (except for small random shocks to the demand curve) and unknown to the participants. We focus on the group experiments with $K = 6$, as in Hommes et al. (2007).

The main question in these experiments was whether agents can learn and coordinate on the unique REE, in a world where consumers and producers act as if they were maximizing utility and profits, but where they do *not* know underlying market equilibrium equations and only observe time series of realized market prices and their own forecasts. Our choice for a nonlinear, S-shaped supply curve enables us to investigate whether agents can avoid systematic forecasting errors, as would, for example, occur along a 2-cycle under naive expectations, or can even learn a REE steady state in a nonlinear cobweb environment.

In their experiment, Hommes et al. (2007) considered a stable and an unstable treatment, which only differ in the parameter λ tuning the nonlinearity of the supply curve as in (1.2). In the stable treatment, if all subjects use naive expectations, prices converge to the RE steady state. In contrast, in the unstable treatment, if all subjects use naive expectations, prices diverge from the RE steady state and converge to the stable 2-cycle, with large and systematic forecasting errors.

Figure 1.5 shows time series of the realized prices in two typical group experiments, one stable and one unstable treatment, that only differ in the magnitude of the parameter λ tuning the nonlinearity of the supply curve. For both treatments, the sample mean of realized prices is very close to the (unknown) RE price. Moreover, in the stable treatment, the sample variance is close to the variance (0.25) of the RE

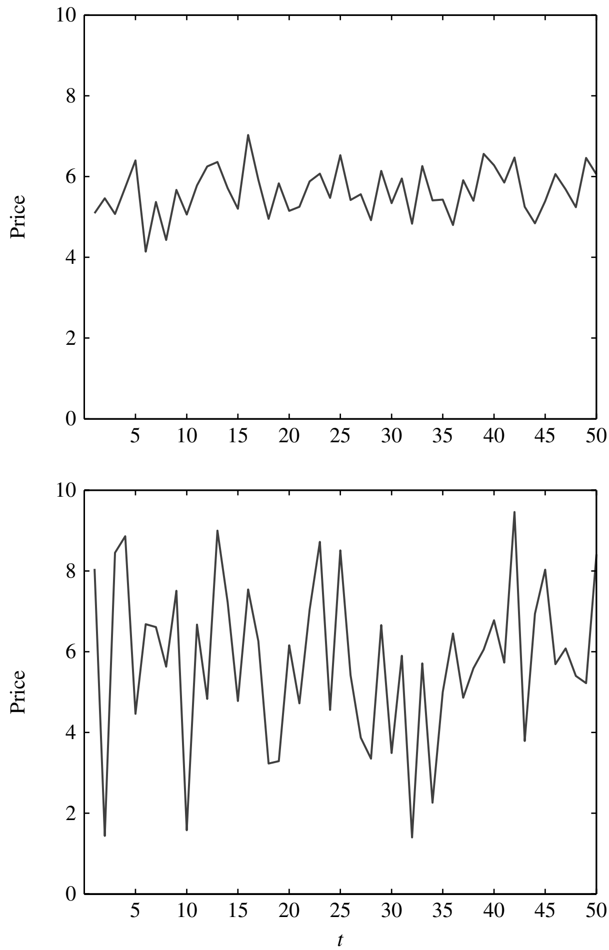

Figure 1.5. Realized market prices in two different cobweb group experiments. In the stable treatment (upper panel) the price quickly converges to the RE price with small random fluctuations, whereas in the unstable treatment (lower panel)) prices do not converge and exhibit excess volatility, with strongly fluctuating prices around the RE price.

benchmark. In contrast, in the unstable treatment the sample variance is significantly higher than the variance (0.25) of the RE benchmark, so that the unstable treatment exhibits excess volatility. Hommes et al. (2007) also look at autocorrelations in realized market prices, and find that there is no statistically significant autocorrelations in realized market prices, for both the stable and the unstable treatments. Apparently, the heterogeneous interactions of individual forecasting rules have washed out all linear predictable structure in realized aggregate market prices.

Hence, in a stable cobweb environment with unknown demand and supply curves, rational expectations may be a reasonable description of long run aggregate price behavior. In an unstable cobweb environment, however, full coordination on rational expectations does not arise and the market exhibits excess volatility. Homogeneous expectation models are – to our best knowledge – unable to explain all laboratory experiments simultaneously, and therefore heterogeneity is a key feature in explaining experiments across different treatments.

1.5 Behavioral rationality and heterogeneous expectations

A theory of bounded rationality and learning must be based on some reasonable consistency between expectations and realizations. Broadly speaking, we have discussed two stories of learning. One story of adaptive learning, where all agents use the same simple rule, the perceived law of motion, and try to learn the optimal parameters of the forecasting heuristic. According to this view, a representative agent optimizes the parameters of his forecasting rule within a given (simple) class of rules. The second story assumes that there are different classes of rules and that evolutionary selection or reinforcement learning determines which classes are more popular. According to this view, agents are heterogeneous and tend to switch to heuristics that have been more successful in the recent past. More sophisticated classes of rules may require higher information-gathering costs. We now combine these two bounded rationality stories of adaptive learning and evolutionary selection into a theory of behavioral rationality and heterogeneous expectations.

As an example, assume that agents can choose between a simple rule, naive expectations, and a more sophisticated SAC learning rule. The SAC rule requires more effort, therefore it is more costly and can only be obtained at per period information costs $C \geq 0$. Hence, agents can choose between two forecasting rules:

$$p^e_{1,t} = \alpha_{t-1} + \rho_{t-1}(p_{t-1} - \alpha_{t-1}), \tag{1.18}$$

$$p^e_{2,t} = p_{t-1}. \tag{1.19}$$

In a cobweb world with SAC learning versus naive expectations, the market equilibrium price is determined by

$$p_t = \frac{a - n_{1,t}sp^e_{1,t} - n_{2,t}sp_{t-1}}{d}, \tag{1.20}$$

where $n_{1,t}$ and $n_{2,t}$ represent the fractions of producers using SAC learning respectively naive expectations.

Evolutionary selection or reinforcement learning determines how many agents will adopt each strategy. The fractions of SAC learners and naive producers depend upon an evolutionary performance measure given by (minus) squared prediction errors, and the fractions of the two types are again represented by a logit model

$$n_{1,t+1} = \frac{e^{-\beta[(p_t - p^e_{1,t})^2 + C]}}{e^{-\beta[(p_t - p^e_{1,t})^2 + C]} + e^{-\beta(p_t - p_{t-1})^2}}, \tag{1.21}$$

$$n_{2,t+1} = \frac{e^{-\beta(p_t - p_{t-1})^2}}{e^{-\beta[(p_t - p^e_{1,t})^2 + C]} + e^{-\beta(p_t - p_{t-1})^2}}. \tag{1.22}$$

This model is very similar to the model with costly rational versus free naive expectations in the previous section, with rational agents replaced by SAC learning. Recall that rational expectations requires perfect information, including knowledge about market equilibrium equations and beliefs of all other agents. SAC learning assumes *no* knowledge about beliefs of other agents, but instead SAC learning tries to extract information (including how prices may be affected by expectations of other agents) from *observable* quantities using a simple, linear model with time-varying parameters in an unknown nonlinear environment.

