



UvA-DARE (Digital Academic Repository)

Evolution of market heuristics

Anufriev, M.; Hommes, C.H.

[Link to publication](#)

Citation for published version (APA):

Anufriev, M., & Hommes, C. (2007). Evolution of market heuristics. (CeNDEF Working Paper Universiteit van Amsterdam; No. 07-06). Amsterdam: Universiteit van Amsterdam.

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <http://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

Evolution of Market Heuristics*

Mikhail Anufriev Cars Hommes[†]

*CeNDEF, Department of Economics, University of Amsterdam,
Roetersstraat 11, NL-1018 WB Amsterdam, Netherlands*

Abstract

The time evolution of aggregate economic variables, such as stock prices, is affected by market expectations of individual investors. Neo-classical economic theory assumes that individuals form expectations rationally, thus enforcing prices to track economic fundamentals and leading to an efficient allocation of resources. However, laboratory experiments with human subjects have shown that individuals do not behave fully rational but instead follow simple heuristics. In laboratory markets prices may show persistent deviations from fundamentals similar to the large swings observed in real stock prices.

Here we show that evolutionary selection among simple forecasting heuristics can explain coordination of individual behavior leading to three different aggregate outcomes observed in recent laboratory market forecasting experiments: slow monotonic price convergence, oscillatory dampened price fluctuations and persistent price oscillations. In our model forecasting strategies are selected every period from a small population of plausible heuristics, such as adaptive expectations and trend following rules. Individuals adapt their strategies over time, based on the relative forecasting performance of the heuristics. As a result, the evolutionary switching mechanism exhibits path dependence and matches individual forecasting behavior as well as aggregate market outcomes in

*This work was supported by the ComplexMarkets E.U. STREP project 516446 under FP6-2003-NEST-PATH-1.

[†]Corresponding author. Tel.: +31-20-5254246; fax: +31-20-5254349; e-mail: C.Hommes@uva.nl.

the experiments. Our results are in line with recent work on agent-based models of interaction and contribute to a behavioral explanation of universal features of financial markets.

JEL codes: E37, G12, D84, C91, C92.

Keywords: Expectations feedback; Experiments; Heuristics; Evolutionary learning; Asset-pricing model.

Introduction

In social systems today's individual decisions depend upon expectations or beliefs about future developments. Think for example of the stock market, where an investor buys stocks today when she expects stock prices to rise in the future. Expectations affect individual behavior and the realized market outcome (e.g. prices and traded quantities) is an aggregation of individual behavior. A market is an *expectations feedback system*: market history shapes individual expectations which, in turn, determine current aggregate market behavior and so on. But how do individuals actually form market expectations, and what is the aggregate outcome of the interaction of individual market forecasts?

Traditional economic theory assumes that all individuals have rational expectations (1–2). In a market model, this means that forecasts coincide with mathematical expectations, conditioned upon available information. In a rational world individual expectations coincide, on average, with market realizations, and markets are efficient with prices fully reflecting economic fundamentals (3). In the traditional view, there is no room for market psychology and “irrational” herding behavior. An important underpinning of the rational approach comes from an early evolutionary argument made by Alchian (4) and Friedman (5), that “irrational” traders will not survive competition and will be driven out of the market by rational traders, who will trade against them and earn higher profits.

However, starting from Simon (6), many economists argue that rationality imposes unrealistically strong informational and computational requirements upon individual behavior and it is more reasonable to model individuals as *boundedly rational*, using simple rules of thumb in decision making. Laboratory experiments indeed showed that individual decisions under uncertainty are at odds with perfectly rational behavior, and can be much better described by simple *heuristics*, which sometimes may lead to persistent biases (7–9). Models of bounded rationality have also been applied to forecasting behavior, and several adaptive learning algorithms have been proposed to describe market expectations. For example, Sargent (10) and Evans and Honkapohja (11) advocate the use of different forms of bounded rationality in the modeling of expectations and decision making in macroeconomics, while Erev and Roth

(12) propose reinforcement learning as an explanation of average behavior in a number of experiments in a game-theoretical setting. It is interesting that in some models (13) adaptive learning enforces convergence to rational expectations, while in others (14–15) learning may not converge at all but instead lead to excess volatility and persistent deviations from rational equilibrium similar to real markets (16–17).

