

Evolutionary Selection of Individual Expectations and Aggregate Outcomes*

Mikhail Anufriev[†] Cars Hommes[‡]

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

July 2009

Abstract

In recent ‘learning to forecast’ experiments with human subjects of Hommes, et al. (2005), three different patterns in aggregate price behavior have been observed: slow monotonic convergence, permanent price oscillations and dampened price fluctuations. These different aggregate outcomes are at odds with a representative agent who is fully rational or employs a single adaptive learning rule. We construct a simple model of individual learning, based on performance based evolutionary selection or reinforcement learning among heterogeneous expectations rules, explaining these different aggregate outcomes. Agents are boundedly rational and choose from a small set of simple price prediction rules, such as naive, adaptive or trend following expectations, similar to individual rules estimated in the experiment. Agents update their active rule by evolutionary selection based upon forecasting performance. Simulations show that after some initial learning phase, coordination on a common rule occurs. Out-of-sample predictive power of our model is higher than of the rational or other homogeneous expectations benchmarks. Our results show that heterogeneity in expectations is crucial to describe individual forecasting behavior as well as aggregate price behavior.

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

*We thank the participants of the “Complexity in Economics and Finance” workshop, Leiden, October 2007, the “Computation in Economics and Finance” conference, Paris, June 2008, the “Learning and Macroeconomic Policy” workshop, Cambridge, September 2008, and seminar participants at the University of Amsterdam, Sant’Anna School of Advanced Studies (Pisa), University of New South Wales, Catholic University of Milan, University of Warwick, Paris School of Economics, and the European University Institute (Florence) 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 and by the EU 7th framework collaborative project “*Monetary, Fiscal and Structural Policies with Heterogeneous Agents (POLHIA)*”, grant no.225408.

[†]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.H.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), Erev and Roth (1998) and Camerer and Ho (1999) propose reinforcement learning as an explanation of average behavior in a number of experiments in a game-theoretical setting. 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 forecasting rules have been proposed, e.g., Brock and Hommes (1997) and Branch and Evans (2006), see Hommes (2006) for an extensive overview.

Laboratory experiments with human subjects, with full control over 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). 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 different aggregate market outcomes. Although our model is very simple it fits the experimental data surprisingly well (see, e.g., Fig. 10).

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. In Section 5 we discuss how our model fits the experimental data. Finally, Section 6 concludes.

2 Learning to Forecast Experiments

In this section we discuss the laboratory experiments. Subsection 2.1 recalls the experimental design, Subsection 2.2 focuses on aggregate price behavior, 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 , while the risky asset pays stochastic dividends, independently identically distributed (IID), with mean \bar{y} . 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, assuming all other underlying model assumptions to be satisfied. *Ceteris paribus* learning to forecast experiment data can be used to test various expectations hypotheses, such as rational expectations or adaptive learning models, see the discussion in Duffy (2008).

Participants in the experiments knew that the actual price realization of the risky asset is determined by market clearing on the basis of the investment strategies of the pension fund. The exact functional form of the strategies and the market 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 there 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 the 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 own earnings $e_{i,t}$ and the fundamental parameters (the risk-free interest rate $r = 0.05$ and the dividend mean $\bar{y} = 3$):

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

Note that, since the price p_t is unknown at the beginning of period t , it is not included in 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 a 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 long lasting bubbles. The fraction of robot traders increased in response to the deviations of the asset 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.

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 never became larger than 0.2. Recently, Hommes, Sonnemans, Tuinstra, and van de Velden (2008) ran experiments without robot traders in which long lasting bubbles occurred.




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.

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

$$e_{i,t} = \begin{cases} 1 - \left(\frac{p_t - p_{i,t}^e}{7}\right)^2 & \text{if } |p_t - p_{i,t}^e| < 7, \\ 0 & \text{otherwise,} \end{cases} \quad (2.4)$$

so that forecasting errors exceeding 7 would result in no reward at a given period. At the end of the session the accumulated earnings of every participant were converted to euros (1 point computed as in (2.4) corresponded to 50 cents) and paid out.

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

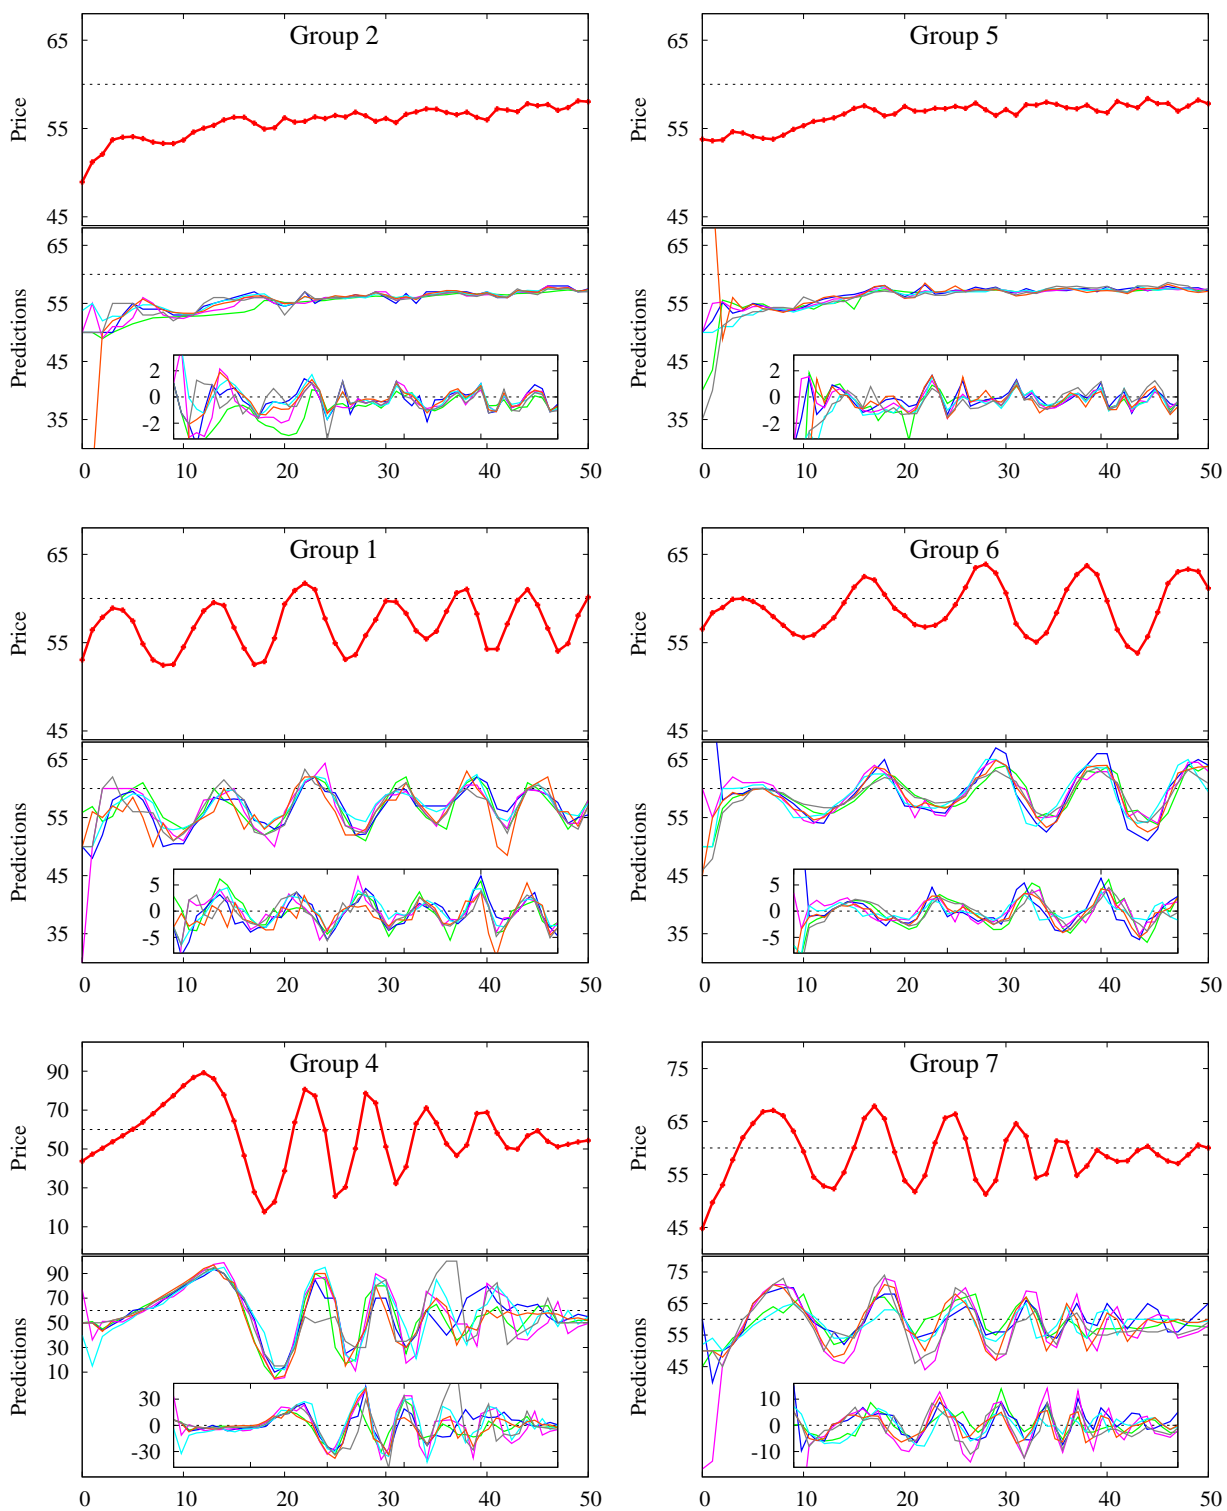


Figure 2: **Three different qualitative outcomes in 6 sessions of the forecasting experiment.** For every panel an *upper part* shows observed prices (red) in comparison with the fundamental level (black); a *lower part* shows individual predictions of 6 participants with forecasting errors in an *inner frame*. In two sessions (**upper panels**) the price converges monotonically to the fundamental value, in two sessions (**middle panels**) price exhibited persistent oscillations, in two sessions (**lower panels**) the price converged through large (observe the change of scale for group 4) but damping oscillations.

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

To summarize, in 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 similar forecasts in every period.

The purpose of this paper is to explain these “stylized facts” simultaneously by a simple model of individual learning behavior.

2.3 Individual Forecasting Rules

Which forecasting rules did individuals use in the learning to forecast experiment? Comparison of the RE benchmark in Fig. 1 with the lab experiments in Fig. 2 suggests that rational expectations is not a good explanation of individual forecasting and aggregate behavior. For the oscillating groups 1 and 6 this is immediately clear. One could perhaps argue that the other groups show a tendency to converge to RE after 50 periods, but in contrast to the RE-benchmark, e.g., in the monotonically converging groups 2 and 5 market prices are consistently below the RE-benchmark 60. Moreover, RE does not explain the slowly dampened oscillating patterns.

To get some intuition for individual behavior it is useful to look at some of the time evolution of individual predictions. Fig. 3 shows time series of some (lagged) individual forecasts together with the realized price. The timing in the figure is important. For every time t on the horizontal axes we show the price p_t together with the individual two-period ahead forecast $p_{i,t+2}^e$ of that price by some participant i . In this way we can infer graphically how the two-period ahead forecast $p_{i,t+2}^e$ uses the last observed price p_t . For example, if they coincide, i.e., $p_{i,t+2}^e = p_t$, it implies naive expectations in period t .

In group 2, subject 5 extrapolates price changes in the early stage of the experiment (see the upper left panel), but, starting from period $t = 6$, uses a simple *naive rule* $p_{t+2}^e = p_t$.

³Price dynamics in group 3 (not shown, but see the concluding remarks) 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 between individual predictions. Even in 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.

