

Evolutionary Selection of Expectations in Positive and Negative Feedback Markets

Mikhail Anufriev · Cars H. Hommes · Raoul
H.S. Philipse

Abstract An economic environment is a feedback system, where the dynamics of aggregate variables depend on individual expectations and *vice versa*. The type of feedback mechanism is crucial for the aggregate outcome. Experiments with human subjects (Heemeijer, Hommes, Sonnemans, and Tuinstra, 2009) have shown that price converges to the fundamental level in a negative feedback environment but fails to do so under positive feedback. We present an explanation of these experimental results by means of a model of evolutionary switching between heuristics. Active heuristics are chosen endogenously, on the basis of their past performance. Under negative feedback an adaptive heuristic dominates explaining fast price convergence, whereas under positive feedback a trend-following heuristic dominates resulting in persistent price deviations and oscillations.

JEL codes: D03, G12, D84, C91, C92.

Keywords: Evolutionary learning; Expectations feedback; Experiments.

This paper is **forthcoming** in the *Journal of Evolutionary Economics*.

M. Anufriev · C. Hommes
CeNDEF, School of Economics, University of Amsterdam,
Roetersstraat 11, NL-1018 WB Amsterdam, the Netherlands
and
Tinbergen Institute,
Gustav Mahlerplein 117, NL-1082 MS Amsterdam, the Netherlands

R. Philipse
University of Oxford, Pembroke College,
OX1 1DW, Oxford, United Kingdom

1 Introduction

Expectations play a major role in both financial and commodity markets affecting individual decisions and having a large impact on aggregate variables, such as prices. Sometimes the influence of expectations seems larger than that of other, more tangible fundamental factors. Recent literature on Heterogeneous Agent Models shows that when non-fundamental expectations are widespread among market participants they can result in price bubbles, see, e.g., the review in Hommes (2006). Dynamics of heterogeneous, non-fundamental expectations might explain different “stylized facts” of financial markets, above all excess volatility of asset prices (Shiller, 1981). For this reason it is important to investigate how individuals form expectations about market variables and what is their aggregate effect. The expectations of the participants in real markets are not easily observable and are also affected by different uncontrolled factors. This makes experiments in a controlled environment particularly attractive for studying expectations (Duffy, 2006). The number of experiments on expectations is growing, see, e.g., Hey (1994), Hommes, Sonnemans, Tuinstra, and van de Velden (2005, 2007, 2008), Adam (2007) and Fehr and Tyran (2008), see Hommes (2011) for a recent survey. By and large these experiments reject the dominating in economic theory Rational Expectations hypothesis. The development of an alternative theory of expectation formation thus becomes an important task.

In this paper we contribute to this task by fitting the Heuristics Switching Model (HSM), introduced in Anufriev and Hommes (2009, 2012), to a recent experiment on expectation formation presented in Heemeijer, Hommes, Sonnemans, and Tuinstra (2009) (HHST, henceforth). This experiment is one of several “learning to forecast” experiments that have been performed in the CREED laboratory of the University of Amsterdam (UvA). Participants of these experiments have a single task to predict next period’s realization of an aggregate variable (price), which depends on the average of all individual predictions. The details of the pricing equation may change in these experiments, see the examples in Hommes, Sonnemans, Tuinstra, and van de Velden (2005), Hommes, Sonnemans, Tuinstra, and van de Velden (2007), Hommes, Sonnemans, Tuinstra, and van de Velden (2008) and Sonnemans and Tuinstra (2010). These early experiments have already pointed out that the type of feedback between expectations and realization is important. In the “hog cycle” model of a market for an agricultural good with a production lag, when the expectations about next period’s price are high, producers produce a large amount of the good and the realized price is low. This is an example of a *negative* feedback system. The opposite situation holds in *positive* feedback systems. For example, if expectations about the future price of a speculative asset are high, the demand for the asset will increase and the realized price will be also high. Experiments in Heemeijer, Hommes, Sonnemans, and Tuinstra (2009) focus on the role of the feedback type. In a perfectly symmetric setting, where the only difference between two treatments is the sign of the average forecast in the pricing equation, two different qualitative patterns have been observed. In the negative feedback treatments, where higher expectations lead to lower price, the realized market price converges to the fundamental level very quickly and remains there for the rest of the experiment. In contrast, in a positive feedback market, where higher expectations lead to higher price realizations, the price exhibits fluctuations and long lasting deviations from the fundamental level. The purpose of this paper is to demonstrate that a *single* model of individual expectations can explain these different aggregate outcomes.

Estimation of individual forecasts in HHST reveals that participants of the experiment tended to use simple behavioral rules in their forecasting activity. These rules strongly depended on the feedback type. This paper takes these results as a starting point to fit the data with a heuristic switching model (HSM) which also explains why different rules are chosen under different feedback. This HSM proposed in Anufriev and Hommes (2012) extends the adaptive belief framework of Brock and Hommes (1997). In the HSM agents’ expectations are governed by several behavioral heuristics, e.g., adaptive or trend-following rules. Agents switch between heuristics learning to use those heuristics which performed better in the past. The impact of different heuristics on the aggregate price is, therefore, changing over time, explaining different aggregate outcome observed in the experiment. For the negative feedback experiment we find that in the initial periods the adaptive heuristic performs much better than the others. As this heuristic attracts enough participants, the price dynamics exhibit convergence to the fundamental level with small oscillations around it due to small noise. Along this

converging time path the adaptive heuristic performs slightly better than the others, supporting this outcome in the long run. In the positive feedback environment the performance of the trend-following heuristic is better. When this heuristic is used by a majority of participants, the price deviates from the fundamental value leading to persistent price trends, and confirming the dominating trend-following heuristic.