Laboratory experiments with human subjects and controlled economic fundamentals are well suited to study how individuals form expectations and how their interaction shapes aggregate market behavior (18–19). But the results from laboratory experiments are also mixed. Early experiments, with various market designs such as double auction trading, show convergence to equilibrium (20–21), while more recent asset pricing experiments exhibit deviations from equilibrium with persistent bubbles and crashes (22–23). A clear explanation of these different market phenomena is still lacking (24). 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.

In recent learning to forecast experiments, described at length in Hommes et al. (23), qualitatively different aggregate outcomes have been observed in the same experimental setting. In a stationary environment participants, during a number of periods, had to predict the price of a risky asset (say a stock) having knowledge of the fundamental parameters (mean dividend and interest rate) and previous price realizations, but without knowing the forecasts of others. If all agents would behave rationally or learn to behave rationally, the market price would quickly converge to a constant fundamental value $p^f = 60$. While in some groups convergence did happen, in other groups prices persistently fluctuated (see Fig. 1 *Left*). What was even more striking, is that in all groups individuals were able to coordinate their forecasts (see Fig. 2 *Left*). In this paper we present a simple model based on *evolutionary selection of simple heuristics* explaining how coordination of individual forecasts can emerge and, ultimately, enforce the different aggregate market outcomes. To our best knowledge, this is the first learning model explaining *different* time series patterns in the *same* laboratory experiments. The only differences between the model simulations in Figs. 1 and 2 are the initial prices and the initial

distribution over the heuristics.

Laboratory experiment

A number of sessions of a computerized learning to forecast experiment have been performed in the CREED laboratory at the University of Amsterdam; see (23) for a detailed description. In each session of the experiment six participants had to predict the price of an asset for 51 periods and have been rewarded according to the accuracy of their predictions. The participants were told that they are advisors to a pension fund and that this pension fund can invest money either in a risk free asset with real interest rate r per period or in shares of an infinitely lived risky asset. In each period the risky asset pays uncertain dividend which is a random variable, independent identically distributed (IID), with mean \bar{y} . The price of the risky asset, p_t , is determined by a market clearing equation on the basis of the investment strategies of the pension fund. The exact functional form of the strategies and the equilibrium equation were unknown to the participants, but they were informed that the higher their forecast is, the larger will be the demand for the risky asset of the pension fund. Participants also knew the values of the parameters $r = 0.05$ and $\bar{y} = 3$, and therefore had enough information to compute the *rational fundamental price* (i.e. the discounted sum of the expected future dividend stream) of the risky asset $p^f = \bar{y}/r = 60$.

Every session of the experiment lasted 51 periods. In every period each of the 6 participants provided a two period ahead forecast for the price of the risky asset, given the available information. This information consisted of past prices (up to two lags) of the risky asset and own past predictions (up to one lag) made by the participant. Participants did not know the predictions of other participants, neither did they know exactly how their own forecast affected the equilibrium price. When all 6 predictions for the price in period $t + 1$ have been submitted, the current market clearing price was computed according to a standard model of asset pricing, see e.g. (25):

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

where \bar{p}_{t+1}^e denotes an (equally weighted) average of the 6 individual forecasts, $r(= 0.05)$ is the risk free interest rate, $\bar{y}(= 3)$ is the mean dividend, ε_t is a stochastic term representing small demand and supply shocks, and n_t stands for a small fraction of “robot” traders who always submit a fundamental forecast p^f . These robot traders were introduced as a “stabilizing force” in the experiment to prevent the occurrence of large bubbles. The fraction of robot traders increased as the price moved away from its fundamental equilibrium level:

$$n_t = 1 - \exp\left(-\frac{1}{200} |p_{t-1} - p^f|\right). \quad (2)$$

This mechanism reflects the feature that in real markets there is more agreement about over- or undervaluation when the deviation from the fundamental is large. At the end of each period every participant was informed about the realized price and her earnings were defined by a quadratic scoring rule

$$e_{t,h} = \max\left(1300 - \frac{1300}{49}(p_t - p_{t,h}^e)^2, 0\right). \quad (3)$$

There were 7 sessions of the experiment. The stochastic shocks ε_t were the same in all sessions (normally distributed, with mean 0 and standard deviation 0.5).