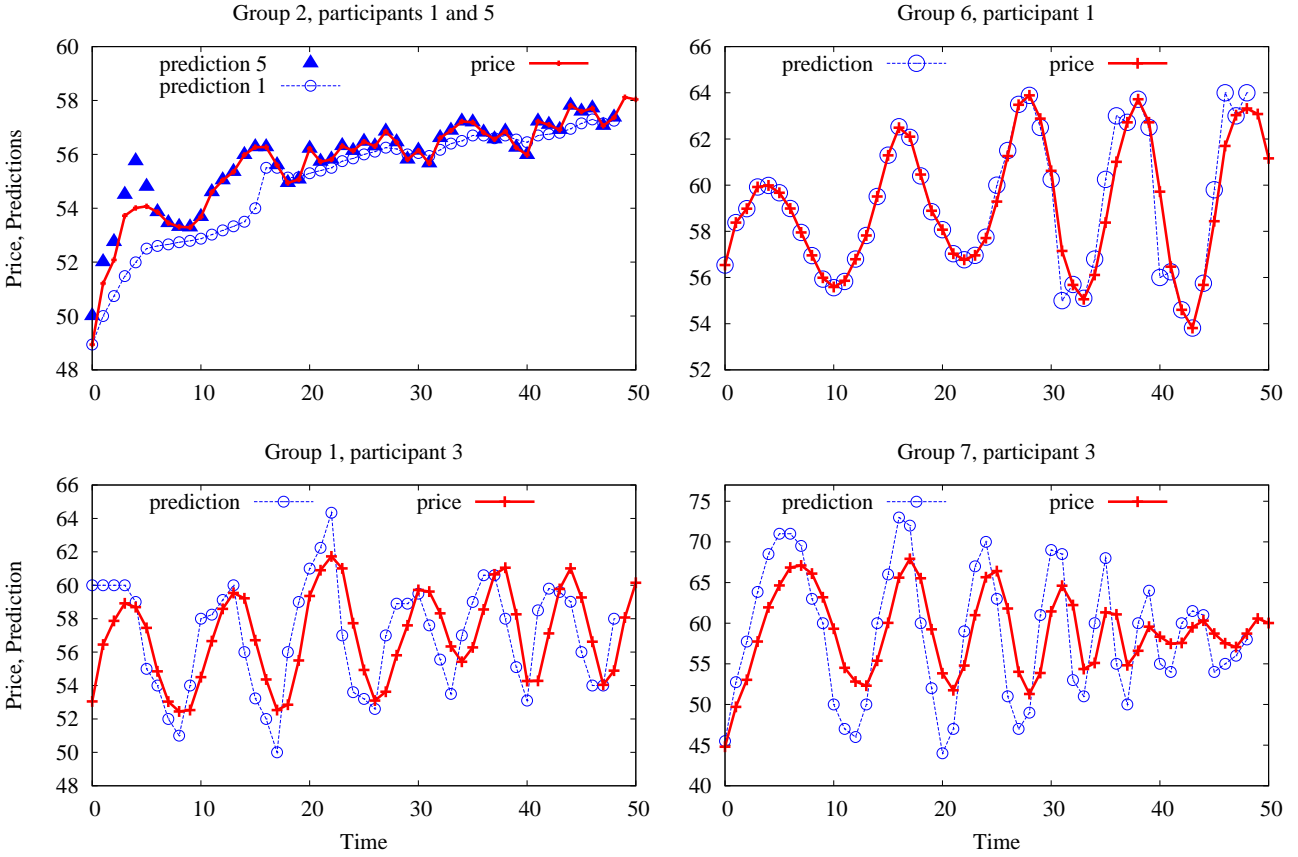


Figure 3: **Switching of the experiment’s participants between simple rules.** For any point on the abscissa, representing time t , the price p_t (red) and the forecast $p_{i,t+2}^e$ (blue) are shown. This forecast, $p_{i,t+2}^e$, was made immediately after the announcement of the price p_t .

In other words, in period 6, subject 5 switched from an extrapolative to a naive forecasting rule. Subject 1 from the same group used a “smoother”, *adaptive* forecasting strategy, always predicting a price between the previous forecast and the previous price realization. These graphs already suggest individual heterogeneity in forecasting strategies.

In the oscillating group 6, subject 1 used naive expectations in the first half of the experiment (until period 24, see the upper right panel). Naive expectations however, lead to prediction errors in an oscillating market, especially when the trend reverses. In period 25 subject 1 switches to a different, trend extrapolating prediction strategy. In what follows, this subject uses a trend extrapolating strategy switching back to the naive rule at periods of expected trend reversal (e.g., in periods 27 and 28, 32 and 33, 37 and 38, 42, 43 and 44, and 47).

Interestingly, participant 3 from another oscillating group 1 starts out predicting the fundamental price, i.e., $p_{t+1}^e = p^f = 60$ in the first four periods of the experiments (see the lower left panel). But since the majority in group 1 predicts a lower price, the realized price is much lower than the fundamental, causing participant 3 to switch to a different, trend extrapolating strategy. Trend extrapolating predictions overshoot the realized market price at the moments of trend reversal. Towards the end of the experiment participant 3 learned to anticipate the trend changes to some extent.

Finally, in group 7 with dampened price oscillations subject 3 started out with a strong

trend extrapolation (see the lower right panel). Despite very large prediction errors (and thus low earnings) at the turning points, this participant stuck to strong trend extrapolation. Perhaps only in the last 4 periods some kind of adaptive expectations strategy was used.

Using the individual experimental data HSTV estimated the forecasting rules for each individual based on the last 40 periods, to allow for some learning phase⁵. Simple linear rules of the form

$$p_{h,t+1}^e = \alpha_h + \beta_h p_{t-1} + \gamma_h (p_{t-1} - p_{t-2}) + \delta_h p_{h,t}^e \quad (2.5)$$

were estimated. These rules are so-called *first order heuristics*, since they only use the last forecast, the last observed price and the last observed price change to predict the future price. Remarkably, individual forecasts are well explained by the first order heuristics: for 33 out of 42 participants (i.e., for 78%) an estimated linear rule falls into this simple class with an R^2 typically higher than 0.80. In fact, within the class of first order heuristics three extremely simple rules came up from the estimation results, characterizing the three different observed aggregate outcome: convergence, permanent oscillations and dampened oscillations.

Participants from the converging groups 2 and 5 often used an *adaptive expectations* rule of the form:

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

with weight $0 \leq w \leq 1$. Note that at the moment when the forecast for the price p_{t+1} is submitted, the price p_t is still unknown (see Eq. 2.2), so that the last observed price is p_{t-1} . At the same time, the last forecast p_t^e is of course known when forecasting p_{t+1} . Notice also that for $w = 1$, we obtain the special case of *naive expectations*.⁶ The individual forecast series shown in the upper left panel of Fig. 3 are examples of estimated rules of the form (2.6), with $w = 1$ for subject 5 and $w \simeq 0.25$ for subject 1.

Especially for the subjects in the permanent and the dampened oscillating groups estimation revealed a simple *trend-following* forecasting rule of the form:

$$p_{t+1}^e = p_{t-1} + \gamma (p_{t-1} - p_{t-2}), \quad (2.7)$$

where $\gamma > 0$. This rule has a simple behavioral interpretation: the forecast uses the last price observation and adjusts in the direction of the last price change. The extrapolation coefficient γ measures the strength of the adjustment. The estimates of this coefficient ranged from relatively small extrapolation values, $\gamma = 0.4$, to quite strong extrapolation values, $\gamma = 1.3$.

Finally, especially in the permanently oscillating groups 1 and 6, a number of participants used slightly more sophisticated AR(2) rules of the form

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

This is an example of an *anchoring and adjustment* rule (Tversky and Kahneman, 1974), since it extrapolates the last price change from the reference point or *anchor* $(p^f + p_{t-1})/2$ describing the “long-run” price level. One could argue that the anchor for this rule, defined as an equally

⁵Ideally one would like to estimate a model allowing for switching between different forecasting rules, consistent with the qualitative observations of switching behavior. However, the data sets are too short to estimate nonlinear switching models on individual data.

⁶The adaptive rule (2.6) was estimated for 5 out of 12 participants in groups 2 and 5, and three among them had w insignificantly different from 1 (i.e., they used naive rule). Four other participants used an AR(1) rule, $p_{t+1}^e = a + b p_{t-1}$, conditioning only on the past price with a coefficient $b < 1$.

weighted average between the last observed price and the fundamental price, was unknown in the experiment, since subjects were not provided explicitly with the fundamental price.⁷ Therefore, in our evolutionary selection model in Section 4 one of the rules will be we (2.8) with the fundamental price p^f replaced by a proxy given by the (observable) sample average of past prices $p_{t-1}^{av} = \sum_{j=0}^{t-1} p_j$, to obtain

$$p_{t+1}^e = 0.5 (p_{t-1}^{av} + p_{t-1}) + (p_{t-1} - p_{t-2}). \quad (2.9)$$

To distinguish these rules, we will refer to the forecasting rule (2.8), with an anchor partly determined by the fixed fundamental price p^f , simply as an anchoring and adjustment (AA) heuristic, and to the more flexible forecasting rule (2.9), with an anchor learned through a sample average of past prices, as the *learning anchoring and adjustment (LAA)* heuristic.

Two important observations follow by the above discussion. First, the subjects in the experiment tended to base their predictions on past observations, using relatively simple and intuitive rules of thumb, such as adaptive expectations or trend extrapolation. Second, it seems that participants tried to learn from past errors, and their learning behavior was in the form of *switching* between different heuristics. These two observations, 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.

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 studying price fluctuations under *homogeneous expectations* in the forecasting experiment.

The model with homogeneous expectations is given by

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

The first equation describes forecasting behavior with a simple first order heuristic f as in (2.5), which can be either adaptive expectations (in which case f does not depend on p_{t-2}) or trend following expectations (in which case f 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 the equilibrium pricing equation used in the experiment, cf. (2.2). We present an analysis of the so-called *deterministic skeleton model*, setting term ε_t in (3.1) to zero, as well as stochastic simulations with the same realizations of the shocks, ε_t , as in

⁷It is remarkable that exactly this anchor and adjustment rule (2.8) was estimated for participant 3 in group 1, who *did* submit the fundamental price forecast in the first four periods of the experiment, see the lower left panel of Fig. 3. For 4 out of 12 participants of groups 1 and 6 the estimated AR(2) rule was very close to (2.8), as it can be seen in the middle panels of Fig. 6.

the experiment, in order to investigate how the noise affects price fluctuations. In terms of *deviations* from the fundamental price the model can be rewritten as

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

Fig. 4 shows example of simulated dynamics for different adaptive, trend-following and anchoring and adjustment rules which have been estimated from individual forecasting experimental data.

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 (2.6). 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.*

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 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 stochastic price series (shown by triangles and squares) still exhibit almost monotonic convergence. Adaptive expectations thus seems a good explanation of the aggregate 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 rule as:

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

This extrapolative rule contains both the trend following 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 (2.7) is obtained, while $\alpha = p^f/2$, $\beta_1 = 1.5$ and $\beta_2 = -1$ correspond to the anchoring and adjustment heuristic (2.8). The rules for which the forecasts are not consistent with realizations will be disregarded by the participants, sooner or later. Therefore, both in the formal analysis and in simulations we confine our attention to the rules satisfying the following simple steady-state consistency requirement:

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

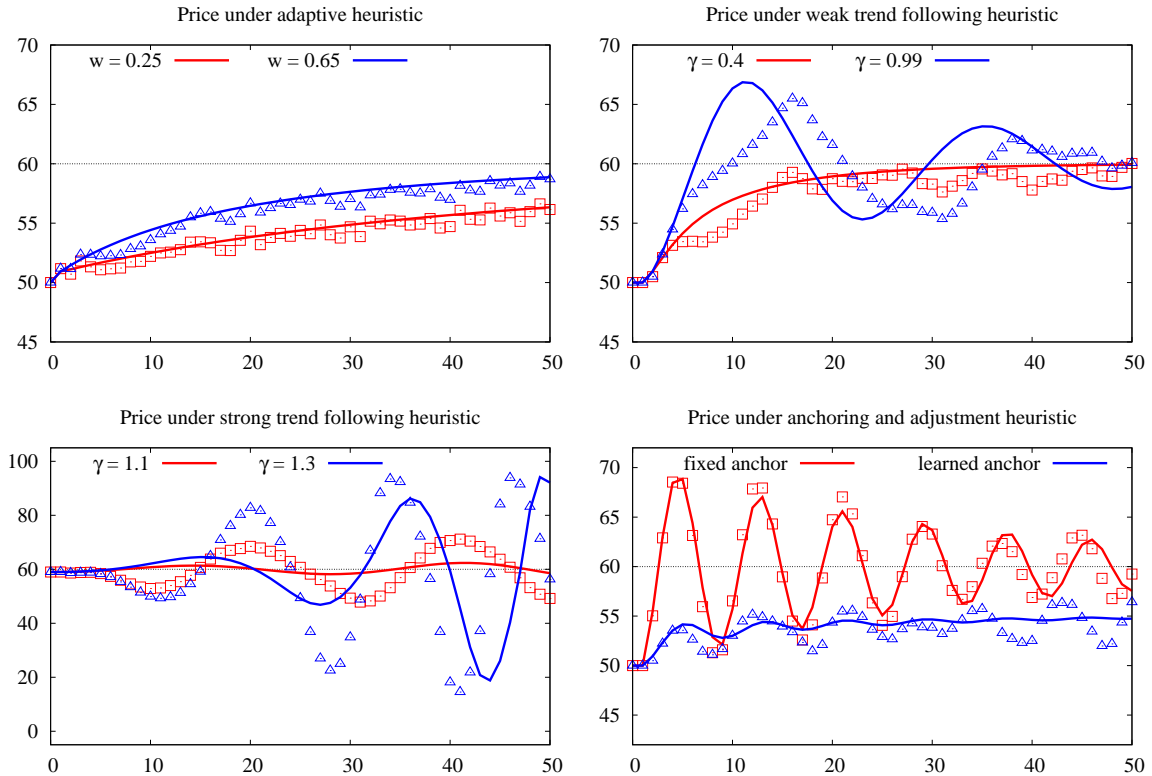


Figure 4: **Model (3.1) with homogeneous expectations.** The trajectories of the deterministic skeleton (the curves) and stochastic simulations with the same realization of shocks as in the experiment (triangles and squares) for different forecasting heuristics. **Upper left panel:** Dynamics with the 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). **Lower left panel:** Dynamics with the strong trend extrapolation oscillates (slowly) around the fundamental steady-state and diverges to a quasi-periodic cycle. **Lower 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 very slowly and almost monotonically.