According to (1.21–1.22), as long as the squared forecasting error $(p_t - p_{t-1})^2$ from naive expectations is smaller than the squared forecasting error $(p_t - p^e_{1,t})^2$ from SAC learning plus the per period costs C, most producers will "free ride" and not bother about statistical learning. When the squared prediction errors for naive expectations become larger, however, most producers will switch prediction strategy and buy the SAC learning forecast.

Figure 1.6 illustrates the dynamics in the cobweb model with SAC learning versus naive expectations. The price dynamics becomes chaotic, with a complicated underlying strange attractor. Prices fluctuate irregularly, but at the same time adaptive learning enforces convergence of the learning parameters, the sample average α_t and the sample autocorrelation ρ_t. The sample average α_t converges to the (unknown) steady state price p^*, where demand and supply intersect, while the sample autocorrelation coefficient converges to a constant of about -0.4. Agents thus *learn to be contrarians*, as $\rho_t \to -0.4$, consistent with the SAC in realized prices. The presence of the SAC learning rule in the ecology of forecasting strategies ensures that much of the strongly negative autocorrelation in realized market prices is "arbitraged away," similar to the learning-to-forecast laboratory experiments.

In the cobweb market with heterogeneous agents, adaptive learning picks up the correct sample average and first-order sample autocorrelation. Interactions and evolutionary switching between these strategies cause complicated dynamical behavior. The bifurcation diagram in Figure 1.6 (bottom panel) illustrates a *rational route to randomness*, that is, complicated dynamics arises when agents become more sensitive to past performance of the strategies. The irregular price fluctuations are caused by the interaction of a simple, destabilizing strategy and a more sophisticated, costly, stabilizing strategy. This example illustrates that, agents are able to learn the optimal linear AR(1) rule in an unknown, complex heterogeneous environment.

When we apply the same 2-type model with SAC learning versus naive expectations to a stable cobweb model, evolutionary selection and adaptive learning enforce convergence to the rational expectations price. Hence, this simple 2-type heterogeneous

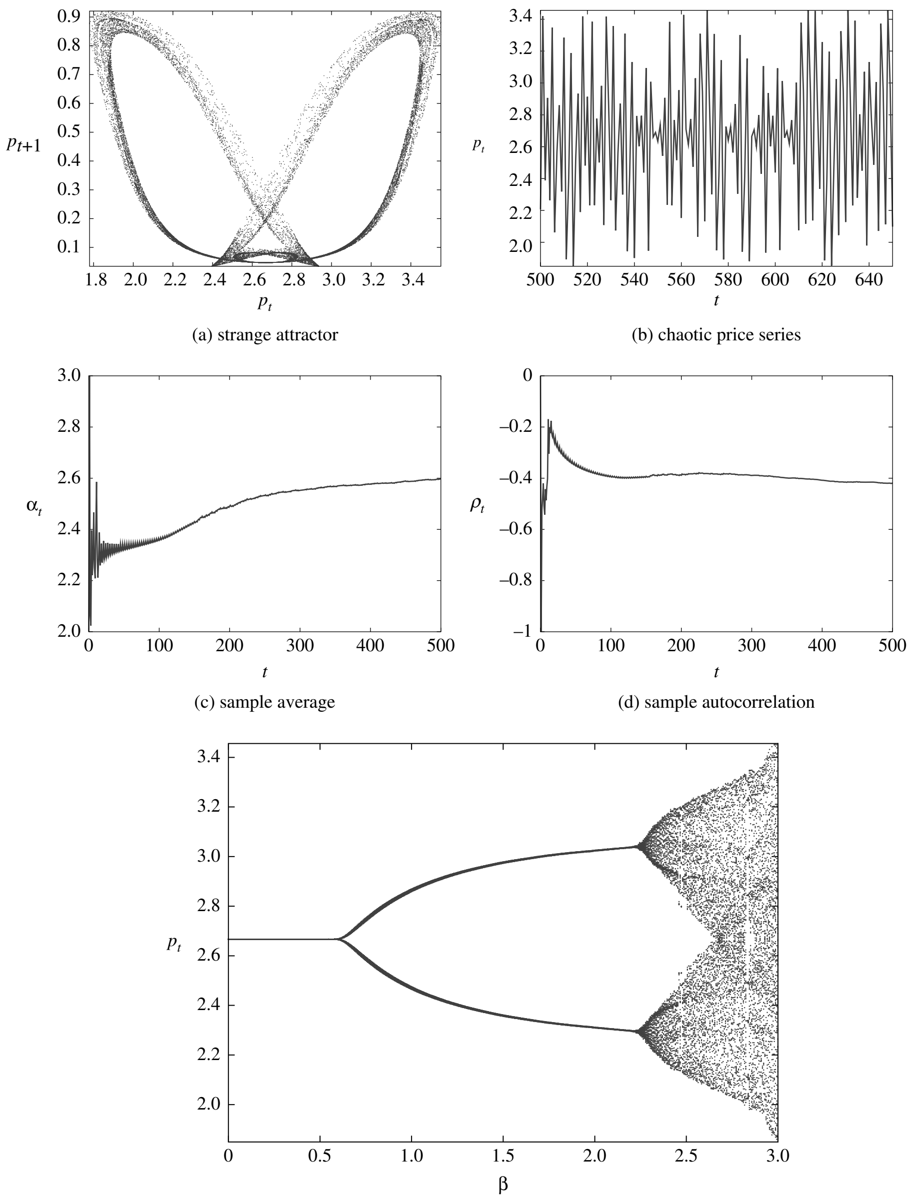

(e) bifurcation diagram

Figure 1.6. SAC learning versus naive expectations. Agents learn to be contrarians, as the first-order autocorrelation coefficient ρ_t approaches -0.4. The bifurcation diagram shows a rational route to randomness, as the intensity of choice β increases. Parameters are $\beta = 3$, $a = 4$, $d = 0.5$, $s = 1$, $C_1 = 1$, $p_0 = 2$, $\alpha_0 = 2$, $\rho_0 = 0$.

expectations model is able to explain both the coordination onto rational expectations in the stable treatment and the excess volatility with erratic price fluctuations in the unstable treatment of the cobweb laboratory experiments.

