

Chapter 7

Complexity, Evolution and Learning

Empirical and Experimental Validation of Heterogeneous Expectations

Cars Hommes

7.1 Introduction

A paradigm shift in economics is taking place. In traditional, neoclassical economics a representative agent who behaves perfectly rational has been the main working hypothesis and mathematical analysis of simple tractable models its main focus. A problem with this approach is that it requires unrealistically strong assumptions about individual behaviour, such as perfect knowledge and information about the economy and extremely high computational abilities to do what is optimal. An advantage of the neoclassical research programme, partly explaining its success, is that rationality imposed through optimizing behaviour and model consistent expectations enforces strong discipline on the modelling framework leaving no room for market psychology and unpredictable, irrational behaviour.

An alternative complexity view is now emerging, based on interaction of many heterogeneous agents, whose behaviour is only *boundedly rational*. In this new behavioural agent-based approach, computer simulation models are the main modelling framework. An advantage is that it becomes possible to describe in detail individual behaviour at the micro level based on realistic assumptions. The Santa Fe conference proceedings Anderson et al. (1988) and Arthur et al. (1997a) contain many contributions within the complexity view. The recent *Handbook of computational economics* (Tesfatsion and Judd 2006) contains many chapters describing the state of the art of agent-based economics. There is however still an important problem with the bounded rationality research programme: it leaves too many degrees of freedom. There is only one way (or perhaps a few ways) one can be right, but there are many ways one can be wrong. To turn the alternative view into a successful research programme, one has to “tame the wilderness of bounded rationality”.

C. Hommes

School of Economics, University of Amsterdam, Amsterdam, The Netherlands

A key feature that distinguishes economics from natural sciences is that market realizations depend on future *expectations* and, at the same time, expectations about future developments are based on current and past realizations. An economy is an expectations feedback system in which beliefs and realizations co-evolve. Agents are “smart” and will adapt their behaviour if it benefits them. If all agents are perfectly rational, in equilibrium individual expectations and realizations must coincide on average, leading to the neoclassical representative rational agent model. But if agents are heterogeneous and only boundedly rational, one needs a convincing theory of heterogeneous expectations. In this chapter we discuss, a simple story of *heterogeneous expectations* and some empirical and experimental validation. Agents can choose from a class of simple heuristics disciplined by *adaptive learning* and *evolutionary selection*. An extensive recent survey of this approach including many references to related work can be found in Hommes (2006).

This chapter is organized as follows. Section 7.2 describes a simple example, an asset pricing model with heterogeneous beliefs, and illustrates how the asset price dynamics may become unstable when expectations are driven by reinforcement learning based on past strategy performance. Section 7.3 discusses the empirical validity of a simple version of the model with two types of traders, fundamentalists and technical analysts, and how it explains the “dot com bubble” in stock prices in the late 1990s. Section 7.4 discusses how this approach matches the stylized facts of learning to forecast laboratory experiments with human subjects. Finally, Sect. 7.5 briefly describes a future perspective.

7.2 An Asset Pricing Model with Heterogeneous Beliefs

As a simple example of a model with heterogeneous expectations we consider the asset pricing model with heterogeneous beliefs of Brock and Hommes (1998). This model may be viewed as a simple stylized version of the Santa Fe artificial stock market model introduced by Arthur et al. (1997b). Agents can invest in a risk free asset that pays a fixed return $1 + r$ or in a risky asset that pays uncertain dividends y_t in each period. The market clearing pricing equation is given by

$$(1 + r)p_t = \sum_{h=1}^H n_{ht} E_{ht}(p_{t+1} + y_{t+1}) + \varepsilon_t, \quad (7.1)$$

where p_t is the price of the risky asset, n_{ht} the (time varying) fraction of trader type h , $E_{ht}(\cdot)$ the belief of type h about next period’s price plus dividend, and ε_t a noise term representing, for example a small fraction noise traders. In the special case when *all* agents are rational the asset price will be equal to the *rational, fundamental benchmark* p_t^* , given by the discounted sum of expected future dividends

$$p_t^* = \frac{E_t[y_{t+1}]}{1+r} + \frac{E_t[y_{t+2}]}{(1+r)^2} + \dots$$

This fundamental benchmark is nested as a special case within the general heterogeneous agent model. In the case of IID dividends with mean \bar{y} , the fundamental price becomes constant,

