

Evolutionary Selection of Individual Expectations and Aggregate Outcomes*

Mikhail Anufriev[†] Cars Hommes[‡]

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

First Version: July 2007
This Version: October 2008

This is PRELIMINARY version of the paper.
Please do not quote it without the authors' permission.

Abstract

We construct a model of evolutionary or reinforcement learning, where boundedly rational agents choose from a few simple rules for price prediction, such as naive, adaptive or trend following expectations. Agents update their active rule by evolutionary selection based upon forecasting errors. Simulations show that after some initial learning phase, coordination on a common rule occurs. Which rule survives evolutionary selection depends on the initial conditions, particularly on the price pattern in the first few periods and the initial shares of agents attached to the rules. Consequently, evolutionary learning exhibits path dependence and different patterns of realized prices are generated, explaining the results of the recent experiment on expectations formation in a standard asset pricing setting (Hommes, Sonnemans, Tuinstra and Van de Velden, 2005). Tuning the parameters, these patterns can be made both qualitatively and quantitatively close to those observed in the experiments. We thus provide an explanation of the experimental findings using a low dimensional nonlinear deterministic model with few parameters.

Keywords: Learning, Heterogeneous Expectations, Expectations Feedback, Experimental Economics

*We thank the participants of the “Computation in Economics and Finance” conference, Montreal, July 2007, the “Complexity in Economics and Finance” workshop, Leiden, October 2007 and the “Learning and Monetary Policy??” workshop, Cambridge, September 2008 for stimulating discussions and helpful comments on earlier versions of this paper. This work was supported by the ComplexMarkets E.U. STREP project 516446 under FP6-2003-NEST-PATH-1.

[†]Tel.: +31-20-5254248; fax: +31-20-5254349; e-mail: M.Anufriev@uva.nl.

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

1 Introduction

In social systems today's individual decisions crucially depend upon expectations or beliefs about future developments. Think for example of the stock market, where an investor buys (sells) stocks today when she expects stock prices to rise (fall) 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 (Muth, 1961; Lucas and Prescott, 1971). 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 (Fama, 1970). 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 (1950) and Friedman (1953), 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, following Simon (1957), 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 have shown 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 (Tversky and Kahneman, 1974; Kahneman, 2003; Camerer and Fehr, 2006). 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 (1993) and Evans and Honkapohja (2001) advocate the use of *adaptive learning* in modeling expectations and decision making in macroeconomics, while Arthur (1991) and Erev and Roth (1998) 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 (Bray and Savin, 1986) adaptive learning enforces convergence to rational expectations, while in others (Bullard, 1994) learning may not converge at all but instead lead to excess volatility and persistent deviations from rational equilibrium similar to real markets (Shiller, 1981; De Bondt and Thaler, 1989). Recently, models with *heterogeneous expectations* and evolutionary selection among the forecasting rules have been proposed, e.g. Brock and Hommes (1997) and ?, see Hommes (2006) for an extensive overview.

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 (Marimon, Spear, and Sunder, 1993; Peterson, 1993). But the results from laboratory experiments are mixed. Early experiments, with various market designs such as double auction trading, show convergence to equilibrium (Smith, 1962), while more recent asset pricing experiments exhibit deviations from equilibrium with persistent bubbles and crashes (Smith, Suchanek, and Williams, 1988; Hommes, Sonnemans, Tuinstra, and Velden, 2005). A clear explanation of these different market phenomena is still lacking (Duffy, 2008). It is particularly challenging to provide a general *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, Sonnemans, Tuinstra, and Velden (2005), three qualitatively different aggregate outcomes have been observed in the same experimental setting. In a stationary environment participants, for 50 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 in the laboratory price convergence did occur, in other groups prices persistently fluctuate (see Fig. 2, upper parts of different panels). Another striking finding in the experiments is that in all groups individuals were able to coordinate on a common predictor (see Fig. 2, lower parts of different panels). The main purpose of this paper is to 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. Although our model is very simple it fits the experimental data surprisingly well (see e.g Fig. 11).

The paper is organized as follows. In Section 2 we review the findings of the laboratory experiment and we look at individual forecasting rules which will form the basis of our evolutionary model. Section 3 is devoted to the analysis of implied price dynamics under *homogeneous* forecasting rules which were identified in the experiment. A learning model based on evolutionary selection between simple forecasting heuristics is presented, analyzed and simulated in Section 4. Finally, Section 5 concludes.

2 Learning to Forecast Experiments

In this section we discuss the laboratory experiments. Subsection 2.1 recalls the experimental design, Subsection 2.2 focusses on aggregate price behaviour, while Subsection 2.3 discusses individual prediction rules.

2.1 Experimental Design

A number of sessions of a computerized learning to forecast experiment in the CREED laboratory at the University of Amsterdam have been presented in Hommes, Sonnemans, Tuinstra, and Velden (2005), henceforth HSTV. In each session human subjects had to predict the price of an asset for 51 periods and have been rewarded for the accuracy of their predictions. Fig. 2 shows the result of the experiment for six different groups. The reader can immediately recognize two striking results of the experiment: different qualitative patterns in aggregate price behavior and high coordination of individual forecasts, even though individuals do not know the forecasts of others. Before starting to develop an explanation for these findings, we briefly describe the experimental design.

Each market consists of six participants, who were told that they are advisors to a pension fund and that this pension fund can invest money either in a risk-free or in a risky asset. In each period the risk-free asset pays a fixed interest rate $r (= 0.05)$, while the risky asset pays stochastic dividends, independently identically distributed (IID), with mean $\bar{y} (= 3)$. Trading in the risky asset had been computerized, using a demand schedule derived from mean-variance maximization, given the subject's individual forecast. Hence, subject's only task in every period was to give a two period ahead point prediction for the price of the

risky asset, and their earnings were inversely related to their prediction errors. An advantage of this approach is that it provides clean data on expectations, which can be used to test e.g. the rational expectations hypothesis or adaptive learning behavior models, see e.g. the discussion in Duffy (2008). Participants knew that the actual price realization of the risky asset is determined by market equilibrium 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. However, they were informed that the higher their own forecast is, the larger will be the demand for the risky asset. Stated differently, they knew that their was positive feedback from individual price forecasts to the realized market price. They were also aware that, ultimately, the demand also depends on the forecasts of other participants, but they did not know their number nor their identity.

More formally the session of the experiment can be presented as follows. At the beginning of every period $t = 0, \dots, 50$ every participant $i = 1, \dots, 6$ provides a forecast for the price of the risky asset in the next period, p_{t+1} , given the available information. An individual forecast, $p_{i,t+1}^e$, can be any number (with two decimals) between 0 and 100. The information set $\mathbb{I}_{i,t}$, at date t , consists of past prices, past own predictions¹, past earnings and the fundamental parameters (the risk-free interest rate ($r = 0.05$) and the dividend mean (\bar{y}):

$$\mathbb{I}_{i,t} = \{p_0, \dots, p_{t-1}; p_{i,0}^e, \dots, p_{i,t}^e; e_{i,0}, \dots, e_{i,t-1}; r, \bar{y}\}. \quad (2.1)$$

Note that, since the price p_t is unknown at the beginning of period t , it is not included into the information set. The same holds for the earnings $e_{i,t}$ in period t , which will depend on the price p_t . Notice also that participants can, in principle, compute the *rational fundamental price* of the risky asset, $p^f = \bar{y}/r = 60$, given by the discounted sum of the expected future dividend stream.

The market clearing price was computed according to a standard mean-variance asset pricing model (Campbell, Lo, and MacKinlay, 1997; Brock and Hommes, 1998):

$$p_t = \frac{1}{1+r} \left((1-n_t) \bar{p}_{t+1}^e + n_t p^f + \bar{y} + \varepsilon_t \right). \quad (2.2)$$

The market price at date t depends on the average of individual predictions, $\bar{p}_{t+1}^e = \sum_i p_{i,t+1}^e / 6$, and the fundamental forecast p^f given by small fraction n_t of “robot” traders. It is also affected by a small stochastic term ε_t , representing e.g. demand or supply shocks. The robot traders were introduced in the experiment as a far from equilibrium stabilizing force to prevent the occurrence of large bubbles. The fraction of robot traders increased in response to the deviations of the last price from its fundamental level:

$$n_t = 1 - \exp \left(- \frac{1}{200} |p_{t-1} - p^f| \right). \quad (2.3)$$

This mechanism reflects the feature that in real markets there is more agreement about over- or undervaluation of an asset when the price deviation from the fundamental level is large².

At the end of the period every participant was informed about the realized price p_t . The earnings were determined by a quadratic scoring rule

$$e_{i,t} = \max \left(1300 - \frac{1300}{49} (p_t - p_{i,t}^e)^2, 0 \right). \quad (2.4)$$

¹Past prices and predictions were visualized on the computer screen both in a graph and a table.

²In the experiments the fraction of robot trader became never larger than $n_t \leq 0.2x$. Recently, ? ran experiments without robot traders, and observed long lasting bubbles.

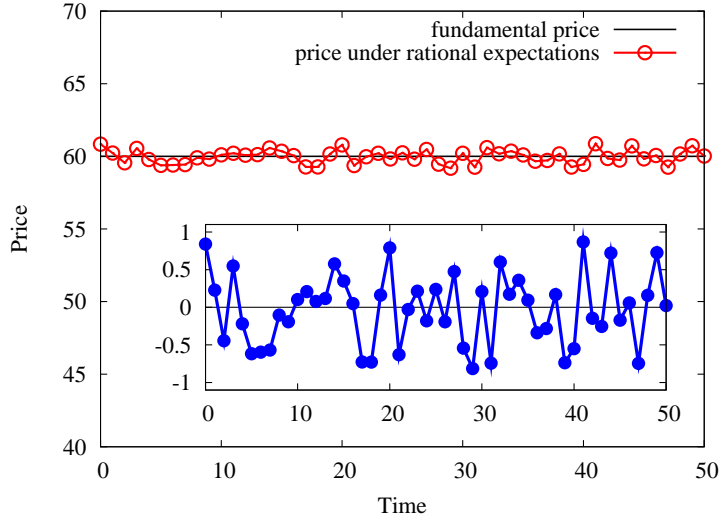


Figure 1: **Price evolution and prediction errors (inner frame) under Rational Fundamental Expectations.** When every participant predicts p^f the realized price (red) fluctuates around fundamental level. Small, non-systematic prediction errors (blue) are due to stochastic shocks.

There were seven sessions of the experiment, each with the same realizations of the stochastic shocks ε_t drawn independently from a normal distribution with mean 0 and standard deviation 0.5. The same stochastic process $\{\varepsilon_t\}_{t=0}^{50}$ will be used in our simulations.

Fig. 1 shows the simulation of realized prices, which would occur when all individuals use the rational, fundamental forecasting rule, $p_{i,t+1}^e = p^f$, for all i and t . Under rational expectations the realized price $p_t = \varepsilon_t / (1 + r)$ randomly fluctuates around the fundamental level $p^f = \bar{y} / r = 60$ with small amplitude. In the experiment, one can not expect rational behavior at the outset, but aggregate prices might converge to their fundamental value through individual learning.

2.2 Aggregate price behavior

Fig. 2 shows time series of prices, individual predictions and forecasting errors in six different sessions of the experiment. A striking feature of aggregate price behavior is that three different qualitative patterns emerge. The prices in groups 2 and 5 converge slowly and almost monotonically to the fundamental price level. In groups 1 and 6 persistent oscillations are observed during the entire experiment. In groups 4 and 7 prices are also fluctuating but their amplitude is decreasing.³

A second striking result concerns individual predictions. In all groups participants were able to coordinate their forecasting activity. The forecasts, as shown in the lower parts of the panels in Fig. 2, are dispersed in the first periods but then become very close to each other in all groups.⁴ The coordination of individual forecasts has been achieved in the absence of

³Price dynamics in group 3 (not shown) is more difficult to classify. Similar to group 1 it started with moderate oscillations, then stabilized at a level below the fundamental, suddenly falling in period $t = 40$, probably due to a typing error of one of the participants.

⁴To quantify the degree of coordination, HSTV analyze the average prediction error over time and across the six participants for each group. It turns out that this error is explained more by the “common” prediction error (measured as the deviation of the average prediction from the realized price), than by the dispersion