The rest of the paper is organized as follows. In Section 2 we set out the background on which this work rests, discuss bounded rationality and recall the HHST experiment. Section 3 presents models with a single heuristic and HSM with several heuristics replicating the experimental data. In Section 4 we compare the performances of different models, optimize the best performing model and discuss the robustness of the model simulations. Finally, Section 5 presents the conclusions drawn from our research.

2 Agents' Rationality and Market Feedback

One of the main theoretical and methodological foundations of this paper is the concept of *bounded rationality* initially proposed by Herbert Simon, see, e.g., Simon (1957). According to Sargent (1993) bounded rationality provides an alternative to the stringent restrictions enforced by the concept of rational expectations (RE). Full rationality imposes two requirements on economic models. The first is that individuals act rationally, maximizing some objective function subject to perceived constraints. The second requirement is that there is consistency of perceptions about these constraints among different individuals, which is guaranteed by REs. REs require agents to have profound knowledge about their surroundings, such as knowledge of the precise equations of the environment and expectations of other agents. This suggests that REs only have limited and mostly theoretical appeal. We may expect from the outset that RE models will perform poorly descriptively, because the assumptions required by the concept are unlikely to be satisfied in most economic environments.

In contrast, bounded rationality does not assume that economic agents behave as all-foreseeing beings who know the complicated structure of the whole economy. Instead it creates a system in which agents try to maximize their outcome by altering their forecasts to eliminate systematic forecast errors. For instance, the adaptive learning approach advocated by Sargent (1993) and extensively used in macroeconomics (e.g., Evans and Honkapohja, 2001) assumes that agents form expectations by means of some underlying economic model whose parameters are unknown. Agents behave like econometricians and estimate these parameters. By explicitly modeling a learning process, models of bounded rationality can be used to describe non-equilibrium dynamics and, hence, are more useful in explaining experimental data. At the same time, experiments can be used to discriminate between plausible learning models. By abstracting from irrelevant details of individual learning processes, a good model should, in particular, explain why aggregate variables, such as price, converge to the RE trajectory in some laboratory experiments, but fail to do so in others. Recent empirical evidence suggests that the type of the environment's feedback can be crucial for the outcome.

2.1 Feedback Mechanisms

In economic, biological and physical systems the type of feedback between output and inputs is important. A *positive* feedback system reinforces a change in input by responding to a perturbation in the same direction. In contrast, a *negative* feedback system reverses a change in input and responds to a perturbation in the opposite direction. Davidovits (2008) points out that the negative feedback inherent to most biological systems allows the system to keep its response to disturbances at a relatively constant level. Examples include the mechanism regulating body temperature by sweating or shivering and mechanism regulating the level of blood glucose whose failure would result in diabetes. Thus, a negative feedback stabilizes the system, as it forces it to return to an equilibrium state after a disequilibrium shock. This principle can be illustrated by a classical example of negative feedback in a physical system, the centrifugal governor. In *Selfish Gene* Dawkins (1976) explains

The Watt governor consists of a pair of balls which are whirled round by a steam engine. Each ball is on the end of a hinged arm. The faster the balls fly round, the

more does centrifugal force push the arms towards a horizontal position, this tendency being resisted by gravity. The arms are connected to the steam valve feeding the engine, in such a way that the steam tends to be shut off when the arms approach the horizontal position. So, if the engine goes too fast, some of its steam will be shut off, and it will tend to slow down. If it slows down too much, more steam will automatically be fed to it by the valve, and it will speed up again.

On the other hand, in physics and economics many systems have positive feedback which may trigger and reinforce instability. The chain of nuclear fission reactions is caused by a single neutron bombarding a nucleus of Uranium 235. As a result of collision, the Uranium 235 nucleus splits into smaller parts, liberates energy and emits new neutrons. These neutrons hit other Uranium 235 nuclei and so on.

In economic markets in general both types of feedback will play a role. On the whole however, negative feedback is usually associated with supply driven commodity markets and positive feedback with demand driven speculative asset markets. In commodity markets, if producers expect an increase in future prices of a certain commodity, they will start producing more of it. According to the law of supply and demand this will then lead to a decrease in price. In a speculative asset market, if all traders expect a stock's value to increase they will start buying the stock. This creates a self-fulfilling prophecy by causing the stock to increase in value. Some economists argue that financial markets can easily lead to a state of speculative bubbles, along which the growth of the asset prices is mainly driven by investors' optimistic expectations, "animal spirit", rather than by some fundamental factors of economy, see, e.g., Akerlof and Shiller (2009) and Kirman (2011).

The distinction between positive expectations feedback and negative expectations feedback is related to the idea of strategic complements and substitutes (Haltiwanger and Waldman, 1989). If the majority of agents predict a higher price level in the market with positive feedback, every other agent is better off adjusting his/her own prediction upwards as well. Thus every agent has an incentive to behave similarly to others. Such behavior is typical for an environment with strategic substitutability. On the other hand, in a market with negative feedback, every agent has an incentive to considerably deviate from the majority, decreasing his/her prediction in response to an increase of the prediction from the majority. This is a characteristic of environment with strategic complementarity.