$$p^* = \bar{y}/r.$$

Assuming that the beliefs about future dividends are correct (for example because they can be inferred from past observations of the exogenous dividend process), the model can be rewritten in deviations

$$x_t = p_t - p^*$$

from the fundamental and simplifies to:

$$(1+r)x_t = \sum_{h=1}^H n_{ht} E_{ht} x_{t+1} + \varepsilon_t. \quad (7.2)$$

Strategy choice follows an *evolutionary selection* principle, that is, “strategies that have performed better attract more followers”. This can be modelled in several ways, but we follow Brock and Hommes (1997) where the fractions of belief type h are determined by the discrete choice model (a random utility model)

$$n_{ht} = \frac{e^{\beta U_{h,t-1}}}{Z_{t-1}}, \quad (7.3)$$

where

$$Z_{t-1} = \sum_j e^{\beta U_{j,t-1}}$$

is normalization factor and $U_{h,t-1}$ measures the *past performance* or *fitness* (for example realized profits, forecasting performance, etc.) of strategy h . The parameter β is the *intensity of choice* measuring the sensitive of agents to differences in strategy performance. In the extreme case $\beta = 0$, agents behave randomly and all fraction types are fixed with equal weights; at the other extreme, $\beta = \infty$, all agents immediately switch to the best predictor (the “neoclassical limit”).

Which ones out of an ocean of possible forecasting rules will agents use? In a real market, it seems unlikely that many agents will coordinate on a very complicated rule. Therefore, we use simple rules, such as linear rules with only one time lag (written in deviations

$$x_{t-1} = p_{t-1} - p^*$$

from the fundamental):

$$f_{ht} = p^* + g_h x_{t-1} + b_h,$$

where g_h is a *trend* parameter and b_h a *bias* parameter. Another simple rule not using any fundamental price information is the *trend extrapolating rule*

$$f_{ht} = p_{t-1} + g_h(p_{t-1} - p_{t-2}),$$

which simply extrapolates the last price change. So far, the parameters in the forecasting rules have been fixed, but one can introduce *adaptive learning* to learn the parameters over time. For example, agents may update forecasting parameters by sample average or by employing a recursive ordinary least squares scheme (OLS-learning), as additional observations become available (see for example Evans and Honkapohja (2001) for an extensive treatment of adaptive learning in macroeconomics and Sargent (2008) for a recent discussion of the importance of learning in macroeconomics and monetary policy).

Figure 7.1 shows simulations of the price fluctuations in an example with four belief types, including fundamentalists and trend followers, and fitness given by last period's realized profits. When the intensity of choice is small, the steady state is typically stable and the asset prices converge to the fundamental benchmark. Intuitively this may be understood by observing that for small intensity of choice, agents are more or less randomly distributed over the different strategies, and as a result the average forecast is close to the fundamental enforcing convergence to the fundamental price. In contrast, when the intensity of choice is large agents typically coordinate on a common strategy and the dynamics destabilizes. In particular, coordination on a trend following strategy may occur, leading to persistent price deviations from fundamental. Indeed the asset pricing dynamics in Fig. 7.1 is characterized by irregular switching between phases of close to the fundamental price fluctuations with fundamentalists dominating the market and phases of temporary bubbles when trend following strategies dominate the market. Excess volatility and temporary bubbles are driven by short run profit opportunities. The noisy simulation illustrates that even in this simple model the start and burst of the temporary bubbles are highly unpredictable.

7.3 Empirical Validation

How relevant are these bubble and crash dynamics to real financial market data? We briefly discuss the estimation of a simple version of the model with two types of agents, using yearly S&P 500 data; see Boswijk et al. (2007) for a detailed analysis.