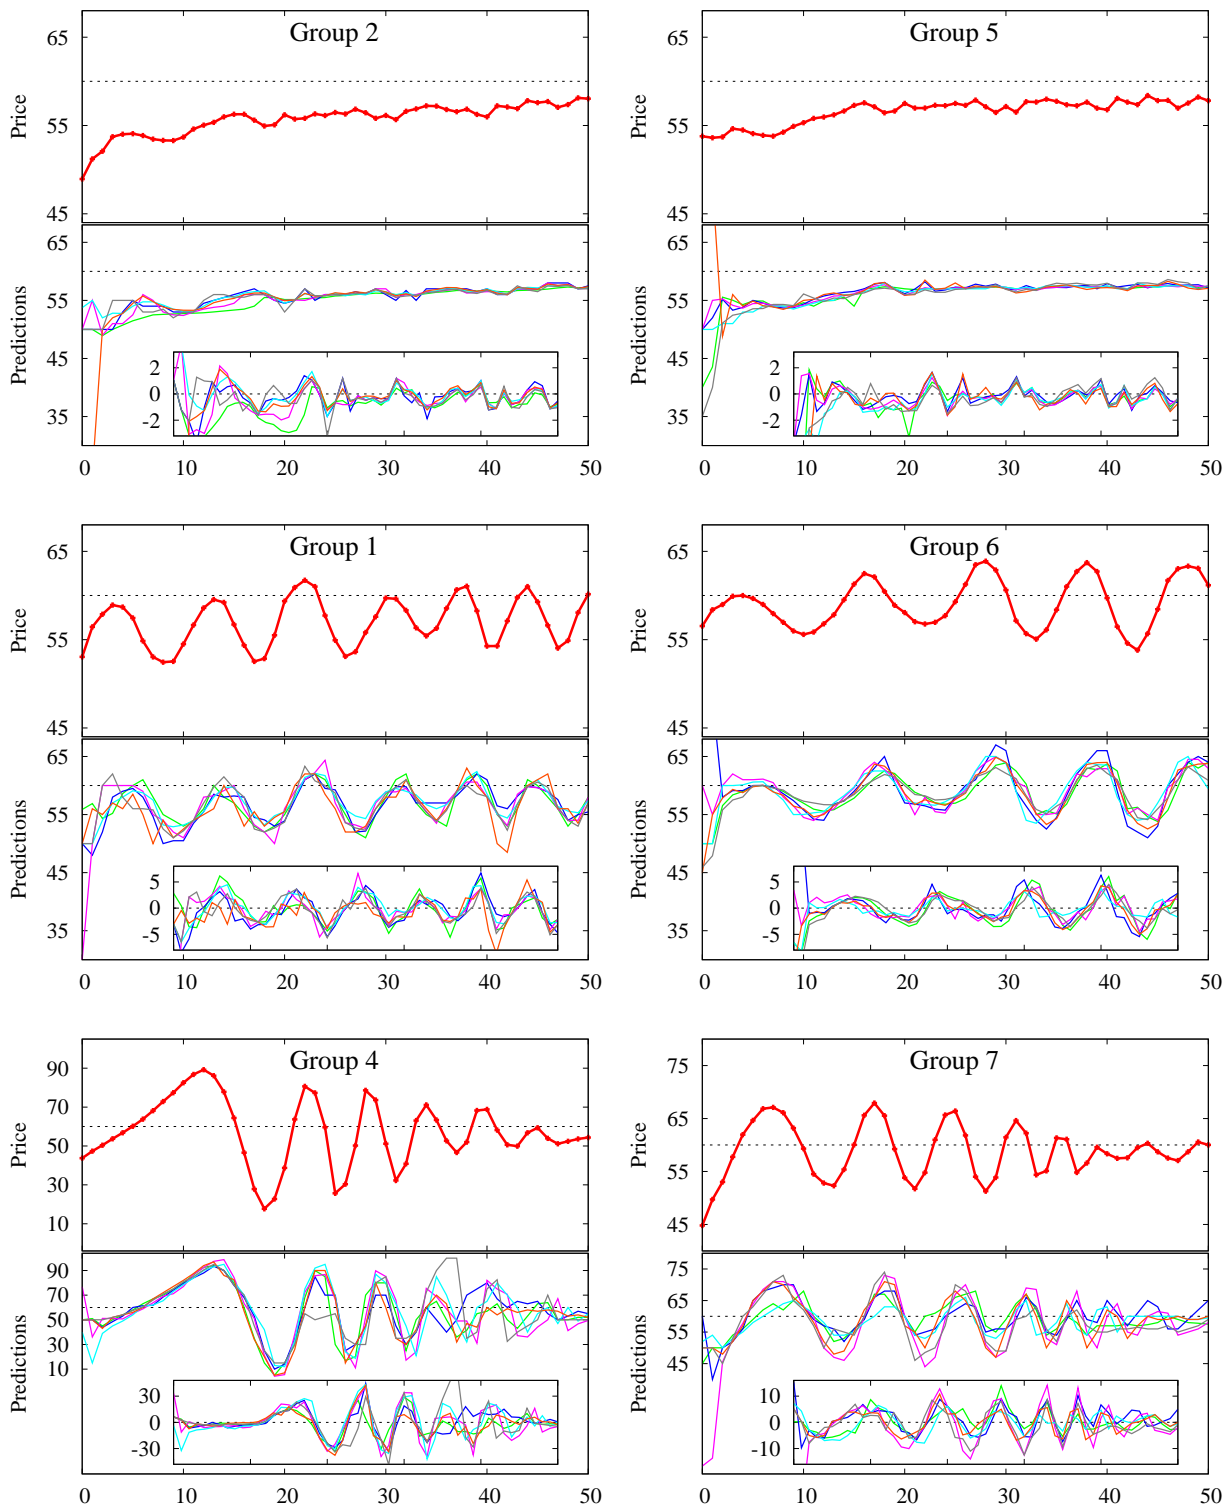


Figure 2: **Results of the forecasting experiment in different sessions.** Upper panels show the groups with monotonic convergence, **middle panels** show the groups with constant oscillations, **lower panels** show the groups with damping oscillations. For every panel **upper part** shows observed prices (red) in comparison with fundamental level (black); **lower part** shows individual predictions; and **inner frame** shows individual forecasting errors.

between individual predictions. Even in the groups 4 and 7 with the lowest coordination, the dispersion between individual predictions accounts only for 29% and 34%, respectively, of the average total prediction error.

any communication between subjects and knowledge of past and present predictions of other participants.

To summarize, the HSTV learning to forecast experiment we have the following:

- participants were unable to learn the rational, fundamental forecasting rule; only in some cases individual predictions moved (slowly) in the direction of the fundamental price towards the end of the experiment;
- three different price patterns were observed: (i) slow, (almost) monotonic convergence, (ii) persistent price oscillations with almost constant amplitude, and (iii) large initial oscillations dampening slowly towards the end of the experiment;
- already after a short transient, participants were able to coordinate their forecasting activity, submitting in every period similar forecasts.

The purpose of this paper is to explain these findings simultaneously by a simple model of individual learning behavior.

2.3 Individual Forecasting Rules

Our model will be based upon some features of individual forecasting behavior detected in the experiment. We therefor first discuss these behavioral aspects.

Structure of the forecasting errors of the participants (see small frames in Fig. 2) gives some mixed evidence concerning the abilities of agents to learn from the past mistakes. Only in groups 2 and 5 with monotonic convergence the forecasting errors decreased to the values less than 1, which is close to the errors under the rational expectations benchmark. Errors in the oscillating groups 1 and 6 persistently decreased and increased but remained within the same range during all the experiment. Finally, relatively small initial errors in groups 4 and 7 were followed by huge errors in the middle of the experiment, which then dropped. Therefore, neither rational expectation theory, nor any learning theory implying that a participant always successfully learns from own past mistakes is not applicable to this experiment.

If the behavior of participants is not consistent with the rational expectations, which other models of expectation formation can be used? And would it be correct to conclude that agents in such groups as 1 or 6 did not learn at all? To answer these questions we consider a few examples of the time evolution of individual predictions. In every panel of Fig. 3 the dynamics of forecasts submitted by a certain participant is shown against the price. The timing in the figure is important. For every time t on the horizontal axes we show price p_t against the individual forecast $p_{i,t+2}^e$ of participant i . Notice that this forecast is given immediately after the corresponding price is announced, so that the dependence of forecasts on the information set (2.1) can be seen clearly.

In group 2 the subject 5 tried to extrapolate price changes in the beginning of the experiment (see the upper left panel). However, starting from period $t = 6$ this participant always used simple *naive rule* $p_{t+1}^e = p_{t-1}$. Subject 1 from the same group used more “smooth”, *adaptive* strategy in forecasting, always predicting a price inside an interval between the previous forecast and previous price realization. In oscillating group 6 the subject 1 used the naive rule in the first half of the experiment (see the upper right panel). Such rule, however, leads to the prediction errors, especially during the periods of trend. Consequently, in the middle of the experiment the prediction strategy has been changed. Now the participant *extrapolated trend* in prices, switching back to the naive rule at the periods of expected trend reversal.

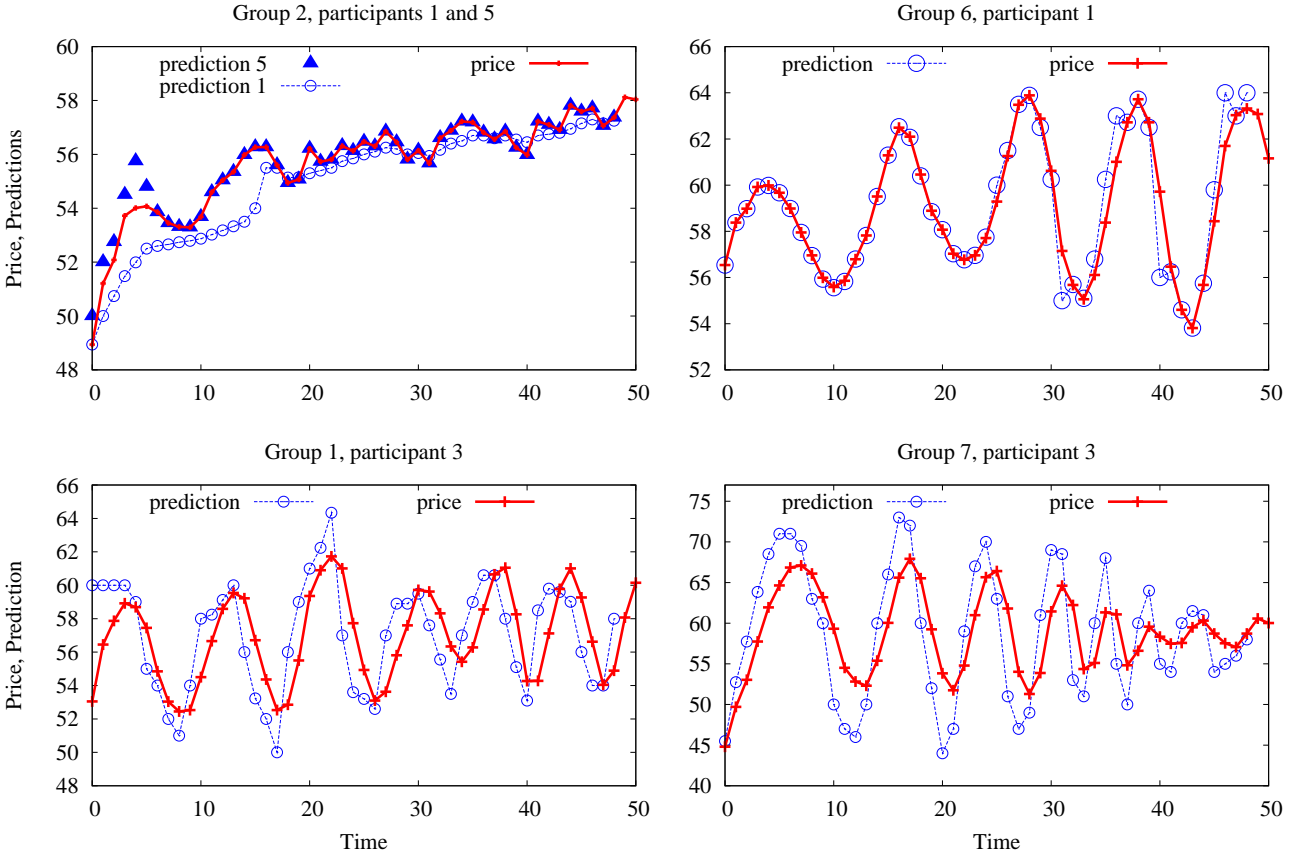


Figure 3: **Switching of the experiment's participants between simple rules.** For any point on the abscissa, representing time t , the price p_t (red) and the forecast $p_{i,t+2}^e$ (blue) are shown. This forecast was made immediately after the announcement of the price.

Participant 3 from another oscillating group 1 made the fundamental prediction $p_{t+1}^e = p^f$ in the first four periods (see the lower left panel). Being, probably, unsatisfied with the rewards from this strategy (which is rational only in the fully rational world), the participant started to extrapolate observable price trends. Such prediction rule often overshoots at the moments of the trend reversal implying low earnings. Towards the end of the experiment the participant learned to anticipate the trend changes to some extent. Finally, in group 7 with damping oscillations the subject 3 started with a strong trend extrapolation (see the lower right panel). Despite very low earnings on the turning points of trend, the participant stuck to this rule until the last 4 periods, when something similar to the adaptive strategy was used.

The availability of the participants' predictions allowed HSTV to estimate the individual forecasting rules directly. The estimation was performed for the predictions over the last 40 periods and many intuitive rules with simple interpretation came up from the estimation. Participants from converging groups often adapted their last forecast in the direction of the last price. This prediction strategy is known as *adaptive expectations*:

$$p_{t+1}^e = w p_{t-1} + (1 - w) p_t^e = p_t^e + w (p_{t-1} - p_t^e), \quad (2.5)$$

with weight $0 \leq w \leq 1$. Note that at the moment when forecasts of price p_{t+1} are submitted, price p_t is still unknown (see Eq. 2.2) and the last observed price is p_{t-1} . At the same time, the last own forecast p_t^e is known when forecasting p_{t+1} . Two rules shown in the upper left

panel of Fig. 3 are examples of such prediction strategy with $w = 1$ (for naive expectations of subject 5) and $w \simeq 0.25$ for subject 1.

Often the predictions were well approximated by the *trend-following* rule:

$$p_{t+1}^e = p_{t-1} + \gamma(p_{t-1} - p_{t-2}), \quad (2.6)$$

where $\gamma > 0$. Under this rule the participant believes that the price is determined by the last observation adjusted in the direction of the last price change. Extrapolation coefficient γ gives the strength of the adjustment. The estimates of this coefficient ranged from relatively small values as $\gamma = 0.4$ to quite high values as $\gamma = 1.3$.

Finally, some participants used more sophisticated rules extrapolating the last price change from a certain reference point describing the “long-run” level of the price process. For instance, for the third participant in group 1 the estimation suggests the following *anchoring and adjustment* rule:

$$p_{t+1}^e = 0.5(p^f + p_{t-1}) + (p_{t-1} - p_{t-2}). \quad (2.7)$$

The reference point or *anchor* for this rule is defined as an equally weighted average between the last observed price and the fundamental price. Since in the experiment subjects were not provided explicitly with the fundamental price, it can be argued that rule (2.7) was not feasible.⁵ Therefore, in our simulations we will use analogous rule replacing the fundamental price p^f in (2.7) by a proxy given by the observed sample average of past prices $p_{t-1}^{av} = \sum_{j=0}^{t-1} p_j$. The forecasting strategy

$$p_{t+1}^e = 0.5(p_{t-1}^{av} + p_{t-1}) + (p_{t-1} - p_{t-2}). \quad (2.8)$$

is the anchoring and adjustment rule with learned anchor.

Two important observations are suggested by the above analysis. First, the subjects in the experiment tend to base their predictions on the past observations, using relatively simple and intuitive rules of thumb, such as naive rule or trend extrapolation. Following seminal paper of Tversky and Kahneman (1974) we will use the term “*heuristics*” for these simple prediction rules. Second, it seems that participants try to learn from the past errors, and their learning behavior takes form of a *switching* between different heuristics (and not, e.g. better estimation of parameters of a particular heuristic). Moreover, this learning not necessarily leads to the improvement in the forecasting performance of the agents. These two ideas, simplicity of the forecasting strategy and (imperfect) evolutionary switching on seemingly more successful rules will form the basis of our learning model in Section 4.

The trend-following and anchoring and adjustment heuristics are examples of the *extrapolation expectation* formation as defined in Williams (1987). Such expectational schemes often come up from the forecasting experiments. For instance, Hey (1994) estimates the individual prediction rules in the experiment with exogenously given process and found that many subjects use extrapolation expectations. In the sessions with stable underlying process, the trend extrapolation started with the past price level, so that the trend-following heuristic (2.6) was employed. Instead, for the unstable processes the extrapolation started with the long-run average as in the anchoring and adjustment heuristic (2.7). Such behavior of the agents seems quite reasonable, and the similar tendency was observed in the HSTV experiments: trend-following behavior prevailed in the groups with converging price (both oscillating

⁵It is remarkable, however, that the rule (2.7) came from the predictions of the participant who *did* submit the fundamental price forecast in the first periods of the experiment, see the lower left panel of Fig. 3.

and not), while the anchoring and adjustment heuristic often used in the groups with constant oscillations. However, in the HSTV experiment, the price process was endogenous, so that the stability/instability of the underlying process was the emerging property. In such circumstances development of the trend-following heuristics in one groups and the anchoring and adjustment heuristics in other groups is non-trivial and unexpected result deserving some explanation.

3 Price Behavior under Homogeneous Expectations

Having a set of estimated individual forecasting rules, one can ask whether these homogeneous expectation rules can generate the qualitatively different patterns observed in the experiments. The experimental evidence about forecasting behavior suggests strong coordination on a common prediction rule. One can therefore suspect that this common rule (which, for whatever reason, turned out to be different in different groups) generates the resulting pattern. In this Section we investigate this conjecture by investigating price fluctuations under *homogeneous expectations* in the forecasting experiment.

The model with homogeneous expectations consists of the following equations:

$$\begin{cases} p_{t+1}^e = f(p_{t-1}, p_{t-2}, p_t^e) \\ n_t = 1 - \exp\left(-\frac{1}{200} |p_{t-1} - p^f|\right) \\ p_t = \frac{1}{1+r} \left((1-n_t)p_{t+1}^e + n_t p^f + \bar{y} + \varepsilon_t \right). \end{cases} \quad (3.1)$$

The first equation describes the forecasting behavior with a simple forecasting heuristic f , which can be either adaptive (in which case it does not depend on p_{t-2}) or extrapolative (in which case it does not depend on p_t^e). The second equation gives the evolution of the share of “robot” traders, identical to the rule used in the experiment. The third equation is an equilibrium condition from the pricing model used in the experiment (cf. (2.2)), with the same stochastic component. We concentrate on an analysis of the so-called *deterministic skeleton model* of (3.1), setting all stochastic term ε_t to zero, and then present stochastic simulations to investigate how the noise affects price fluctuations. In terms of *deviations* from the fundamental price the model can be rewritten as

$$p_t - p^f = \frac{1}{1+r} \left((1-n_t)p_{t+1}^e + n_t p^f - p^f \right) = \frac{1-n_t}{1+r} (p_{t+1}^e - p^f), \quad (3.2)$$

Fig. 4 shows example of simulated dynamics for different adaptive, trend-following and anchoring and adjustment rules.