2.2 Experimental Evidence

A number of experiments focusing on the type of feedback have recently been performed, see, e.g., Fehr and Tyran (2008) and Sutan and Willinger (2009). This paper builds upon the learning to forecast experiment of Heemeijer, Hommes, Sonnemans, and Tuinstra (2009), HHST, henceforth.¹ The experiment is designed in such a way that the only difference between two treatments is whether the feedback is positive or negative. In both treatments the price formation process is of the form

$$p_t = f(\bar{p}_t^e),$$

where \bar{p}_t^e is the average price forecast of all the participants. The functional form of f is linear but its slope depends on the type of market. In the negative feedback market the (unknown) price generating law of motion is given by

$$p_t = 60 - \frac{20}{21} (\bar{p}_t^e - 60) + \varepsilon_t, \quad (1)$$

while in the positive feedback market the law of motion is given by

$$p_t = 60 + \frac{20}{21} (\bar{p}_t^e - 60) + \varepsilon_t, \quad (2)$$

where $\varepsilon_t \sim N(0, 1/4)$ is a random sequence, the same in all experimental sessions. For both price-generating processes the rational expectations equilibrium, which require $\bar{p}_t^e = p_t$, lead to the same RE outcome, $p_t = 60 + \varepsilon_t$. Thus under both negative and positive feedback the

¹ The data of all the learning to forecast experiments, including the experiments we discuss in this paper, are freely available at the personal web page of Cars Hommes, <http://www1.fee.uva.nl/cendef>.

same equilibrium strategy prevails, when all participants predict 60. Further, both processes have the same absolute slope coefficient of $20/21$ with respect to the average forecast. The only difference between the two processes is the sign of the slope coefficient.²

Thirteen experimental sessions have been performed in the CREED laboratory at the University of Amsterdam, six in the negative feedback treatment with the price given by (1), and seven in the positive feedback treatment with the price given by (2). In each session 6 participants had to predict the price during 50 periods (with $t = 0, \dots, 49$). Let $p_{i,t}^e$ denote the forecast for the price p_t submitted by participant i at the beginning of period t . The average price forecast in (1) and (2) is computed as

$$\bar{p}_t^e = \frac{1}{6} \sum_{i=1}^6 p_{i,t}^e. \quad (3)$$

Participants were paid according to accumulated rewards, with the reward per period computed (in euros) as

$$e_{i,t} = \frac{1}{2} \max \left(1 - \left(\frac{p_t - p_{i,t}^e}{7} \right)^2, 0 \right). \quad (4)$$

Therefore, the accuracy of their prediction was the only concern of participants. Participants were not informed explicitly about the price generating process. Before the experiment, they were given qualitative information on how the market functions. They knew that the price depends on the aggregate price prediction of a number of players and they knew the type of feedback, positive or negative. They did not know the identity and the number of other participants. Moreover, participants had no information about forecasts of others. In each period, after the individual predictions were submitted, the aggregate price has been computed from market clearing and shown to the participants, graphically and in table form, together with their own last and previous forecasts and rewards.

Experimental Outcome: Aggregate Price Dynamics

Fig. 1 shows the market price observed in the positive and negative feedback experimental markets. The price in the negative feedback treatments fluctuates heavily in the first 5 periods and then quickly converges towards the value 60 to which we refer as “the equilibrium price”. The price in the positive feedback treatments does not converge to the equilibrium value and in several groups exhibits slow oscillatory movement with relatively high amplitude.³

Another interesting characteristic of the experimental outcomes is that the price development seems to depend on the price in the first period. In the experiment this initial price p_0 is a result of more or less random guesses made by the participants. Indeed, at time 0 participants do not have any knowledge of past prices. They are only given the range $[0, 100]$ where the initial forecast has to lie. Apparently, the initial price realization had a big influence on the rest of the dynamics. In the negative feedback market shown in Fig. 1(a), prices fluctuated more strongly in the beginning, when the initial price significantly differed from the equilibrium level. In the positive feedback market a more striking difference could be observed, see Fig. 1(b). In all three groups where the standard deviation in the last 30 periods was less than 1, the initial price was in the range of 20 from the equilibrium level. On the other hand, in both groups where the initial price differed from the equilibrium level by a value larger than 20, strong oscillations were observed during the experiment.

Experimental Outcome: Individual Predictions

HHST investigated the prediction behavior of individuals by estimating parameters c , o_l and s_l of a linear heuristic

$$p_{i,t}^e = c + \sum_{l=1}^3 o_l p_{t-l} + \sum_{l=1}^3 s_l p_{i,t-l}^e,$$

² Bao, Hommes, Sonnemans, and Tuinstra (2010) ran similar learning to forecast experiment with a time varying fundamental price subject to large unanticipated shocks.

³ Ignoring the first 5 periods, the standard deviation of the price in 6 groups with negative feedback were 0.63, 0.93, 2.97, 0.57, 1.55 and 0.84. In all these cases the interval of oscillations contained the price 60. Ignoring the first 20 periods, the standard deviation of the price in 7 groups with positive feedback was 3.35, 0.89, 0.82, 7.01, 2.63, 5.96 and 0.80. In all three groups where the standard deviation was less than 1, the equilibrium price 60 *did not* belong to the interval of oscillations.




Fig. 1 Price dynamics in the HHST experiment. (a) The price quickly converges to the equilibrium level 60 under negative feedback in all 6 experimental groups. (b) The price never converges to the equilibrium level 60 under positive feedback and exhibits strong oscillations over 50 periods of the experiment in 4 of the 7 groups.

where p_t is the price at time t and $p_{i,t}^e$ is the expectations of individual i for the price p_t . They found that such a simple, linear rule can be fitted to 71 out of 78 participants (without autocorrelations in the residuals). Furthermore, the forecasts given by 40 of the 78 participants can be described by an even simpler heuristic in the form of

$$p_{i,t}^e = \alpha_1 p_{t-1} + \alpha_2 p_{i,t-1}^e + (1 - \alpha_1 - \alpha_2)60 + \gamma(p_{t-1} - p_{t-2}). \quad (5)$$

This *first-order* heuristic depends only on the last observed price, the last own forecast and the last price change. A key difference between the two treatments has been found in the estimation of the trend parameter, γ . In the sessions with negative feedback, γ was usually non-significant, while in the sessions with positive feedback, γ was, in most cases, significantly positive.⁴ The use of simple forecasting rule is consistent with the finding of the behavioral literature that agents often replace optimizing rules with relatively simple rules of thumb, see, e.g., Conlisk (1996).