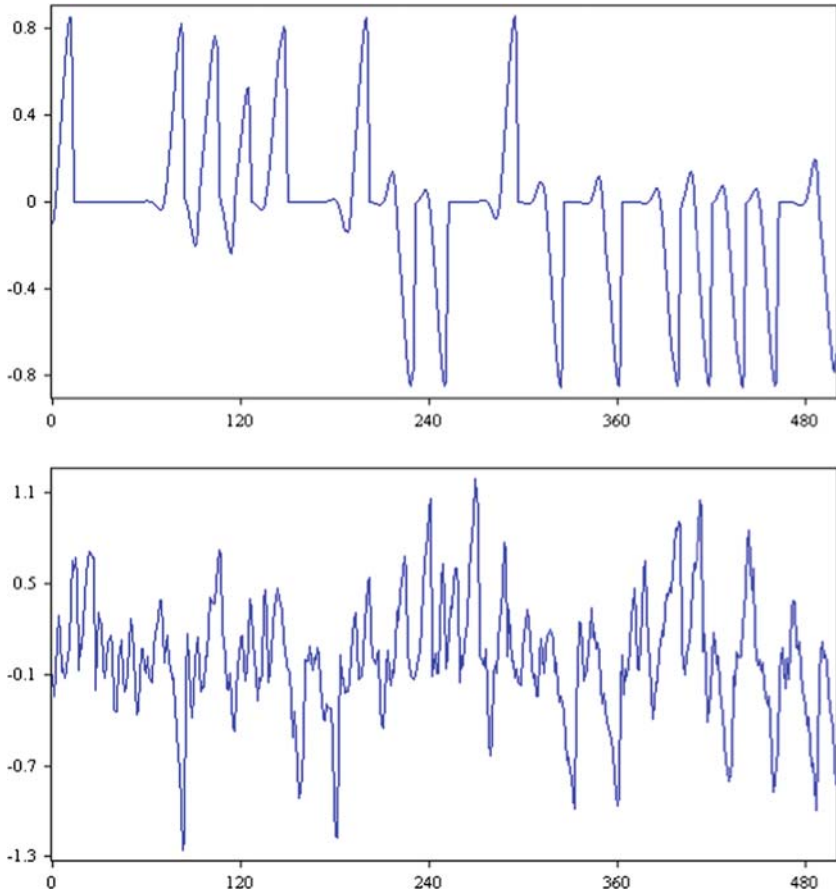


Fig. 7.1 Chaotic (*top*) and noisy chaotic (*bottom*) time series of asset prices (in deviations from the fundamental price) in an example with four trader types. Prices fluctuate irregularly around the benchmark fundamental price (which corresponds to 0). Parameters are: $g_1 = 0$, $b_1 = 0$; $g_2 = 1.1$, $b_2 = 0.2$; $g_3 = 0.9$, $b_3 = -0.2$ and $g_4 = 1.21$, $b_4 = 0$, $r = 0.1$ and $\beta = 90$

Figure 7.2 shows the logs of yearly S&P 500, 1871–2003, and a benchmark fundamental price based on dividends with a constant growth rate g . This is the standard Gordon model and the fundamental benchmark is given by

$$p_t^* = \frac{1+r}{r-g} y_t,$$

where g is the growth rate of dividends and r is the required rate of return for investors to hold the risky asset (given by the sum of the risk premium to hold

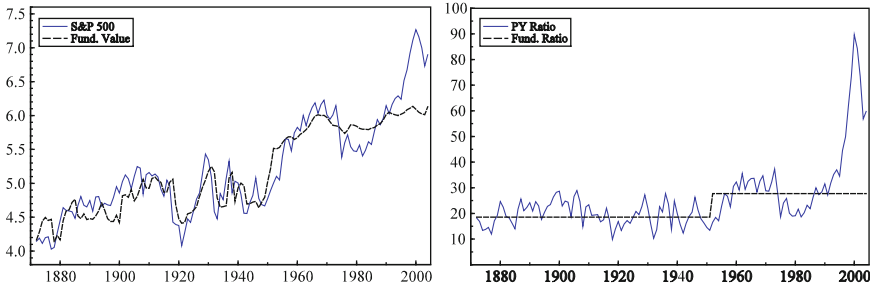


Fig. 7.2 Time series of the log of S&P 500, 1871–2003 and the benchmark fundamental for a dividend process with constant growth rate

stocks and the risk free interest rate). The corresponding fundamental price to cash flow ratio

$$\delta_t^* = \frac{p_t^*}{y_t} = \frac{1+r}{r-g} = m$$

is constant along the fundamental (the right plot in Fig. 7.2 allows for one jump in the fundamental in 1950, due to a jump in the risk premium; see for example Fama and French (2002)). Figure 7.2 shows that the realized price–dividend ratio shows large swings around the fundamental benchmark, fluctuating between 10 and 30 for more than a century, rising to unprecedented values of almost 90 in the 1990s, and coming down to values below 60 in recent years.