3.1 Adaptive Heuristic

Assume that all participants use the same adaptive heuristic $p_{t+1}^e = w p_{t-1} + (1-w) p_t^e$ in their forecasting activity. Notice that naive expectations is obtained as a special case, for $w = 1$. The following result describes the behavior of system (3.1) in this case.

Proposition 3.1. *Consider the deterministic skeleton of (3.1) with the adaptive prediction rule (2.5). This system has a unique steady-state with price equal to fundamental price, i.e. $p^* = p^f$. The steady-state is globally stable for $0 < w \leq 1$, with a real eigenvalue λ , $0 < \lambda < 1$, so that the convergence is monotonic.*

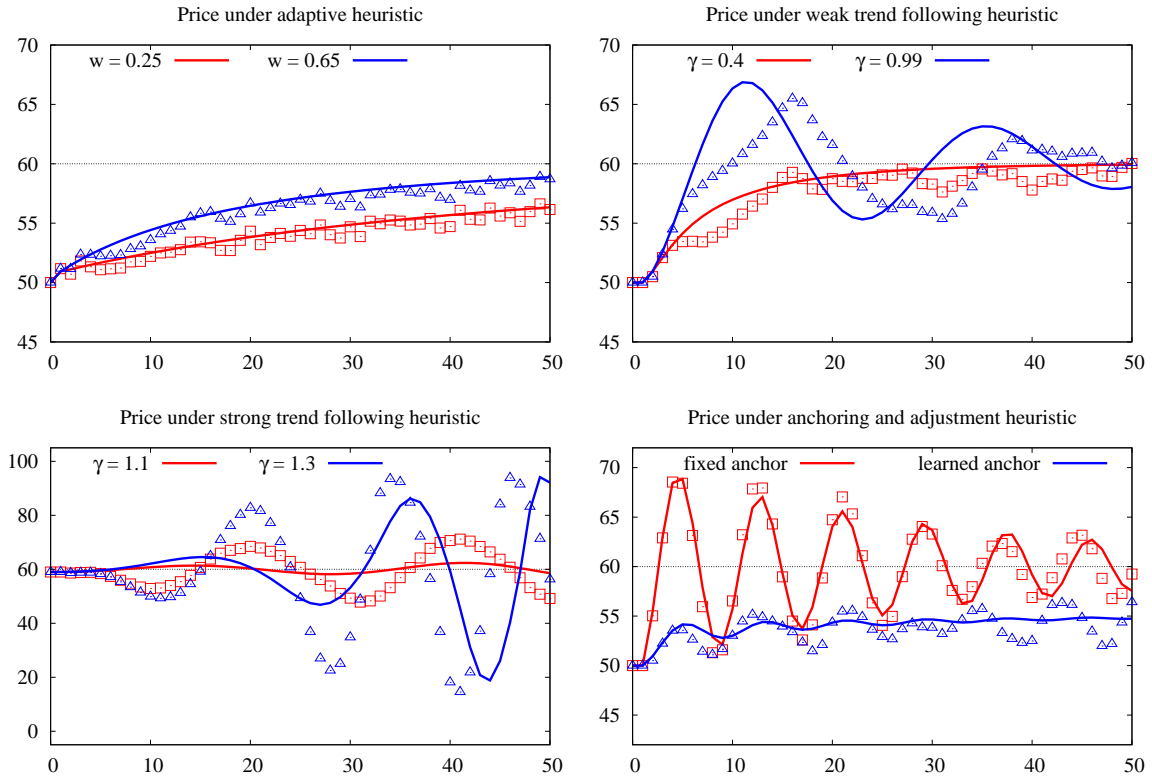


Figure 4: **Model (3.1) with homogeneous expectations.** The trajectories of the deterministic skeleton (the curves) and stochastic simulations with noise (triangles and squares) are shown for different forecasting heuristics. **Upper left panel:** Dynamics with the adaptive forecast converge to the fundamental steady-state. **Upper right panel:** Dynamics with weak trend extrapolation converges to the fundamental steady-state. Convergence can be either monotonic (for small extrapolation coefficients), or oscillating (for high coefficients). **Middle left panel:** Dynamics with the strong trend extrapolation oscillates (slowly) around the fundamental steady-state and diverges to a quasi-periodic cycle. **Middle right panel:** Dynamics with the anchoring and adjustment heuristic oscillates around the steady-state and (ultimately) converges. The same heuristic with learned anchor generates small amplitude oscillations around its current long-run estimation, which converges extremely slowly and almost monotonically.

Proof. See Appendix A. □

The dynamics with the adaptive forecasting heuristic is illustrated in the upper left panel of Fig. 4 for two different values of the weight w assigned to the past price. When the weight is relatively low, e.g. $w = 0.25$ as for participant 1 in group 2, the error correction is small, and the price only converges slowly to the fundamental steady-state. In the case of larger weight, e.g. $w = 0.65$ as estimated for subject 4 of group 5, convergence is somewhat faster. In the case of adaptive expectations, the role of stochastic shocks is minimal. Shocks slightly perturb the system, but the price trajectory (shown by triangles and squares) still exhibits almost monotonic convergence. Adaptive expectations thus seems a good explanation of the price pattern observed in the experimental groups 2 and 5.

3.2 Extrapolative Rules

Consider now the dynamics with homogeneous extrapolative expectations. For the sake of generality we write the extrapolative forecasting rules as follows:

$$p_{t+1}^e = \alpha + \beta_1 p_{t-1} + \beta_2 p_{t-2}. \quad (3.3)$$

This extrapolative rule contains both the trendfollowing and the anchor and adjustment heuristic as special cases. Indeed, setting $\alpha = 0$, $\beta_1 = 1 + \gamma$ and $\beta_2 = -\gamma$, the trend-following heuristics (2.6) is obtained, while $\alpha = p^f/2$, $\beta_1 = 1.5$ and $\beta_2 = -1$ correspond to the anchoring and adjustment heuristic (2.7).

The rules for which the forecasts are not consistent with realizations will be disregarded by the participants, sooner or later. Therefore, both in formal analysis and in simulations we confine our attention to the rules satisfying the following definition.

Definition 3.1. The extrapolative rule (3.3) is called *consistent* in the steady-state p^* , if it predicts p^* in this steady-state.

In other words, consistent rules give unbiased predictions at the steady-state. Obviously, the extrapolative rule is consistent in p^* if and only if $\alpha = (1 - \beta_1 - \beta_2)p^*$. For instance, the trend-following heuristics are consistent in the steady-state with any price, while the anchoring and adjustment heuristic (2.7) is consistent only when $p^* = p^f$.

The following result describes all possible steady-states of the asset-pricing dynamics with consistent extrapolative heuristic, as well as their local stability.

Proposition 3.2. *Consider the dynamics of the deterministic skeleton of (3.1) with extrapolative prediction rule (3.3).*

There exists a unique steady-state in which the rule is consistent. In this steady-state, $p^ = p^f$ and $n^* = 0$. Such “fundamental” steady-state is locally stable if the following three conditions are met*

$$\beta_2 < (1 + r) - \beta_1, \quad \beta_2 < (1 + r) + \beta_1, \quad \beta_2 > -(1 + r). \quad (3.4)$$

The steady-state generically exhibits pitch-fork, period-doubling or Neimark-Sacker bifurcation, if the first, second or third inequality in (3.4) turns to equality, respectively. The dynamics is oscillating when $\beta_1^2 + 4\beta_2(1 + r) < 0$.

Proof. See Appendix B. □

Dynamical system (3.1) with homogeneous extrapolative expectations (3.3) may have multiple steady-states. Proposition 3.2 asserts, however, that the extrapolative rule is consistent in only one of them, with fundamental price p^f . The stability conditions (3.4) are illustrated in Fig. 5 in the parameter space (β_1, β_2) . The two filled regions contain all the rules for which the extrapolative heuristic (3.3) generates stable dynamics. For the pairs lying below the parabolic curve, the dynamics are oscillating. The loss of stability happens when the pair (β_1, β_2) leaves the stability area. The dynamics immediately after the bifurcation are determined by the type of the bifurcation through which the stability is lost. For instance, after the pitchfork bifurcation the price diverges from the fundamental level but converges to one of two new steady-states. The Neimark-Sacker bifurcation implies an existence of the stable, quasi-periodic dynamics right after the bifurcation.⁶

⁶To be precise, the stable quasi-periodic dynamics arises after the *supercritical* Neimark-Sacker bifurcation. This is the case for our system, as simulations suggest.




Figure 5: **Stability of the fundamental steady-state in an asset-pricing model with homogeneous extrapolative expectations.** The dynamics (3.1) with expectations (3.3) converges to the steady-state if the coefficient pair (β_1, β_2) belongs to the union of light and dark grey regions. The dynamics oscillates if the pair lies below the parabolic curve.

Three points shown in Fig. 5 correspond to the extrapolative forecasting rules estimated in the experiment. Two trend-following heuristic (2.6) with different values of the extrapolation coefficient γ are labeled as $\gamma = 0.4$ and $\gamma = 1.3$. The anchoring and adjustment heuristic (2.7) is labeled as AAA.

Trend-following heuristic. These results imply that the price can converge as well as diverge under the trend-following rule $p_{t+1}^e = p_{t-1} + \gamma(p_{t-1} - p_{t-2})$. To distinguish between these two cases we will use the terms *weak* and *strong* trend extrapolation, respectively.

The dynamics with the weak trend extrapolation is illustrated in the upper right panel of Fig. 4. When extrapolative coefficient is sufficiently small, the convergence is monotone. For large γ convergence becomes oscillatory, as in case $\gamma = 0.99$. Notice that the dynamics do not resemble the damping oscillations observed in the experiment. Also the estimation of individual strategies did not reveal any trend-following rule which would generate such converging oscillations. The case of the strong trend extrapolation is illustrated in the middle panels of Fig. 4. The dynamics diverge from the steady state through oscillations of increasing amplitude (left panel), and settle down to a quasi-periodic attractor shown in the right panel in coordinates (p_t, p_{t-1}) . The speed of divergence and amplitude of the long run fluctuations increase with γ , as shown by comparison of cases with $\gamma = 1.1$ and $\gamma = 1.3$.

Anchoring and adjustment heuristic. Applying Proposition 3.2 to the anchoring and adjustment rule (2.7) we conclude that the price dynamics is converging. Since parameters of such rule are very close to the Neimark-Sacker bifurcation, the convergence to the fundamental steady-state is oscillatory and slow. Indeed, as the lower left panel of Fig. 4 shows, after first 50 periods the price still oscillates with relatively large amplitude. The dynamics of the anchoring and adjustment rule with learned anchor (2.8) is shown in the same panel. Notice that the dynamics with such heuristic is described by a non-autonomous system, whose formal analysis is complicated. Simulations with rule (2.8) converge to the same fundamental steady-state

as with rule (2.7), but much slower and with less pronounced oscillations. In the presence of noise, the oscillations of both time series look qualitatively similar to the dynamics of groups 1 and 6 of the experiment.

Finally, in the lower right panel of Fig. 4 we show the attractor for the anchoring and adjustment heuristic $p_{t+1}^e = 0.5(p^f + p_{t-1}) + 1.1(p_{t-1} - p_{t-2})$, which differs from our standard representation (2.7) by stronger extrapolation of the past price trend. In this case the stability is lost through the Neimark-Sacker bifurcation. The dynamics converge to the 8-cycle when the noise is not present, and it is quasi-periodic otherwise.

In this Section we analyzed the dynamics underlying the HSTV experiment in case, when a single heuristic is used. Even if the patterns of monotonic convergence, constant oscillations and also damping oscillations can be reproduced, we believe that the model with homogeneous forecasting rule (3.1) gives unsatisfactory explanation for the experiment. First, the model has difficulties in reproducing some patterns. For example, the damping oscillations, which can be generated only by the strong trend following rule, look hardly similar to the dynamics in the groups 4 and 7. Second, on the conceptual level the explanation with homogeneous rules would leave unanswered the question *why* different patterns self-emerged in the experiment.

Finally, the assumption of homogeneous forecasting rule is not as plausible as one may think. Even if the HSTV experiment revealed a large degree of similarity between individual forecasts in every period, it did not demonstrate a coordination on a common *prediction rule*. In Fig. 6 we plot the coefficients of the estimated individual extrapolative rules on the stability region.⁷ The dispersion of the forecasting rules is clear. Nevertheless, there are some important tendencies. As one may expect the majority of the rules in the converging groups belong to the region of monotonic convergence. In every of the groups with constant oscillations, there are couple of rules lying very close to the locus of the Neimark-Sacker bifurcation. Finally, the groups with damping oscillations are characterized by the presence of both stable and unstable rules.

OLD TEXT FROM BEGINNING SECTION

We show that within the HSTV experimental setting (i) the adaptive heuristics would translate into (almost) monotonic convergence of prices, (ii) different trend-following heuristics would generate all three different patterns: monotonic convergence, constant oscillations and damping oscillations, and (iii) the anchoring and adjustment heuristic would be consistent with slowly damping oscillations. We will stress, however, some discrepancies between experiment results and the model of homogeneous expectations and, at the end, will reject the model.

4 Model of Evolutionary Switching of Heuristics

In this Section we demonstrate that the learning process in the form of switching between simple heuristics can explain different qualitative scenarios observed in the experiment. Before describing a model, let us recall the most important “stylized facts” which we found during the analysis of individual prediction strategies.

- participants tend to base their predictions on past observations, following simple routines;
- learning of people has a form of switching from one routine to another;

⁷Of course, in every group there were rules which cannot be represented by the extrapolative prediction (3.3), e.g. adaptive heuristic.