In other words, consistent rules give unbiased predictions at the steady-state. Obviously, the extrapolative rule is consistent in p^* if and only if $\alpha = (1 - \beta_1 - \beta_2)p^*$. Notice that, the trend-following heuristic (2.7) is consistent at any steady-state, while the anchoring and adjustment heuristic (2.8) is consistent only at the steady state with $p^* = p^f$.⁸

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 the fraction of robot traders $n^* = 0$. The “fundamental” steady-state is locally*

⁸In related learning-to-forecasting experiments Heemeijer, Hommes, Sonnemans, and Tuinstra (2009) find that the estimated linear forecasting rules for many subjects are consistent in the steady state $p^* = p^f$, both in market environments with positive and negative expectations feedback.




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 fundamental steady-state if the pair of coefficients (β_1, β_2) belongs to the union of light and dark grey regions. The edges of the triangle at the border of the stability region correspond to pitchfork, period-doubling and Neimark-Sacker bifurcations respectively. The price dynamics is oscillating if the pair (β_1, β_2) lies below the parabolic curve. The three dots correspond to two trend-following heuristics (labeled $\gamma = 0.4$ and $\gamma = 1.3$) and an anchoring and adjustment heuristic (labeled AA), which will be used in the learning model in Section 4.

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 a pitch-fork, period-doubling or Neimark-Sacker bifurcation, if the first, second or third inequality in (3.4) turns into an equality, respectively. Moreover, the dynamics is oscillating (i.e., the eigenvalues of the linearized system are complex) when $\beta_1^2 + 4\beta_2(1 + r) < 0$.

Proof. See Appendix B. □

In general, the 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 only in the fundamental steady state p^f . The stability conditions (3.4) are illustrated in Fig. 5 in the parameter space (β_1, β_2) . The dark regions contain all rules for which the extrapolative heuristic (3.3) generates stable dynamics. For the pairs lying below the parabolic curve, the dynamics are oscillating. A loss of (local) stability occurs when the pair (β_1, β_2) leaves the stability area and crosses the boundary formed by the triangle. The dynamics immediately after the bifurcation are determined by the type of bifurcation through which stability is lost. For instance, after the pitchfork bifurcation the price diverges from its fundamental level and converges to one of two new stable steady-states. The Neimark-Sacker bifurcation implies existence of (quasi-)periodic price fluctuations right after the bifurcation.

The three dots shown in Fig. 5 correspond to three extrapolative forecasting rules estimated from individual experimental data. Two trend-following heuristic (2.7) with different values of the extrapolation coefficient γ are labeled as $\gamma = 0.4$ and $\gamma = 1.3$. The anchoring and adjustment heuristic (2.8) is labeled as AA.

Trend-following heuristic. These results imply that the price may either converge or diverge under the trend-following rule (2.7), depending upon the parameter γ . 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 the extrapolative coefficient is sufficiently small (e.g., $\gamma = 0.4$), convergence is monotonic; for larger γ -values (e.g., $\gamma = 0.99$) convergence becomes oscillatory. Notice however that these stable oscillations are different (e.g., of a lower frequency) compared to the dampened oscillations observed in groups 4 and 7 in the experiments. Indeed, the estimation of individual strategies in groups 4 and 7 did not reveal a trend-following rule which would generate such converging oscillations. The case of the strong trend extrapolation is illustrated in the bottom left panel of Fig. 4. The price dynamics diverges from the fundamental steady state through oscillations of increasing amplitude. 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 (2.8) we conclude that the implied price dynamics is converging. Since the parameters of the anchoring and adjustment rule are very close to the Neimark-Sacker bifurcation, the convergence to the fundamental steady-state is oscillatory and slow (cf. the bottom right panel of Fig. 4). Interestingly, for the stochastic simulation the convergence is even slower and the amplitude of the price fluctuations remains more or less constant in the last 20 periods, with an amplitude ranging from 55 to 65 comparable to that of the permanently oscillatory group 6 in the experiments. The small shocks ε_t added in the experimental design, to mimic (small) shocks in a real market, thus seem to be important to keep the price oscillations alive.

The bottom right panel of Fig. 4 also shows the price dynamics of the learning anchoring and adjustment (LAA) rule (2.9). Notice that the dynamics with a LAA heuristic is described by a non-autonomous system, whose formal analysis is complicated. Simulations under homogeneous expectations given by the LAA heuristic (2.9) converge to the same fundamental steady-state as with the anchoring and adjustment heuristic (2.8), but much slower and with less pronounced oscillations. In the presence of noise, the price oscillations under the LAA heuristic are qualitatively similar to the price fluctuations in the permanently oscillatory group 1 of the experiment (with prices fluctuating below the fundamental most of the time, between 52 and 62).

Homogeneity versus heterogeneity

In this Section we analyzed the price dynamics underlying the experiment, under the assumption that expectations are homogeneous and all individuals use the same forecasting heuristic. Qualitatively, all three observed patterns in the experiments, monotonic convergence, constant oscillations and dampened oscillations can be reproduced. However, the dampened oscillations have a different amplitude and frequency compared to the experimental groups 4 and 7. Moreover, a model with homogeneous expectations leaves open the question *why* different patterns in aggregate behavior emerged in different experimental groups.

The dots in Fig. 6 represent the coefficients of the estimated individual extrapolative rules

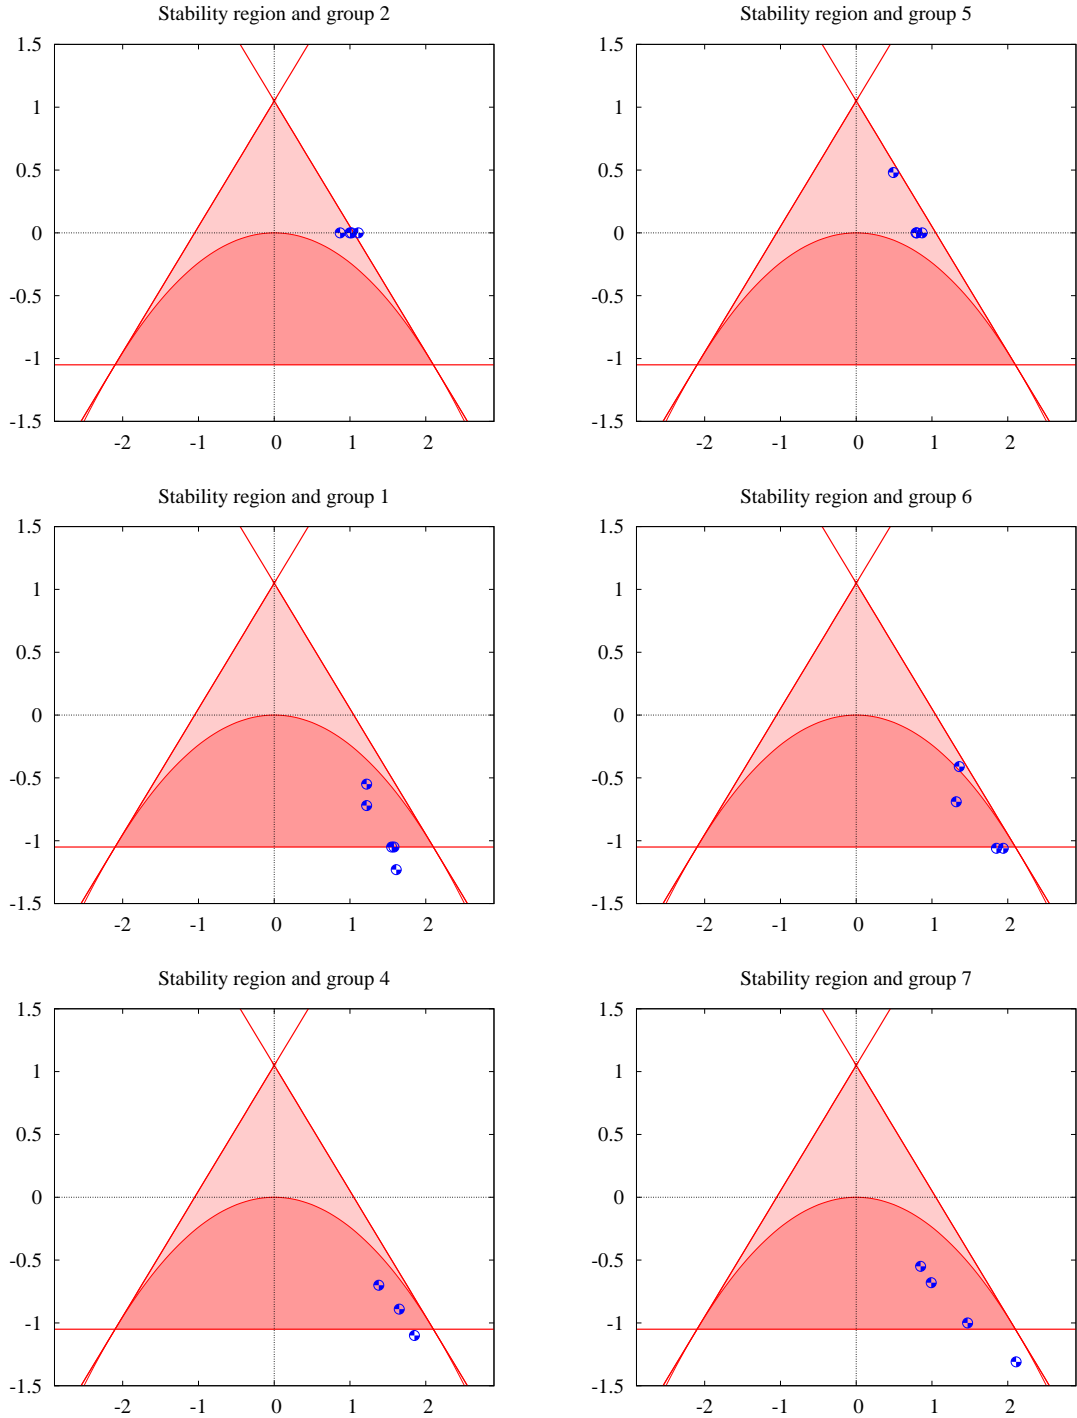


Figure 6: **Stability of a model with homogeneous extrapolative rules estimated in the experiment.** **Upper panel:** In both groups with converging price, all rules generate stable monotonic dynamics. **Middle panel:** In both groups with oscillating price, there were two rules on the stability border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

(3.3).⁹ While the dispersion of individual forecasting rules is clear, the figure suggests some

⁹In every group there were rules which cannot be represented by the extrapolative prediction (3.3), e.g., an

regularities. In the converging groups 2 and 5, the majority of rules belong to the region of monotonic convergence. In contrast, in the oscillating groups almost all individual rules lie in the oscillatory region (i.e., the linear forecasting rule has complex eigenvalues). Furthermore, in groups 1 and 6 with constant price oscillations, at least two individual rules in every group are very close to the locus of the Neimark-Sacker bifurcation (i.e., are close to complex unit roots), while in groups 4 and 7 with dampened oscillations there is at least one (strongly) unstable individual forecasting rule. These results suggest that *heterogeneous expectations* may be a key element in order to explain the learning to forecast experiments.

4 Evolutionary Selection of Forecasting Heuristics

In this Section we demonstrate that evolutionary selection between simple forecasting heuristics can explain the different qualitative scenarios observed in the experiment. Before describing a model, let us recall the most important “stylized facts” which we found in the individual and aggregate experimental data:

- participants tend to base their predictions on past observations following simple forecasting heuristics;
- individual learning has a form of switching from one heuristic to another;
- in every group some form of coordination of individual forecasts occurs; the rule on which individuals coordinate may be different in different groups;
- coordination of individual forecasting rules is not perfect and some heterogeneity of the applied rules remains at every time period.