1.6 Financial markets as complex adaptive systems

The cobweb commodity market model discussed in the previous two sections is the simplest nonlinear framework in which the role of heterogeneous expectations has been studied extensively. In the last two decades, several structural heterogeneous agent models have been introduced in the finance literature; see LeBaron (2006) and Hommes (2006) for extensive surveys. In most of these heterogeneous agent models different groups of traders, having different beliefs or expectations about future asset prices, coexist. Two broad classes of traders can be distinguished. The first are *fundamentalists*, believing that the price of an asset is determined by underlying economic fundamentals, as measured, for example, by the expected future dividend stream. Fundamentalists predict that the asset price will move in the direction of its fundamental value and buy (sell) the asset when the price is below (above) its fundamental value. The second typical trader type are *chartists* or *technical analysts*, believing that asset prices are not determined by fundamentals only, but that they can be predicted by simple technical trading rules based upon observed patterns in past prices, such as trends or cycles.

An important critique from "rational expectations finance" upon behavioral finance with boundedly rational agents is that "irrational" traders will *not* survive in the market, because they will on average lose money and therefore they will be driven out of the market by rational investors, who will trade against them and drive prices back to fundamentals. According to this view it can be assumed, at least in the long run, that all agents behave "as if" they are all rational (Friedman, 1953). This "Friedman hypothesis" is essentially an evolutionary argument, suggesting that wealth- or profit-based reinforcement learning will drive out irrational investors.

At the Santa Fe Institute (SFI), Arthur et al. (1997b) and LeBaron et al. (1999) have built an early artificial stock market, where traders select their forecasting rules and trading strategies from a large population of trading rules, based upon an evolutionary "fitness measure," such as past realized profits or squared prediction errors. Strategies with higher fitness have a bigger chance of being adopted by individual traders. Computer simulations with genetic algorithms of this artificial stock market are characterized by two different regimes: close to the fundamental fluctuations, where the efficient market hypothesis (EMH) holds, and periods of persistent deviations from fundamentals and excess volatility, where the market is dominated by technical trading. Asset prices switch irregularly between these different regimes, creating stylized facts, such as time-varying, clustered volatility (GARCH effects) and fat tails, similar to those observed in real financial data. Here we briefly discuss the asset pricing model with heterogeneous beliefs of Brock and Hommes (1998), which may be viewed as a stylized, more tractable version of the SFI artificial stock market. We only illustrate the main features of the model by a simple 4-type example; an extensive treatment of the model is given in Chapter 6.

Agents can either buy an infinitely lived risky asset that pays an uncertain dividend y_t, or invest in a risk free asset that pays a fixed rate of return r. Dividends follow an exogenous stochastic process, known to all agents. In a perfectly rational world, all traders expect the future price of the risky asset to follow the fundamental price p_t^*, given by the discounted sum of expected future dividends. Boundedly rational traders, however, believe that in a heterogeneous world prices can in general *deviate* from their fundamental value. Let

$$x_t = p_t - p_t^* \tag{1.23}$$

denote the price deviation from the fundamental value. In the asset pricing model with heterogeneous beliefs, with different trader types h, $1 \leq h \leq H$, the market clearing price deviation from the fundamental benchmark is determined by

$$(1+r)x_t = \sum_{h=1}^{H} n_{ht} f_{ht}, \tag{1.24}$$

where n_{ht} is the time-varying fraction of trader type h in period t and f_{ht} is the forecast of type h at time t. Each forecasting rule f_h may be viewed as a "model of the market" of type h according to which prices will deviate from the fundamental price. For example, a forecasting strategy f_h may correspond to a technical trading rule, based upon short run or long run moving averages, or a trading range break strategy of the type used in real markets.

A convenient feature of our model formulation in terms of deviations from a benchmark fundamental is that it can be used for empirical and experimental testing of the theory. In this general setup, the benchmark rational expectations asset pricing model will be *nested* as a special case, with all forecasting strategies $f_h \equiv 0$.

An evolutionary selection mechanism describes how the fractions of different trader types will be updated over time. Fractions are updated according to an evolutionary fitness measure U_{ht}, given, e.g., by past realized profits or forecasting performance. Without discussing it in detail here, we give the expression for *realized profits* in deviations from the fundamental[8]:

$$U_{ht} = \big(x_t - (1+r)x_{t-1}\big)\left(\frac{f_{h,t-1} - Rx_{t-1}}{a\sigma^2}\right). \tag{1.25}$$

We assume here that there are no information-gathering costs, so that all forecasting strategies are freely available to all agents. Fractions of each type are given by a discrete

[8] Realized profits for type h in (1.25) are obtained by multiplying the realized excess return of the risky asset over the risk free asset times the demand for the risky asset by traders of type h. The demand is derived from mean-variance maximization of expected next period's wealth. The parameter a represents risk aversion while σ^2 represents constant conditional belief about the variance; see Chapter 6, Section 6.3 for details.

choice or multinomial logit model

$$n_{ht} = \frac{e^{\beta U_{h,t-1}}}{Z_{t-1}}, \qquad Z_{t-1} = \sum_{h=1}^{H} e^{\beta U_{h,t-1}}, \tag{1.26}$$

where Z_{t-1} is a normalization factor, so that all fractions n_{ht} add up to one. The crucial feature of (1.26) is that the higher the fitness of trading strategy h, the more traders will select strategy h. As before, the parameter β is the intensity of choice, measuring how sensitive the mass of traders is to selecting the optimal prediction strategy. The financial market with heterogeneous traders is represented by the market equilibrium equation (1.24) coupled with an evolutionary selection of strategies (1.25–1.26). Prices and beliefs thus coevolve over time.

What do price fluctuations in this stylized asset pricing model with heterogeneous beliefs look like? Brock and Hommes (1998) present an analysis of the evolutionary dynamics in the case of simple, *linear* forecasting rules

$$f_{ht} = g_h x_{t-1} + b_h, \tag{1.27}$$

where the parameter g_h represents a *trend* and the parameter b_h represents an upward or downward *bias* in prices. These very simple linear predictors were chosen as the simplest class of rules and to keep the analysis of the dynamical behavior tractable. It turns out that a rational route to randomness, that is, a bifurcation route to strange attractors arises as the intensity of choice to switch prediction or trading strategies becomes high, even when there are no information-gathering costs and all fundamental and technical trading strategies are freely available to all agents. Figure 1.7 shows chaotic price fluctuations on a strange attractor in a typical example, with four different types, fundamentalists versus three different classes of chartists. The chaotic price fluctuations are characterized by an irregular switching between phases of close-to-the-EMH-fundamental-price fluctuations, phases of "optimism" with prices following an upward trend, and phases of "pessimism," with (small) sudden market crashes. In fact, one could say that prices exhibit evolutionary switching between the fundamental value and temporary speculative bubbles.