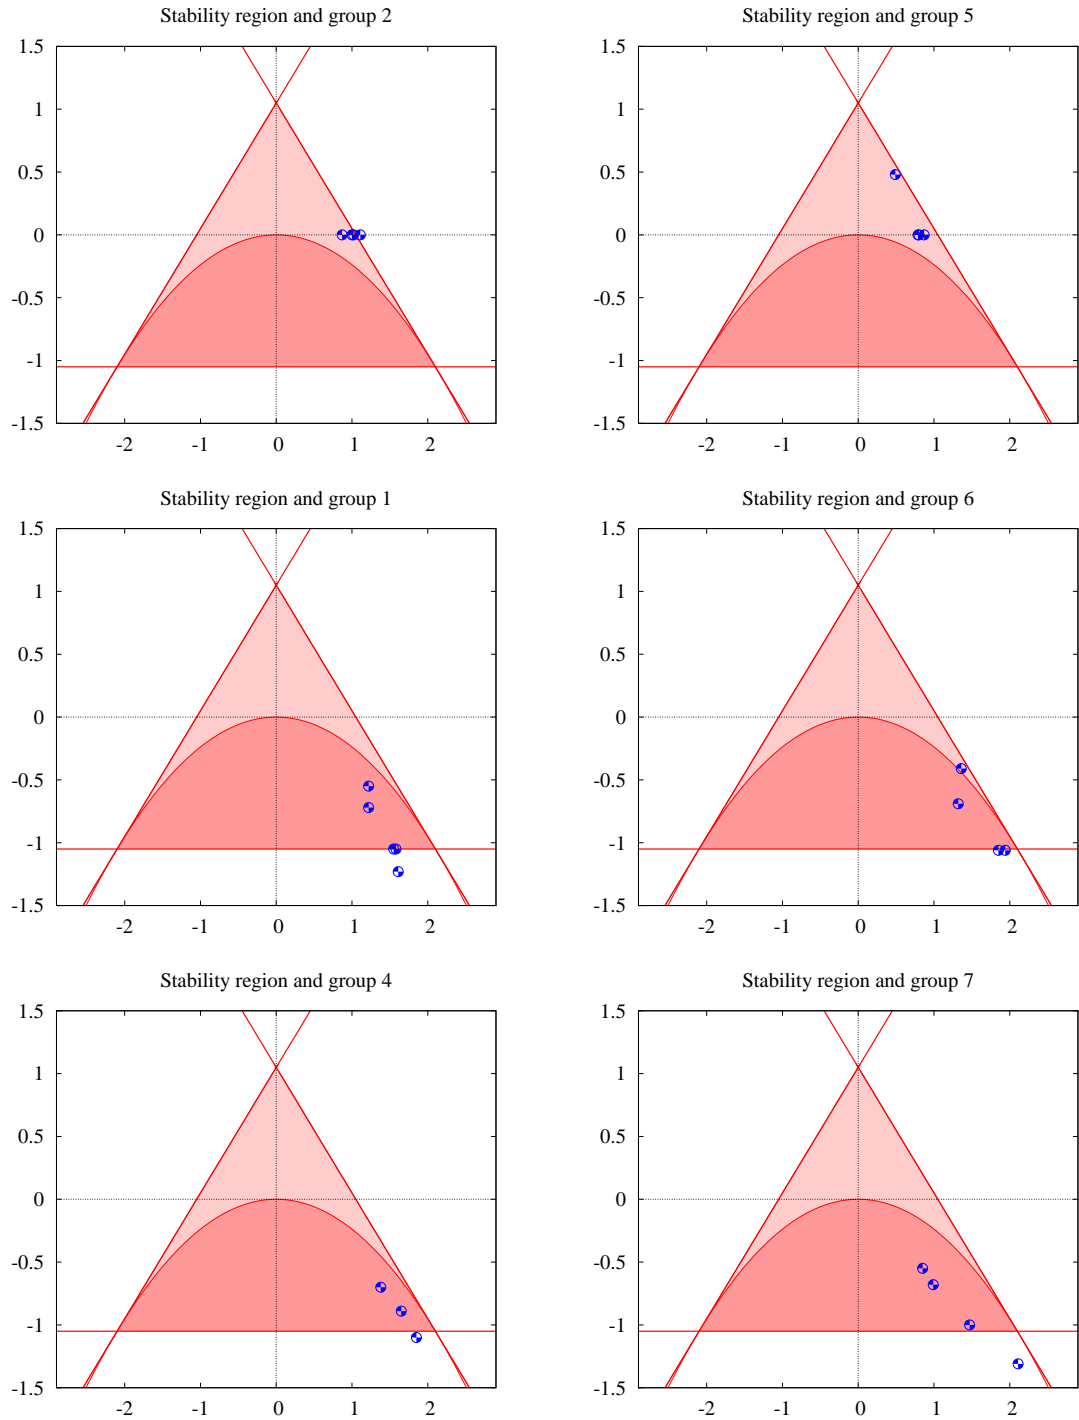


Figure 6: **Stability of a model with homogeneous extrapolative rules estimated in the experiment.** **Upper panel:** In both groups with converging price, all rules generate stable monotonic dynamics. **Middle panel:** In both groups with oscillating price, there were two rules on the stability border of the Neimark-Sacker bifurcation. **Lower panel:** In both groups with damping oscillations, both stable and unstable rules were present.

- in every group, at every period of time, forecasts are very close to each other;
- some heterogeneity of the applied rules remains at every time period.

The idea of the model is the following. Assume that there exists a pool of simple prediction rules (e.g. adaptive or trend-following heuristics) commonly available to the participants of the experiment. At every time period these heuristics deliver forecasts for the next period price, and the current price is then computed. However, the *impacts* of heuristics in the price determination are changing because of the participants' learning. The better a heuristic performed in the past, the higher impact it gets in the price determination. As a result, the price and heuristics' impacts co-evolve in a dynamical process feeding back each other.

It turns out that this process can lead to a coordination on the successful rule, implying certain aggregate price pattern. However, the rule on which agents coordinate can depend on the impacts of heuristics in the initial stage of the experiment and first price realizations. These initial conditions were, indeed, casual in the experiment. It could happen, for example, that in one session participants with adaptive expectations dominated during the first few rounds, while in other session, say, the weak trend-followers had a majority. The initial prices, in turn, depended on the forecasts made when participants had only few variables in the information set.⁸

Below we show that this natural idea can be formalized as a model with necessary path-dependence structure.

4.1 Formal Model

Let \mathcal{H} denote a set of H heuristics which participants can use for the prediction of price. In the beginning of time t any rule $h \in \mathcal{H}$ gives a two-period ahead point prediction for price p_{t+1} . The prediction is described by deterministic function f_h of available information

$$p_{h,t+1}^e = f_h(p_{t-1}, p_{t-2}, \dots; p_{h,t}^e, p_{h,t-1}^e, \dots). \quad (4.1)$$

The price in period t is computed on the base of these predictions as in (2.2):

$$p_t = \frac{1}{1+r} \left((1-n_t) \bar{p}_{t+1}^e + n_t p^f + \bar{y} + \varepsilon_t \right), \quad (4.2)$$

where \bar{p}_{t+1}^e is the average predicted price, r is the risk free interest rate, \bar{y} is the mean dividend, and ε_t is a noise term. Finally, n_t is the share of robot traders evolving as in the experiment (cf. (2.3)) according to

$$n_t = 1 - \exp \left(- \frac{1}{200} |p_{t-1} - p^f| \right). \quad (4.3)$$

In further simulations we use the same values of parameters and the same noise realization as in the HSTV experiment. In particular, the fundamental price always predicted by robots is set to $p^f = \bar{y}/r = 0.05/3 = 60$.

An important difference with respect to the experiment is that in our evolutionary model, the average \bar{p}_{t+1}^e in (4.2) is taken with respect to the predictions given by different heuristics

$$\bar{p}_{t+1}^e = \sum_{h=1}^H n_{h,t} p_{h,t+1}^e, \quad (4.4)$$

⁸In the HSTV experiment the forecasts of the first period price were around 50, probably because it was the middle point of the available range. Some participants started with fundamental prediction.

with $p_{h,t+1}^e$ defined in (4.1). The weight $n_{h,t}$ assigned to the heuristic h is called *impact* of this heuristic. Impact is evolving in time and depends on the past relative performance of all H heuristics, with more successful in the past heuristics attracting a higher number of followers.

Let the performance measure of a heuristic in a given period be the squared forecasting error generated by this heuristic. This definition is consistent with the reward structure of the experiment. Taking the weighted average of the past squared forecasting errors, we define

$$U_{h,t-1} = -(p_{t-1} - p_{h,t-1}^e)^2 + \eta U_{h,t-2} \quad (4.5)$$

as a *performance measure* of the heuristic h up to (and including) time $t - 1$. The parameter $0 \leq \eta \leq 1$ represents the memory strength. It measures relative weight agents give to past errors of heuristic h . 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 impact of rule h is updated according to a *discrete choice model with asynchronous updating*

$$n_{h,t} = \delta n_{h,t-1} + (1 - \delta) \frac{\exp(\beta U_{h,t-1})}{Z_{t-1}}, \quad (4.6)$$

where $Z_{t-1} = \sum_{h=1}^H \exp(\beta U_{h,t-1})$ is a normalization factor. There are two important parameters in (4.6). Parameter $0 \leq \delta \leq 1$ gives some persistence or inertia in the impact of 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 impacts of the rules never change, no matter what their past performance was. If $0 < \delta \leq 1$, in each period a fraction $1 - \delta$ of participants update their rule according to the well known *the discrete choice model* used for example in Brock and Hommes (1997). The parameter $\beta \geq 0$ represents the intensity of choice measuring how sensitive individuals are to differences in strategy performance. The higher intensity of choice β , the faster individuals will switch to more successful rules. In the extreme case $\beta = 0$, the impacts in (4.6) move to an equal distribution independent of their past performance. At the other extreme $\beta = \infty$, *all* agents who update their heuristic (i.e. fraction $1 - \delta$) switch to the most successful predictor.

Initialization. The model is initialized by providing (i) a sequence $\{p_0, p_1, \dots, p_{\text{in}}\}$ of initial prices, long enough to allow any forecasting rule in \mathcal{H} to generate its prediction, and (ii) an initial distribution $\{n_{h,\text{in}}\}, 1 \leq h \leq H$ of heuristics' impacts (summing to 1). Additionally, the initial share of robot traders and initial performances of all H heuristics set to 0.

Given initial prices, the heuristic forecasts can be computed and, using the initial impacts of the heuristics, the price $p_{\text{in}+1}$ 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 impacts $n_{h,\text{in}}$ for the individual rules are used, since past performance is not well defined yet. The price $p_{\text{in}+2}$ is computed and the initialization stage is finished. Starting from this moment the evolution according to (4.2) is well defined: first the performance measure in (4.5) is updated, then, the new impacts of the heuristics are computed according to (4.6), and the new prediction of the heuristics are obtained according to (4.1). Finally, the new average forecast (4.4) and the new fraction of robot traders (4.3) are computed, and a new price is determined by (4.2).

4.1.1 Example with Four Heuristics

The evolutionary model can be simulated with an arbitrary set of heuristics.⁹ For the definitiveness we will deal with a case when set \mathcal{H} contains only four forecasting rules. These rules, referred as ADA, WTR, STR and AAA are given in Table 1. In analysis we typically substitute the LAA heuristic by the similar AAA heuristic.

ADA	adaptive heuristic	$p_{1,t+1}^e = 0.65 p_{t-1} + 0.35 p_{1,t}^e$
WTR	weak trend-following rule	$p_{2,t+1}^e = p_{t-1} + 0.4 (p_{t-1} - p_{t-2})$
STR	strong trend-following rule	$p_{3,t+1}^e = p_{t-1} + 1.3 (p_{t-1} - p_{t-2})$
LAA	anchoring and adjustment rule with learned anchor	$p_{4,t+1}^e = 0.5 (p_{t-1}^{av} + p_{t-1}) + (p_{t-1} - p_{t-2})$
AAA	anchoring and adjustment rule with fixed anchor	$p_{4,t+1}^e = 0.5 (p^f + p_{t-1}) + (p_{t-1} - p_{t-2})$

Table 1: Heuristics used in an evolutionary model. In simulations in Figs. 7–8 the first four heuristics are used.

Our choice of these four rules is based on two considerations. First, in pool \mathcal{H} we included only those rules that were estimated in the experiment, slightly modifying them to obtain consistent rules with simple behavior interpretation. Second, we include the rules which, in a model with homogeneous expectations discussed in Section 3, generate qualitatively different dynamics. It allows one to get some non-trivial interaction between different heuristics, so that qualitatively different patterns can be obtained. Our experimentations with alternative choice of heuristics suggest that the main result of simulations (possibility to reproduce qualitatively different price patterns) will not change as soon as the second condition is satisfied.

The complete evolutionary model is given by the following system:

$$\left\{ \begin{array}{l}
 p_{1,t+1}^e = 0.65 p_{t-1} + 0.35 p_{1,t}^e \\
 p_{2,t+1}^e = p_{t-1} + 0.4 (p_{t-1} - p_{t-2}) \\
 p_{3,t+1}^e = p_{t-1} + 1.3 (p_{t-1} - p_{t-2}) \\
 p_{4,t+1}^e = 0.5 p_{t-1}^{av} + 1.5 p_{t-1} - p_{t-2} \\
 n_t = 1 - \exp \left(- \frac{1}{200} |p_{t-1} - p^f| \right) \\
 U_{h,t-1} = - (p_{t-1} - p_{h,t-1}^e)^2 + \eta U_{h,t-2} \quad 1 \leq h \leq 4 \\
 n_{h,t} = \delta n_{h,t-1} + (1 - \delta) \frac{\exp(\beta U_{h,t-1})}{Z_{t-1}} \quad 1 \leq h \leq 4 \\
 p_t = \frac{1}{1+r} \left((1 - n_t) (n_{1,t} p_{1,t+1}^e + \dots + n_{4,t} p_{4,t+1}^e) + n_t p^f + \bar{y} + \varepsilon_t \right),
 \end{array} \right. \quad (4.7)$$

where as before p_{t-1}^{av} stands for the average of all the past prices up to p_{t-1} . Notice that in the deterministic skeleton the last equation can be rewritten in deviations as

$$p_t - p^f = \frac{1 - n_t}{1 + r} \sum_{h=1}^4 n_{h,t} (p_{h,t+1}^e - p^f). \quad (4.8)$$

⁹The software for simulations *evexex* is freely available at <http://www.cafed.eu/evexex> together with brief documentation and configuration settings used for the simulations reported below.

	Experiment		Simulations								
	Observed Prices $t = 0$	$t = 1$	Initial Prices p_0 p_1		Heuristics' Initial Impacts ADA WTR STR LAA				Parameters β η δ		
Group 2	48.94	51.21	49	50.5	0.25	0.35	0.15	0.25	0.4	0.7	0.9
Group 5	53.78	53.61	54	53.5	0.25	0.35	0.15	0.25	0.4	0.7	0.9
Group 1	53.05	56.45	51	54	0.15	0.35	0.35	0.15	0.4	0.7	0.9
Group 6	56.54	58.38	56	58	0.1	0.3	0.4	0.2	0.4	0.7	0.9
Group 4	43.72	47.33	42	47	0	0.1	0.9	0	0.4	0.7	0.9
Group 7	44.81	49.71	44	48	0	0.17	0.66	0.17	0.4	0.7	0.9

Table 2: Initial conditions for simulation of different qualitative scenarios reported in Figs. 7–8. Initial prices in simulations are chosen close to the prices observed in the experiment (second and third columns).

Comparing it with (3.2) we observe that our model can be seen as a generalization of homogeneous expectation model with *weighted average* forecast of four heuristics on the case of *endogenous* weights.

Due to presence of the demand/supply shocks ε_t in the pricing equation, system (4.7) is, in general, stochastic. In the next section we simulate this system for 50 periods with *the same* noise process $\{\varepsilon_t\}_{t=0}^{50}$ as was employed in the experiment. Resulting (deterministic) dynamics will be referred as “simulated path” and will be compared with the experimental data. Section 4.3 is devoted to the analysis of the deterministic skeleton of system (4.7) when the noise term ε_t is absent. Finally, Section 4.4 presents another type of simulations, so-called “one step ahead predicted path” when the experimental data are used at each step to correct the prediction errors made by the model.

4.2 Simulated Path

As soon as four heuristics are fixed, there are only three free “learning” parameters in the model: β , η and δ . Provided that these parameters are given, system (4.7) is initialized with two initial prices, p_0 and p_1 , and four initial impacts $n_{h,\text{in}}$ used at periods $t = 2$ and $t = 3$.

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.

Six simulations reported in Figs. 7–8 aim to imitate the price dynamics observed in six experimental groups. Thus, we use the same realizations $\{\varepsilon_t\}_{t=0}^{50}$ as in the experiment. To stress path-dependence, for all simulations we have fixed the parameter values as $\beta = 0.4$, $\eta = 0.7$, $\delta = 0.9$. The simulations only differ in initial conditions, which are reported in Table 2. We stress that no fitting exercise has been performed at this stage. All the plots have been easily obtained through a trial-and-error experimentation with different initial conditions and parameters. In particular, we experimented with initial prices $\{p_0, p_1\}$ close to the prices observed in the first two rounds of the corresponding experimental group (shown in the second and third columns of Table 2). At the same time initial distribution of agents over the pool of heuristics, i.e. initial impacts $\{n_{1,\text{in}}, n_{2,\text{in}}, n_{3,\text{in}}, n_{4,\text{in}}\}$, were chosen in such a way to match

the experiment dynamics during the first few periods.¹⁰ For replication of groups 2 and 5, the initial impacts of heuristics are distributed almost uniformly, with slight dominance of WTR heuristics to produce small initial trend in prices. For groups 1 and 6, where initial trend was even stronger, both trend heuristics WTR and STR were initialized with high weights. Finally, in groups 4 and 7 with the strongest trend in price in initial periods, the rule STR got very large impact.