The main idea of the model is as follows. 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 next period’s price, and the realized market price depends upon these individual forecasts. However, the *impacts* of different forecasting heuristics upon the realized prices are changing over time because the participants are learning based on evolutionary selection: the better a heuristic performed in the past, the higher its impact is in determining next period’s price. As a result, the realized market price and impact of the forecasting heuristics co-evolve in a dynamical process with mutual feedback. It turns out that this evolutionary model exhibits *path dependence* thus explaining coordination on different forecasting heuristics leading to different aggregate price behavior.

4.1 The Model

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

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

adaptive heuristic or linear rules with three lags.

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 the 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 all simulations we use the same parameter values and the same realization of stochastic shocks ε_t as in the experiment. In particular, the fundamental price as predicted by robots is set to $p^f = \bar{y}/r = 0.05/3 = 60$.

In our evolutionary model, the average \bar{p}_{t+1}^e in (4.2) is a population weighted average of the different forecasting heuristics

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

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

Similar to the incentive structure in the experiment, the performance measure of a forecasting heuristic in a given period is based on its squared forecasting error. More precisely, the *performance measure* of heuristic h up to (and including) time $t-1$ is given by

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

The parameter $0 \leq \eta \leq 1$ represents the *memory*, measuring the relative weight agents give to past errors of heuristic h . In the special case $\eta = 0$, the impact of each heuristic is completely determined by the most recent forecasting error; for $0 < \eta \leq 1$ all past prediction errors, with exponentially declining weights, affect the impact of the heuristics.

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. In the special case $\delta = 0$, (4.6) reduces to the *the discrete choice model* with synchronous updating used in Brock and Hommes (1997) to describe endogenous selection of expectations. The more general case, $0 \leq \delta \leq 1$, gives some persistence or inertia in the impact of rule h , reflecting the fact (consistent with the experimental data) that not all the participants update their rule in every period or at the same time (see Hommes, Huang, and Wang (2005) and Diks and van der Weide (2005)). Hence, δ may be interpreted as the average per period 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 discrete choice model. The parameter $\beta \geq 0$

represents the intensity of choice measuring how sensitive individuals are to differences in strategy performance. The higher the intensity of choice β , the faster individuals will switch to more successful rules. In the extreme case $\beta = 0$, the 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., a fraction $1 - \delta$) switch to the most successful predictor.

Initialization. The model is initialized by providing 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, as well as an initial distribution $\{n_{h,\text{in}}\}, 1 \leq h \leq H$ of the impacts of different heuristic (summing to 1). Additionally, the initial share of robot traders and initial performances of all H heuristics are set to 0.

Given initial prices, the heuristic’s 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. Thereafter, the price $p_{\text{in}+2}$ is computed and the initialization stage is finished. After this initialization state 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.¹⁰ Since one of our goals is to explain the three different observed patterns in aggregate price behavior – monotonic convergence, permanent oscillations and dampened oscillations – we keep the number of heuristics as small as possible and consider a model with only four forecasting rules. These rules, referred to as ADA, WTR, STR and LAA and given in Table 1, were obtained as simple descriptions of typical individual forecasting behavior observed and estimated in the experiments, as discussed in Section 3.

Table 1: **Heuristics used in the evolutionary model.** In simulations in Figs. 7–8 the first four heuristics are used. The LAA heuristic is obtained from the simpler AA heuristic, by replacing the (unknown) fundamental price p^f by the sample average p_{t-1}^{av} .

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	$p_{4,t+1}^e = 0.5 (p_{t-1}^{av} + p_{t-1}) + (p_{t-1} - p_{t-2})$
AA	anchoring and adjustment rule with fixed anchor	$p_{4,t+1}^e = 0.5 (p^f + p_{t-1}) + (p_{t-1} - p_{t-2})$

¹⁰The software for simulations *evexex* is freely available at <http://www.cafed.eu/evexex> together with a brief documentation and configuration settings used for the simulations reported below.

The evolutionary model with 4 forecasting heuristics is given by:

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

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

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

Comparing with (3.2) we observe that our model can be seen as a generalization of the homogeneous expectation model to the case of heterogeneous expectations with the forecast p_{t+1}^e replaced by a *weighted average*, with endogenous, time varying weights, of the individual forecasts of the four heuristics.

Due to the presence of demand/supply shocks ε_t in the pricing equation, the model (4.7) is, in general, stochastic. In the next section we simulate this system for 50 periods with *the same* realisation of the noise process $\{\varepsilon_t\}_{t=0}^{50}$ as in the experiment. The resulting dynamics will be referred to as “simulated path” and will be compared with the experimental data. Section 4.3 is devoted to the (local stability) analysis of the deterministic skeleton of system (4.7) when the noise term ε_t is absent. Section 5 is devoted to investigation of the “one step ahead” forecasting errors of our stochastic nonlinear evolutionary selection model.

4.2 Simulated Path

As soon as the four heuristics matched by experimental forecasting data 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 in periods $t = 2$ and $t = 3$.

Figs. 7–8 illustrate six typical simulated time series of aggregate prices, individual forecasts and fractions of each of the four heuristics together with the aggregate price series in the experiments. In all simulations we have fixed the parameter values as $\beta = 0.4$, $\eta = 0.7$ and $\delta = 0.9$;¹¹ the simulations only differ in initial conditions, as reported in Table 2. For all simulations, we use the same realizations $\{\varepsilon_t\}_{t=0}^{50}$ of the (small) noise as used in the experiment. These simulations should thus be viewed as 50-period ahead forecasts of the patterns of aggregate price behavior and underlying individual forecasting behavior.

¹¹These parameters were obtained after some simple trial and error simulations for groups 1, 2, and 7 in Anufriev and Hommes (2009).

Table 2: **Initial conditions for simulation.** Initial prices and impacts for 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).

	Experiment		Simulations								
	Observed Prices		Initial Prices		Heuristics' Initial Impacts				Parameters		
	$t = 0$	$t = 1$	p_0	p_1	ADA	WTR	STR	LAA	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

Figs. 7 shows that the switching heuristics model indeed exhibits *path-dependence* and is capable of reproducing all three qualitatively different price patterns observed in the experiments, that is, monotonically converging prices, permanent oscillations or dampened oscillations. This path-dependent feature of the model 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.

We stress that, at this stage, no fitting exercise has been performed. All plots have been easily obtained through some 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). The initial distribution of agents over the four heuristics, i.e., initial impacts $\{n_{1,0}, n_{2,0}, n_{3,0}, n_{4,0}\}$, turn out to be important in order to replicate the aggregate price patterns in the 50-periods ahead forecasting simulations.¹² For replication of the monotonic convergence in groups 2 and 5, the initial impacts of heuristics are distributed almost uniformly, with a slight dominance of the WTR heuristics to produce a small initial trend in prices. For the oscillating groups 1 and 6, where the initial trend was stronger, both trend heuristics WTR and STR were initialized with somewhat higher weights. Finally, in the dampened oscillating groups 4 and 7, with the strongest trend in price in initial periods, the STR rule has a large initial impact.

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 close during the simulations, causing slow (almost) monotonic convergence of the price to its fundamental equilibrium $p^f = 60$. For group 2 the increase of the price together with a series of subsequent positive shocks ε_t leads to a temporary slight domination of 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 its fundamental level.

In the 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

¹²Ideally, one would like to estimate the distribution of participants over the four heuristics in the initial phase of the experiment, but this seems impossible due to insufficient data and the fact that especially in the initial phase learning takes place.

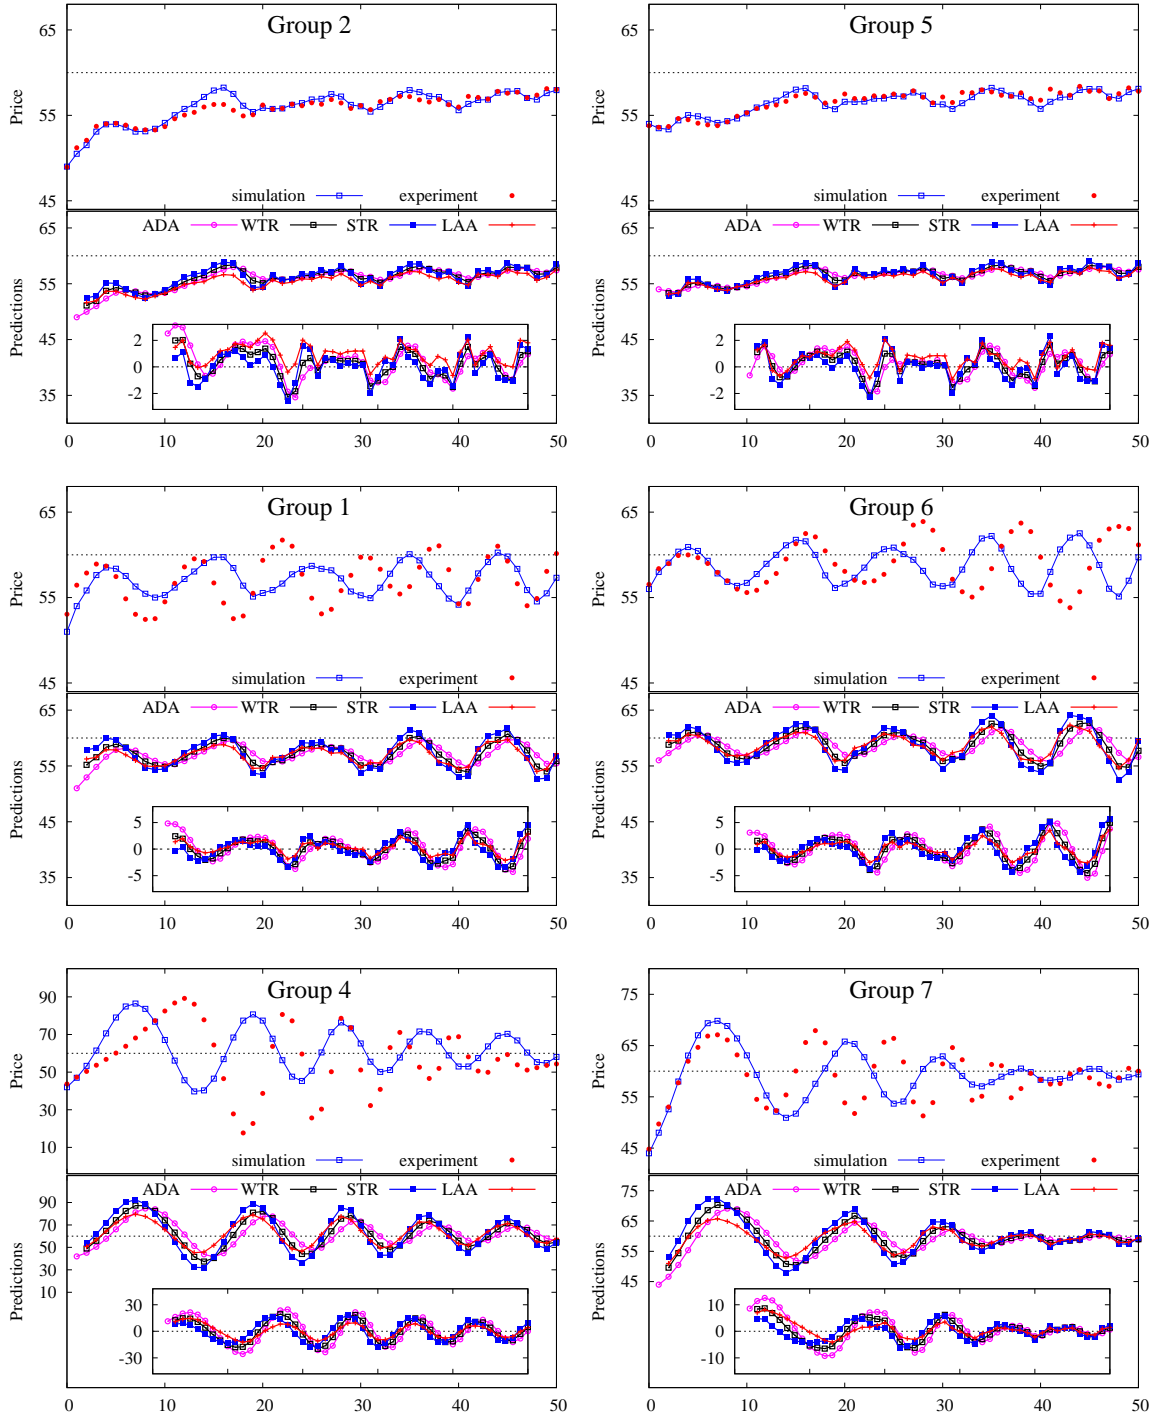


Figure 7: **Laboratory experiments and heuristics switching model simulations.** **Upper parts of panels** show prices for laboratory experiments in different groups (red) with corresponding simulations of the evolutionary model (blue). All three different aggregate market outcomes are reproduced: monotonic 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).