Findings of the experiment

The main findings of the experiment are as follows. First, realized asset prices differed significantly from the rational fundamental price in all sessions. Comparison of the experiment with prediction of the rational expectations model shows that on average the asset was undervalued. Furthermore, prices exhibited excess volatility, with much larger swings than the rational expectations model.

Second, three different price patterns were observed, see Fig. 1 *Left*. In group 2 (*Top*) and group 5 (not shown) the price of the asset slowly converged, almost monotonically, to the fundamental price. In group 4 (not shown) and group 7 (*Middle*) large initial fluctuations were observed, dampening slowly towards the end of the experiment. In group 1 (*Bottom*) and group 6 (not shown) the price oscillates around the fundamental price with an (almost)

constant amplitude. (Price dynamics in group 3 (not shown) was more difficult to classify, somewhere between oscillations and convergence).

Third, analysis of the individual price predictions reveals that during each session the participants were able to coordinate on a common prediction strategy, as illustrated in Fig. 2 *Left*. Finally, estimation of the individual predictions (based on the last 40 observations, to allow for a short learning phase) showed that participants had a tendency to use simple, linear forecasting rules, such as naive expectations (i.e. the forecast is simply the last observed price) and adaptive expectations (a weighted average of the last observed price and the last forecast). Many participants only used the two most recently observed prices, for example in a simple linear trend extrapolation forecasting rule.

Methods

In our simulation model agents will select rules from a population of simple forecasting rules or heuristics. The choice of heuristics will be governed by an *evolutionary selection* mechanism, based on the principle that more successful strategies will attract more followers. Strategy performance is measured by accumulated (negative) squared prediction errors, in line with the payment incentives in the laboratory experiments.

Forecasting Heuristics

To keep our model as simple as possible, but rich enough to explain the different observed price patterns, we have chosen only 4 heuristics which are intuitively simple and were among the rules estimated on the individual forecasts in the experiment. A behavioral interpretation underlies each heuristic. The first heuristic is an *adaptive expectations* (ADA) rule, using a weighted average between the last observed market price and the last individual forecast. Note that at the moment when forecasts of price p_{t+1} are submitted, price p_t is still unknown (see Eq. 1) and the last observed price is p_{t-1} . At the same time, the last own forecast $p_{t,1}^e$ is

known when forecasting p_{t+1} . We have chosen the following ADA rule:

$$p_{t+1,1}^e = 0.65 p_{t-1} + 0.35 p_{t,1}^e. \quad (4)$$

The second and third heuristics are *trend following* rules extrapolating a weak or a strong trend respectively. They simply predict the last observed *price level* plus a multiple of the last observed *price change*, and only differ in the magnitude of the extrapolation factor. In the case of *weak trend rule* (WTR) the factor is small and equal to 0.4, so that the rule is

$$p_{t+1,2}^e = p_{t-1} + 0.4 (p_{t-1} - p_{t-2}). \quad (5)$$

The *strong trend rule* (STR) has a larger extrapolation factor 1.3 and is given by

$$p_{t+1,3}^e = p_{t-1} + 1.3 (p_{t-1} - p_{t-2}). \quad (6)$$

The fourth heuristic is slightly more complicated. It combines an average prediction of the last observed price and an estimate of the long run equilibrium price level with an extrapolation of the last price change. More precisely, the rule is given by

$$p_{t+1,4}^e = 0.5 (p_{t-1}^{av} + p_{t-1}) + (p_{t-1} - p_{t-2}), \quad (7)$$

where p_{t-1}^{av} is the sample average of all past prices, that is, $p_{t-1}^{av} = \sum_{j=0}^{t-1} p_j$. This rule is an *anchoring and adjustment heuristic* (A&A), since it uses a (time varying) *anchor* or reference point, $0.5 (p_{t-1}^{av} + p_{t-1})$, defined as an (equally weighted) average between the last observed price and the sample mean of all past prices, and extrapolates the last price change ($p_{t-1} - p_{t-2}$) from there. Tversky and Kahneman (7) have shown that people often rely on such anchoring and adjustment heuristics. The A&A rule has been obtained from a related linear forecasting rule $p^e = 30 + 1.5 p_{t-1} - p_{t-2} = 0.5(p^f + p_{t-1}) + (p_{t-1} - p_{t-2})$, used by some individuals in the experiment. In the experiment however, subjects did *not* know the fundamental price p^f explicitly, but were able to learn an anchor $0.5(p^f + p_{t-1})$ and extrapolate price changes from there. Therefore, we replaced p^f in the rule by a proxy, given by the observed sample average of prices.