2.3 Evolutionary Model of Heuristics

Which model of expectations would fit the results of the HHST experiment? As discussed above, the RE model implies that individuals immediately predict the equilibrium price 60. The results of the experiment show strong and persistent deviations from the equilibrium especially in the positive feedback treatment. This suggests that the more flexible concept of bounded rationality can be more relevant. Marcet and Sargent (1989) introduced *adaptive learning* models, where agents behave as econometricians, applying some linear model for price prediction and re-estimating the coefficients of the model as more and more data are available. The models of adaptive learning generate convergence, although with different speed, to the equilibrium price for both feedback dynamics, (1) and (2), considered in the HHST experiment (Evans and Honkapohja, 2001). This outcome, however, is at odds with the experimental evidence. In fact, it is not very surprising that RE or adaptive learning dynamics do not fit experiments well. Both approaches require a relatively high amount of knowledge from the participants about the environment where they operate: under REs agents have a full knowledge of the equations governing the system, while under adaptive learning agents know the functional form of these equations.

The “learning to forecast” experiments are made in such a way that participants have limited knowledge about the environment. Every time period of the experiment can be thought of as a one-shot game in which every participant has many strategies (every positive number up to two decimals), and the payoffs are given by (4). In this interpretation of the experiment,

⁴ The value of γ was significantly different from 0 in 3 out of 19 first-order heuristics (5) estimated for the participants in the negative feedback experiments and in 15 out of 21 first-order heuristics estimated for the participants in the positive feedback experiments. In the former case values of γ were estimated as -0.44 , -0.38 and 0.06 . In the latter case values of γ in the range from 0.28 to 0.97 were observed.

participants are not explained that they are playing a game against each other and do not know the exact payoff structure, since the equation for realized price is outside of their information set. Such an experimental setting is attractive, because it corresponds to many real world situations in which individuals make decisions without detailed information and understanding of the underlying (complex) environmental structure. The behavior of agents in such situation is nicely described by Mailath (1998), p. 1349-1350:

The typical agent is not like Gary Kasparov, the world champion chess player who knows the rules of chess, but also knows that he doesn't know the winning strategy. In most situations, people do not know they are playing a game. Rather, people have some (perhaps imprecise) notion of the environment they are in, their possible opponents, the actions they and their opponents have available, and the possible payoff implications of different actions. These people use heuristics and rules of thumb (generated from experience) to guide behavior; sometimes these heuristics work well and sometimes they don't. These heuristics can generate behavior that is inconsistent with straightforward maximization.

To explain the experimental results we use the model of Anufriev and Hommes (2012) where predictions are made by simple behavioral rules, heuristics, as suggested by the above quote. To explain the behavior of agents in *both* positive and negative feedback experiments, the forecasting heuristics are combined in a so-called switching model, which is an extension of Brock and Hommes (1997). More than one forecasting heuristic is tried by agents and when a heuristic performs well it attracts more followers. Thus, the relative weights attached to different heuristics are time-varying.

3 Model and Simulations

In this section we, first, present two different forecasting heuristics and study the dynamics of the corresponding models with homogeneous expectations. Then, we combine the heuristics in the switching model.

We evaluate the explanatory power of different models in two ways. In this section we apply a first test. We compare a 50-periods ahead model simulation (so-called *simulated path*) with the experimental outcomes. The price dynamics, generated by a model, must have some of the characteristics of the price developments observed in the experiments. In the negative feedback market the price must fluctuate heavily in the first few periods and then converge quickly to the equilibrium value, see Fig. 1(a). In the positive feedback market the price must oscillate slowly and either gradually converge towards the constant level or not converge at all, see Fig. 1(b). After identifying a model that satisfies these requirements for the simulated path of 50 periods, we apply a second test in Section 4. There the model performance in one-period-ahead forecasting is investigated numerically, the models are optimized over parameters and tested for robustness.

3.1 First-order heuristics

Before defining the switching model a small number of forecasting heuristics has to be chosen. A prediction heuristic must fulfill a certain number of conditions to be relevant. It should be informationally accessible to the experiment's participants, and it should be simple and intuitive for them to use. We define the following two heuristics. The first heuristic is an *adaptive heuristic* given by

$$p_{t+1}^e = wp_t + (1 - w)p_t^e, \quad (6)$$

where an agent's expectation of the price in the next period, p_{t+1}^e , depends on the price of today and agents' expectation about the price of today, with weights w and $(1 - w)$, respectively.

The second prediction heuristic is a *trend heuristic* of the form

$$p_{t+1}^e = p_t + \gamma(p_t - p_{t-1}), \quad (7)$$

where the expectation for the price in the next period depends on the last price plus γ times the last price *change*. This heuristic is interpreted as people expecting a constant trend in price developments.