Under homogeneous, rational expectations and a constant mean dividend process, the asset price dynamics is extremely simple: one constant price equal to the fundamental at all dates. Under the hypothesis of *heterogeneous expectations* among traders, the situation changes dramatically, and an extremely rich dynamics of asset prices and returns emerges, with bifurcation routes to strange attractors. In contrast to Friedman's hypothesis, in the evolutionary competition driven by (short run) realized profits, fundamentalists *cannot* drive out chartists trading strategies. The market is characterized by an irregular switching between periods where fundamental analysis dominates and other periods where technical trading is more profitable. Short run profit opportunities lead boundedly rational agents to adopt trend-following strategies, causing persistent price deviations from fundamentals. In empirical work, e.g., Brock et al. (1992), it has

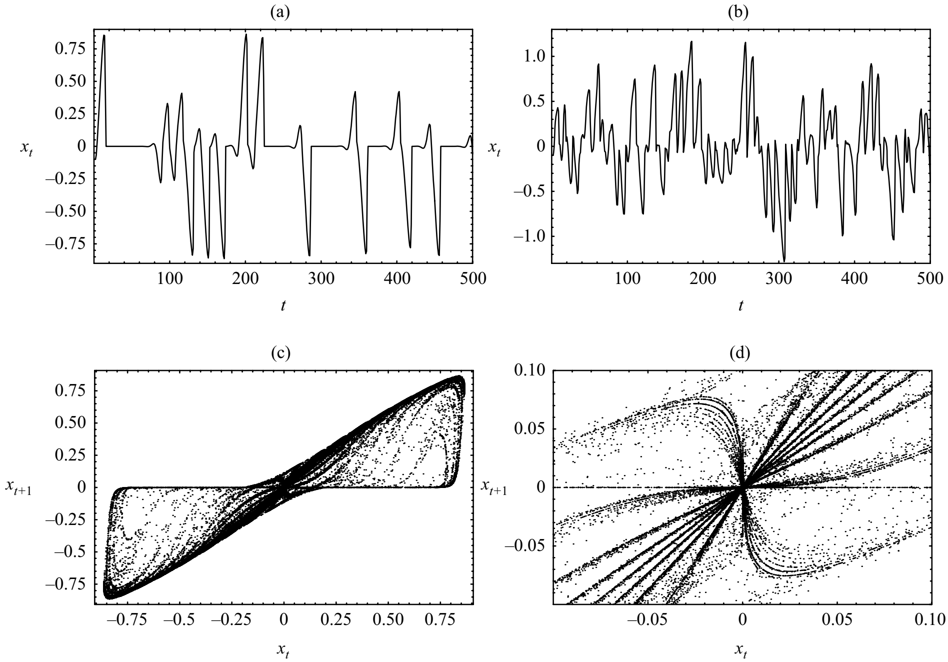

Figure 1.7. Chaotic (top left) and noisy chaotic (top right) price time series of asset pricing model with heterogeneous beliefs with four trader types. Strange attractor (bottom left) and enlargement of strange attractor (bottom right). Belief parameters are $g_1 = 0$, $b_1 = 0$; $g_2 = 0.9$, $b_2 = 0.2$; $g_3 = 0.9$, $b_3 = -0.2$ and $g_4 = 1 + r = 1.01$, $b_4 = 0$; other parameters are $r = 0.01$ and $\beta = 90.5$.

been shown that simple technical trading rules applied to real data such as the Dow Jones Index can indeed yield positive returns.

1.6.1 Estimation of a model with fundamentalists versus chartists

From a *qualitative* viewpoint, the chaotic price fluctuations in the asset pricing model with heterogeneous beliefs bear a close resemblance to observed fluctuations in real markets. But do the endogenous irregular fluctuations explain a statistically significant part of stock price movements? Here, we briefly discuss an estimation of a heterogeneous agent model, with fundamentalists versus trend followers, using yearly S&P 500 stock market data; see Chapter 7 for a more detailed discussion.

Figure 1.8 shows time series of yearly log prices of the S&P 500 stock market index, 1880–2003, around a benchmark fundamental (top left panel) and the corresponding price-to-earnings (PE-)ratio (top right panel). The fundamental price is a nonstationary stochastic process, following an exogenous stochastic earnings process with constant mean growth rate. The S&P 500 shows large swings around this RE benchmark fundamental. These large swings become even more pronounced from the PE-ratio plots. If the asset price would perfectly track its fundamental value, the PE-ratio would be constant at 17.5, as indicated by the horizontal line (right panel). For the S&P 500,