The first three rules (Eqs. 4–6) 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 heuristic (Eq. 7) combines adaptive learning of the price level and trend extrapolation.

Evolutionary Switching

Which forecasting heuristics from the population should agents choose? Our simulation model is based upon evolutionary switching between the four forecasting heuristics, driven by the past relative performance of the heuristics. 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 , $1 \leq h \leq 4$, up to (and including) time period t is given by

$$U_{t,h} = -(p_t - p_{t,h}^e)^2 + \eta U_{t-1,h}. \quad (8)$$

The parameter $0 \leq \eta \leq 1$ measures 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 rule h is updated according to a *discrete choice model with asynchronous updating* (26-27)

$$n_{t,h} = \delta n_{t-1,h} + (1 - \delta) \frac{\exp(\beta U_{t-1,h})}{Z_{t-1}}, \quad (9)$$

where $Z_{t-1} = \sum_{h=1}^4 \exp(\beta U_{t-1,h})$ is a normalization factor. There are two important parameters in Eq. 9. 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 updates their rule according to the well known *discrete choice model* used for example in Brock and Hommes (15). The parameter $\beta \geq 0$ 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 Eq. 9 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.

In the *evolutionary heuristics switching model* the price p_t in period t is computed as

$$p_t = \frac{1}{1+r} \left((1-n_t)(n_{t,1}p_{t+1,1}^e + n_{t,2}p_{t+1,2}^e + n_{t,3}p_{t+1,3}^e + n_{t,4}p_{t+1,4}^e) + n_t p^f + \bar{y} + \varepsilon_t \right), \quad (10)$$

where $p_{t+1,1}^e, \dots, p_{t+1,4}^e$ are the predictions for period $t+1$ according to the 4 heuristics in Eqs. 4–7, $n_{t,1}, \dots, n_{t,4}$ are the fractions using these heuristics described by Eqs. 8–9, n_t stands for a small fraction of “robot” traders described by Eq. 2, $r(= 0.05)$ is the risk free interest rate, $\bar{y}(= 3)$ is the mean dividend, $p^f(= 60)$ is the fundamental price and ε_t is the stochastic term representing small demand and supply shocks (taken to be the same as in the experiment).

Initialization of the model

The model is initialized by two initial prices, p_0 and p_1 , and initial weights $n_{1,h}$, $1 \leq h \leq 4$ (summing to 1; the initial share of robot traders $n_1 = 0$). Given p_0 and p_1 , the heuristics forecasts can be computed and, using the initial weights of the heuristics, the price p_2 can be computed. In the next period, the forecasts of the heuristics are updated, the fraction of “robot” traders is computed, while the same initial weights $n_{1,h}$ for individual rules are used (past performance is not well defined yet in period 3). The price p_3 is computed and the initialization stage is finished. Starting from period 4 the evolution according to Eq. 10 is well defined: first the performance measure in Eq. 8 is updated, then, the new weights of the heuristics are computed according to Eq. 9 and finally a new price is determined by Eq. 10.

Results and Discussion

Our evolutionary selection mechanism contains three parameters, β , η and δ , measuring respectively (i) how sensitive individuals are with respect to differences in strategy performance, (ii) how much relative weight they give to the most recent errors, and (iii) how strongly an individual sticks to her previous strategy. We have performed numerous simulations and found that the *path-dependence* feature of the model, in particular the capability to produce both

persistent oscillating and converging patterns, remains valid for a large range of parameters. Qualitatively, the simulation results are robust with respect to the parameters, but some quantitative features, such as the speed of convergence, the amplitude and frequency of oscillations and the stability of long run equilibrium, may change when parameters are varied.