4.2.1 Explaining different qualitative scenario

Fig. 7 can be directly compared with Fig. 2 (we use the same scale in all corresponding panels of these two figures). Upper parts of panels in Fig. 7 show realized prices for both the experiments and the heuristic switching model. The heuristic switching model qualitatively matches all three different patterns, slow monotonic convergence to the fundamental price, persistent price oscillations and dampened oscillatory price movements, found in the laboratory experiments. Lower parts of panels show predictions of four heuristics over the simulation, which looks quite close, especially for monotonic convergence case. In inner frames the forecasting errors of different heuristics are shown. The errors do not disappear to the end of simulations, in line with individual forecasts observed in the experiment.

Fig. 8 plots the corresponding transition paths of the impacts of each of the four forecasting heuristics. In the case of monotonic convergence (see the upper panels), the impacts of all four heuristics remain relatively similar during the simulations causing slow (almost) monotonic convergence of the price to the fundamental equilibrium $p^f = 60$. For group 2 the increase in price together with a series of subsequent positive shocks ε_t leads to a temporary domination of the dynamics by the STR 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 level.

In two simulations for the groups with constant oscillations (see the middle panels), the weak and strong trend followers represent the largest proportions in the initial distribution of heuristics, and prices rise. However already after 5 periods the impact of the LAA heuristic starts to increase, because it predicts better than the static strong and weak trend followers, who either overestimate or underestimate the price trend. The impact of the anchoring adjustment heuristic gradually increases, so it dominates the market within 10 periods, rising to more than 70% after 40 periods. It explains coordination of individual forecasts as well as persistent price oscillations around the long run equilibrium level.

Finally, in the last two simulations (see lower panel) a large initial impact of (strong) trend followers leads to an extreme rise of market prices in the first 7 periods, followed by large price oscillations. Relatively fast (after period 10 in case of group 7), however, the impact of STR rule decreases, while the impact of the LAA heuristic, rises to more than 80% after 30 periods. Once again, the flexible anchoring and adjustment heuristic forecasts better than the static strong trend following rule, which overestimates the price trend. In simulation for group 7 after 40 periods the impact of the anchoring adjustment heuristic starts slowly decreasing, and consequently the price oscillations slowly stabilize. The decline of the impact of the LAA heuristic implies also smaller coordination between individual predictions during the last 10 periods, which is also consistent with experimental data.

¹⁰This is done with a purpose to imitate initial distribution of four heuristics over participants in the experiment. Direct estimation of this distribution in every group seems to be difficult due to obvious insufficiency of data.

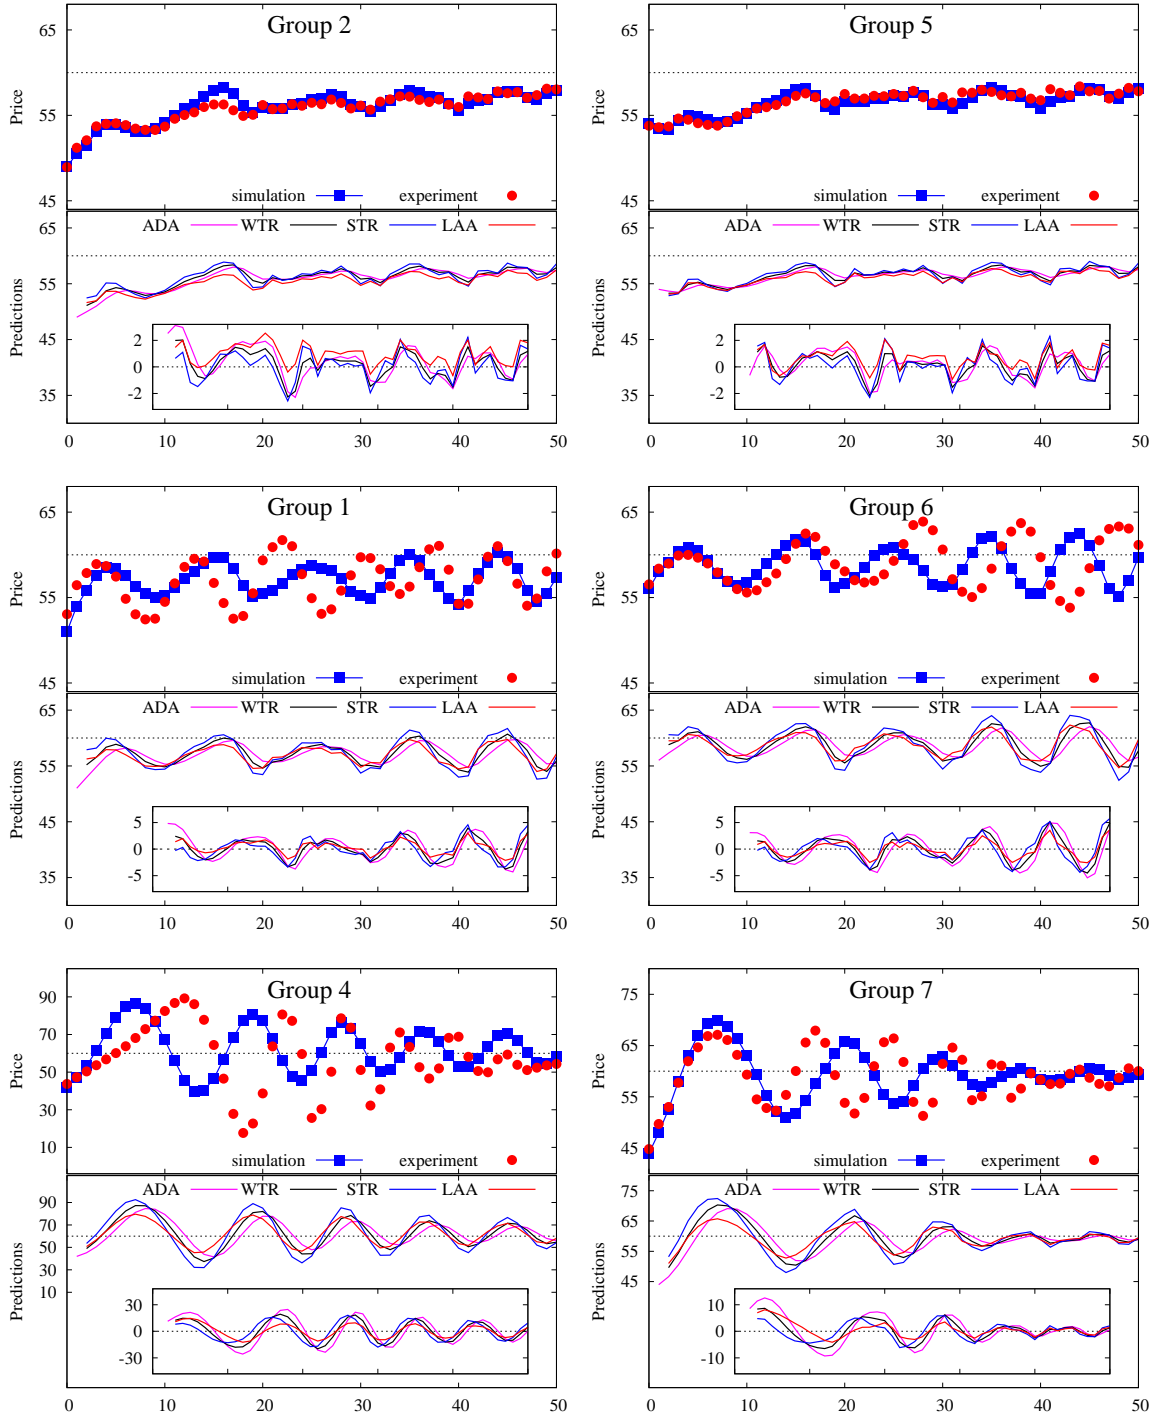


Figure 7: **Laboratory experiments and heuristics switching model simulations.** **Upper parts of panels** show prices for laboratory experiments in different groups (red) with corresponding simulations of the evolutionary model (blue). All three different aggregate market outcomes are reproduced: monotonic convergence to equilibrium (**top panels**), permanent oscillations (**middle panels**), and oscillatory convergence (**bottom panels**). **Lower parts of panels** show predictions and forecasting errors (inner frames) of four heuristics: adaptive expectations (ADA, purple), weak trend followers (WTR, black), strong trend followers (STR, blue) and anchoring adjustment heuristic (LAA, red).

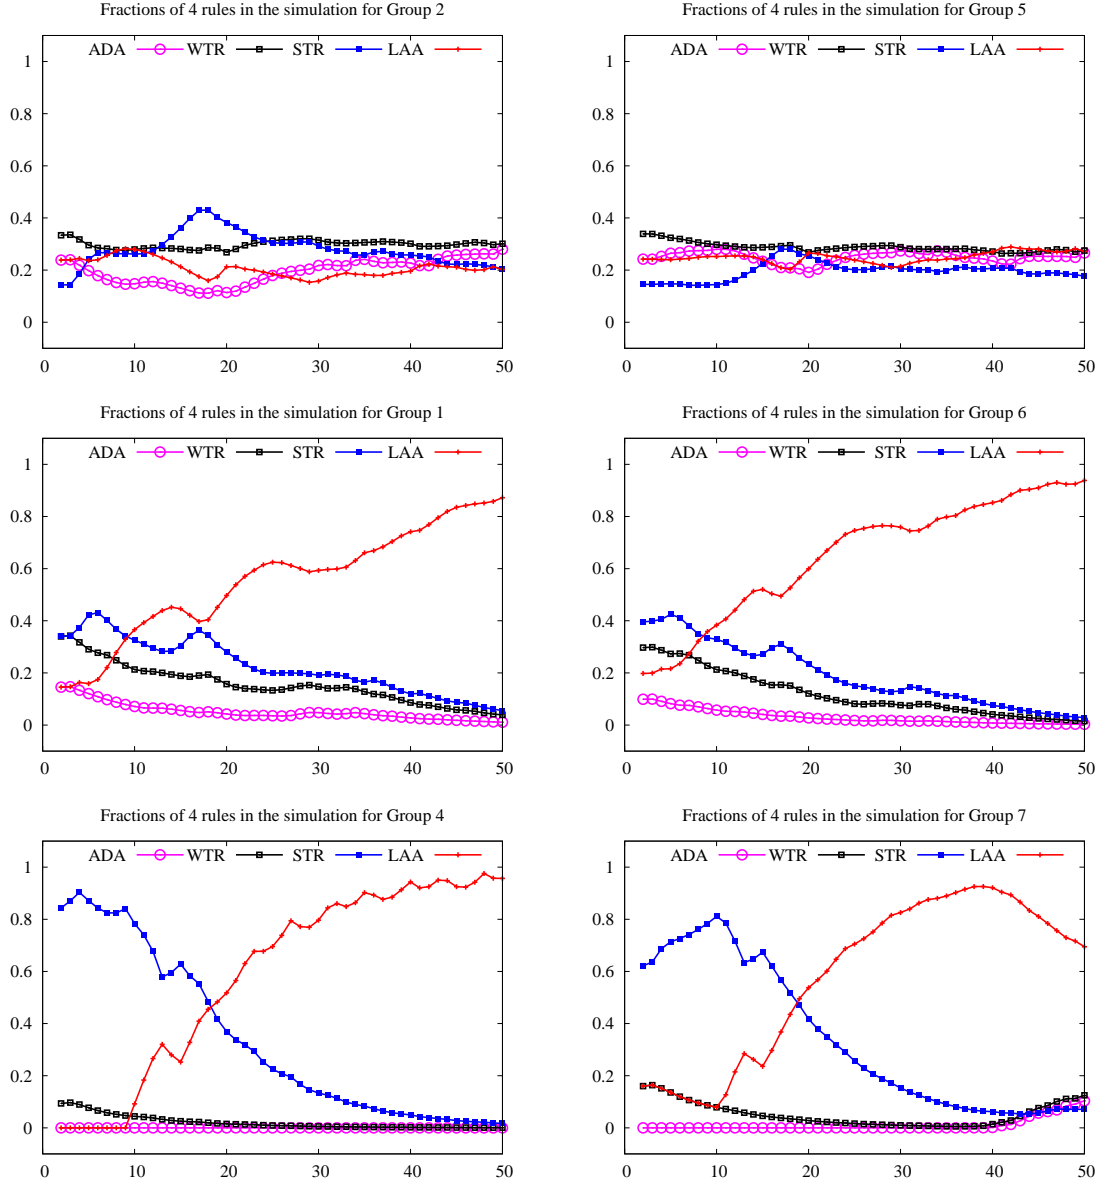


Figure 8: **Simulated fractions of the forecasting rules in the heuristics switching model.** Fractions of four forecasting heuristics: adaptive expectations (ADA, purple), weak trend followers (WTR, black), strong trend followers (STR, blue) and anchoring adjustment heuristic (LAA, red). Coordination of individual forecasts explains three different aggregate market outcomes reported in Fig. 7.

4.2.2 Fitting the simulated path of the model

Now we turn to the question of how good our explanation of the experiment is and, in particular, whether the model with four heuristic is capable to explain the experiment better than the homogeneous expectation model.

In Table 3 we show a mean squared error (MSE) generated by the deterministic models with different heuristics while fitting the experimental price data. It means that for any period t we compute a squared deviation of the price observed in the experiment, p_t^{exp} , from the price

Table 3: MSE for different groups and different specifications.

Specification	Group 2	Group 5	Group 1	Group 6	Group 4	Group 7
Fundamental Prediction	18.037	11.797	15.226	8.959	291.376	22.047
ADA – exp prices	0.841	0.200	7.676	8.401	330.101	51.526
WTR – exp prices	4.419	1.983	8.868	6.252	308.549	30.298
STR – exp prices	585.789	478.525	638.344	509.266	1231.064	698.361
AAA – exp prices	39.308	21.760	17.933	17.345	289.134	87.878
LAA – exp prices	5.475	3.534	5.405	14.404	307.605	69.749
ADA – fitted prices	0.514	0.199	6.832	7.431	312.564	36.436
WTR – fitted prices	4.222	1.844	8.670	6.228	292.150	19.764
STR – fitted prices	413.435	42.488	182.284	29.200	580.543	579.141
AAA – fitted prices	26.507	13.228	11.117	13.981	258.010	63.777
LAA – fitted prices	2.055	1.859	4.236	13.433	284.880	45.153
4 heuristics (Table 2)	0.449	0.302	8.627	14.755	526.417	29.520
4 heuristics (fitted)	0.313	0.245	7.227	7.679	235.900	18.662

generated by simulation, p_t^{sim} , and average these deviations over time

$$MSE = \frac{1}{49} \sum_{t=2}^{50} (p_t^{exp} - p_t^{sim})^2.$$

Notice that we skip two time periods corresponding to the initialization stage of our simulations.

Second line of Table 3 shows the MSE for a model where participants submit rational, fundamental prediction. Next ten lines show the MSE when the model from Section 3 using one of the five heuristics defined in Table 1 (see the first column) is used. Such model is initialized by the prices in the first two periods, and we consider two possible initializations. First, we initialize the model with prices observed in the corresponding experimental session; second, we vary these experimental prices in order to get the best fit. Finally, the last two lines show the MSE for the four heuristics switching model both with initializations as in Table 2 and with fitted initial prices (so that initial impacts are not changed).

First of all, notice that the fundamental model is extremely poor compared with actual experiment realizations. This is also clear if one compares experimental data with time series for fundamental homogeneous predictions shown in Fig. 1. As expected, in the groups 2 and 5 with monotonic convergence the ADA heuristic performs extremely well giving small values of the MSE. All other heuristics, especially the STR, are much worse in fitting the experiment. However, the model with switching can generate even better fit than any of the four heuristics. It is remarkable that this happens despite the fact that over all 50 periods of simulation these four heuristics had quite similar impacts. In the groups 1 and 6 with constant oscillations the LAA, ADA and WTR heuristics generate the smallest MSE. The model with four heuristics does not improve the best fit of the homogeneous expectations model, but its MSE is comparable with those of the best heuristics. Similarly, in the groups 4 and 7 with damping oscillations the LAA and WTR heuristics perform better than other. The model now improves the results for the best heuristics, even if the overall fit is not as good as for the other groups.