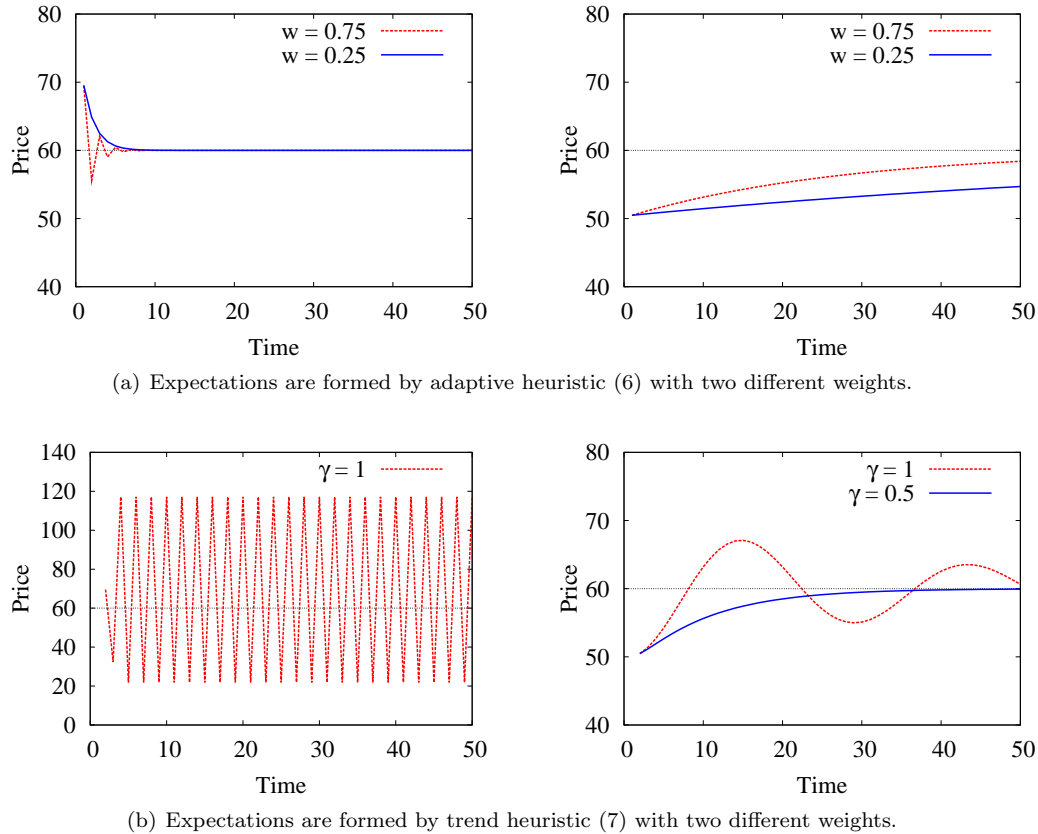


Fig. 2 50 periods ahead forecast obtained by the models with homogeneous expectations in the negative (*left* panels) and positive (*right* panels) feedback environments.

3.2 Dynamics generated by single heuristics

What kind of aggregate behavior does each of these heuristics imply for the negative and positive feedback environments? Fig. 2 shows the deterministic simulated paths with the adaptive and trend heuristics, i.e., 50-periods ahead simulations by the model with homogeneous expectations. The first forecasts generated by the heuristics are set to 50 in this simulation.

The adaptive heuristic (6) with $w = 0.75$ generates oscillatory converging price dynamics in the negative and monotonically converging price dynamics in the positive feedback markets, see Fig. 2(a). This outcome is similar to the dynamics observed in the negative feedback sessions, but is very different from the oscillations observed in the positive feedback sessions. Notice that a decrease of the weight w in (6) leads to less oscillating and, eventually, to monotonic convergence under negative feedback and very slow convergence under positive feedback, as illustrated by example with $w = 0.25$.

Straightforward stability analysis reveals that the trend heuristic (7) produces converging prices in the negative feedback market when $-21/20 < \gamma < 1/40$, and converging prices in the positive feedback market when $-41/40 < \gamma < 21/20$. For γ outside of the interval, the price dynamics under corresponding feedback diverge. Interestingly, all the values of γ estimated in the positive feedback sessions of the HHST experiment fall within the interval of convergence under positive feedback. But the dynamics with these γ 's would diverge under negative feedback. The dynamics of the model where expectations are generated by the trend heuristic are illustrated in Fig. 2(b). For $\gamma = 1$ the dynamics under negative feedback (left panel) does not converge to the equilibrium but remain bounded because the forecasts in the simulations are limited to the interval $[0, 100]$. When γ is changing to 0.5 the initial oscillations are less wild but eventually converge to the same 2-cycle. The right panel of Fig. 2(b) shows that for $\gamma = 1$ the price oscillates on the positive feedback market, resembling the experimental outcome. When γ decreases the oscillations are less pronounced and convergence eventually become monotonic, as shown for $\gamma = 0.5$.

3.3 Heuristic Switching Model

In the model of heterogeneous expectations, different heuristics can be used. Consequently, the average price expectations in (1) and (2) is given by

$$\bar{p}_t^e = \sum_{h=1}^H n_{h,t} p_{h,t}^e, \quad (8)$$

where H is the number of heuristics and $n_{h,t}$ is the impact of heuristic h at time t , which depends on the past performance of the heuristic.

The performance of the heuristics is measured by squared forecasting errors, consistently with incentives given to the participants in the experiment. The performance of heuristic h at time t is given by

$$U_{h,t} = -(p_t - p_{h,t}^e)^2 + \eta U_{h,t-1}, \quad (9)$$

where p_t is the realized price obtained by applying (1) or (2), respectively. The parameter $0 \leq \eta \leq 1$ represents the weight that agents attribute to past forecasting errors. The impact of heuristic h changes according to a discrete choice model with asynchronous updating (Hommes, Huang, and Wang, 2005; Diks and van der Weide, 2005)

$$n_{h,t+1} = \delta n_{h,t} + (1 - \delta) \frac{\exp(\beta U_{h,t})}{Z_t}, \quad (10)$$

where $Z_t = \sum_{h=1}^H \exp(\beta U_{h,t})$ is a normalization factor.

Two parameters are important in (10). The parameter δ is inversely related to the frequency with which every agent updates “active” forecasting heuristic. Positive values of δ capture the tendency of people to stick to their previously chosen rule despite the evidence that an alternative rule performs better. Such inertia is widely reported in experiments (Kahneman, 2003). In a large population, δ is also the average percentage of agents who do not update their heuristic in every period. The parameter $\beta \geq 0$ determines how strongly those agents who update their heuristic react to a difference in performance between heuristics. If $\beta = 0$ agents will not consider the differences in the performance of the heuristics at all; all heuristics will be given equal impacts. If, on the contrary, the value of β is very large, agents who update their forecasting heuristic will all switch to the best performing heuristic.