To stress the path-dependence, for all reported simulations we have fixed the parameter values at $\beta = 0.4$, $\eta = 0.7$, $\delta = 0.9$. The simulations thus only differ in initial conditions, that is, in the two initial prices $\{p_0, p_1\}$ and in the initial distribution of agents over the population of heuristics, i.e. initial weights $\{n_{1,1}, n_{1,2}, n_{1,3}, n_{1,4}\}$. For the three simulations in Fig. 1 the initial conditions are as follows:

- for group 2 with monotonic convergence: initial prices: $p_0 = 49, p_1 = 50.5$; initial shares $n_{1,1} = n_{1,4} = 0.25$, $n_{1,2} = 0.35$ and $n_{1,3} = 0.15$;
- for group 7 with dampened oscillations: initial prices: $p_0 = 44, p_1 = 48$; initial shares $n_{1,1} = 0$, $n_{1,2} = n_{1,4} = 0.17$ and $n_{1,3} = 0.66$;
- for group 1 with persistent oscillations: initial prices: $p_0 = 51, p_1 = 54$; initial shares $n_{1,1} = n_{1,4} = 0.15$, $n_{1,2} = n_{1,3} = 0.35$.

Fig. 1 shows realized prices (*Left*) for both the experiments and the heuristics switching model, as well as the shares (*Right*) of the 4 heuristics. Fig. 2 shows the individual forecasts in the experiments (*Left*) as well as the forecasts of the 4 heuristics (*Right*). Similar to the experiments, in the simulation model coordination of forecasts arises.

Fig. 1 *Left* shows that the heuristics switching model closely matches all three different patterns, slow monotonic convergence to the fundamental price, dampened oscillatory price movements and persistent price oscillations, in the laboratory experiments. The model 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. Fig. 1 *Right* plots the corresponding transition paths of the fractions of each of the four forecasting heuristics. In the case of monotonic convergence (*Top*), the four fractions (and the individual forecasts) remain relatively close together during the simulation

causing slow (almost) monotonic convergence of the price to the fundamental equilibrium 60. The increase in price causes a temporary domination of the dynamics by the strong trend heuristic between periods 13 and 23. However, this rule overestimates the price trend so that, ultimately, the adaptive heuristic takes the lead, and price converges to fundamental. In the second simulation (*Middle*), 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 however, the fraction of strong trend followers decreases, while the fraction of the fourth rule, the anchoring and adjustment heuristic, rises to more than 80% after 30 periods. The flexible anchoring and adjustment heuristic forecasts better than the static strong trend following rule, which overestimates the price trend. After 40 periods the fraction of the anchoring adjustment heuristic starts slowly decreasing, and consequently the price oscillations slowly stabilize. In the third simulation (*Bottom*) weak and strong trend followers each represent 35% 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 anchoring adjustment heuristic starts to increase, because once again it predicts better than the static strong and weak trend followers, who either overestimate or underestimate the price trend. The fraction of the anchoring adjustment heuristic gradually increases and dominates the market within 10 periods, rising to more than 70% 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, path dependent aggregate market outcomes. This explanation is consistent with recent work on agent-based models of interaction explaining emergent phenomena in financial markets, including fat tails, clustered volatility, temporary bubbles and crashes and scaling laws (28–32). In real markets small price movements triggered by random news about market fundamentals may be reinforced by trend following strategies causing excessively volatile markets. Evolution of market heuristics along the lines sketched here provides a simple, universal, behavioral

explanation of these emergent phenomena.

Materials

The simulation program for the model described in this paper together with brief documentation and configuration settings used for the reported simulations is freely available at <http://www.cafed.eu/evexex>.