Recalling the simulations of Fig. 7 in the groups with oscillation, the following problem with the MSE as a measure of fit becomes clear. Namely, even if our model can generate qualitatively similar oscillations, they always have different frequencies from those which were

Table 4: Squared distance between coefficients of AR-2 estimation

Specification	Group 2	Group 5	Group 1	Group 6	Group 4	Group 7
Fundamental Prediction	0.946	0.671	2.673	3.610	2.311	2.002
ADA – exp prices	0.239	0.006	2.182	2.898	1.691	1.494
WTR – exp prices	0.066	0.529	0.383	0.627	0.203	0.165
STR – exp prices	1.494	2.583	0.112	0.020	0.240	0.342
AAA – exp prices	1.095	1.848	0.010	0.038	0.045	0.094
LAA – exp prices	0.747	1.544	0.003	0.050	0.003	0.013
ADA – fitted prices	0.100	0.000	1.584	2.159	1.385	1.157
WTR – fitted prices	0.068	0.343	0.262	0.435	0.174	0.139
STR – fitted prices	1.358	2.192	0.078	0.001	0.147	0.242
AAA – fitted prices	1.036	1.755	0.005	0.029	0.038	0.083
LAA – fitted prices	0.640	1.277	0.000	0.033	0.000	0.004
4 heuristics (Table 2)	0.383	0.744	0.011	0.008	0.157	0.239
4 heuristics (fitted)	0.144	0.499	0.009	0.003	0.121	0.048

observed in the experiment.¹¹ Consequently, big errors will be generated at the periods when oscillations in experiment and simulations are in different phases. To deal with this problem, we use indirect inference technique. On the first stage, we estimate an AR(2) econometric model both on the experimental and on the simulational data. On the second stage we compute the Euclidean distance between estimators. Results of such statistics are reported in Table 4.

Our main focus on those groups where the MSE was not a good measure, i.e. on the groups 1, 6, 4 and 7. Notice that again the fundamental strategy performs extremely bad. Also the model with ADA heuristics, which was leading in two converging groups, generates large deviation from the underlying experimental estimates. In fact, in groups with constant oscillations the STR and LAA heuristics perform better than others. The switching model does not improve their fit, but generates similar results, which are at any rate better than the results of the homogeneous model with other heuristics. In the groups with damping oscillations, the LAA heuristic gives the best fit, but the model with four heuristics is better than the second best.

To summarize, even if different heuristics can be better in fitting the experimental data of different sessions, the model with four heuristics always performs not worse than the second best heuristic and in some cases even improves the fit. Notice that these results are obtained without fit of parameters and initial impacts. The main advantage of the model with four heuristics is, of course, that this model can be applied to all experimental sessions.

4.3 Analysis of the Deterministic Skeleton of the Model

Previous simulations pose a number of interesting theoretical questions concerning the dynamics produced by the model: How general is its path-dependent property? Are fluctuations only short-run phenomena or they do not die out in the long run as well? Is the role of noise crucial for generating fluctuations?, etc. To address these questions, in this Section we consider the deterministic skeleton of system (4.7), letting $\varepsilon_t = 0$, and we analyze its dynamic properties. To deal with non-autonomous system, in our analysis we substitute the LAA rule

¹¹Simulations of the model show that generated frequencies are more affected by the choice of heuristics than by the learning parameters. With other extrapolative coefficients in the heuristics or with additional heuristics, the quantitative fit of the model can be improved. Recall, however, that our choice of heuristics was driven by the estimation of the experimental data and simplicity of the model.

by analogous AAA heuristic with fixed anchor p^f instead of time varying anchor p_{t-1}^{av} .

4.3.1 Local stability of four heuristic model of evolutionary switching

For the sake of generality let us introduce the following four heuristics,¹² one of which is adaptive and three other are extrapolative, consistent in p^f , rules:

$$\begin{aligned} p_{1,t+1}^e &= w p_{t-1} + (1-w) p_{1,t}^e, \\ p_{h,t+1}^e &= (1 - \beta_{h,1} - \beta_{h,2}) p^f + \beta_{h,1} p_{t-1} + \beta_{h,2} p_{t-2} \quad \text{for } h = 2, 3, 4. \end{aligned} \quad (4.9)$$

The resulting dynamics can be described by a multi-dimensional system. As following result demonstrates, the price dynamics can be constant only on the fundamental level.¹³

Proposition 4.1. *Assume that dynamics of an asset pricing model with evolutionary switching (4.7) in deterministic skeleton generates constant price level p^* . Then price is on the fundamental level, $p^* = p^f$. Furthermore, the share of robots is also fixed and equal to zero, and all heuristics with non-zero weights give fundamental forecast.*

Proof. See Appendix C. □

Turning now to the local stability analysis of the fundamental steady-state, notice from (4.8) that, analogously to the homogeneous expectation model in Section 3, local stability of price dynamics at the fundamental steady-state is not affected by the dynamics of robot traders. Eliminating robot traders from the dynamics we obtain differentiable system. Standard analysis of its Jacobian leads to the following

Proposition 4.2. *The fundamental steady-state of an asset pricing model with evolutionary switching (4.7) (in deterministic skeleton and with fixed anchor in the anchoring and adjustment heuristic) is locally stable if (i) parameters η and δ are less than 1, and (ii) all the roots of polynomial*

$$P(\mu) = \mu^2 \frac{w}{4(1+r)} + (1-w-\mu) \left(\mu^2 - \mu \frac{\beta_{2,1} + \beta_{3,1} + \beta_{4,1}}{4(1+r)} - \frac{\beta_{2,2} + \beta_{3,2} + \beta_{4,2}}{4(1+r)} \right). \quad (4.10)$$

lie inside the unit circle.

The fundamental steady-state is unstable, if at least one of the roots of polynomial (4.10) is outside the unit circle.

Proof. See appendix D where straight-forward computations show that Jacobian of the system has eigenvalues 0, η and δ (of multiplicity 4), as well as three other eigenvalues which are roots of polynomial (4.10). □

When the heuristic coefficients are specified, the roots of the third-order polynomial $P(\mu)$ can be computed. Notice that in general the local stability does not depend on the intensity of choice β . Furthermore, its dependence on the two other parameters of learning process, η and δ , is also limited. As soon as $\delta \neq 1$, i.e. impacts of heuristics are not “frozen” over time, and $\eta < 1$, i.e. agents discount their past performances, the local stability conditions are completely determined by polynomial (4.10) and depend only on the actual heuristics.

¹²To get a model of the previous session, one sets $\alpha_2 = 0, \beta_{2,1} = 1.4, \beta_{2,2} = -0.4$ in the rule $h = 2$, $\alpha_3 = 0, \beta_{3,1} = 2.3, \beta_{3,2} = -1.3$ in the rule $h = 3$ and $\alpha_4 = 30, \beta_{4,1} = 1.5, \beta_{4,2} = -1$ in the rule $h = 4$.

¹³This result holds also for the non-autonomous system with LAA heuristic.

Parameters η and δ , being the eigenvalues of the Jacobian matrix, affect, however, the speed of convergence.

It does not seem to be possible to derive the roots of polynomial (4.10) directly. Therefore we will compute them using numerical methods for four heuristics ADA, WTR, ATR and AAA defined in Table 1. For these rules the roots of polynomial $P(\mu)$ are given as

$$\mu_1 = 0.473668, \quad \mu_2 = 0.634594 - 0.268898i, \quad \mu_3 = 0.634594 + 0.268898i.$$

The modulus of the complex eigenvalues is equal to 0.689214. Thus, the fundamental steady-state is locally stable. Consistent with this result, all simulations which we performed with *deterministic version* of four heuristics model presented in the previous section did converge to the fundamental steady-state.

At this point we can conjecture that a small amount of noise ε_t representing in the experiment demand/supply shocks was crucial in generating oscillating time series in groups 1 and 6. Two caveats are necessary, however. First, the fundamental steady state can be locally unstable under different pool of heuristics. We shall see in Section 4.3.3 that two heuristics, unstable under homogeneous expectations STR and stable under homogeneous expectations AAA, would be sufficient to generate non-converging dynamics in the evolutionary model. Second, even if the fundamental steady-state is stable, the other attractors may co-exist.

4.3.2 Local stability of four heuristics model with constant impacts

Which rules have to be included to the heuristic pool to generate non-converging dynamics in the model? It is intuitively clear that dynamics would not converge if high enough impact is given to the “unstable” heuristics, i.e. those which generate unstable dynamics under homogeneous expectations.

To justify this intuition let us consider an auxiliary version of our model, when impacts of different heuristics on the price are not changing over time. Formally, such *constant impacts model* corresponds to the special case of our evolutionary model with $\delta = 1$. Assuming that the forecasts are given in general formulation (4.9), price evolution in the constant impacts model is described by

$$p_t - p^f = \frac{1}{1+r} \exp\left(-\frac{1}{200}|p_{t-1} - p^f|\right) \sum_{h=1}^4 n_h (p_{h,t+1}^e - p^f),$$

where impacts of heuristics n_h are arbitrary constants summed up to 1. It is straight-forward to find out that the dynamics of such model is stable when all the roots of the polynomial

$$P_1(\mu) = \mu^2 \frac{n_1 w}{1+r} + (1-w-\mu) \left(\mu^2 - \frac{\mu}{1+r} \sum_{h=2}^4 n_h \beta_{h,1} - \frac{1}{1+r} \sum_{h=2}^4 n_h \beta_{h,2} \right) \quad (4.11)$$

lie inside the unit circle. Comparing $P_1(\mu)$ with stability polynomial in Proposition 4.2, we observe that local stability of evolutionary model is governed by the local stability of the constant impacts model with *equal* impacts. This result is not surprising. Indeed, evolutionary model tends to choose the best performed heuristic at any step, and assigns the impacts to the forecasting rules according to their performances. In the fundamental steady-state all four heuristics perform equally well, so if dynamics converge to this steady-state all heuristics will have equal impacts.

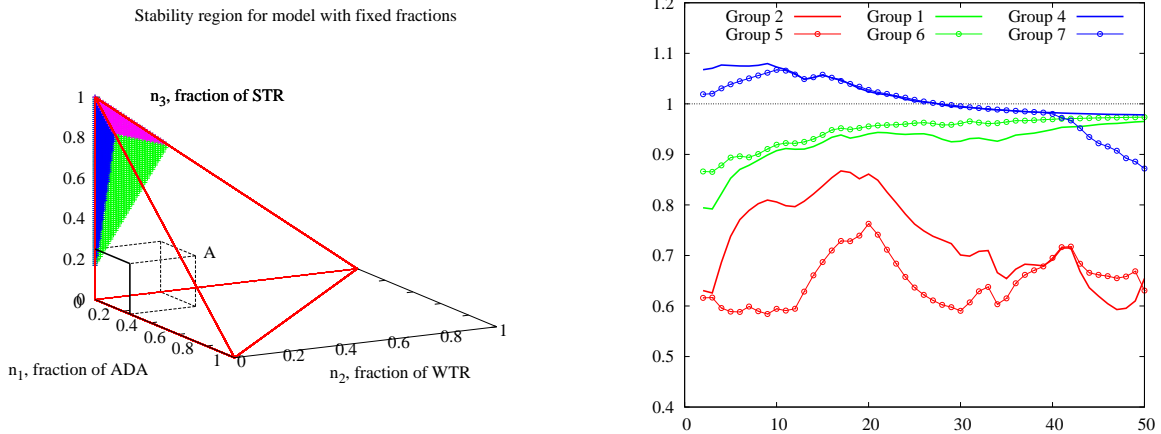


Figure 9: **Local stability of fundamental steady-state in the fixed fraction model.** **Left panel:** The fundamental steady-state is unstable in a model with four heuristics, when fixed impacts (n_1, n_2, n_3) belong to the filled conic region of the unit simplex. **Right panel:** Evolution of the modulus of the largest eigenvalue of polynomial (4.11) during simulations of evolutionary model.

However, the dynamics of the constant impacts model can be unstable if the distribution of heuristics is not uniform. In the right panel of Fig. 9 we show simplex

$$\Delta_4 = \left\{ (n_1, n_2, n_3, n_4) : \sum_{h=1}^4 n_h = 1, n_h \geq 0 \quad \forall h \right\}$$

of all possible impacts. Inside this simplex we draw a region containing all the points where the system with fixed impacts and four heuristics defined in Table 1 is unstable. This instability region was obtained numerically by evaluating the roots of polynomial $P_1(\mu)$ in (4.11) for different values of impacts n_h . The instability region has conic shape concentrating close to the upper left vertex of the simplex where the STR rule has the highest impact. Notice that among four heuristic this is the only rule which generates unstable dynamics when used alone. Consequently, if the impact of STR is relatively high and the impacts of remaining three heuristics are low, the dynamics of fixed fraction model are unstable.

The simplex shown in Fig. 9 also illustrates that the point of equal distribution of impacts (i.e. point A with $n_1 = n_2 = n_3 = n_4 = 0.25$) does not belong to the region of instability. As we said above, it implies that the fundamental steady-state of evolutionary switching model is locally stable for four chosen heuristics. It is still, however, interesting to investigate whether the evolving distribution of impacts generated by such a model and illustrated in Fig. 8 corresponds to the stability or instability in the model with fixed impacts. The right panel of Fig. 9 shows an evolution of the largest eigenvalue of polynomial $P_1(\mu)$ for simulations discussed in Section 4.2. As expected in the converging groups 2 and 5 the distribution of impacts is always such that the constant impact model would be stable. In oscillatory groups 1 and 6 the impacts are evolving to the state where the constant fraction model is close to bifurcation. Finally, in groups 4 and 7 the initial impacts correspond to the instability but over time the model is stabilizing.

4.3.3 Persistent fluctuations in an evolutionary model

Observing the simplex in Fig. 9, one can easily understand that the evolutionary switching model with two heuristics, STR and AAA, will generate the dynamics with *unstable funda-*




Figure 10: **Bifurcation Diagrams in the model with two Heuristics.** 100 points after 600 transitory steps are shown for the evolutionary switching model with STR and AAA heuristics without noise. Initial impacts of the heuristics are equal. Benchmark parameters are $\beta = 0.4$, $\eta = 0.7$ and $\delta = 0.9$, while extrapolation parameter in STR heuristic $\gamma = 1.3$. **Left panel:** Bifurcation diagram with respect to the extrapolation parameter γ . Dynamics converges to the fundamental steady-state for low values of γ and to the quasi-cycle for high values of γ . **Right panel:** Bifurcation diagram with respect to the memory parameter η . Dynamics oscillates wildly for low values of η and converges to a quasi-cycle with a small period for high values of η .

mental steady-state. Indeed, on the vertical edge of simplex $n_1 = n_2 = 0$, so it corresponds to a model with constant fractions of these two heuristics. Stability of the fundamental steady-state in evolutionary model then depends on the middle point of this edge, which belongs to the unstable region.

Two panels of Fig. 10 show the bifurcation diagram of the model with two heuristics with respect to the extrapolation coefficient γ in the STR heuristic (left) and with respect to the memory parameter η (right). The autonomous system in the case of competing STR and AAA heuristics undergoes Neimark-Sacker bifurcation when the coefficient of extrapolation of the STR, γ , becomes large. Fundamental steady-state loses its stability and endogenous fluctuations are generated. When parameter $\gamma = 1.3$ the fundamental equilibrium is unstable. Small values of memory parameter η imply, then, that the agents forget the previous performances of both heuristics quite fast. Since the STR is typically self-reinforcing on the short time scale, this heuristic will often dominate despite those errors which it does when trend is reverting. Consequently, oscillations are especially large for small η . When memory increases the model tends to produce smaller fluctuations. This is because the STR heuristic has quite low performance and is not used very often.

4.4 One-step ahead predicted time-series

All our simulations so far represented the “simulated paths” when the system (4.7) governs the initialized simulations and all the model forecasting errors are ignored. Due to path-dependence feature of the model, these forecasting errors are accumulated, causing sometimes a big discrepancy between the simulated and observed time series. For example, the model generated incorrect phase in the oscillatory groups. In practice, the forecasting errors made by the dynamical model at each time step have to be used to improve its future predictions.