In order to simulate the model, one should

- choose H different forecasting heuristics;
- fix three learning parameters, β , η and δ ;
- initialize prices in order for the heuristics to yield the initial forecasts;
- initialize the impacts for all heuristics so that the initial forecasts are combined to determine the average price forecast.

Given these initializations, the model works as follows. For every time t , first, the forecasts $p_{h,t}^e$ of H heuristics are computed on the basis of past prices. Second, they are combined using (8) to provide the average price forecast. Third, the price predicted by the model at time t is computed using (1) for the negative feedback market or (2) for the positive feedback market. This price is denoted simply as p_t . Fourth, the performance of every heuristic $U_{h,t}$ is calculated using (9) on the basis of the realized price p_t . Finally, the relative impacts of heuristic for the next period are computed using (10). The same steps are then repeated for time $t + 1$, and so on.

3.4 Dynamics of the Heuristics Switching Model

Let us apply the heuristic switching model (HSM) given by (9) and (10) to the experimental results of HHST. The parameters of the heuristics are chosen as $w = 0.75$ for the adaptive heuristic (6), and $\gamma = 1$ for the trend heuristic (7). Recall that the heuristics with these parameters describe the two markets relatively well: the negative feedback market is well described by the adaptive heuristic, while the positive feedback market is well described by the trend heuristic. After some trial and error simulations, we set the parameters of the HSM




Fig. 3 Heuristic switching model with $H = 2$ heuristics, adaptive and trend. The parameters are $w = 0.75$, $\gamma = 1$, $\beta = 1.5$, $\delta = 0.1$ and $\eta = 0.1$. *Left*: Simulated path of price without noise (lines) and with experimental noise (circles). *Right*: Evolution of the heuristics' impacts.

to $\beta = 1.5$, $\delta = 0.1$ and $\eta = 0.1$. We also choose equal initial impacts of both heuristics and $p_0 = 50$ as the initial price.

Figs. 3(a) and 3(b) show the outcome of the model's simulations in, respectively, the negative and positive feedback environments. In the left panels the two types of dynamics of the simulated path are shown. The lines show the simulation without noise in the laws of motion (1) and (2), while the circles correspond to the simulation with the same noise realization ε_t as in the experiments. The right panels show the evolution of the heuristics' impacts for the simulated path without noise. We observe striking difference in the dynamics between the negative and positive feedback environments. Indeed, the price dynamics of the heuristic switching model do adhere to the characteristics of the experimental outcomes in both treatments. In the negative feedback market the price oscillates heavily in the first periods and then quickly converges. In the positive feedback market the price slowly and (when augmented by the experimental noise) persistently oscillates around the equilibrium.

It is particularly informative to analyze the evolution of the heuristics' impacts. When the feedback is negative, the impact of the trend heuristic immediately falls to almost 0 and increases later on only at the stage when the price has already converged to the equilibrium level, i.e., when the predictions of both heuristics are similar. When the feedback is positive, the opposite phenomenon takes place with the trend heuristic dominating from the beginning of the simulations. The intuition of this result is as follows. The trend heuristic performs well during the long phases of the trends and performs poorly during the periods with frequent fluctuations around the constant price. At the same time, an extensive use of the trend heuristic results in the trends under the positive feedback and in oscillations under the negative feedback, see Fig. 2(b). Thus, under the positive feedback, the success of the trend heuristic reinforces its use, which makes the trend in prices sustainable. The adaptive heuristic performs relatively poorly during the trend phases and loses its impact. On the other hand, under the negative feedback, the trend heuristic generates oscillatory dynamics on which it performs very poorly,




Fig. 4 Positive feedback environment. Dynamics of the heuristic switching model with $H = 2$ heuristics for different initial prices. *Left*: No noise is added to the law of motion (2). *Right*: Experimental noise is added to the law of motion (2).

much worse than adaptive heuristic. Coordination on the adaptive heuristic leads to fast convergence through initial oscillations, as shown in Fig. 2(a).

The initial 10 periods of the simulations are explained now, but what happens next? In the positive feedback market agents attach a higher impact to the trend heuristic at the trend phases and decrease their impacts when the price development changes direction. Even when all subjects use the trend-following heuristic, the trend cannot be sustained forever and, at a certain moment, the trend will be reverted and the impact of the adaptive heuristic will grow. This occurs in the periods 15 – 17 of the simulations. Afterwards the downward trend reinforces the use of the trend heuristic, but since the price is already close to the equilibrium, the relative impacts of the heuristics are similar. Notice that the model generates such convergence to the equilibrium only in the absence of noise, see the left panel of Fig. 3(b). In the negative feedback market, the price is stabilized at the level close to the equilibrium during the periods 10 – 15. However, the steady-state dynamics with price at the equilibrium level is not stable in the model with switching. Indeed, when the price stabilizes both heuristics give the same predictions and their impacts are the same. But the trend heuristic reinforces a trend and leads to the overshooting of the equilibrium level. As a result, dynamics converge to the 2-period cycle with price being very close to the equilibrium level but jumping around it.⁵ At the cycle, the forecast of the trend heuristic is worse than the forecast of the adaptive heuristic, which results in their different impacts: around 80% of the adaptive heuristic and around 20% of the trend heuristic.

The heuristic switching model is also able to reproduce the same pattern of dependence on the initial condition, which was observed in the HHST experiment. Fig. 4 shows that the different aggregate outcomes (convergence and oscillations) within the same environment of the positive feedback can be attributed to the path-dependent property of the HSM. Depending on the initial price level, the model produces qualitatively different outcomes during 50 periods. Both without (the left panel) and with noise (the right panel), the price generated by the model stays closer to the fundamental level during all simulation, when the initial price, p_0 , is closer to the fundamental level.