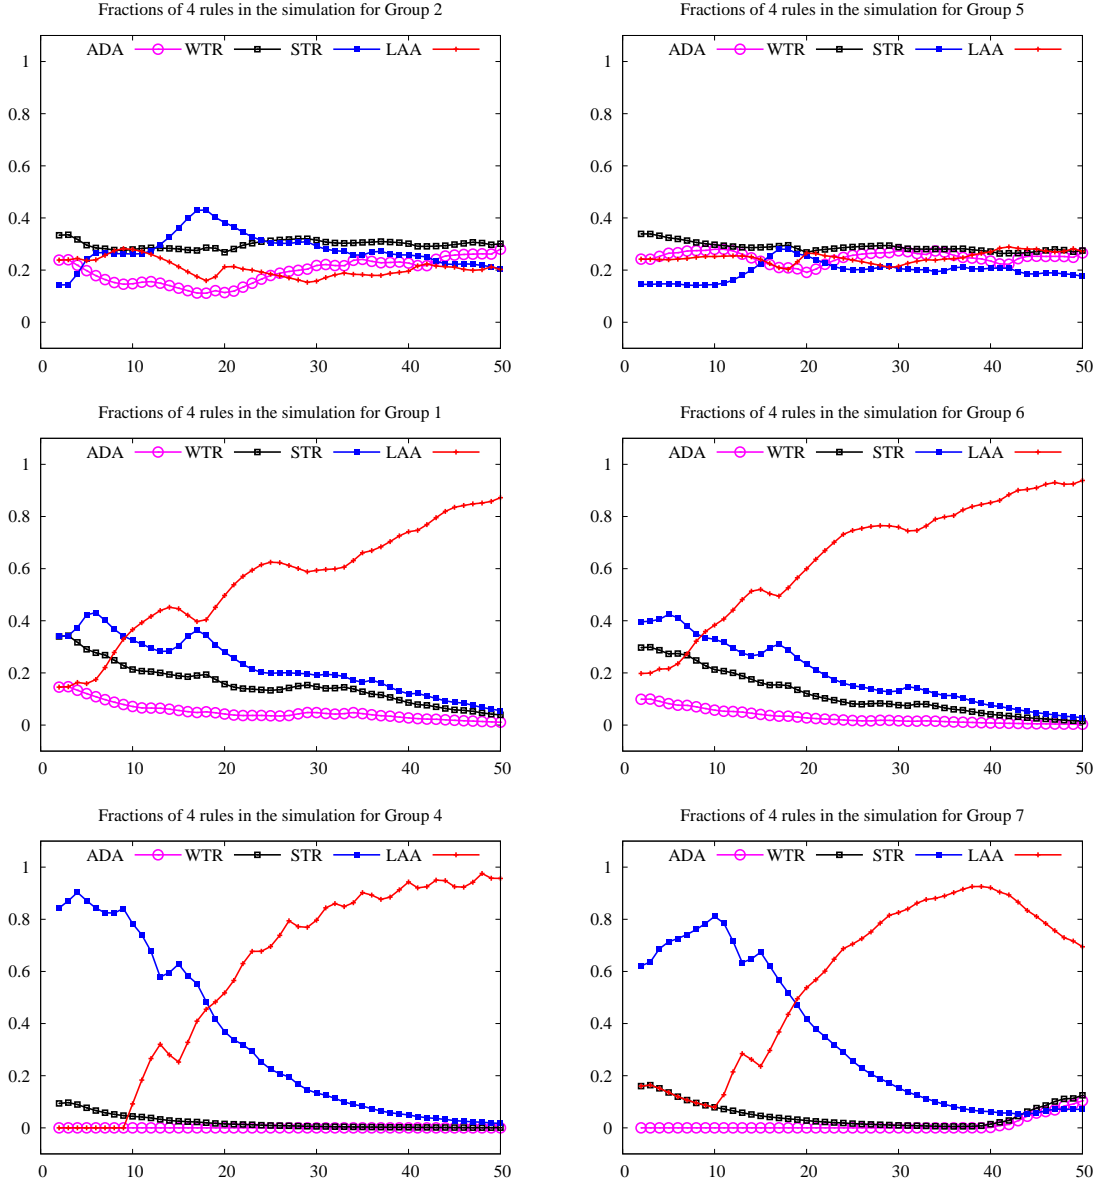


Figure 8: **Simulated fractions of the forecasting rules in the heuristics switching model.** Fractions of four forecasting heuristics: adaptive expectations (ADA, purple), weak trend followers (WTR, black), strong trend followers (STR, blue) and anchoring adjustment heuristic (LAA, red). Coordination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

heuristics, and aggregate prices rise. However already after 5 periods the impact of the learning anchor and adjustment (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, it dominates the market within 10 periods, rising to more than 70% after 40 periods. Our switching heuristics model thus explains coordination of individual forecasts as well as persistent price oscillations around the long run equilibrium level.

Finally, in the last two simulations (see the lower panel) a large initial impact of (strong)

trend followers leads to a strong rise of market prices in the first 7 periods, followed by large price oscillations. Already relatively quickly (e.g. after period 10 in group 7), however, the impact of the 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 weaker coordination between individual predictions during the last 10 periods, which is also consistent with experimental data.

4.3 Analysis of the Deterministic Skeleton of the Model

The simulations pose a number of interesting theoretical questions concerning the dynamics produced by the model: How general is the path-dependent property? Are fluctuations only short-run phenomena or are fluctuations persistent in the long run as well? Are the fluctuations endogenously generated or are exogenous shocks crucial for sustained fluctuations? 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 keep the stability analysis tractable and avoid a non-autonomous system, we replace the LAA rule by an analogous AA heuristic with fixed anchor p^f instead of time varying anchor p_{t-1}^{av} .¹³

4.3.1 Local stability of four heuristic evolutionary switching model

For the sake of generality let us introduce the following four heuristics,¹⁴ one of which is adaptive and three other are extrapolative rules (consistent in p^f according to Def. 3.1):

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

The resulting dynamics can be described by a multi-dimensional system. The following result shows that if the price converges, it must converge to the fundamental price level.¹⁵

Proposition 4.1. *Assume that the price dynamics of the deterministic skeleton asset pricing model with evolutionary switching (4.7) converges to a constant price level p^* . Then the price converges to its fundamental level, i.e., $p^* = p^f$. Furthermore, the share of robots is fixed and equal to 0, and all heuristics with non-zero weights predict the fundamental price.*

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, the 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 a differentiable system. Standard analysis of its Jacobian leads to the following

¹³Recall from 2.3 that a number of participants used an AR(2) forecasting rule close to the AA heuristic, cf. footnote 7.

¹⁴To get the 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 $h = 4$.

¹⁵This result holds also for the non-autonomous system with the LAA heuristic.

Proposition 4.2. *The fundamental steady-state of the asset pricing model with evolutionary switching (4.7) (i.e., the deterministic skeleton and with fixed anchor in the anchoring and adjustment heuristic) is locally stable if (i) $\eta < 1$ and $\delta < 1$, and (ii) all roots of the 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 a straight-forward computation shows that the 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). \square

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 other two parameters of the learning process, η and δ , is also limited. As soon as $\delta \neq 1$, i.e., the 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 only depend on the coefficients of the forecasting heuristics. The parameters η and δ , being eigenvalues of the Jacobian matrix, affect, however, the speed of convergence.

A simple analytical expression of the roots of polynomial (4.10) is not available. Using numerical methods to compute the roots for the four heuristics ADA, WTR, ATR and AA defined in Table 1 yields

$$\mu_1 \simeq 0.47, \quad \mu_2 \simeq 0.63 - 0.27i, \quad \mu_3 \simeq 0.63 + 0.27i.$$

The modulus of the complex eigenvalues is approximately equal to 0.69. Thus, the fundamental steady-state is locally stable. Consistent with this result, all simulations which we performed with the *deterministic version* of the 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 demand/supply shocks in the experiment, was crucial in generating persistently oscillating time series in groups 1 and 6. Two caveats are necessary, however. First, the fundamental steady state can be locally unstable under a different pool of heuristics. For example, a switching model with only two of the four heuristics, the unstable STR under homogeneous expectations and the stable AA rule under homogeneous expectations, will generate non-converging dynamics. Second, even if the fundamental steady-state is stable, then other attractors, such as a stable cycle, may co-exist.

4.3.2 Local stability of four heuristics model with constant impacts

Which combination of rules might generate non-converging dynamics? It is intuitively clear that the price dynamics will not converge if the impact given to “unstable” heuristics, i.e., to those heuristics which generate unstable dynamics under homogeneous expectations, is high enough. Therefore, consider an auxiliary version of the model, where the impacts of different heuristics are not changing over time. Formally, such a *constant impacts model* corresponds

to a special case within our evolutionary model with $\delta = 1$. Assuming that the forecasts are as in (4.9), the 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 the impacts of heuristics n_h are arbitrary constants summing up to 1. The dynamics of the constant impacts model is locally stable when all 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 the polynomial in (4.10), we observe that the local stability of the evolutionary switching model is governed by the local stability of the constant impacts model with *equal* impacts, i.e., with $n_h = 0.25$, $h = 1, \dots, 4$. This result is not surprising. Indeed, the evolutionary model tends to choose the best performing heuristic at any time 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.

However, the dynamics of the constant impacts model can be unstable if the distribution of heuristics is not uniform. In the left panel of Fig. 9 we show the 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. The dark region in this simplex contains all points where the constant impacts model with four heuristics defined in Table 1 is *unstable*. This instability region was obtained numerically by evaluating the roots of the polynomial $P_1(\mu)$ in (4.11) for different values of the impacts n_h . The instability region has a conic shape connected to the upper left vertex of the simplex, where the STR rule has the highest impact. Notice that among the four heuristic this is the only rule which generates unstable dynamics under homogeneous expectations. Consequently, if the impact of STR is relatively high and the impacts of the remaining three heuristics are relatively low, the dynamics of fixed fractions 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 noted above, this implies that the fundamental steady-state of the evolutionary switching model with our four heuristics is locally stable.

The right panel of Fig. 9 illustrates how the evolving distribution of impacts generated by the switching model affects the (in)stability in the model with fixed impacts, by showing the evolution of the largest eigenvalue of the polynomial $P_1(\mu)$ for the 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 is stable. In the oscillatory groups 1 and 6, in the early stage (say the first 20 periods) the system is stable, while towards the end of the simulations the impacts are evolving to a state where the constant impact model is close to bifurcation. For the oscillatory groups, the system thus evolves from a stable process to a near unit root process. Finally, in groups 4 and 7 the initial impacts correspond to instability, because of a relatively large impact of the STR, but over time the system becomes stable.

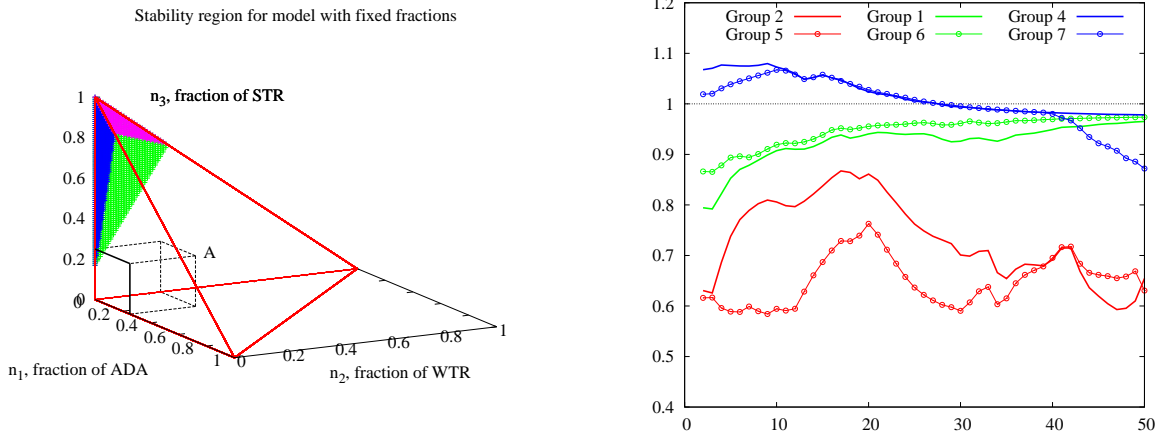


Figure 9: **Local stability of fundamental steady-state in the constant impacts model.** **Left panel:** The fundamental steady-state is unstable in a model with four heuristics, when fixed impacts (n_1, n_2, n_3) belong to the dark conic region of the unit simplex. **Right panel:** Evolution of the modulus of the largest eigenvalue of polynomial (4.11) corresponding to the simulations of the evolutionary model in Figures 7 and 8.