There are two competing views concerning the explanation of swings in price-to-cash flows. Some attribute them to rational responses to macroeconomic fundamentals, while others judge that irrational swings in investor sentiment play a significant role. Shiller (2000) gives a lucid description of both views, stressing the relevance of psychological factors.

Boswijk et al. (2007) estimated a simple two-type model of the form

$$R^*x_t = n_t\phi_1x_{t-1} + (1 - n_t)\phi_2x_{t-1} + \varepsilon_t, \tag{7.4}$$

where

$$R^* = (1+r)/(1+g), \quad x_t = \delta_t - m$$

is the deviation of the price-to-cash flow from the fundamental, n_t and $(1 - n_t)$ are the fractions of the two types (depending on past realized profits) and ϕ_hx_{t-1} , $h = 1, 2$, are the forecasts of the two types of next period’s deviation from the fundamental (only the first lag was significant). The estimation results yield significant estimates $\phi_1 = 0.76$ and $\phi_2 = 1.14$, implying that type 1 are fundamentalists believing in *mean reversion* of the price towards its fundamental value, while type 2 are trend followers, believing that the price bubble will continue. Figure 7.3

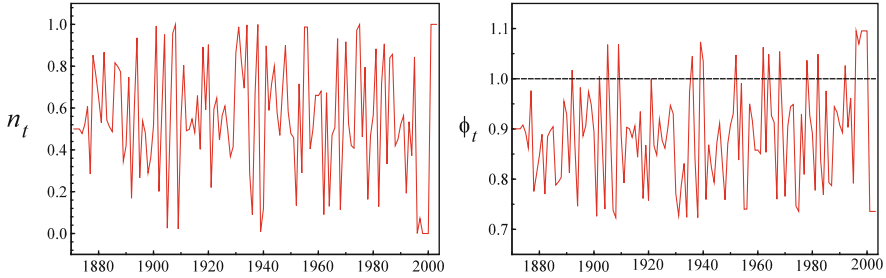


Fig. 7.3 Time series of the estimated fraction n_t of fundamentalists (*left panel*) and average market sentiment $\phi_t = \{n_t\phi_1 + (1 - n_t)\phi_2\}/R^*$ (*right panel*)

shows the time variation of the estimated fraction n_t of fundamentalists. Significant heterogeneity with strategy switching and large fluctuations in the fractions of both types occur. In particular, one observes a low fraction of fundamentalists for 5 or 6 subsequent years in the late 1990s. The average coefficient

$$\phi_t = \{n_t\phi_1 + (1 - n_t)\phi_2\}/R^*$$

in Fig. 7.3 shows that market sentiment fluctuates considerably over the years, with average traders believing in explosive asset prices in the late 1990s. This simple model explains the “dot com bubble” in the late nineties as being *triggered by fundamentals*, in the form of good news (a new technology) about the economy, and *strongly amplified by trend following strategies* based on reinforcement learning driven by short run profits.

The PD-ratio has come down to values below 60 in recent years, and one may ask the question: *Will the bubble resume?* Figure 7.4 shows prediction of both the nonlinear model with strategy switching and the linear model with a representative fundamentalist, believing in average mean reversion. Clearly the nonlinear switching model predicts much larger swings in price-to-cash flow fluctuations of asset prices than its linear, representative agent counterpart.

7.4 Learning to Forecast Experiments

Laboratory experiments with human subjects are well suited to discipline the “wilderness of bounded rationality”. In laboratory markets with carefully controlled market fundamentals one can investigate which behavioural rules are more likely to be used by human subjects in different market environments. In this section we briefly discuss laboratory experiments on expectations formation in the asset pricing framework of the previous sections. We address the following questions:

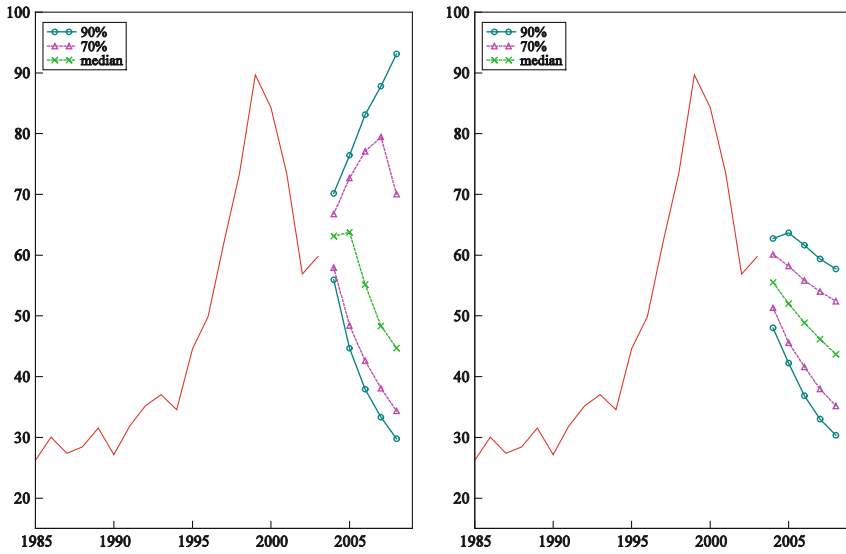


Fig. 7.4 Quantiles of 2,000 simulated predictions of the PD-ratio for the nonlinear evolutionary switching model (left) and the linear, representative agent model (right). Both models are estimated using data until 2003 and then predict up to 5 years ahead

- How do *boundedly rational* agents form expectations and how do they learn in a *heterogeneous* world?
- How do individual forecasting rules *interact* and what is the *aggregate outcome* of individual interaction?
- Will *coordination* occur, even when there is limited market information?
- Does *learning* enforce convergence to rationality?

Hommes et al. (2005) performed learning to forecast experiments in an asset pricing framework similar to that used in Sects. 7.2 and 7.3. Six human subjects have to forecast the price of a risky asset for 50 periods, and their payment is inversely related to their forecasting errors. There is expectations feedback, since the realized market price is determined by aggregation of individual forecasts. After all individual make a forecast, the computer computes a market clearing price derived from standard mean–variance maximization demand functions using the individuals forecasts as inputs. Since subjects only forecast and trading is completely computerized, agents may be viewed as rational optimizers, given their individual forecasts. Such a laboratory environment thus produces “clean data” on expectations, and one can test various expectations hypotheses. Except for the six subjects, there is a seventh robot trader in the market, who always predicts the fundamental price and whose weight increases (from 0 to at most 0.2) when prices deviate more from fundamental.

Subjects thus have *limited information* about the market. They are told that they are advisors to a pension fund, which will invest more in the risky asset, when the

subject makes a higher forecast. They also know that the asset price is determined by market clearing. From the qualitative market information, subjects should be able to understand that the asset market exhibits *positive feedback*, that is, higher forecasts lead to higher realized market prices. Subjects also know the interest rate $r = 0.05$ and the mean dividend $\bar{y} = 3$, and could use these to compute the fundamental price

$$p^f = \bar{y}/r = 60.$$

Furthermore in forecasting p_{t+1} , they know past realized prices (up to p_{t-1}), their own past forecasts (up to $p_{t,h}^e$) and their own earnings (up to $e_{t-1,h}$). However, subjects do *not* know market equilibrium equations, the forecasts of others and the number of pension funds in the market. The information in these experiments is therefore similar to what is often assumed in models with boundedly rational traders.

The (unknown) price generating process is given by

$$p_t = \frac{1}{1+r} \left((1-n_t) \frac{p_{t+1,1}^e + \dots + p_{t+1,6}^e}{6} + n_t p^f + \bar{y} + \varepsilon_t \right), \quad (7.5)$$

where n_t is the share of robot traders given by

$$n_t = 1 - \exp \left(-\frac{1}{200} |p_{t-1} - p^f| \right), \quad (7.6)$$

$p_{t+1,h}^e$, $1 \leq h \leq 6$, are the individual forecasts and ε_t is a small noise term. If all subjects would forecast rationally and use the fundamental price of 60 as their individual forecast, the realized market prices would be close to 60 with small random fluctuations around it. This is perhaps not what one would expect right from the start in a market with limited information, but an interesting question is whether the market price will at least converge to the fundamental price. It is useful to briefly mention two other homogeneous agents benchmarks. If all subjects would use *naive expectations*, that is, use the last price observation to forecast

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

starting say with an initial forecast of 50, then realized market prices will converge monotonically to the fundamental price 60. If on the other hand all subjects use a simple trend extrapolation rule

$$p_{t+1}^e = (p_{t-1} + 60)/2 + (p_{t-1} - p_{t-2}), \quad (7.7)$$

then prices will fluctuate around the fundamental for 50 periods (about six oscillations). One may wonder how individual subjects would arrive at such a rule, but