3.5 Discussion

The first simulations of the heuristic switching model point to a *behavioral* explanation of the difference in the experimental outcomes between positive and negative feedback markets. When people, as subjects in the experiment or agents in the model, cannot make strategic decisions due to the absence of full knowledge of the environment they are operating in, they rely on behavioral rules of thumb. In the learning to forecast experiment different rules are possible, and some of them provide better forecasts than others. The *learning* of agents then takes a form of evaluating different forecasting possibilities and switching to those which

⁵ This cycle is almost invisible at Fig. 3(a). The prices are 59.74418 and 60.2559 along the cycle. The impacts are 0.8125 and 0.1875 for the adaptive and trend heuristics, respectively.

performed better in the past. Agents in the HSM learn individually (not socially, through the interaction with others) by applying a counterfactual analysis of alternative forecasting rules on the basis of past data. As a result, the population of agents switch to more successful rules and the aggregate dynamics may change its properties (e.g., the trend in prices may revert). Then, *via* the re-evaluation of performances, dynamics feed back to the distribution of the rules' impacts. Three parameters of the model allow to capture the behavioral characteristics such as imperfect switching behavior and, consequently, heterogeneity (especially when β is small), inertia in switching (when δ is close to 1), and short memory of past performances (when η is close to 0).

While the behavioral assumptions underlying our model are known from the behavioral literature,⁶ our model does not aim to provide a precise description of the behavior of subjects in the experiment.⁷ Instead, the aim of the model is to outline a mechanism explaining both negative and positive feedback markets at the same time. Indeed, in the simulations discussed in Section 3.4 the same heuristics and values of the learning parameters have been used. Our model is, essentially, a parsimonious version of the numerous computational learning models based on genetic algorithms (see, e.g., Arifovic (1996) and Hommes and Lux (2011)) or its modifications such as Individual Evolutionary Learning (see Arifovic and Ledyard, 2007 and Anufriev, Arifovic, Ledyard, and Panchenko, 2011).

Our model also stresses importance of heterogeneity in the explanation of the experiments. According to the results of Section 3.2, the homogeneous expectations model with simple first-order heuristics we considered can not explain negative and positive feedback experimental data simultaneously. The results of the experiments can be explained, however, by assuming that agents learn to change their forecasting methods. The heuristic switching model can be simulated with many heuristics. For example, the model analyzed in Anufriev and Hommes (2012) had 4 different heuristics. We found, however, that the two heuristics are sufficient to reproduce the result of the HHST experiment qualitatively, and preferred such parsimonious version of the model over other possible specifications.⁸

4 Validation and Robustness

In this section we further examine the performance of the heuristic switching model with the adaptive and trend forecasting heuristics. First, we investigate the one-period-ahead forecasting performance of the HSM, compare it with performances of the other models, and also fit the HSM to the data.⁹ Second, we investigate robustness of the model with respect to the change in parameter values.

4.1 One-period ahead model predictions

We now apply the heuristic switching model in a different way than before. The difference is that at each period of the simulation the past *experimental* prices are used to compute the heuristics' forecasts and impacts. Every time the heuristic switching model will produce a one-period ahead forecast of the price. This forecast is based upon the same information as subjects had in the experiment. Formally, it means that in (6), (7) and (9) we substitute the price predicted by the model, p_t , by the price observed in one of the experimental groups,

⁶ HHST and other learning to forecast experiments showed that agents use first-order heuristics for forecasting. Indeed, the heuristics we have chosen for the simulations were used by some subjects in the experiment. Recent learning to forecast experiments provide evidence of *switching* between simple heuristics, see illustration in Anufriev and Hommes (2009). Anufriev, Bao, and Tuinstra (2011) estimate switching parameters in a multiple choice experiment.

⁷ For example, we ignore the fact that actual impacts in the experiment could be a multiple of 1/6, proceeding instead with arbitrary values of impacts, which is valid only under a continuum of agents. An alternative interpretation is that every agent combines two forecasts, adaptive and trend, with time varying weights.

⁸ We simulated the model with 3 heuristics, adding to the adaptive and trend heuristic a, so-called, anchoring and adjustment heuristic obtained from (5) by replacing the (unknown to the participants) fundamental price 60 by the average of all past prices. The anchoring and adjustment heuristic combines some features of the trend and adaptive heuristics. It is, therefore, not surprising that the dynamics of the model with 3 heuristics (not shown in the paper) is not very different from the dynamics shown in Fig. 3.

⁹ For other examples of fitting the similar learning models to the experimental or survey data, see Colucci and Valori (2006) and Branch (2004).




Fig. 5 Dynamics of the HHST experiment as compared with the one-period ahead predictions by the heuristic switching model, augmented by experimental noise. The parameters are $w = 0.75$, $\gamma = 1$, $\beta = 1.5$, $\delta = 0.1$ and $\eta = 0.1$. *Left*: Price dynamics in the experiment (points) and forecasts generated by the model (line). *Right*: Evolution of impacts of two heuristics.

$p_t^{\text{Gr X}}$. In particular, the impacts of the heuristics are updated now as in (10) on the basis of the performance given by

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

After the forecasts and impacts of all heuristics are computed, the corresponding price equation generates the one-period ahead prediction of the price by the HSM.