5 Empirical Validation

In this section we address the issue of how well the nonlinear stochastic switching model with four forecasting heuristics fits with the experimental data. This section is divided in three parts. We, first, illustrate the one-period ahead forecasts of the model visually, then turn to the rigorous evaluation of the in-sample performance of the model, and, finally, look at out-of-sample forecasts made by the model.

5.1 One-period ahead simulations

Fig. 10 compares the experimental data with the one-step ahead predictions made by our model, using the same benchmark parameters $\beta = 0.4$, $\eta = 0.7$ and $\delta = 0.9$ as before. In these simulations the initial prices coincide with the initial prices in the first two periods in the corresponding experimental group, while the initial impacts of all heuristics are equal to 0.25. Fig. 10 suggests that the switching model with four heuristics fits the experimental data quite nicely.

It is useful to briefly discuss the differences of the stochastic simulations in Fig. 10 with the earlier “simulated paths” in Fig. 7. At each time step, the simulated path only uses *simulated price data* as inputs to compute the heuristics’ forecasts and to update their impacts. Therefore, the simulated paths are essentially 50 period ahead forecasts generated by the nonlinear switching model. These simulated paths already showed that the nonlinear switching model (augmented by the same small noise as in the experiment) is capable of generating all three different patterns observed in the experimental data. In contrast, the one-step ahead predictions of the nonlinear switching model in Fig. 10 use past *experimental price data*, i.e., exactly the same information that was available to participants in the experiments, as inputs for the forecasting rules and the updating of fractions. An immediate observations by comparing these simulations is that the one-period ahead forecasts can follow more easily the sustained oscillations as well as the dampened oscillatory patterns. While the simulated

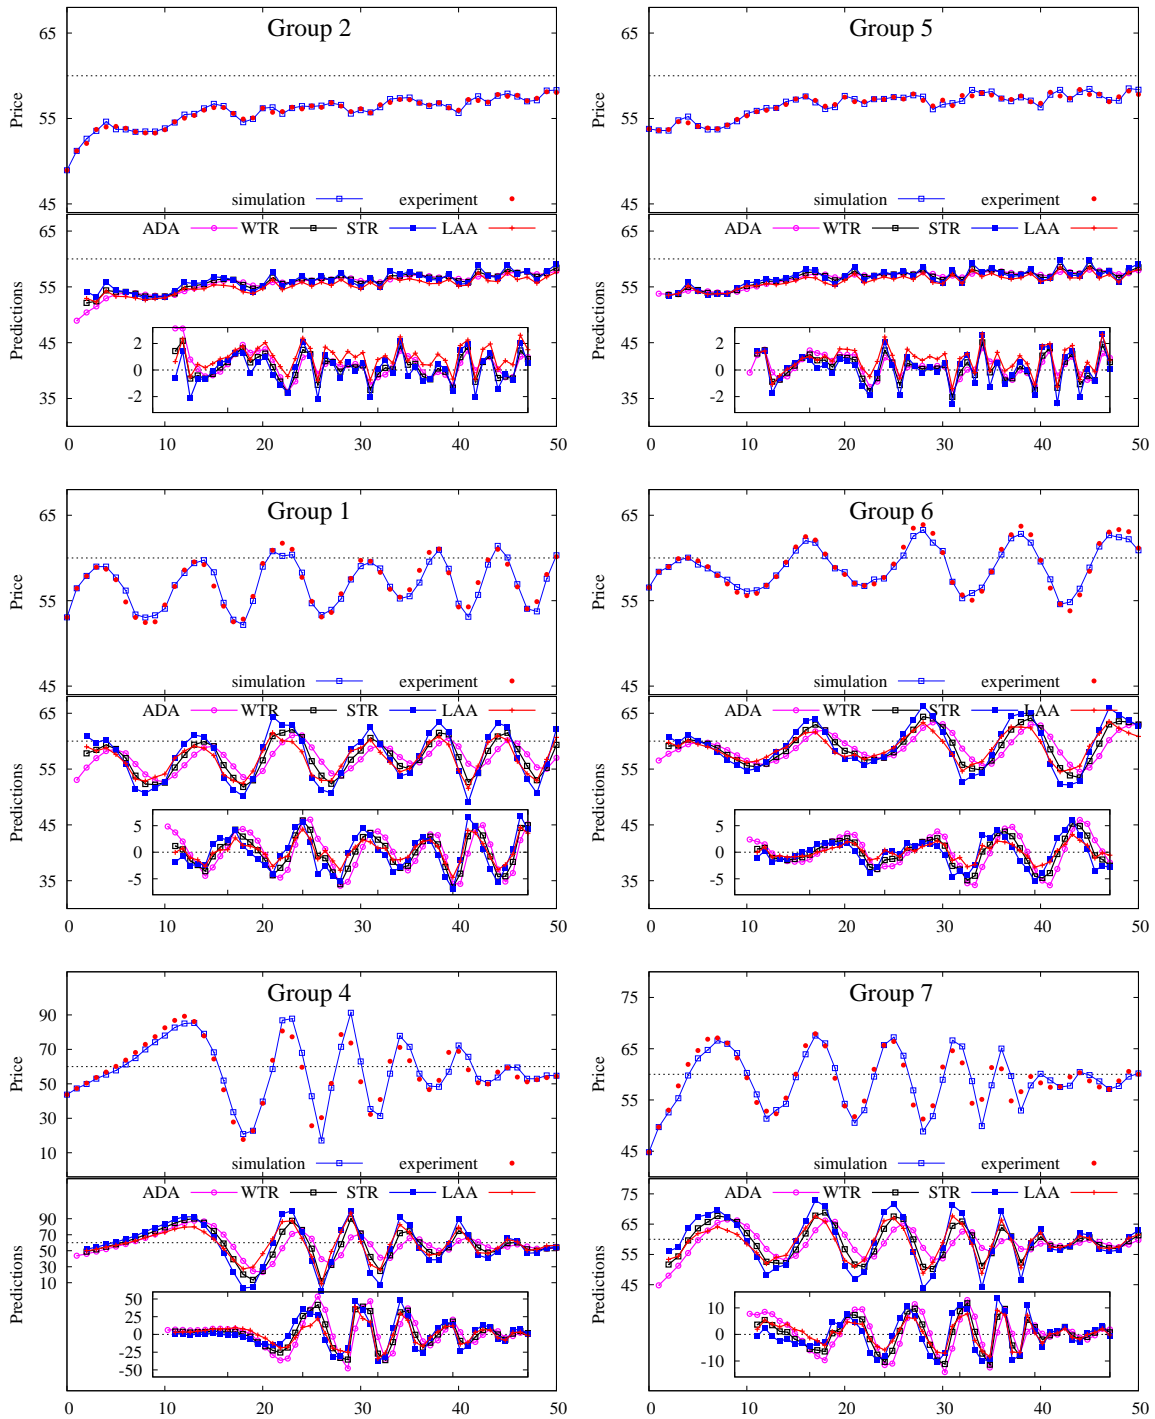


Figure 10: **Laboratory experiments and one-step ahead predictions of the evolutionary model.** Upper parts of panels show prices for laboratory experiments in different groups (red) with corresponding one-step ahead predictions of the evolutionary model (blue). Lower parts of panels show predictions 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).

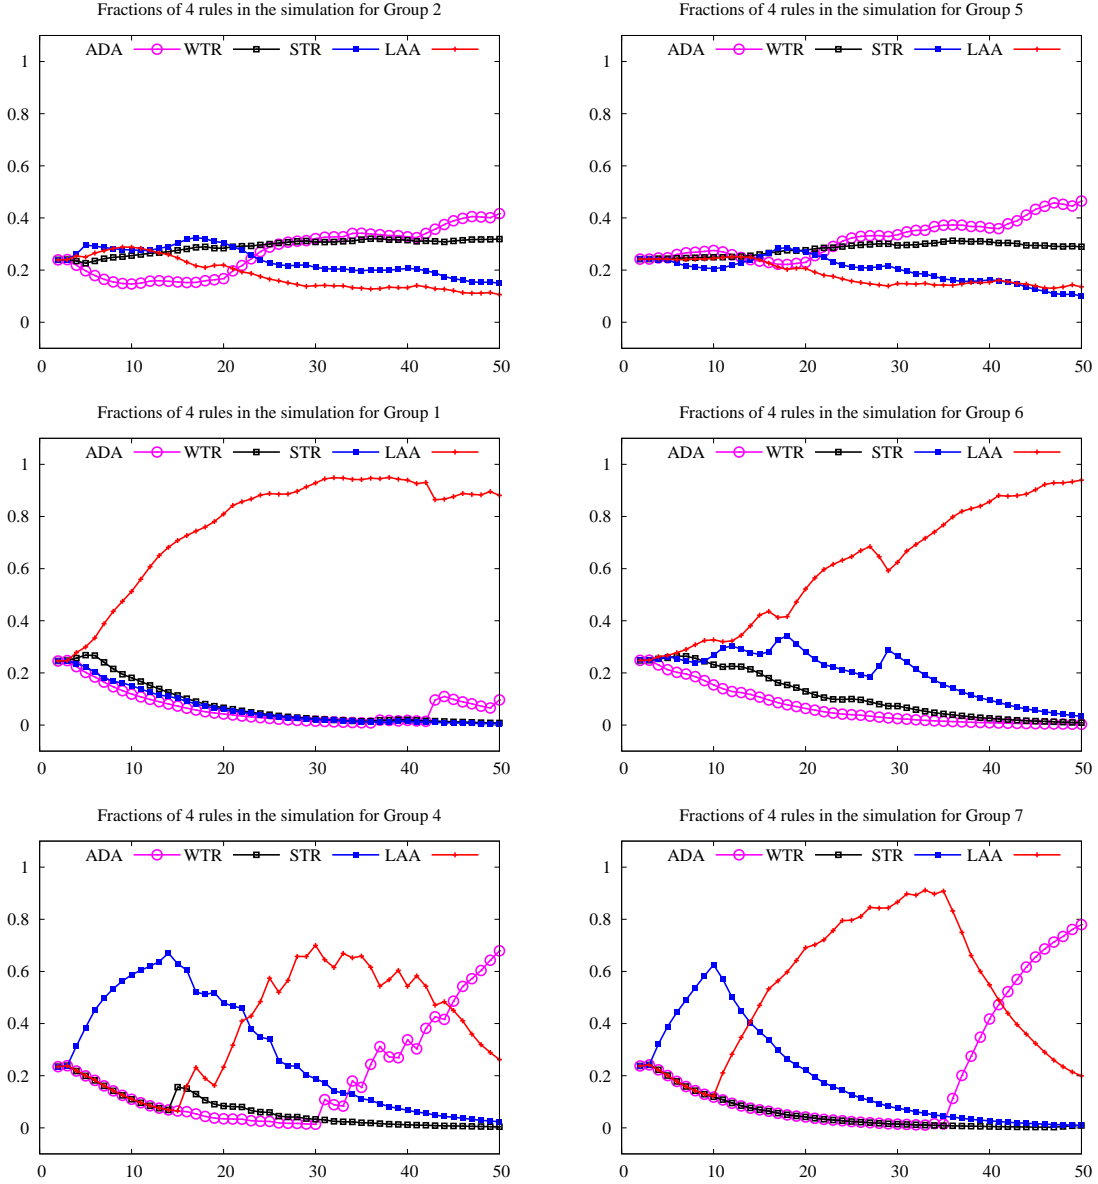


Figure 11: **Evolution of heuristic impacts during the one-step ahead predictions of the model.** Fractions of four forecasting heuristics: adaptive expectations (ADA, purple), weak trend followers (WTR, black), strong trend followers (STR, blue) and anchoring adjustment heuristic (LAA, red).

paths could only reproduce the (dampened) oscillatory patterns for 50 periods if the initial impact of strong trend followers is sufficiently large, the one-step ahead forecasts easily follow oscillatory patterns starting from a uniform initial distribution of forecasting rules, i.e., $n_{1,0} = n_{2,0} = n_{3,0} = n_{4,0} = 0.25$.