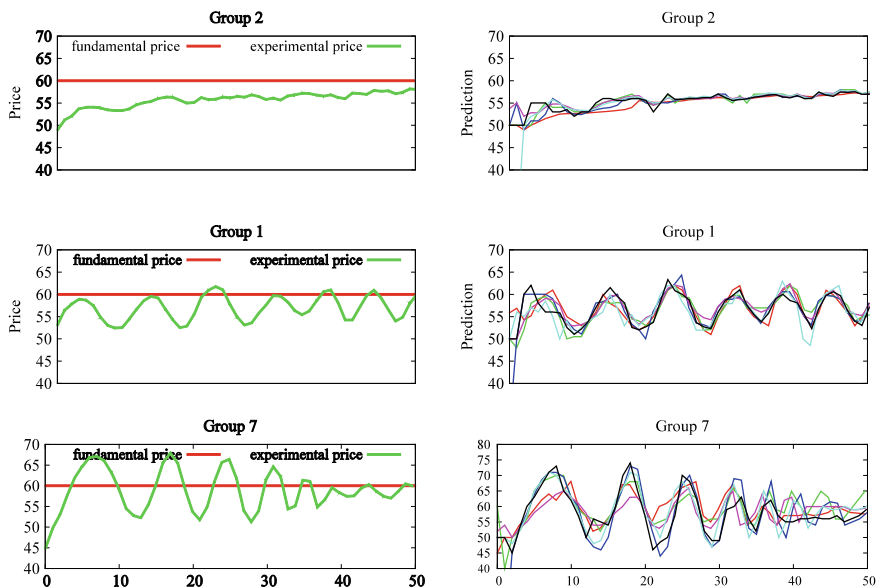


Fig. 7.5 Three typical outcomes of realized prices (*left panel*) and individual forecasts (*right panel*) of the learning to forecast laboratory experiments: (1) monotonic convergence, (2) permanent oscillations, and (3) dampened oscillations

remarkably estimation of the forecasting rules showed that a number of individuals use a rule very similar to (7.7).

Figure 7.5 shows some typical outcomes of realized prices (*left panel*) and individual forecasts (*right panel*). Three qualitatively different patterns are observed: (1) monotonic convergence, (2) permanent oscillations, and (3) dampened oscillations. Monotonic convergence is very similar to what would happen if all subjects use a naive forecast. The permanent oscillations are similar to what would happen if all subjects use a simple linear AR2 rule such as (7.7). In the third case of dampened oscillations a strong price trend emerges in the beginning of the experiment, but the strong trend gets weaker and reverses when prices deviate too much from their fundamental value. Another striking feature of the experiment is that in all cases there is strong *coordination* on a common prediction rule, as illustrated in the right panel of Fig. 7.5. Coordination however is *path dependent*, since different qualitative outcomes are observed in different markets.

Estimation of individual prediction rules shows that for most subjects (more than 90%) forecasting is well explained by a simple linear model with no more than three lags in prices and individual forecasts. In fact, for a majority of subjects (more than 50%) a simple rule with only one or two lags fits the forecasting behaviour very well. Some simple rules that have been estimated include:

- *Adaptive expectations*

$$p_{t+1}^e = wp_{t-1} + (1-w)p_t^e$$

(in converging groups)

- *Linear rules*

$$p_{t+1}^e = \alpha + \beta_1 p_{t-1} + \beta_2 p_{t-2}$$

(in oscillating groups)

- *Trend-extrapolating rules*

$$p_{t+1}^e = p_{t-1} + \gamma(p_{t-1} - p_{t-2})$$

(in oscillating groups)

In order to explain these experiments Anufriev and Hommes (2009) recently developed a *heuristics switching model*. There are a number of *simple heuristics* and in the beginning agents choose heuristics *randomly*. Agents evaluate the *past performance* of these heuristics based on forecasting accuracy, and subsequently tend to *switch* to more successful heuristics. Figure 7.6 shows simulations of the heuristics switching model reproducing the three different patterns observed in the laboratory experiments.