There are two places in the dynamical system (4.7) where “correct” (i.e., observed in the experiment) price, p_t^{exp} , can be employed instead of the predicted price (i.e., generated by the

Table 5: MSE of the one-step ahead forecast for different groups and different specifications.

Specification	Group 2	Group 5	Group 1	Group 6	Group 4	Group 7
Fundamental Prediction	18.037	11.797	15.226	8.959	291.376	22.047
naive	0.060	0.062	3.397	2.292	126.162	12.652
AAA	5.537	3.447	2.930	0.863	60.751	5.647
ADA	0.126	0.050	5.440	4.303	185.591	18.825
WTR	0.081	0.132	1.902	1.038	86.254	8.674
STR	0.556	0.612	2.792	0.767	81.523	13.663
LAA	0.433	0.434	0.427	0.603	60.025	5.564
4 heuristics ($\delta = 1$)	0.082	0.158	1.128	0.605	62.865	6.683
4 heuristics (Figs. 11-12)	0.066	0.103	0.426	0.266	40.766^a	4.148
4 heuristics (best fit)	0.057	0.035	0.405	0.188	33.653 ^a	2.8151
$\beta \in [0, 1]$	0.99	0.99	0.1	0.99	0.13	0.23
$\eta \in [0, 1]$	0.63	0.98	0.99	0.1	0.82	0.45
$\delta \in [0, 1]$	0.80	0.00	0.45	0.78	0.60	0.44

^a Computed for $\beta = 0.1$, $\eta = 0.7$ and $\delta = 0.9$.

model) p_t . First, the forecasts of the future price by all four heuristics can make use of the observed prices. Second, the square deviation in the performance measure can be computed on the basis of the observed price. When the experimental prices feed back to the model in this way, the initial values of prices and heuristics' impacts which we use in the simulations are not important, since after a few periods the actual price dynamics will mostly govern the system.

Fig. 11 compares the experimental data with the one-step ahead predictions made by our model. For these simulations we chose the benchmark parameters $\beta = 0.4$, $\eta = 0.7$ and $\delta = 0.9$, with an exception of the group 4 for which $\beta = 0.1$.¹⁴ In all the simulations the initial values are such that the prices in the first two periods coincide with the prices observed in the corresponding experimental group, while the initial impacts of heuristics are all equal to 0.25. We again observe that the model can easily reproduce three different qualitative dynamics. However now, when the model forecasting errors are taken into account, the mismatching in the oscillatory groups disappear. Fig. 12 shows how in different groups different heuristics are taking the lead after starting with identical uniform distribution. It is interesting to observe how the groups with damping oscillations go through three different phases where the STR, the LAA and the ADA heuristics subsequently dominate.

Table 5 compares the mean squared error of such one-step ahead prediction for different models: fundamental predictions, six homogeneous expectations models (including also naive expectations), heterogeneous model with 4 heuristics and fixed fractions (corresponding to $\delta = 1$), heterogeneous model with 4 heuristics and benchmark parameters ($\beta = 0.4$, $\eta = 0.7$ and $\delta = 1$), and, finally, the heterogeneous model with 4 heuristics fitted by means of the grid search in the parameter space (the last three lines show the corresponding values of parameters for which the best fit is achieved). Comparing different MSE we observe that the evolutionary model with 4 heuristics generates in average better one-step ahead predictions than the fundamental forecasting, the homogeneous models (with the exception of group 5; the MSE for the best among the four models is shown bold for each group) and the heterogeneous model without evolutionary learning. We also observe that even if the fitting of the model can improve the MSE, this improvement is relatively small if compared with improvement which

¹⁴Relatively large prediction errors, which our heuristics are making for this group, cause overflow of the numerical computations for high values of β and/or η .

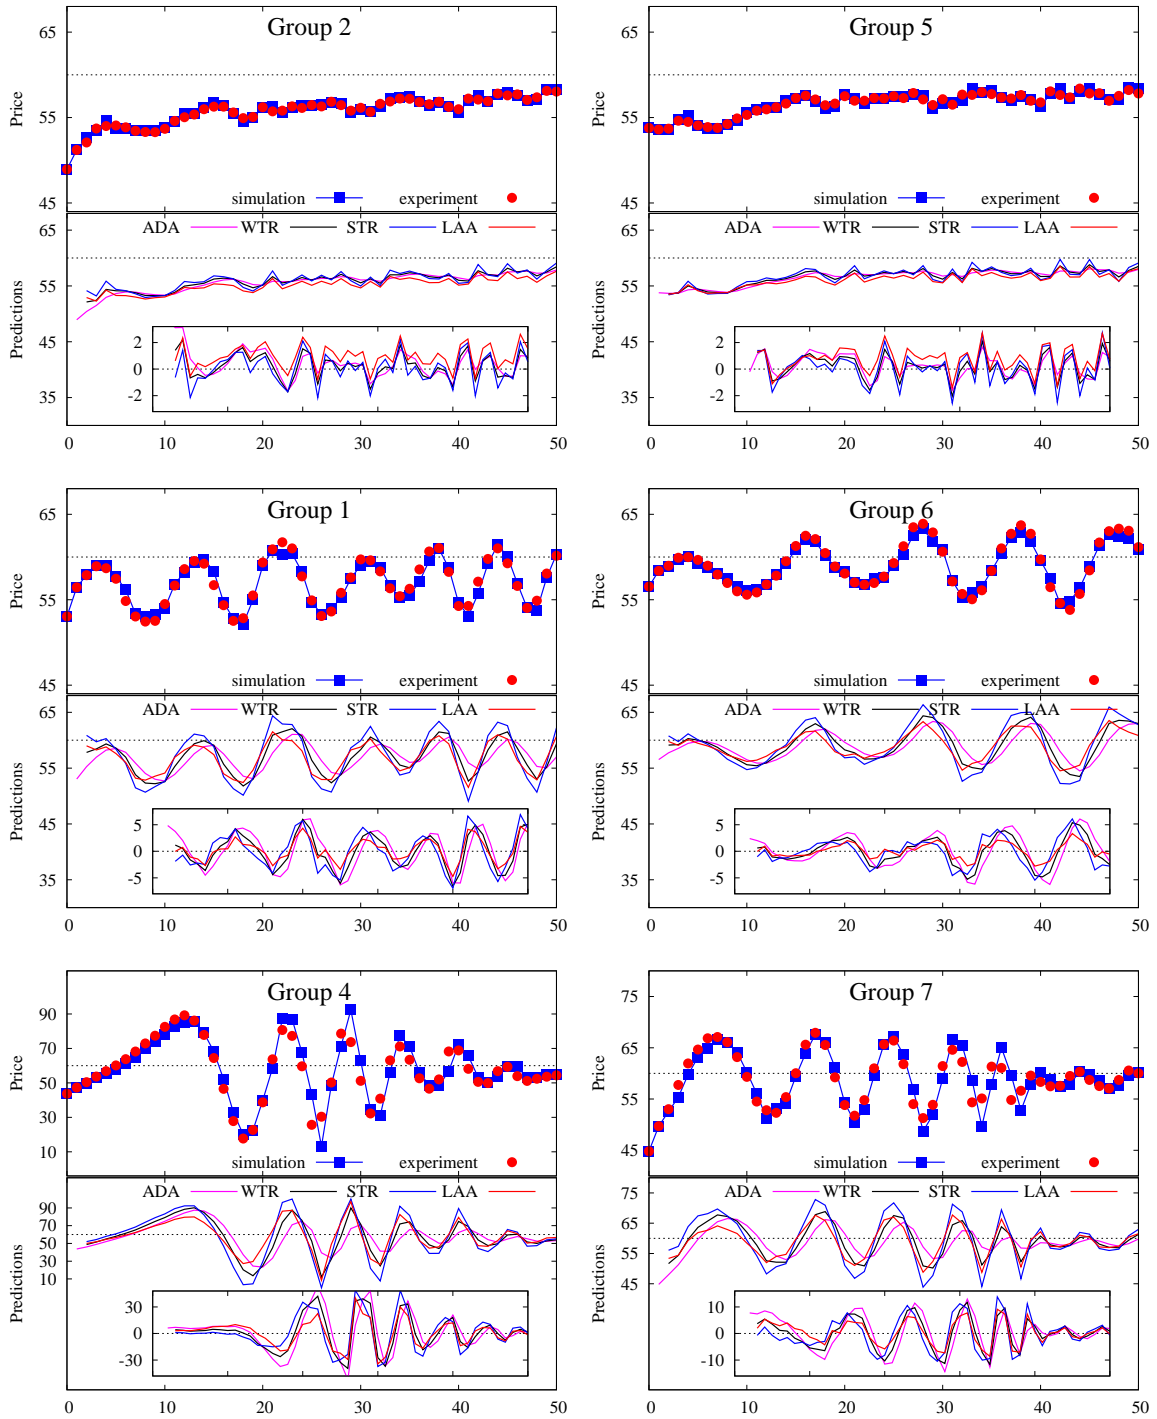


Figure 11: **Laboratory experiments and one-step ahead predictions of the evolutionary model.** Upper parts of panels show prices for laboratory experiments in different groups (red) with corresponding one-step ahead predictions of the evolutionary model (blue). **Lower parts of panels** show predictions and forecasting errors (inner frames) of four heuristics: adaptive expectations (ADA, purple), weak trend followers (WTR, black), strong trend followers (STR, blue) and anchoring adjustment heuristic (LAA, red).

the heterogeneous model makes with respect to the homogeneous models.

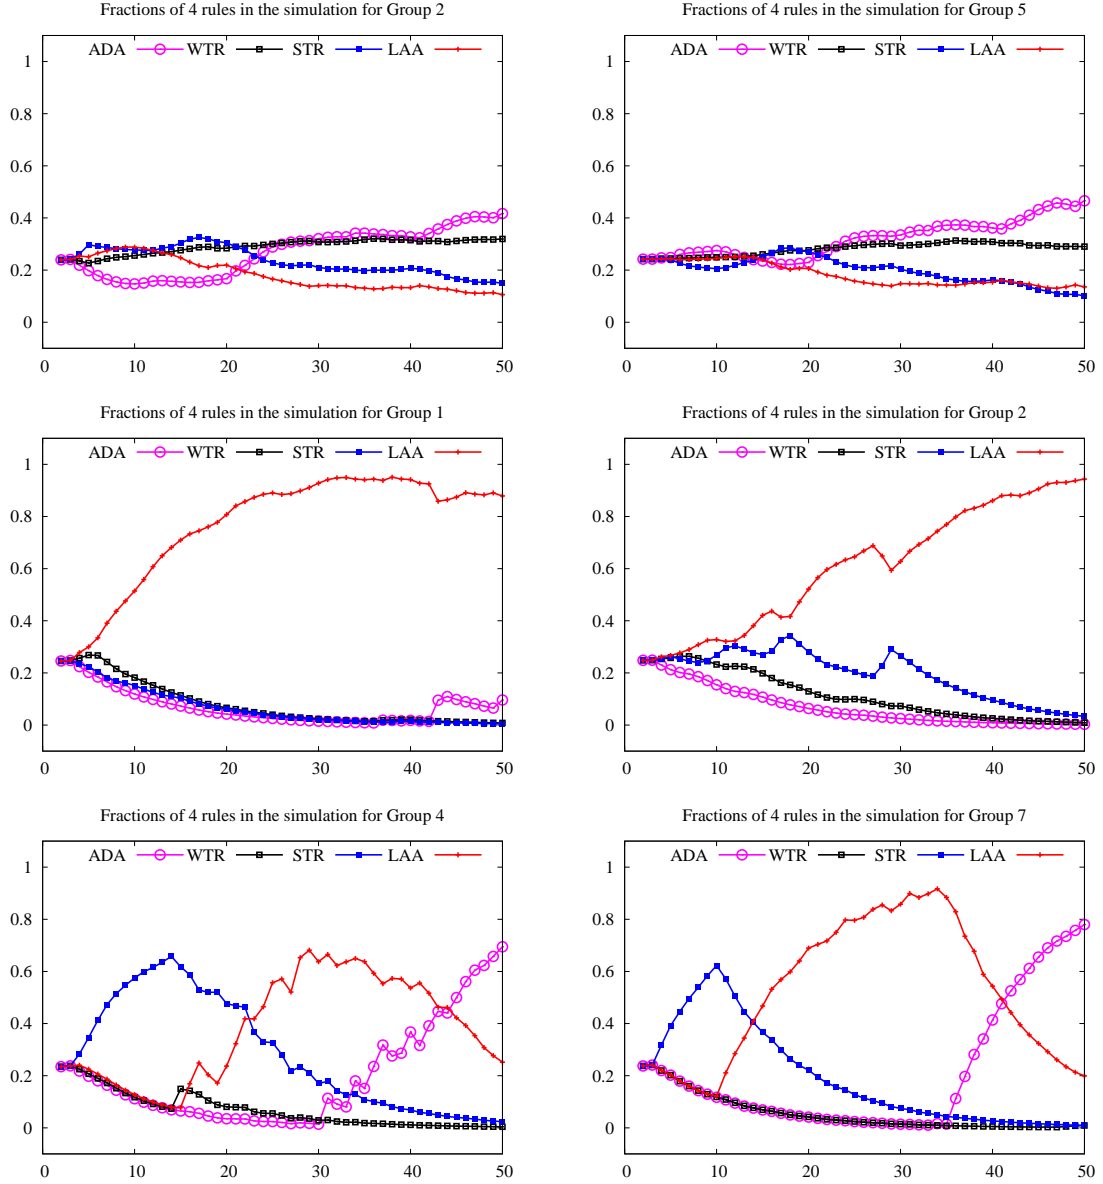


Figure 12: **Evolution of heuristic impacts during the one-step ahead predictions of the model.** Fractions of four forecasting heuristics: adaptive expectations (ADA, purple), weak trend followers (WTR, black), strong trend followers (STR, blue) and anchoring adjustment heuristic (LAA, red).

5 Conclusion

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 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.

Our approach is similar to other models of reinforcement learning: Arthur (1991), Arthur (1993), Erev and Roth (1998) and ?. However, our model is built in a different environment from those which are studied in the standard game theory. Namely, agents in our framework do not have a well defined strategies and they do not know the payoff matrix, which is, in addition, is changing over time in a path-depended manner. To our best knowledge, the model presented in this paper is the first learning model explaining *different* time series patterns in the *same* laboratory experiments.

References

- ALCHIAN, A. A. (1950): "Uncertainty, Evolution, and Economic Theory," *The Journal of Political Economy*, 58(3), 211–221.
- ARTHUR, W. (1991): "Designing Economic Agents that Act like Human Agents: A Behavioral Approach to Bounded Rationality," *American Economic Review*, 81, 353–359.
- (1993): "Economic Agents that Behave like Human Agents," *Journal of Evolutionary Economics*, 3, 1–22.
- BRAY, M., AND N. SAVIN (1986): "Rational Expectations Equilibria, Learning, and Model Specification," *Econometrica*, 54(5), 1129–1160.
- BROCK, W. A., AND C. H. HOMMES (1997): "A Rational Route to Randomness," *Econometrica*, 65(5), 1059–1095.
- (1998): "Heterogeneous Beliefs and Routes to Chaos in a Simple Asset Pricing Model," *Journal of Economic Dynamics and Control*, 22, 1235–1274.
- BULLARD, J. B. (1994): "Learning equilibria," *Journal of Economic Theory*, 64, 468–485.
- CAMERER, C. F., AND E. FEHR (2006): "When Does 'Economic Man' Dominate Social Behavior?," *Science*, 311(5757), 47–52.
- CAMPBELL, J., A. LO, AND A. MACKINLAY (1997): *The Econometrics of Financial Markets*. Princeton University Press, Princeton.
- DE BONDT, W. F., AND R. THALER (1989): "Does the Stock Market Overreact?," *Journal of Finance*, 40(3), 793–805.