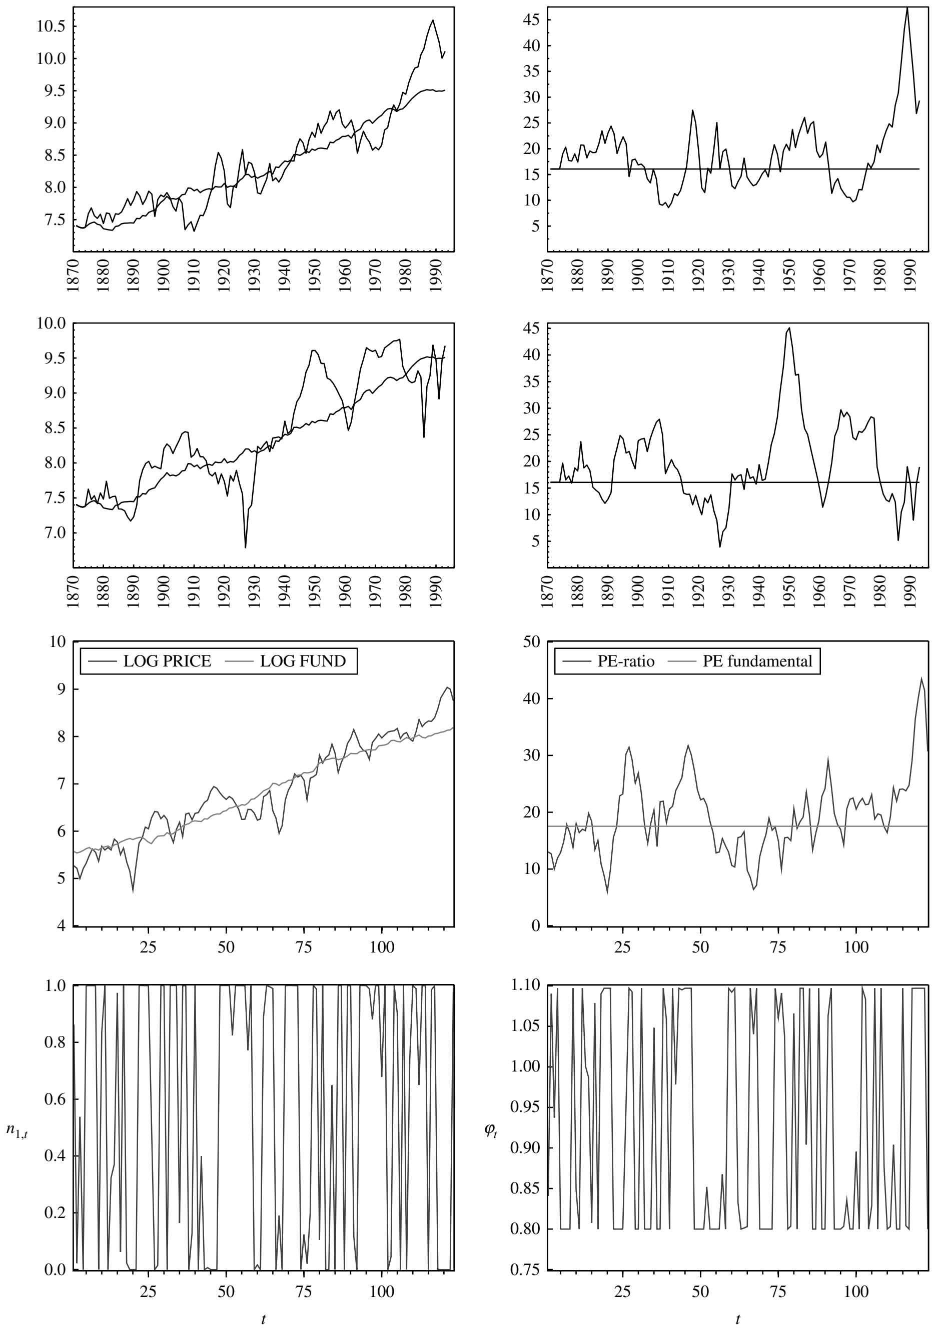

Figure 1.8. Simulated log price series (left panel) and PE-ratios (right panel) for estimated 2-type model with fundamentalists versus trend followers. Parameters: $g_1 = 0.80$, $g_2 = 1.097$ and $\beta = 7.54$. Top-panel: logarithm S&P 500 and fundamental (left) and PE ratio (right). Second panel: fitted model with reshuffled residuals. Third panel: simulated model with normally distributed shocks with mean 0 and $\sigma = 2.975$. Bottom panel: estimated fraction of fundamentalists (left) and average extrapolation coefficient (right).

however, the PE-ratio is characterized by long swings and persistent deviations from the benchmark fundamental ratio. The PE-ratio fluctuates between 8 and 28 for about 100 years, but in the late 1990s, the PE-ratio “explodes” to unanticipated high values of more than 45.

Boswijk et al. (2007) estimated a 2-type model with evolutionary strategy switching for the S&P 500 PE-ratio. The two forecasting functions were $f_{1t} = g_1 x_{t-1}$ and $f_{2t} = g_2 x_{t-1}$, with estimated parameter values $g_1 = 0.80$ and $g_2 = 1.097$, implying that type 1 behaves as a fundamentalist expecting mean-reversion of the PE-ratio toward the fundamental benchmark, while type 2 are trend followers expecting the trend to continue and the price to deviate further from fundamental value. Figure 1.8 shows plots of the fitted model time series and simulations of the estimated model, with shuffled estimated residuals (second panel) as well as normally distributed shocks with mean zero and the same variance (third panel). The simulation with reshuffled residuals shows temporary bubbles similar to the original data series, except that the timing of the large bubble is different (due to the reshuffling). A typical model simulation with normally distributed shocks of the same variance will show occasionally large deviations, up to 45 or even higher, of the PE-ratio from its fundamental benchmark. The time series of the fraction of fundamentalists (bottom left panel) is characterized by irregular switching between periods where almost all agents are fundamentalists or trend followers respectively. Finally, a simulated time series of the average market sentiment, that is, the extrapolation coefficient averaged over the population of traders,

$$\varphi_t = n_{1t} g_1 + n_{2t} g_2, \tag{1.28}$$

is shown (right bottom panel). The average extrapolation coefficient switches irregularly, and occasionally exceeds 1, causing phases of strong trend extrapolation.

These simulations show that endogenous speculative dynamic of a simple asset pricing model with two different belief types, fundamentalists versus chartists, around a benchmark fundamental may explain a significant part of observed stock price fluctuations in real markets. According to our model temporary price bubbles are *triggered by news* about fundamentals, but may become strongly amplified by trend-following strategies. For example, positive news about the economy during a number of consecutive periods may trigger a rise in stock prices, which then may become strongly reinforced by trend-following trading behavior. This may explain the strong rise in stock prices worldwide in the late 1990s, when a new internet technology provided “good news” for the growth of the economy, triggering a rise in stock prices. Our estimated model suggests that, driven by short run profit opportunities, trend-following strategies strongly amplified the rise in stock prices in the late 1990s, thus contributing significantly to the subsequent excessive rise in stocks and the “dot-com” bubble.

1.7 Learning-to-forecast experiments

In the nonlinear economic models discussed so far, expectations play an important role. But how do individuals in complex markets actually form expectations, and what is the aggregate outcome at the macro level of the interactions of individual forecasts at the micro level? Laboratory experiments with human subjects, where economic fundamentals are under control of the experimenter, are well suited to study how individuals form expectations and how their interaction shapes aggregate market behavior.

The results from laboratory experiments are somewhat mixed, however. Early experiments, with various market designs such as double auction trading, show convergence to equilibrium (Smith, 1962, Plott and Sunder, 1982), while more recent asset pricing experiments exhibit persistent deviations from equilibrium with temporary bubbles and sudden crashes (Smith et al., 1988). A clear explanation of these different market phenomena is still lacking (e.g., Duffy, 2008a,b) and this is an important challenge for experimental macroeconomics. It is particularly challenging to provide a universal theory of learning which is able to explain both the possibilities of convergence and persistent deviations from equilibrium. It is intuitively plausible that such a theory needs to be based on heterogeneous expectations and learning.

Here we briefly discuss some recent results from Anufriev and Hommes (2012a,b) to fit a heuristics switching model to laboratory experiments on expectation formation; see Chapter 8 for more details. In the *learning-to-forecast experiments* of Hommes et al. (2005), three different outcomes have been observed in the same experimental setting. Individuals had to make a two-period-ahead forecast of the price of a risky asset, say a stock. These individual forecasts determine aggregate demand and supply, leading to a market clearing price. The equilibrium price was in fact determined in exactly the same way as the asset pricing model with heterogeneous beliefs, as discussed above. Based upon the realized market price and without knowledge of the forecasts of others, individuals were then asked to make their next forecasts, and so on. The experiment lasted 50 periods. The environment in this experiment is stationary and if all agents would behave rationally or learn to behave rationally, the market price would be equal to (or quickly converge to) a constant fundamental price $p^f = 60$. In the experiment coordination of individual forecasts occurred, but three different aggregate market outcomes have been observed (see Figure 1.9, left panels):

(a) slow, monotonic convergence to the constant fundamental price level;
(b) slowly converging oscillatory movements around the fundamental price; and
(c) persistent oscillatory fluctuations around the fundamental.