Fig. 11 shows how, for the one step ahead forecasts, in different groups different heuristics are taking the lead after starting from a uniform distribution. The reader should compare these to the corresponding Fig. 8 for the simulated paths. In the monotonically converging groups, the impact of the different rules stays more or less equal, although the impact of adaptive expectations gradually increases and slightly dominates the other rules in the last

20-25 periods. The oscillatory groups yield similar results as before, with the LAA rule dominating the market early and its impact increasing to about 90% towards the end of the experiment. Notice, that the domination of the LAA rule happens much faster in simulations for group 1, than for group 6. (For instance, 80% impact is reached by the LAA rule after 20 periods for group 1 and after 40 periods for group 6.) This difference reflects the fact that the *frequency* of oscillations in the two experimental groups were not the same. During the experiment we observe about 6 “cycles” in group 1, but only about 4 and a half “cycles” in group 6. The deterministic path could not match the precise frequencies and therefore could not discriminate between the two groups. In contrast, the stochastic one-step ahead simulations match the oscillations closely and produce clear difference in the evolution of impacts. Our model thus explains oscillatory behavior by coordination on the LAA rule by most subjects, and gives higher relative weight to the LAA when the oscillations are more frequent.

Finally, for the groups with the dampened oscillations, one step ahead forecast produces somewhat different and richer evolutionary selection dynamics than the simulated paths; cf. the bottom panels of 8 and 11. The groups with dampened oscillations go through three different phases where the STR, the LAA and the ADA heuristics subsequently dominate. The STR dominates during the initial phase of a strong trend in prices, but starts declining after it misses the first turning point of the trend. The LAA does a better job in predicting the trend reversal and its impact starts increasing. The LAA takes the lead in the second phase of the experiment, with oscillating prices. But the oscillations slowly dampen and therefore, between periods 30-35, the impact of adaptive expectations, which has been the worst performing rule until that point, starts increasing and adaptive expectations dominates the groups in the last 7-9 periods.

5.2 Forecasting performance

Table 3 compares the mean squared error (MSE) of the one-step ahead prediction for 10 different models: the RE fundamental prediction, six homogeneous expectations models (naive expectations, the fixed anchor and adjustment (AA) rule, and each of the four heuristics of the switching model), and three heterogeneous expectations models with 4 heuristics, namely,

Table 3: MSE over 47 periods of the one-step ahead forecast for different groups and 10 different model specifications.

Specification	Group 2	Group 5	Group 1	Group 6	Group 4	Group 7
RE Fundamental Prediction	16.6231	10.8238	15.7581	9.3245	300.9936	21.9123
naive	<i>0.0388</i>	0.0514	3.5415	2.4494	141.0558	13.2453
AA	5.1259	3.4323	2.9309	0.888	65.2296	5.0594
ADA	0.0712	0.0378	5.6734	4.6095	210.3313	19.5158
WTR	0.0862	0.1419	2.0905	1.1339	92.2163	9.2932
STR	0.5001	0.6605	2.9071	0.8131	124.3494	14.7224
LAA	0.4588	0.4756	0.456	0.6591	66.2637	5.8635
4 heuristics ($\delta = 1$)	0.0814	0.1698	1.2417	0.6618	70.8516	7.0956
4 heuristics (Figs. 10-11)	0.0646	0.1108	0.4672	0.2917	47.2492	4.3154
4 heuristics (best fit)	0.0493	<i>0.0353</i>	<i>0.4423</i>	<i>0.1655</i>	<i>34.4932</i>	<i>2.9358</i>
$\beta \in [0, 10]$	10	10	0.1	10	3	0.2
$\eta \in [0, 1]$	0.4	0.9	1	0.1	0.8	0.5
$\delta \in [0, 1]$	0.9	0.6	0.5	0.7	0.6	0.4

the model with fixed fractions (corresponding to $\delta = 1$), the switching model with benchmark parameters $\beta = 0.4$, $\eta = 0.7$ and $\delta = 0.9$, and, finally, the “best” switching model fitted by means of a grid search in the parameter space (the last three lines in Table 3 show the corresponding optimal parameter values). The MSEs for the benchmark switching model are shown in bold and, for comparison, for each group the MSEs for the best among the four heuristics are also shown in bold. The best among 10 models for each group is shown in italic.¹⁶

An immediate observation from Table 3 is that, for all groups, the fundamental prediction rule is by far the worst. This is due to the fact that in the experiment realized prices deviate persistently from the fundamental benchmark. Another observation is that, all models explain the monotonically converging groups very well, with very low MSE.¹⁷ The homogeneous expectations models with naive, adaptive or WTR expectations fit the monotonic converging groups particularly well, some of them slightly better than the benchmark switching model. For the permanent as well as the dampened oscillatory groups, the flexible LAA rule is the best homogeneous expectations benchmark, but the benchmark switching model has an even smaller MSE especially in the permanently oscillating group 6 and the dampened oscillatory groups 4 and 7.

To summarize, the evolutionary learning model is able to make the best out of different heuristics. Indeed, none of the homogeneous expectations models fits *all* different observed patterns, but for each group in the experiment, the lowest MSE is achieved by the best fit switching model. The only exception is group 2 with monotonic convergence, where the errors are very low, and the difference between the best heuristic (naive expectations) and the switching model is minor. Finally, the benchmark switching model also yields low MSEs for all groups, only slightly above the best fit switching model. This suggests that the good fit of the switching model is fairly robust w.r.t. the model parameters.

5.3 Out-of-sample forecasting

Let us now turn to the out of sample validation of the model. In order to evaluate the out-of-sample forecasting performance of the model, we first perform a grid search to find the parameters of the model minimizing the MSE for periods $t = 4, \dots, 43$. Then, the squared forecasting errors of the “best” model are computed for the next 7 periods. The results are reported in the upper part of Table 4 for each group. In the middle part of the table, we also report the corresponding squared prediction errors for the switching model with benchmark parameters. Finally, we compare our structural learning model with a simple non-structural model with three parameters. To this purpose we estimate an AR(2) model to the data up to period $t = 43$ and show in the bottom part of Table 4 the in- and out-of-sample squared prediction errors.

For the converging groups 2 and 5 the squared prediction errors typically increase with horizon but remain very low and comparable with the MSEs computed in-sample. This is not surprising given that the qualitative property of the data (i.e., monotonic convergence) does not change in the last periods, and that the adaptive heuristic, which generates such

¹⁶We evaluate the MSE over 47 periods, for $t = 4, \dots, 50$. This minimizes the impact of the initial conditions (i.e., the initial impacts of the heuristics) for the switching model, since $t = 4$ is the first period when the prediction is computed with both the heuristics forecasts and the heuristics impacts being updated based on the experimental data. For comparison, in all other models we compute errors also from $t = 4$.

¹⁷The only exception is the AA rule, which performs relatively well in the oscillatory groups, but not so well in the monotonically converging groups.

Table 4: **Out-of-sample performance of the heuristic switching model and the AR(2) model.** The in-sample (for $t = 0, \dots, 43$) MSE, the model parameters and 1–, ..., 7– periods ahead out of sample squared forecasting errors are shown. The *top* part: the best fitted in-sample model. The *middle* part: the benchmark model. The *bottom* part: AR(2) model estimated in-sample.

	Group 2	Group 5	Group 1	Group 6	Group 4	Group 7
Best Fit Switching Model						
MSE ₄₀	0.0567	0.035	0.4419	0.1748	39.0888	3.3207
(β, η, δ)	(10,0.4,0.9)	(10,0.9,0.8)	(0.1,1,0.5)	(10,0.1,0.7)	(3.3,0.8,0.6)	(0.2,0.5,0.4)
1 p ahead	0.0073	0.0033	0.4047	0.0305	12.9584	0.227
2 p ahead	0.0016	0.0224	2.7536	0.3976	34.0954	0.5575
3 p ahead	0.0004	0.0968	3.2215	0.2927	0.0001	3.4523
4 p ahead	0.0041	0.2536	0.4972	0.0004	5.5048	2.3656
5 p ahead	0.0033	0.195	3.8239	1.5924	3.4314	0.0225
6 p ahead	0.0054	0.2558	13.4333	6.6856	2.5055	1.3214
7 p ahead	0.0219	0.5148	7.8308	7.8088	0.6532	0.057
Benchmark Switching Model						
MSE ₄₀	0.0681	0.1016	0.4878	0.277	54.3656	4.9562
(β, η, δ)	(0.4,0.7,0.9)	(0.4,0.7,0.9)	(0.4,0.7,0.9)	(0.4,0.7,0.9)	(0.4,0.7,0.9)	(0.4,0.7,0.9)
1 p ahead	0.0237	0.0868	0.163	0.5207	9.9201	0.0136
2 p ahead	0.0194	0.0936	1.7627	2.1952	9.6271	1.5339
3 p ahead	0.0056	0.0958	3.1846	1.3149	11.8802	4.7197
4 p ahead	0.0001	0.2029	1.6712	0.0082	33.93	3.2143
5 p ahead	0.0782	0.0029	0.8245	4.2719	23.2122	0.1401
6 p ahead	0.0282	0.0174	7.5393	14.1611	21.2178	0.4239
7 p ahead	0.0314	0.5498	6.8731	14.0603	16.5857	0.0388
AR(2) Model						
MSE ₄₄	0.2303	0.2315	0.7385	0.2672	65.3395	5.3652
$(\beta_1, \beta_2, \beta_3)$	(8.94, 0.91, -0.07)	(6.77, 0.75, 0.13)	(27.06, 1.39, -0.87)	(18.88, 1.64, -0.96)	(26.36, 1.30, -0.75)	(29.39, 1.20, -0.70)
1 p ahead	1.0385	0.9997	0.0985	1.0364	13.3249	0.1057
2 p ahead	0.7992	0.1967	0.1547	2.6497	2.2298	3.9212
3 p ahead	1.1925	0.2304	0.6581	4.7738	58.221	5.913
4 p ahead	0.2565	0.1628	1.3953	1.3237	135.6244	3.91
5 p ahead	0.7563	0.0333	0.6525	0.0253	85.4903	0.0767
6 p ahead	2.7673	0.7527	13.3626	0.0072	31.3737	4.8219
7 p ahead	2.6237	0.223	18.3608	0.113	6.62	1.7284

convergence, takes a lead already around period 40 (cf. Fig. 11). In the oscillating groups 1 and 6 the out-of-sample errors generated by the switching model are varying with the time horizon. The errors are especially large for the 6– and 7–periods ahead forecasts. The ultimate reason for the relatively large prediction errors is the oscillating behavior observed in the experiment. Even if the switching model with leading LAA heuristic captures this phenomenon qualitatively and also produces oscillations, the error can become large when the two types of oscillations have different frequencies and the prediction goes out of phase. Notice that the same forecasting problem arises also for the AR(2) model in group 1. The prediction errors for the groups 4 and 7 with damping oscillations are not very high, when compared with the in-sample MSE. This is because, towards the end of the experiment, in both treatments convergence was observed. At the end of the in-sample period, i.e., at $t = 43$, the switching model does not yet clearly select the ADA heuristic for group 4, but does select it for group 7. That explains why the forecasting errors in group 4 are larger than in group 7.

Comparing the forecasting errors of the switching and AR(2) models we conclude that the former model is better than the latter, on average. More specifically, in the converging group 5 and oscillating group 6 the out-of-sample performances are very similar, but in the groups

Table 5: **Out-of-sample performance of the heuristic learning model and of the AR(2) model.** The averaged in-sample MSE (for $t = j, \dots, 43 + j$ with $j = 0, \dots, 5$) and averaged 1- and 2- periods ahead squared forecasting errors are shown. The *top* part: the best fitted in-sample switching model. The *middle* part: the benchmark switching model. The *bottom* part: AR(2) model estimated in-sample.

	Group 2	Group 5	Group 1	Group 6	Group 4	Group 7
Best Fit Switching Model						
average MSE ₄₀	0.0422	0.0383	0.4519	0.1702	42.6314	3.6141
1 p ahead	<i>0.0122</i>	<i>0.0321</i>	0.479	<i>0.1921</i>	15.0395	0.7857
2 p ahead	<i>0.0122</i>	<i>0.0901</i>	1.8599	<i>1.0792</i>	57.5144	<i>1.5543</i>
Benchmark Switching Model						
average MSE ₄₀	0.0641	0.1037	0.5242	0.2869	58.2751	5.0055
1 p ahead	0.0417	0.1456	<i>0.4186</i>	0.4097	<i>7.694</i>	0.7922
2 p ahead	0.0801	0.2304	<i>1.5871</i>	1.9035	<i>16.2213</i>	1.9338
AR(2) Model						
average MSE ₄₀	0.2338	0.2407	0.6897	0.2514	65.7466	4.7867
1 p ahead	0.3732	0.4431	0.9981	0.5568	13.4616	<i>0.6682</i>
2 p ahead	0.5052	0.4045	3.5823	1.4944	44.6453	2.0098

1, 2, 4 and 7 the different variations of the switching model outperform the AR(2) model out-of-sample. Furthermore, in group 4, for 3-, 4- and 5- periods ahead predictions, the AR(2) model produces errors larger than 7 in absolute value, which would lead to 0 earnings in the experiment, see (2.4). The switching model never has such large errors.