The four forecasting heuristics used in the simulations are adaptive expectations, a weak and a strong trend-extrapolating rule and an *anchoring and adjustment heuristic*

$$p_{4,t+1}^e = 0.5p_{t-1}^{\text{av}} + 0.5p_{t-1} + (p_{t-1} - p_{t-2}), \quad (7.8)$$

where p^{av} is the sample average of all past prices. This rule uses an anchor (the average of the last observed price and the sample average) and extrapolates a trend from there. Following the terminology of Tversky and Kahneman (1974), it may be viewed as a forecasting anchoring and adjustment heuristic.

The price dynamics in the heuristics switching model is given by

$$p_t = \frac{1}{1+r_f} ((n_{1,t}p_{1,t+1}^e + n_{2,t}p_{2,t+1}^e + n_{3,t}p_{3,t+1}^e + n_{4,t}p_{4,t+1}^e) \times (1-n_t) + p^f n_t + \bar{y} + \varepsilon_t). \quad (7.9)$$

The fractions $n_{h,t}$, $1 \leq h \leq 4$, of the four heuristics are determined by a discrete choice model with *asynchronous updating*

$$n_{i,t} = \delta n_{i,t-1} + (1-\delta) \frac{\exp(\beta U_{i,t-1})}{\sum_{i=1}^4 \exp(\beta U_{i,t-1})}, \quad (7.10)$$

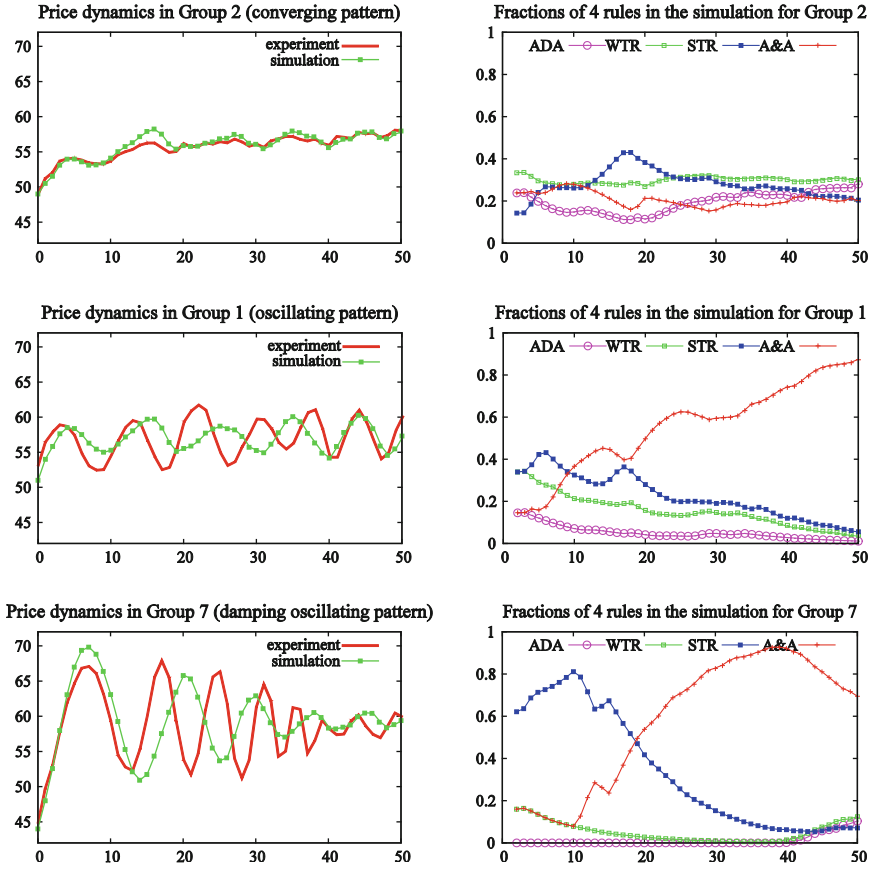


Fig. 7.6 Simulations of the heuristics switching model. Prices for laboratory experiments and evolutionary model. Fractions of four forecasting heuristics: adaptive expectations (ADA), weak trend followers (WTR), strong trend followers (STR), and anchoring adjustment heuristic (A&A). The simulations only differ in initial price forecasts and initial distribution of strategies

with the fitness measure given by (minus) squared prediction errors, that is,

$$U_{i,t-1} = -(p_{t-1} - p_{i,t-1}^e)^2 + \eta U_{i,t-2}. \tag{7.11}$$

The parameter $\eta \in [0, 1]$ represents the *memory strength* in the fitness measure, the parameter $\delta \in [0, 1]$ represents the inertia of traders' switching behaviour (in each period, only a fraction $1 - \delta$ of traders will switch strategy) and the parameter $\beta \geq 0$ is the intensity of choice as before. The fraction n_i of robot traders evolves according to (7.6), as in the experiment.