A simple model based on *evolutionary selection* of *forecasting heuristics* explains how coordination of individual forecasts arises leading to these different aggregate market outcomes. The nonlinear switching model exhibits *path dependence*, since the only differences between the model simulations in Figure 1.9 are the initial states (i.e., initial prices and initial distribution over the heuristics).

The model works as follows (see Chapter 8 for more details). Agents select rules from a population of simple forecasting rules or heuristics. To keep the model as simple as possible, but rich enough to explain the different observed price patterns in the experiments, only four heuristics have been chosen. These heuristics are intuitively

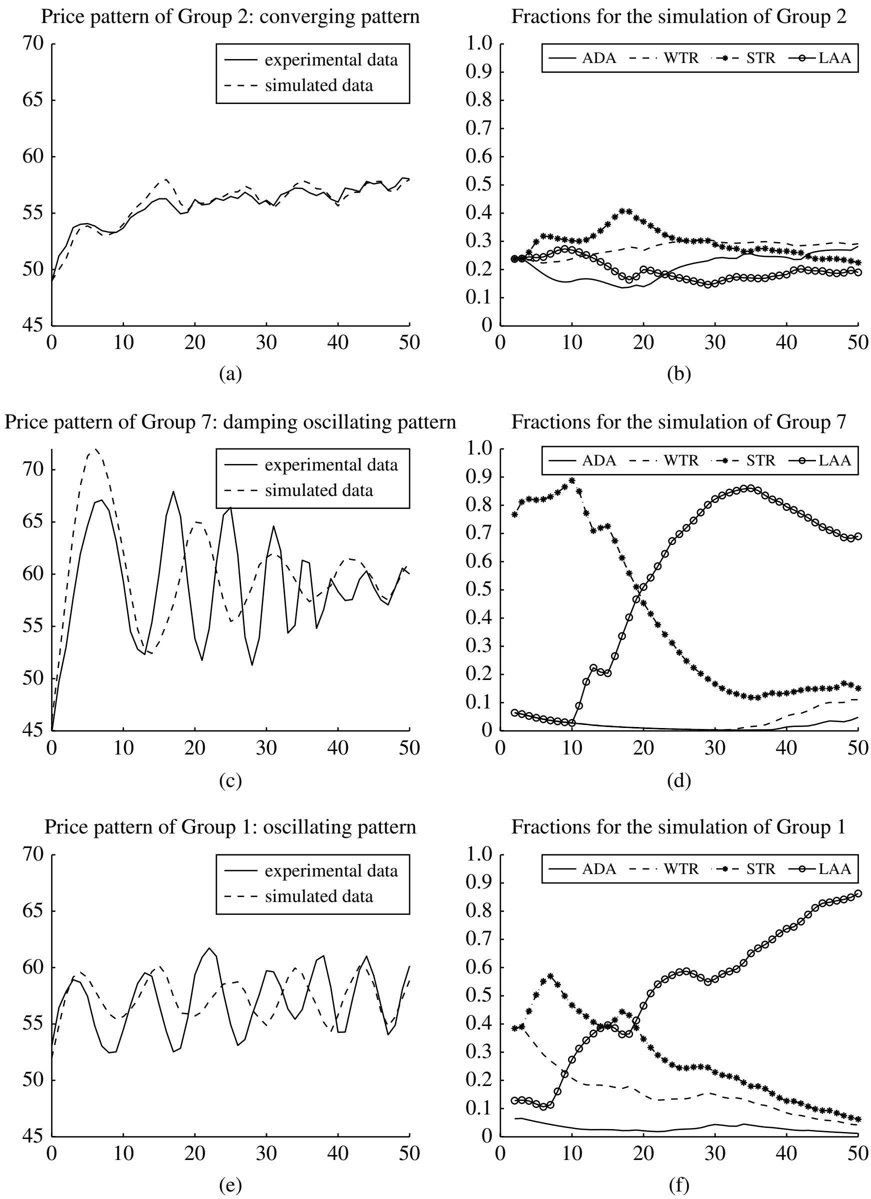

Figure 1.9. Left panels: prices for laboratory experiments (bold) and heuristics switching model (dotted). Right panels: fractions of four forecasting heuristics: adaptive expectations (ADA), weak trend followers (WTR), strong trend followers (STR) and learning anchor and adjustment rule (LAA). Coordination of individual forecasts explains three different aggregate market outcomes: monotonic convergence to equilibrium (top panel), oscillatory convergence (middle panel) and permanent oscillations (bottom panel).

plausible and were among the rules estimated for the individual forecasts in the experiment (these estimations were based on the last 40 observations, to allow for a short learning phase). The following four heuristics have been used in the simulations:

$$p^e_{1,t+1} = 0.65p_{t-1} + 0.35p^e_{1,t}, \tag{1.29}$$

$$p^e_{2,t+1} = p_{t-1} + 0.4(p_{t-1} - p_{t-2}), \tag{1.30}$$

$$p^e_{3,t+1} = p_{t-1} + 1.3(p_{t-1} - p_{t-2}), \tag{1.31}$$

$$p^e_{4,t+1} = \frac{1}{2}(p_{t-1} + \bar{p}) + (p_{t-1} - p_{t-2}), \tag{1.32}$$

where $\bar{p}$ in (1.32) is the sample average of realized prices, i.e., $\bar{p} = \sum_{j=0}^{t-1} p_j$.

Adaptive expectations (ADA) in (1.29) means that the price forecast for period $t+1$ is a weighted average of the last observed price p_{t-1} (weight 0.65) and the last own forecast $p^e_{1,t}$ (weight 0.35). Note that the last observed price has two lags (participants had to make a two-period-ahead forecast), while the last own forecast has only one lag. The *weak trend rule* (WTR) takes the last observed price level p_{t-1} as an *anchor* or *reference point* and extrapolates the last observed price change $p_{t-1} - p_{t-2}$ by a (small) factor 0.4. The *strong trend rule* (STR) in (1.31) is the same as the WTR, except that is has a larger extrapolation factor 1.3. Finally, the *learning anchor and adjustment heuristic* (LAA) uses a (time-varying) *anchor* or *reference point*, defined as an (equally weighted) average between the last observed price and the sample average of all past prices, and extrapolates the last price change from there. The LAA rule has been obtained from a related, simpler AR(2) rule $p^e_{t+1} = \frac{1}{2}(p_{t-1} + 60) + (p_{t-1} - p_{t-2})$, after replacing the (unknown) fundamental price 60 by the observable sample average $\bar{p}$. Such an AR(2) rule (or similar ones) has been estimated for a number of individual forecasts in the experiments. The first three rules are first-order heuristics in the sense that they only use the last observed price level, the last forecast and/or the last observed price change. The fourth rule combines adaptive learning of the price level and trend extrapolation, and therefore we refer to it as a learning anchor and adjustment heuristic (LAA).