- DUFFY, J. (2008): “Experimental Macroeconomics,” in *New Palgrave Dictionary of Economics*, ed. by S. Durlauf, and L. Blume. New York: Palgrave Macmillan, forthcoming.
- EREV, I., AND A. E. ROTH (1998): “Prediction how people play games: Reinforcement learning in games with unique strategy equilibrium,” *American Economic Review*, 88, 848–881.
- EVANS, G. W., AND S. HONKAPOHJA (2001): *Learning and Expectations in Macroeconomics*. Princeton University Press.
- FAMA, E. F. (1970): “Efficient Capital Markets: a Review of Theory and Empirical Work,” *Journal of Finance*, 25, 383–417.
- FRIEDMAN, M. (1953): *Essays in Positive Economics*. University Of Chicago Press.
- HEY, J. D. (1994): “Expectation formation: Rational or adaptive or ...?,” *Journal of Economic Behavior and Organization*, 25, 329–349.
- HOMMES, C. (2006): “Heterogeneous Agent Models in Economics and Finance,” in *Handbook of Computational Economics Vol. 2: Agent-Based Computational Economics*, ed. by K. Judd, and L. Tesfatsion. Elsevier/North-Holland (Handbooks in Economics Series).
- HOMMES, C., J. SONNEMANS, J. TUINSTRA, AND H. V. D. VELDEN (2005): “Coordination of Expectations in Asset Pricing Experiments,” *Review of Financial Studies*, 18(3), 955–980.
- KAHNEMAN, D. (2003): “Maps of Bounded Rationality: Psychology for Behavioral Economics,” *American Economic Review*, 93, 1449–1475.
- LUCAS, R. E., AND E. C. PRESCOTT (1971): “Investment Under Uncertainty,” *Econometrica*, 39(5), 659–681.
- MARIMON, R., S. E. SPEAR, AND S. SUNDER (1993): “Expectationally Driven Market Volatility: An Experimental Study,” *Journal of Economic Theory*, 61(1), 74–103.
- MUTH, J. F. (1961): “Rational Expectations and the Theory of Price Movements,” *Econometrica*, 29(3), 315–335.
- PETERSON, S. P. (1993): “Forecasting dynamics and convergence to market fundamentals: Evidence from experimental asset markets,” *Journal of Economic Behavior and Organization*, 22(3), 269–284.
- SARGENT, T. J. (1993): *Bounded Rationality in Macroeconomics*. Oxford University Press.
- SHILLER, R. (1981): “Do Stock Prices Move Too Much to be Justified by Subsequent Changes in Dividends?,” *American Economic Review*, 71(3), 421–436.
- SIMON, H. A. (1957): *Models of Man: Social and Rational*. John Wiley, New York.
- SMITH, V. L. (1962): “An Experimental Study of Competitive Market Behavior,” *The Journal of Political Economy*, 70(2), 111–137.
- SMITH, V. L., G. L. SUCHANEK, AND A. W. WILLIAMS (1988): “Bubbles, Crashes, and Endogenous Expectations in Experimental Spot Asset Markets,” *Econometrica*, 56(5), 1119–1151.

TVERSKY, A., AND D. KAHNEMAN (1974): “Judgement under uncertainty: Heuristics and biases,” *Science*, 185, 1124–1130.

WILLIAMS, A. W. (1987): “The Formation of Price Forecasts in Experimental Markets,” *Journal of Money, Credit and Banking*, 19(1), 1–18.

APPENDIX

A Proof of Proposition 3.1

From the general relation (3.2) it follows that

$$p_{t+1}^e - p^f = w(p_{t-1} - p^f) + (1-w)(p_t^e - p^f) = (p_t^e - p^f) \left(w \frac{1-n_{t-1}}{1+r} + 1-w \right).$$

The expression in the last parenthesis is a convex combination of 1 and $(1-n_{t-1})/(1+r) < 1$. For positive weight w such combination is always less than 1. Therefore, the dynamical system defines a *contraction* of expectations, which then must globally converge to p^f . The price realization in this point is uniquely defined from (2.5) as $p^* = p^f$. Finally, the evolution of robot traders implies that $n^* = 0$ in this fixed-point.

B Proof of Proposition 3.2

Consider the steady-state (p^*, n^*) with consistent forecasting rule, and notice that (3.2) implies that either $p^* = p^f$ or $1-n^* = 1+r$. The second case is impossible, so $p^* = p^f$ and, therefore, $n^* = 0$.

Using (3.2), the dynamics in deviations is given by

$$p_t - p^f = \exp\left(-|p_{t-1} - p^f|/200\right) \left(\frac{\beta_1}{1+r}(p_{t-1} - p^f) + \frac{\beta_2}{1+r}(p_{t-2} - p^f) \right). \quad (\text{B.1})$$

The first term in the right hand-side is never greater than 1. Thus, dynamics of (B.1) is a superposition of contraction with linear process of second order

$$(1+r)x_t = \beta_1 x_{t-1} + \beta_2 x_{t-2} \quad (\text{B.2})$$

with $x_t = p_t - p^f$. If the latter dynamics is locally stable, the steady-state p^f of original dynamics (B.1) will be also locally stable. Furthermore, since the exponential term in (B.1) is equal to 1 in the steady-state, the linear parts of the dynamics of the last two processes are the same. Thus, processes (B.1) and (B.2) lose stability simultaneously and through the same bifurcation type.

The Jacobian matrix of (B.2) in the steady-state is given by

$$J = \left\| \begin{array}{cc} \frac{\beta_1}{1+r} & \frac{\beta_2}{1+r} \\ 1 & 0 \end{array} \right\|.$$

The standard conditions for the stability can be expressed through the trace $\text{Tr}(J)$ and the determinant $\text{Det}(J)$ of this matrix, and are given by

$$\text{Tr}(J) < \text{Det}(J) + 1, \quad \text{Tr}(J) > -1 - \text{Det}(J), \quad \text{Det}(J) < 1. \quad (\text{B.3})$$

Furthermore, the dynamics is oscillatory if $\text{Tr}(J)^2 - 4\text{Det}(J) = 0$. The substitution of the values of trace and determinant gives inequalities (3.4) and condition $\beta_1^2 + 4\beta_2(1+r_f) < 0$ for the oscillations.

The bifurcation types can be determined from (B.3), since when one of these inequalities turns to equality, the unit circle is crossed by a corresponding eigenvalue of the system. The second inequality is violated when an eigenvalue becomes equal to -1 , which implies the period-doubling bifurcation. The violation of the last inequality in (B.3) implies that two complex eigenvalues cross the unit circle. This happens under the Neimark-Sacker bifurcation. Finally, consider the first inequality, which is violated when one eigenvalue becomes equal to 1 . It turns out that at this occasion two new steady-states are emerging, which implies that the system exhibits pitchfork bifurcation. Indeed, any steady-state (p^*, n^*) should satisfy to

$$\begin{aligned} (1+r)p^* &= (1-n^*)((1-\beta_1-\beta_2)p^f + (\beta_1+\beta_2)p^*) + n^*p^f + \bar{y} \\ &\Downarrow \\ (1+r)p^* &= p^f + (\beta_1+\beta_2)(p^* - p^f)(1-n^*) + \bar{y} \\ &\Downarrow \\ (1+r)(p^* - p^f) &= (\beta_1+\beta_2)(p^* - p^f)(1-n^*) \end{aligned}$$

Thus, in any non-fundamental steady-state, the fractions of robots $n^* = 1 - (1+r)/(\beta_1+\beta_2)$. Only if this fraction belongs to the interval $(0, 1)$ two other steady-states exist with

$$p_{\pm}^* = p^f \pm 200 \log(1-n^*).$$

(The prediction rule is, of course, inconsistent in both steady-states.) When the first inequality in (B.3) is satisfied, these two steady-state do not exist, but they appear at the moment when the inequality changes its sign.

C Proof of Proposition 4.1

In the steady-state with fixed price p^* , the past price sample average will also be equal to p^* . The dynamics (4.8) in the steady-state with fixed price p^* then is given as

$$\begin{aligned} (1+r)(p^* - p^f) &= (1-n_t) \left(n_{1,t}(w(p^* - p^f) + (1-w)(p_{1,t}^e - p^f)) + \right. \\ &\quad \left. + n_{2,t}(p^* - p^f) + n_{3,t}(p^* - p^f) + n_{4,t}(p^* - p^f) \right). \end{aligned} \tag{C.1}$$

In the state with constant price $p^* = p^f$, the fraction of robots $n_t = 0$, so that the above condition simplifies to $0 = (1-w)n_{1,t}(p_{1,t}^e - p^f)$. Then the ADA rule (if it is in actual use) gives fundamental forecast.

If $p^* \neq p^f$, take the limit of $t \rightarrow \infty$ in (C.1). Since adaptive expectations converge to p^* in such limit, we obtain that $1+r \leq 1$, which is impossible.

D Stability of Evolutionary Model

In the body of the paper we have obtained a model (4.7) describing the dynamics of price and other variables under the evolutionary learning over 4 heuristics. We will write the dynamical system using the general notation for four heuristics introduced in (4.9). Recall that all extrapolative rules are assumed to be consistent in p^f . The dynamics below is written in deviations from fundamental price, both in prices and in forecasts. The variables are introduced as follows

$$x_{1,t}^e = p_t^e - p^f, \quad y_t^e = x_{1,t-1}^e, \quad x_{1,t} = p_t - p^f, \quad x_{2,t} = x_{1,t-1}, \quad x_{3,t} = x_{1,t-1}, \quad x_{4,t} = x_{1,t-3}.$$

The following 14-dimensional system of the first order equations describes the dynamics. It consists of 4 equations describing the evolution of performance measures, 4 variables represent the fractions of different forecasting rules, 1 equation describes the price dynamics, which we will write in deviations, and other 3 equations are needed to take lags of price deviations into account, and finally two equations describe the evolution of adaptive expectation rule.

$$\begin{aligned}
x_{t+1}^e &= w x_{1,t-1} + (1-w) x_t^e \\
y_{t+1}^e &= x_t^e \\
U_{1,t-1} &= -(x_{1,t-1} - y_t^e)^2 + \eta U_{1,t-2} \\
U_{h,t-1} &= -(x_{1,t-1} - \beta_{h,1} x_{3,t-1} - \beta_{h,2} x_{4,t-1})^2 + \eta U_{h,t-2} \quad 2 \leq h \leq 4 \\
n_{1,t} &= \delta n_{1,t-1} + \frac{1-\delta}{Z_{t-1}} \exp\left(\beta \left[-(x_{1,t-1} - y_t^e)^2 + \eta U_{1,t-2} \right]\right) \\
n_{h,t} &= \delta n_{h,t-1} + \frac{1-\delta}{Z_{t-1}} \exp\left(\beta \left[-(x_{1,t-1} - \beta_{h,1} x_{3,t-1} - \beta_{h,2} x_{4,t-1})^2 + \eta U_{h,t-2} \right]\right) \quad 2 \leq h \leq 4 \\
x_{1,t} &= \exp\left(-\frac{1}{200}|x_{t-1}|\right) \frac{1}{1+r} \left(\left[\delta n_{1,t-1} + \frac{1-\delta}{Z_{t-1}} \exp\left(\beta \left[-(x_{1,t-1} - y_t^e)^2 + \eta U_{1,t-2} \right]\right) \right] (w x_{1,t-1} + (1-w) x_t^e) + \right. \\
&\quad \left[\delta n_{2,t-1} + \frac{1-\delta}{Z_{t-1}} \exp\left(\beta \left[-(x_{1,t-1} - \beta_{2,1} x_{3,t-1} - \beta_{2,2} x_{4,t-1})^2 + \eta U_{2,t-2} \right]\right) \right] (\beta_{2,1} x_{1,t-1} + \beta_{2,2} x_{2,t-1}) + \\
&\quad \left[\delta n_{3,t-1} + \frac{1-\delta}{Z_{t-1}} \exp\left(\beta \left[-(x_{1,t-1} - \beta_{3,1} x_{3,t-1} - \beta_{3,2} x_{4,t-1})^2 + \eta U_{3,t-2} \right]\right) \right] (\beta_{3,1} x_{1,t-1} + \beta_{3,2} x_{2,t-1}) + \\
&\quad \left. \left[\delta n_{4,t-1} + \frac{1-\delta}{Z_{t-1}} \exp\left(\beta \left[-(x_{1,t-1} - \beta_{4,1} x_{3,t-1} - \beta_{4,2} x_{4,t-1})^2 + \eta U_{4,t-2} \right]\right) \right] (\beta_{4,1} x_{1,t-1} + \beta_{4,2} x_{2,t-1}) \right) \\
x_{2,t} &= x_{1,t-1} \\
x_{3,t} &= x_{2,t-1} \\
x_{4,t} &= x_{3,t-1}
\end{aligned}$$

where

$$Z_{t-1} = \exp\left(\beta \left[-(x_{1,t-1} - y_t^e)^2 + \eta U_{1,t-2} \right]\right) + \sum_{h=2}^4 \exp\left(\beta \left[-(x_{1,t-1} - \beta_{h,1} x_{3,t-1} - \beta_{h,2} x_{4,t-1})^2 + \eta U_{h,t-2} \right]\right)$$

We are interested in stability of this system near the fixed point with price equal to p^f and zero fraction of “robots”. First of all, recall that the term $\exp(-|x_{t-1}|/200)$ in the equation for price deviations can be ignored, since its first-order approximation in this fixed point is 1. The Jacobian matrix \mathbf{J} of the remaining system is given by

$$\begin{vmatrix}
1-w & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & w & 0 & 0 & 0 \\
1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & \eta & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & \eta & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & \eta & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & \eta & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & \frac{9\beta\eta(1-\delta)}{16} & -\frac{\beta\eta(1-\delta)}{16} & -\frac{\beta\eta(1-\delta)}{16} & -\frac{\beta\eta(1-\delta)}{16} & \delta & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & -\frac{\beta\eta(1-\delta)}{16} & \frac{9\beta\eta(1-\delta)}{16} & -\frac{\beta\eta(1-\delta)}{16} & -\frac{\beta\eta(1-\delta)}{16} & 0 & \delta & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & -\frac{\beta\eta(1-\delta)}{16} & -\frac{\beta\eta(1-\delta)}{16} & \frac{9\beta\eta(1-\delta)}{16} & -\frac{\beta\eta(1-\delta)}{16} & 0 & 0 & \delta & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & -\frac{\beta\eta(1-\delta)}{16} & -\frac{\beta\eta(1-\delta)}{16} & -\frac{\beta\eta(1-\delta)}{16} & \frac{9\beta\eta(1-\delta)}{16} & 0 & 0 & 0 & \delta & 0 & 0 & 0 & 0 \\
\frac{1-w}{4(1+r)} & 0 & 0 & 0 & 0 & 0 & \frac{\delta p^f}{1+r} & \frac{\delta p^f}{1+r} & \frac{\delta p^f}{1+r} & \frac{\delta p^f}{1+r} & \frac{w+\beta_{2,1}+\beta_{3,1}+\beta_{4,1}}{4(1+r)} & \frac{\beta_{2,2}+\beta_{3,2}+\beta_{4,2}}{4(1+r)} & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0
\end{vmatrix}$$

It is straight-forward to check that this Jacobian has multipliers equal to 0 (of multiplicity 3) and η and δ (both of multiplicity 4). The remaining three multipliers are the roots of characteristic polynomial for matrix

$$\mathbf{J}_r = \begin{vmatrix}
1-w & w & 0 \\
\frac{1-w}{4(1+r)} & \frac{w+\beta_{2,1}+\beta_{3,1}+\beta_{4,1}}{4(1+r)} & \frac{\beta_{2,2}+\beta_{3,2}+\beta_{4,2}}{4(1+r)} \\
0 & 1 & 0
\end{vmatrix},$$

This characteristics polynomial is given in (4.10). □