References

1. Muth JF (1961) *Econometrica* 29:315-335.
2. Lucas RE, Prescott EC (1972) *Econometrica* 39:659-681.
3. Fama EF (1970) *J Finance* 25:383-417.
4. Alchian AA (1950) *J Pol Econ* 58:211-221.
5. Friedman M (1953) *Essays in Positive Economics* (Univ Chicago Press, Chicago).
6. Simon HA (1957) *Models of Man* (Wiley, New York).
7. Tversky A, Kahneman D (1974) *Science* 185:1124-1131.
8. Kahneman D (2003) *Amer Econ Rev* 93:1449-1475.
9. Camerer CF, Fehr E (2006) *Science* 311:47-52.
10. Sargent TJ (1993) *Bounded Rationality in Macroeconomics* (Clarendon Press, Oxford).
11. Evans GW, Honkapohja S (2001) *Learning and Expectations in Macroeconomics* (Princeton Univ Press, Princeton)
12. Erev I, Roth AE (1999) *Amer Econ Rev* 88:848-881.
13. Bray MM, Savin NE (1986) *Econometrica* 54:1129-1160.
14. Bullard J (1994) *J Econ Th* 64:468-485.
15. Brock WA, Hommes CH (1997) *Econometrica* 65:1059-1095.
16. Shiller RJ (1981) *Amer Econ Rev* 71: 421-436.
17. De Bondt WFM, Thaler R (1985) *J Finance* 40:793-805.
18. Marimon R, Spear SE, Sunder S (1993) *J Econ Th* 61:74-103.
19. Peterson SP (1993) *J Econ Beh and Organ* 22:269-284.

20. Smith VL (1962) *J Pol Econ* 70:111-137.
21. Plot CR, Sunder S (1982) *J Pol Econ* 90:663-698.
22. Smith VL, Suchanek GL, Williams AW (1988) *Econometrica* 56:1119-1151.
23. Hommes CH, Sonnemans JH, Tuinstra J, van de Velden H (2005) *Rev Fin Stud* 18:955-980.
24. Duffy J (in press) in *New Palgrave Dictionary of Economics*, eds. Blume L, Durlauf S (Macmillan).
25. Brock WA, Hommes CH (1998) *J Econ Dyn Ctr* 22:1235-1274.
26. Hommes CH, Huang H, Wang D (2005) *J Econ Dyn Ctr* 29:1043-1072.
27. Diks C, van der Weide R (2005) *J Econ Dyn Ctr* 29:741- 763.
28. Lux T, Marchesi M (1999) *Nature* 397:498-500.
29. Mantegna RN, Stanley HE (1995) *Nature* 376:46-49.
30. Farmer JD, Lo AW (1999) *Proc. Natl. Acad. Sci.* 96:9991-9992.
31. Hommes CH (2006) in *Handbook of Computational Economics, Volume 2: Agent-Based Computational Economics*, eds. Tesfatsion L, Judd KL (Elsevier Science), pp 1109-1186.
32. LeBaron B (2006) in *Handbook of Computational Economics, Volume 2: Agent-Based Computational Economics*, eds. Tesfatsion L, Judd KL (Elsevier Science), pp 1187-1233.