The simulation model is based upon evolutionary switching or *reinforcement learning* between the four forecasting heuristics, driven by their past relative performance. Heuristics that have been more successful in the past will attract more followers. The performance measure is (minus) squared forecasting errors, similar to the financial rewards in the experiment. The performance of heuristic h is given by

$$U_{ht} = -(p_t - p^e_{h,t})^2 + \eta U_{h,t-1}. \tag{1.33}$$

The parameter η measures the relative weight agents give to past errors and thus represents their *memory strength*. When $\eta = 0$, only the performance of the last period plays a role in the updating of the shares assigned to the different rules. For $0 < \eta \leq 1$, all past prediction errors affect the heuristic's performance.

Given the performance measure, the weight assigned to rules is updated according to a *discrete choice model with asynchronous updating*

$$n_{h,t+1} = \delta n_{ht} + (1-\delta)\frac{e^{\beta}U_{ht}}{\sum_{h=1}^{4} e^{\beta U_{ht}}}. \tag{1.34}$$

There are two important parameters in (1.34). The parameter $0 \leq \delta \leq 1$ gives some persistence or inertia in the weight assigned to rule h, reflecting the fact that not all the participants are willing to update their rule in every period. Hence, δ may be interpreted as the fraction of individuals who stick to their previous strategy. In the extreme case $\delta = 1$ the initial weights assigned to the rules never change, no matter what their past performance is. If $0 \leq \delta < 1$, in each period a fraction $1-\delta$ of participants update their rule according to the well known discrete choice model. The parameter β represents the intensity of choice, measuring how sensitive individuals are to differences in strategy performance. The higher the intensity of choice β, the faster individuals will switch to more successful rules. In the extreme case $\beta = 0$, the fractions in (1.34) move to an equal distribution independent of their past performance. At the other extreme $\beta = \infty$, all agents who update their heuristic (i.e., a fraction $1-\delta$) switch to the most successful predictor.

The left panel of Figure 1.9 shows that the heuristics switching model matches all three different patterns, slow monotonic convergence to the fundamental price, dampened oscillatory price movements and persistent price oscillations, in the laboratory experiments. In all simulations in Figure 1.9, the parameters have been fixed at the same values, and the simulations only differ in the initial states, that is, the initial prices and the initial distribution of agents over the population of heuristics. The nonlinear heuristics switching model therefore exhibits *path dependence*, since the simulations only differ in initial states. In particular, the initial distribution over the population of heuristics is important in determining which pattern is more likely to emerge. The right panels of Figure 1.9 plot the corresponding transition paths of the fractions of each of the four forecasting heuristics. In the case of monotonic convergence (top panel), agents start uniformly distributed over the heuristics and the four fractions (and the individual forecasts) remain relatively close together, causing slow (almost) monotonic convergence of the price to the fundamental equilibrium 60. In the second simulation (middle panel), a large initial fraction of (strong) trend followers leads to a strong rise of market prices in the first 7 periods, followed by large price oscillations. After period 10 the fraction of strong trend followers decreases, while the fraction of the fourth rule, the learning anchor and adjustment heuristic, rises to more than 80% after 30 periods. At turning points, the flexible LAA heuristic predicts better than the static STR rule, which overestimates the price trend. After 35 periods the fraction of the LAA heuristic starts slowly decreasing, and consequently the price oscillations slowly stabilize. In the third simulation (bottom panel) weak and strong trend followers each represent 40% of the initial distribution of heuristics, causing a rise in prices which, due to the presence of weak trend followers, is less sharp than in the previous case. However, already after 5 periods the fraction of the LAA rule starts to increase, because once again at turning points it predicts better than the static STR and LTR rules, which either overestimate or underestimate the price trend at turning points. The fraction of the LAA heuristic gradually increases and dominates the market within 20 periods, rising to more than 80% after 40 periods, explaining coordination of individual forecasts as well as persistent price oscillations around the long run equilibrium level.

These simulations illustrate how the interaction and evolutionary selection of individual forecasting heuristics may lead to coordination of individual behavior upon different price patterns and enforce path-dependent aggregate market outcomes. Individuals are behaviorally rational and use simple heuristics consistent with recent observations. Evolutionary learning leads to switching between simple forecasting heuristics based upon recent performance and different types of aggregate behavior may emerge.

1.8 Simple complex systems

The economy is a complex system, with many interacting consumers, firms, investors, banks, etc. But how complex should a model be to describe economic complexity? One could think of a detailed *agent-based model* (ABM) using a "bottom up" approach to model agents' interactions at the micro level and study its aggregate macro behavior. ABMs are becoming increasingly popular in finance and in macro; see, e.g., the collection of papers in the *Handbook of Computational Economics, Volume 2: Agent-Based Computational Economics* (Tesfatsion and Judd, Eds., 2006) and the *Handbook of Financial Markets. Dynamics and Evolution* (Hens and Schenk-Hoppé, Eds., 2009). For ABMs in macro; see, e.g., the monographs of Aoki (2002) and Delli-Gatti et al. (2008); DeGrauwe (2010a,b) contains a stimulating discussion of a "bottom up" approach of ABMs in behavioral macroeconomics.

While detailed ABMs present an important challenge and promising approach in economic modeling, this book emphasizes *simple complex system models* as complementary tools to gain insights in nonlinear interaction mechanisms. The key features of these models are that they are *nonlinear* and that there is some form of *heterogeneity* and endogenous switching between heterogeneous decision rules. Since the economy is inherently uncertain, the "law of motion" of the economy is stochastic. A simple complex system with heterogeneous agents typically is of the form

$$X_{t+1} = F(X_t; n_{1t}, ..., n_{Ht}; \lambda; \delta_t; \epsilon_t), \tag{1.35}$$

where F is a *nonlinear* mapping, X_t is a vector of state variables, say prices (or lagged prices), n_{jt} is the fraction or weight of agents of strategy type h, $1 \leq h \leq H$, λ is a vector of parameters and δ_t and ϵ_t are (vectors of) noise terms. There are (at least) two types of uncertainty relevant for economic modeling, *intrinsic noise* and *model approximation errors*. Intrinsic noise refers to intrinsic uncertainty about economic fundamentals (preferences, technology, future earnings, future growth, etc.) in the economy. The noise term δ_t then represents unexpected random shocks, "news" about economic fundamentals. The second type of noise, *model approximation errors*, represents the fact that a model can only be an approximation of the real world and that part of the economy remains not modeled. Approximation errors will also be present in a physics model, although the magnitude may be smaller. In a financial market, one may for example have a (small) fraction of "noise traders" who trade randomly and whose behavior therefore is uncertain. The law of motion of the economy is then a nonlinear stochastic system as in (1.35).