To further investigate one- and two-period ahead out-of-sample forecasting performance, we fit the model on a moving sample and compute an average of the corresponding squared prediction errors. The results of this exercise for the same three models (best fit, benchmark, and AR(2)) are presented in Table 5. The smallest forecasting errors within the same class (e.g., one-period ahead in group 4) are shown in italic. It turns out that our nonlinear switching model predicts the experimental data better than the simple AR(2) model for both one- and two-period ahead predictions in all groups except for groups 4 and group 1 for the one-period ahead prediction (for group 4, the average out-of-sample prediction of the benchmark switching model however outperforms the AR2 model). In many cases the difference in performance is substantial.

6 Conclusion

The time evolution of aggregate economic variables, such as stock prices, is affected by market expectations of individual agents. 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.

Our results show that performance based 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, e.g. Arthur (1991), Arthur (1993), Erev and Roth (1998) and Camerer and Ho (1999). However, our model is built in a different environment from those which are studied in the standard game theory. For example, agents in our framework do not have a well defined strategy (except than predicting a number between 0 and 100, in two decimals) and they do not know the payoff matrix, which is moreover changing over time in a path-dependent 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.

How robust are these results? While it is beyond the scope of the current paper to study in detail the performance of our nonlinear switching model in other market environments, Fig. 12 illustrates that our switching model also fits some other learning to forecast asset pricing experimental groups in Hommes, Sonnemans, Tuinstra, and Velden (2005) quite nicely. Price behaviour in group 3 (top panel) starts with moderate oscillations, then stabilizes at a level below the fundamental price 60, but suddenly falls in period $t = 40$, probably due to a typing error of one of the participants¹⁸. Our nonlinear switching model of course misses the typing error, but nevertheless matches the overall pattern before and after the unexpected drop quite nicely. Group 8 (middle panel) is an example with a fundamental price of 40 (instead of 60), whose strong oscillations are captured by an initially dominating STR rule, until the LAA rule becomes dominating after period 40. Group 12 (bottom panel) is an example *without* robot traders (i.e. its fraction n_t in the price generating law of motion (2.2) is fixed to 0). Large amplitude price oscillations, with prices almost reaching their maximum 100 and minimum 0, arise, but prices stabilize towards the end of the experiment. The evolution of the impacts of the four forecasting heuristics is similar to the case of dampening price oscillations considered before, with the strong trend rule (STR) initially dominating, the market taken over by the anchor and adjustment rule (LAA) around period 25 and the adaptive expectations (AA) rule finally dominating towards the end. The fact that our nonlinear switching model nicely captures these different patterns shows that the model is robust concerning small in the market setting.

In a related paper, Hommes and Lux (2009) present a heterogeneous expectations model, where individual agents use a genetic algorithm, to explain learning to forecast experiments in a cobweb (hog-cycle) market environment. The GA-learning model captures all stylized facts of the cobweb markets, both at the individual and at the aggregate level, quite nicely across different experimental treatments. In the cobweb markets, aggregate price behavior is

¹⁸The sudden fall of the asset price in group 3 from 55.10 in period 40 to 46.93 in period 41 is due to the fact that one of the participants predicts 5.25 for period 42. It is likely that this corresponds to a typing error (perhaps his/her intention was to type 55.25), since this participants five previous predictions all were between 55.00 and 55.40.

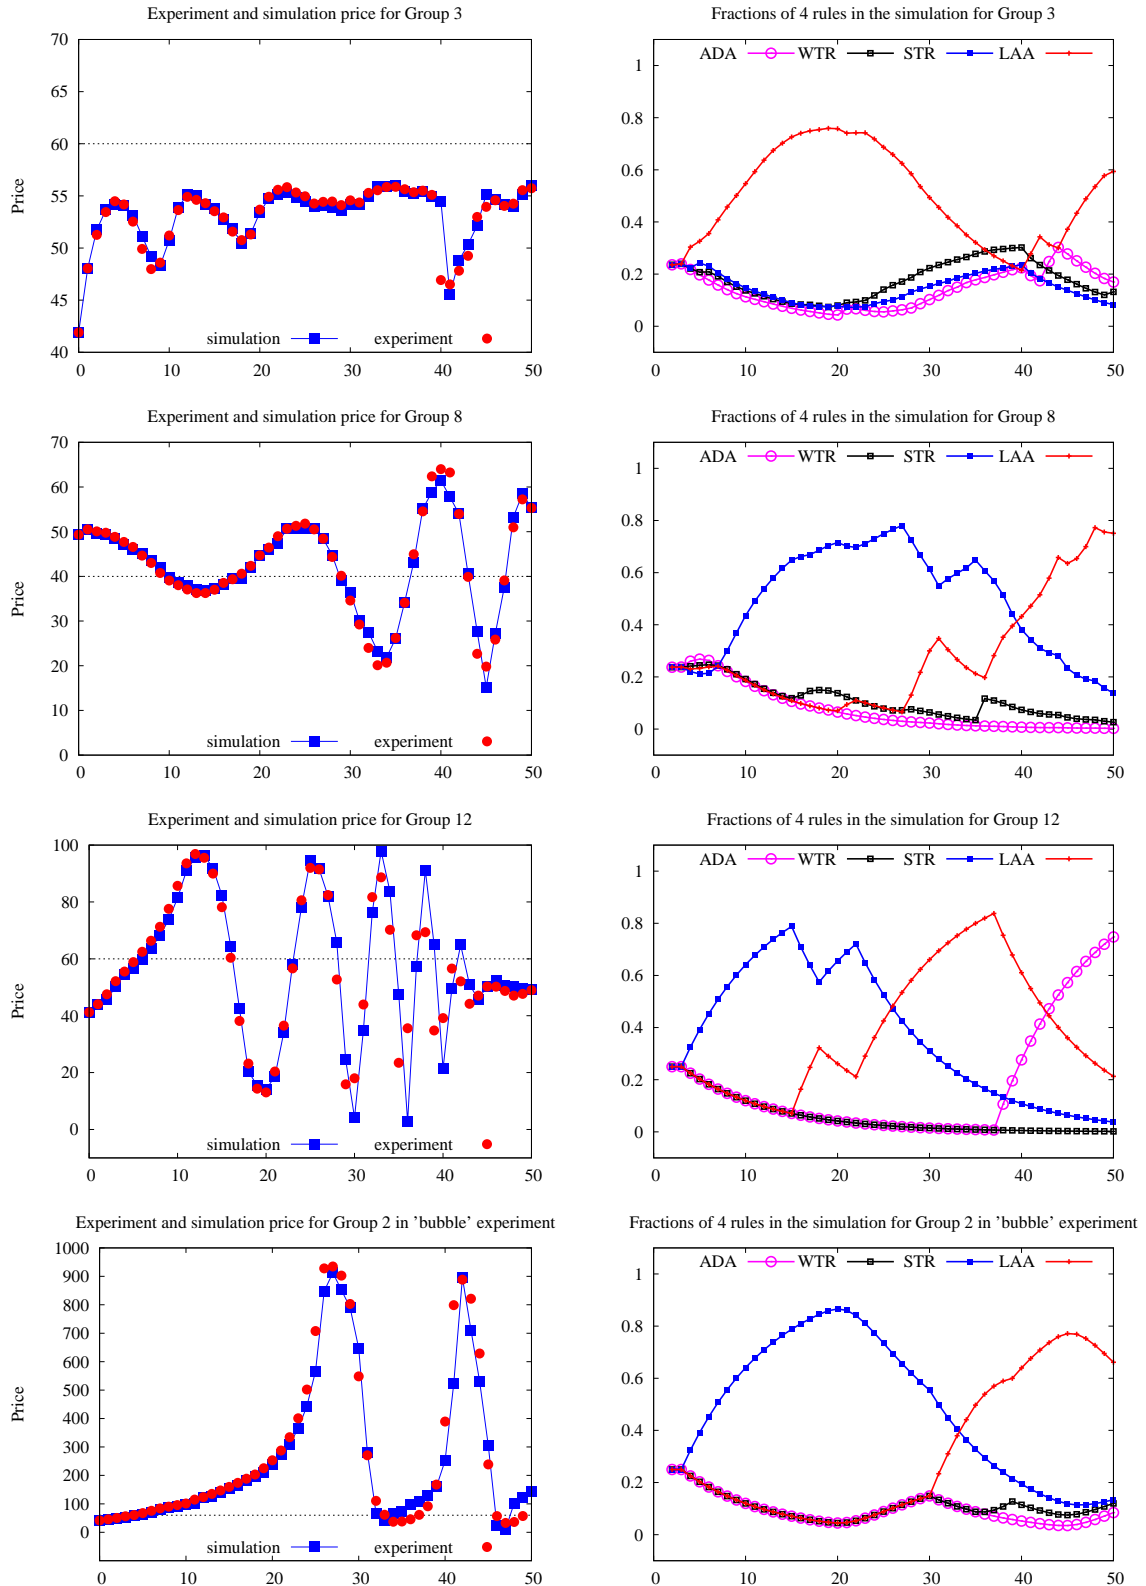


Figure 12: **One-step ahead predictions of the evolutionary switching model for four other experimental groups: (from top to bottom) group with typing error, group with fundamental 40, group without robot traders, group with possibility to submit price forecasts up to 1000. Left panels show prices in experiments (red) with corresponding one-step ahead predictions of the switching model (blue). Right panels show corresponding fractions of the four heuristics: adaptive expectations (ADA, purple), weak trend followers (WTR, black), strong trend followers (STR, blue) and anchoring and adjustment heuristic (LAA, red).**

characterized by random looking (i.e. zero-autocorrelations) price fluctuations around the RE benchmark price. The amplitude of these price fluctuations depends critically upon underlying parameters, especially upon the slopes of demand and supply curves. In the case when the cobweb market is stable under naive expectations, the laboratory experiments as well as the GA-simulations, nicely converge to the RE benchmark price. On the other hand, when the cobweb is unstable under naive expectations, experimental as well as GA-simulated prices are characterized by excess price volatility with zero autocorrelations. Our performance based switching heuristics model can explain the case of convergence to RE, but has more difficulty in capturing simultaneously the high price volatility with randomly (no autocorrelations) looking prices. Apparently, the random looking part of strongly fluctuating prices is better captured by the partly by noise driven GA-simulations than by our simple heuristics switching model. We conjecture that our simple heuristics switching model works well in market environments with at least some structure and persistence in price fluctuations, e.g. possibly irregular, but at least temporarily persistent price trends. A more systematic investigation in which market environments our nonlinear heuristics switching model can explain individual forecasting behavior as well as aggregate price behavior is left for future research.

References

- ALCHIAN, A. A. (1950): "Uncertainty, Evolution, and Economic Theory," *The Journal of Political Economy*, 58(3), 211–221.
- ANUFRIEV, M., AND C. HOMMES (2009): "Evolution of Market Heuristics," *Knowledge Engineering Review*, forthcoming.
- 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.
- BRANCH, W., AND G. EVANS (2006): "Intrinsic heterogeneity in expectation formation," *Journal of Economic Theory*, 127(1), 264–295.
- 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.
- CAMERER, C. F., AND T.-H. HO (1999): "Experienced-Weighted Attraction Learning in Normal Form Games," *Econometrica*, 67(4), 827–874.

- 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.
- DIKS, C., AND R. VAN DER WEIDE (2005): “Herding, a-synchronous updating and heterogeneity in memory in a CBS,” *Journal of Economic Dynamics and Control*, 29, 741–763.
- 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.
- HEEMEIJER, P., C. HOMMES, J. SONNEMANS, AND J. TUINSTRA (2009): “Price stability and volatility in markets with positive and negative expectations feedback: An experimental investigation,” *Journal of Economic Dynamics and Control*, 33, 1052–1072.
- 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., H. HUANG, AND D. WANG (2005): “A robust rational route to randomness in a simple asset pricing model,” *Journal of Economic Dynamics and Control*, 29, 1043–1072.
- HOMMES, C., AND T. LUX (2009): “Individual Expectations and Aggregate Behavior in Learning to Forecast Experiments,” Discussion paper, CeNDEF Working Paper.
- HOMMES, C., J. SONNEMANS, J. TUINSTRA, AND H. VAN DE VELDEN (2008): “Expectations and bubbles in asset pricing experiments,” *Journal of Economic Behavior and Organization*, 67(1), 116–133.
- 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.

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 (2.6) 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

$$J = \begin{vmatrix} \frac{\beta_1}{1+r} & \frac{\beta_2}{1+r} \\ 1 & 0 \end{vmatrix}.$$

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

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

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

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

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

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

This characteristics polynomial is given in (4.10). □