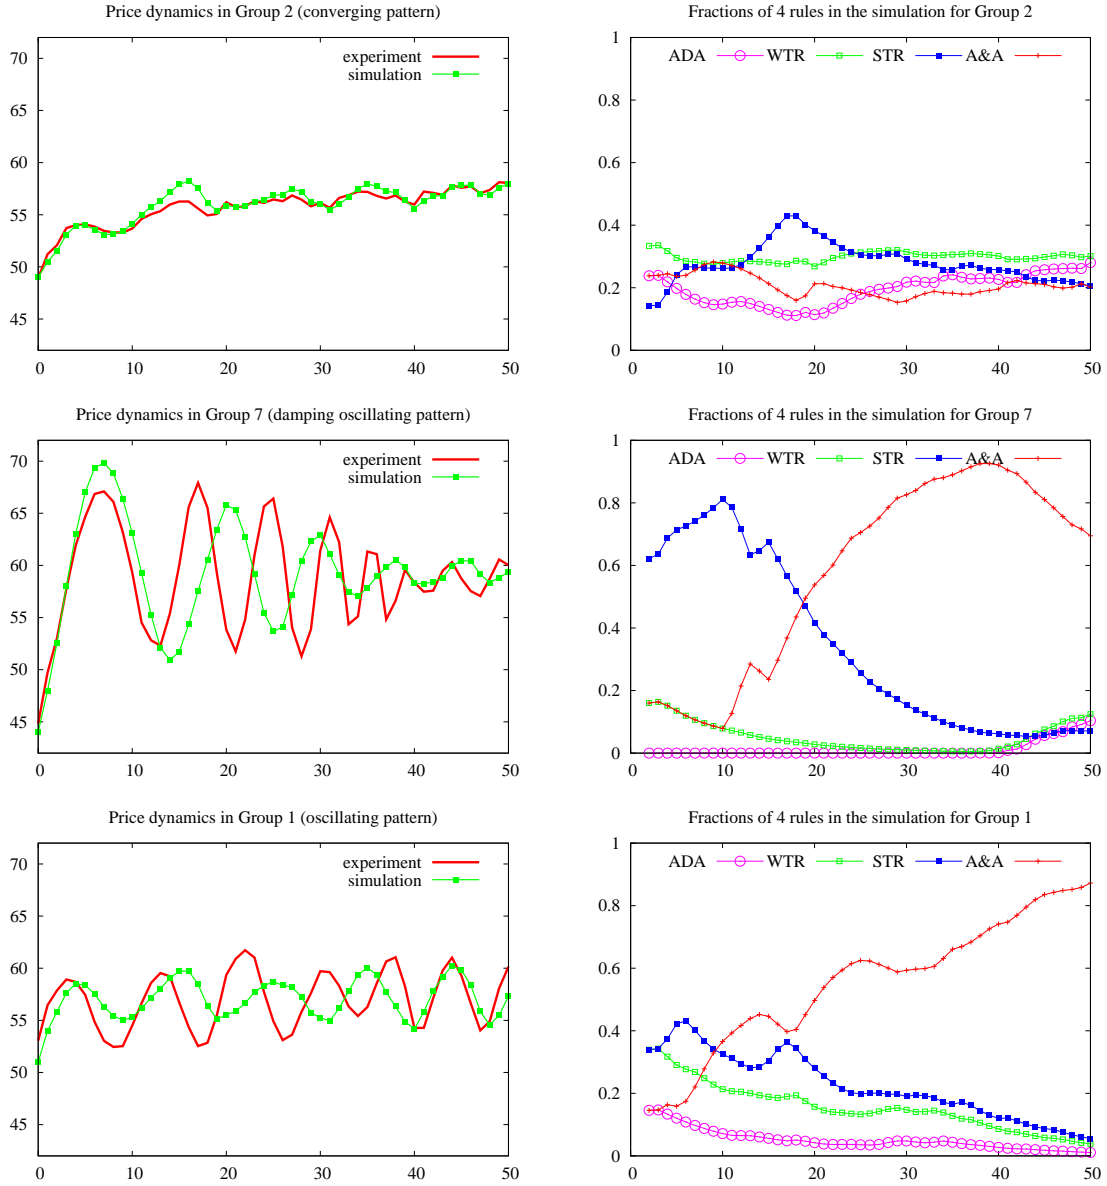


Figure 1: **Laboratory experiments and heuristics switching model simulations.** Prices (*Left*) for laboratory experiments (red) and evolutionary model (green). Fractions (*Right*) of four forecasting heuristics: adaptive expectations (ADA, purple), weak trend followers (WTR, green), strong trend followers (STR, blue) and anchoring adjustment heuristic (A&A, red). Coordination of individual forecasts explains three different aggregate market outcomes: monotonic convergence to equilibrium (*Top*), oscillatory convergence (*Middle*) and permanent oscillations (*Bottom*). Oscillations may be triggered by initial prices, are reinforced when the initial fraction of weak and strong trend heuristics is relatively large and may be sustained by the anchoring adjustment heuristic.




Figure 2: **Coordination in laboratory experiments and model simulations.** Individual predictions of 6 participants in the laboratory experiments (*Left*) and predictions of 4 forecasting heuristics in evolutionary heuristics switching model (*Right*). Heuristics are: adaptive expectations (ADA, purple), weak trend followers (WTR, green), strong trend followers (STR, blue) and anchoring adjustment heuristic (A&A, red). Coordination of individual forecasts arises both in the experiment and in the simulation model in all observed aggregate outcomes: monotonic convergence to equilibrium (*Top*), oscillatory convergence (*Middle*) and permanent oscillations (*Bottom*).