The only differences in the simulations of Fig. 7.6 are the initial price forecasts and the initial distribution over the four heuristics. Trends in realized market prices

are more likely when the initial fractions of the weak and strong trend followers are sufficiently large. Interestingly, the anchoring and adjustment heuristic is important in keeping the fluctuations alive, since in both the permanent and the dampened oscillatory cases their fractions becomes large (more than 80%). Coordination of individual forecasts on simple forecasting heuristics thus explains the three different observed aggregate market outcomes. Oscillations may be triggered by initial prices and small random shocks, are reinforced when the initial fraction of weak and strong trend heuristics is relatively large and may be sustained by the anchoring adjustment heuristic.

7.5 Concluding Remarks

We have discussed a simple theory of heterogeneous market expectations, in which bounded rationality is disciplined through simple heuristics, adaptive learning and evolutionary selection. This theory matches important stylized facts in financial market, such as excess volatility and (temporary) bubbles and crashes. In particular, coordination on trend following strategies, driven by experience based reinforcement learning, may strongly amplify a rise or decline in asset prices triggered by fundamental news. As we have seen, the theory matches for example the “dot com” bubble in stock prices in the late 1990s. The theory is also consistent with learning to forecast laboratory experiments with human subjects and explains observed path-dependent stable and unstable outcomes. In particular, laboratory experiments confirm that coordination on simple trend following strategies may occur and lead to persistent deviations from fundamental and fluctuations in asset prices.

In future work the theory should be tested in different market environments. Complexity models in economics are often based on heterogeneous expectations, and a satisfactory theory of heterogeneous expectations is therefore necessary for a successful research programme on bounded rationality, complexity, agent based economics and evolution.

References

- Anderson PW, Arrow KJ, Pines D (eds) (1988) *The economy as an evolving complex system II*. Addison-Wesley, Reading, MA
- Anufriev M, Hommes CH (2009) Evolution of market heuristics. *Knowl Eng Rev*, forthcoming
- Arthur WB, Durlauf SN, Lane DA (eds) (1997a) *The economy as an evolving complex system II*. Addison-Wesley, Reading, MA.
- Arthur WB, Holland JH, LeBaron B et al. (1997b) Asset pricing under endogenous expectations in an artificial stock market. In: Arthur W, Lane D, Durlauf S (eds) *The economy as an evolving complex system II*. Addison-Wesley, Reading, MA, pp 15–44

- Boswijk HP, Hommes CH, Manzan S (2007) Behavioral heterogeneity in stock prices. *J Econ Dyn Control* 31:1938–1970
- Brock WA, Hommes CH (1997) A rational route to randomness. *Econometrica* 65:1059–1095
- Brock WA, Hommes CH (1998) Heterogeneous beliefs and routes to chaos in a simple asset pricing model. *J Econ Dyn Control* 22:1235–1274
- Evans GW, Honkapohja S (2001) *Learning and expectations in macroeconomics*. Princeton University Press, Princeton
- Fama EF, French KR (2002) The equity premium. *J Finan* 57:637–659
- Hommes CH (2006) Heterogeneous agent models in economics and finance. In: *Handbook computational economics, vol 2: Agent-based computational economics*. Elsevier, Amsterdam, pp 1109–1186
- Hommes CH, Sonnemans J, Tuinstra J, et al. (2005) Coordination of expectations in asset pricing experiments. *Rev Financial Stud* 18:955–980
- Sargent TJ (2008) Evolution and intelligent design. *Am Econ Rev* 98:5–37
- Shiller RJ (2000) *Irrational exuberance*. Princeton University Press, Princeton
- Tesfatsion L, Judd KL (eds) (2006) *Handbook of computational economics, vol 2: Agent-based computational economics*. Elsevier Science, Amsterdam
- Tversky A, Kahneman D (1974) Judgment under uncertainty: heuristics and biases. *Science* 185:1124–1131