In general, the nonlinear stochastic model can be a detailed ABM, a highly nonlinear complex system with many different agent types and of high dimension. For a detailed ABM it will be difficult to use analytical tools and one will mainly have to resort to numerical simulations. In this book, we will emphasize *simple complex systems*, where the dimension of the system is relatively low, the number of different agent types is relatively small, and the system is simple enough to be studied, at least partially, by analytical tools. Since the dynamical behavior of simple complex systems is rich, *a simple nonlinear system model buffeted with noise as in (1.35) may explain a significant part of observed fluctuations and stylized facts in economic and financial markets*. An important goal in simple complex systems modeling is to match the statistical regularities of empirical data both at the macro level and at the micro level. Hence, one would like to match both individual behavioral decision rules, e.g., calibrating them with laboratory experiments, and, at the same time, match aggregate macro behavior and time series properties.

A special case of the nonlinear stochastic system (1.35) arises when all noise terms are set equal to their unconditional mean. We will refer to this system as the *(deterministic) skeleton* denoted by

$$X_{t+1} = F(X_t; n_{1t}, ..., n_{Ht}; \lambda). \tag{1.36}$$

In order to understand the properties of the general stochastic model (1.35) it is important to understand the properties of the nonlinear deterministic skeleton. In particular, one would like to impose as little structure on the noise process as possible, and relate the individual decision rules as well as the aggregate stylized facts of the general stochastic model (1.35) directly to generic properties of the underlying deterministic skeleton. This naturally leads to the study of the dynamics of simple nonlinear dynamical systems.

1.9 Purpose and summary of the book

This book serves three important purposes. First, it presents simple examples of complex systems applications in economics and finance. The simplicity of these examples should help the reader to grasp the essential features of nonlinear complex systems. Second, the methodological part (Chapters 2 and 3) serves as a primer to nonlinear dynamics, introducing the key tools in the analysis of simple nonlinear systems that should be part of the toolbox of any quantitative economist. Third, our main focus is on bounded rationality and heterogeneous expectations in simple complex adaptive economic systems. In particular, we extensively discuss a theory of behavioral rationality, heterogeneous expectations and learning in complex economic systems and confront this theory, both at the macro and the micro level, with empirical time series and laboratory experimental data.

The book is organized in seven chapters following this introduction. Two methodological chapters, Chapters 2 and 3, give an introduction to the mathematical tools of nonlinear, discrete time dynamical systems. Chapter 2 deals with one-dimensional systems, discusses the (in)stability of steady states, introduces elementary bifurcations (tangent, period-doubling, pitchfork, transcritical), defines the notion of chaos and Lyapunov exponents and discusses the period-doubling bifurcation route to chaos. Chapter 3 deals with two- and higher-dimensional systems, discusses the saddle-node and the Hopf (or Neimark–Sacker) bifurcations, introduces the "breaking of an invariant circle" bifurcation route to chaos and strange attractors and introduces the key notions related to chaotic dynamics, such as horseshoes, homoclinic orbits, homoclinic bifurcations, Lyapunov exponents and strange attractors. We have made an attempt to provide an introduction to nonlinear dynamics for non-specialists and a general audience of economists, emphasizing the most important concepts to be used in economic applications and adding references to more advanced mathematical treatments whenever appropriate.

The second part of the book, Chapters 4, 5 and 6, contains simple examples of complex systems modeling in economics and finance. Chapter 4 discusses the *nonlinear cobweb model* with various benchmark *homogeneous expectations*: naive expectations, rational expectations, adaptive expectations and linear backward-looking expectations. In a nonlinear cobweb economy with monotonic demand and supply curves, all of these simple adaptive and backward-looking expectations may lead to chaotic price fluctuations. These benchmark cases provide simple didactic examples of stylized complex dynamics applications to economics, illustrating how expectations feedback may generate complicated dynamics in a *nonlinear* environment. Chapter 4 also discusses the notion of *consistent expectations equilibrium*, where behaviorally rational agents learn the "optimal" linear AR(1) rule in a complex, nonlinear environment. Chapter 5 discusses the cobweb model with *heterogeneous expectations*, focusing on 2-type examples with a sophisticated but costly expectations rule – rational expectations, fundamentalist forecast, a contrarian rule or adaptive learning – competing against a simple, freely available forecasting heuristic such as naive expectations. A common finding is a *rational route to randomness*, i.e., a bifurcation route to chaos, as agents become more sensitive to differences in evolutionary fitness. Moreover, the complexity of the price dynamics increases as the learning detects and exploits more structure in price fluctuations.

Chapter 6 discusses a standard financial market *asset pricing model* with *heterogeneous beliefs*. A number of simple examples, with two, three and four trader types – fundamentalists versus chartists – is discussed. A rational route to randomness arises, even when there are no information-gathering costs for more sophisticated strategies. Hence, simple technical trading rules are not driven out of the market, but survive evolutionary competition driven by (short run) realized profits. Another simple 2-type example, fundamentalists versus conditional trend followers, exhibits coexistence of a stable fundamental steady state and a stable limit cycle. Hence, there is path dependence:

the market may either converge to the stable fundamental price or it may perpetually oscillate around the fundamental price with the fractions of fundamentalists and chartists changing over time. In the presence of noise, the system exhibits clustered volatility, with the market switching irregularly between close to the fundamental price fluctuations and large swings and excess volatility in asset prices. Finally, the case with many different expectations types is discussed. The model with many different trader types is well approximated by the so-called *large type limit* (LTL), a tool that can be used to study markets with many different trader types.

The final part of the book, Chapters 7 and 8, discusses the *empirical validity* of heterogeneous expectations models. Chapter 7 discusses the estimation of a simple 2-type asset pricing model with fundamentalists versus trend followers on yearly data of the S&P 500 stock market index. Behavioral heterogeneity is statistically significant with large swings in the fractions of both types of traders. In particular, the heuristics switching model explains the dot-com bubble as being triggered by good news about economic fundamentals – a new internet technology – strongly amplified by technical trading.

Finally, Chapter 8 discusses the empirical validity of heterogeneous expectations models by laboratory experiments with human subjects. *Learning-to-forecast experiments*, where subjects' only task is to forecast prices in an expectations feedback environment, in the cobweb and the asset pricing frameworks are discussed. Coordination on different types of aggregate market behavior – stable convergence or persistent price oscillations – arises. A simple heuristics switching model – exhibiting path dependence – can explain coordination on these different aggregate outcomes. Another striking finding is that negative expectations feedback markets, such as commodity prices in a cobweb framework where high price expectations yield high production and thus low realized market prices, are rather stable, while positive feedback markets, such as speculative asset markets, tend to oscillate around the fundamental price. The simple heuristics switching model explains both types of aggregate behavior as well as individual behavior. In the negative feedback market, adaptive expectations dominates evolutionary competition, because the trend-following heuristics perform poorly. In contrast, in positive feedback markets trend-following rules perform relatively well and amplify price oscillations, possibly leading to bubbles and crashes.