The one-period ahead predictions of the switching model with adaptive and trend heuristics are shown in the left panels of Fig. 5 for one of the groups from the negative feedback treatment and for two groups from the positive feedback treatment. In the right panels the dynamics of the impacts of the two heuristics are shown. In all these simulations the parameters of the model are as before, $\beta = 1.5$, $\delta = 0.1$ and $\eta = 0.1$. Fig. 5(a) shows the price behavior in the negative feedback market. The prices predicted by the HSM take slightly longer to converge than the price outcomes of the experiments. After an initial phase the adaptive

Model	Negative Feedback	Positive Feedback	12 groups
fundamental	2.5712	46.8344	24.7028
adaptive	2.3001	2.9992	2.6497
trend	21.1112	0.9260	11.0186
mixed	7.9798	1.0518	4.5158
HSM	3.2967	0.9065	2.1016

Table 1 MSEs of 5 different models over 47 time periods of the experiment. The results are averaged over all 6 experimental groups of the negative feedback treatment (2nd column), and over 6 out of 7 experimental groups in the positive feedback treatment (3rd column), see footnote 11. The last column displays the joint result for 12 experimental groups.

heuristic starts to clearly dominate the trend heuristic, explaining the price convergence to the fundamental level. Fig. 5(b) shows the price dynamics in the positive feedback market with a large deviation of the initial price from the fundamental value. The price deviates from the fundamental value persistently, oscillates and does not show any sign of convergence. The trend heuristic dominates during this simulation. Finally, Fig. 5(c) shows a different outcome in the positive feedback market, when the initial price is close to the fundamental value. The price dynamics now stays consistently above but close to the fundamental value and the impacts of the two heuristics are more or less balanced during the simulation.

4.2 Model optimization

The one-period ahead predictions by the heuristic switching model can be compared with similar predictions obtained by the other models. As a fitness measure we will use the mean squared error (MSE) of the model's predictions over k periods. For experimental group the MSE is given by

$$\text{MSE}_k = \frac{1}{k} \sum_{t=50-k}^{49} (p_t^{\text{Gr X}} - p_t^{\text{Mod}})^2,$$

where $p_t^{\text{Gr X}}$ is the realized price in group X of the experiment and p_t^{Mod} is the one-period ahead prediction of the model Mod, taken all information up to and including period $t - 1$ into account. When these MSEs are computed over the data in G different groups and then averaged, the result is denoted as MSE_{kG} .

Table 2 shows the MSEs of five different models over the last $k = 47$ periods (i.e., for $t = 3, \dots, 49$).¹⁰ The second column shows the MSE computed over 6 groups of the negative feedback experiment, the third column shows the MSE computed over 6 groups of the positive feedback experiment¹¹, and the last column gives the MSE over all these 12 groups. We compare the following five models: 'fundamental' model, corresponding to the RE, according to which the prediction is 60 every period; two models with homogeneous expectations, 'adaptive' and 'trend', when expectations generated by, respectively, (6) and (7) are directly plugged into the laws of motion; heterogeneous expectations model where these two heuristics are combined with fixed impacts equal to 50% each period; and the heuristic switching model described in Section 3.4. The benchmark values of parameters have been used in these models, $w = 0.75$, $\gamma = 1$, $\beta = 1.5$, $\delta = 0.1$ and $\eta = 0.1$. The MSEs for the best among the five models are shown in bold.

As expected from the previous analysis, in the negative feedback environment the fundamental model works well, and the adaptive model is better than the trend model. In contrast, in the positive feedback environment the trend model works very well, the adaptive model is

¹⁰ The errors of the first three periods are not taken into account for the following reason. Both heuristics of the HSM require two prices to be initialized (for adaptive heuristic we set $p_1^e = p_0^{\text{Gr X}}$). Therefore, at time $t = 2$ the first performances of the heuristics are obtained when their first predictions, $p_{h,2}^e$, are compared with experimental price $p_2^{\text{Gr X}}$. The performances determine the impacts and only from period $t = 3$ the HSM is able to generate price prediction. For comparison, in all other models we compute errors also from $t = 3$. We also have computed similar MSEs on the data starting from periods $t = 4, 5$ and 6 and obtained very similar results.

¹¹ We eliminated one of the groups of the positive feedback experiment as an outlier. In this group at time $t = 7$ one of the participants predicted 5250, making, presumably, a typing error (in previous three period the predictions of this participant were 50.0, 52.0 and 51.0). As a result, the realized price was 877.8 and was above 100 for another four periods.

Model			Negative Feedback	Positive Feedback	12 groups
adaptive	benchmark	MSE	2.3001	2.9992	2.6497
		w	0.75	0.75	0.75
		AIC	2070.3458	2220.0313	2150.1399
	optimized	BIC	2070.3458	2220.0313	2150.1399
		MSE	2.3001	1.8218	2.3680
		\hat{w}	0.75	1.00	0.91
trend	benchmark	AIC	2072.3451	1940.8575	2088.7528
		BIC	2076.6802	1945.1926	2093.0879
		MSE	21.1112	0.9260	11.0186
	optimized	γ	1.00	1.00	1.00
		AIC	3320.6515	1557.2206	2953.9288
		BIC	3320.6515	1557.2206	2953.9288
mixed	benchmark	MSE	1.9237	0.7480	2.2940
		$\hat{\gamma}$	-0.33	0.71	-0.16
		AIC	1971.5531	1438.7973	2070.8502
	optimized	BIC	1975.8881	1443.1323	2075.1852
		MSE	7.9798	1.0518	4.5158
		(w, γ, n)	(0.75, 1.00, 0.50)	(0.75, 1.00, 0.50)	(0.75, 1.00, 0.50)
HSM	benchmark	AIC	2771.9429	1629.0202	2450.8370
		BIC	2771.9429	1629.0202	2450.8370
		MSE	2.3001	0.7568	2.6276
	optimized	(w, γ, \hat{n})	(0.75, 1.00, 0.76)	(0.75, 1.00, 0.38)	(0.75, 1.00, 1.00)
		AIC	2072.3458	1445.4272	2147.4263
		BIC	2076.6809	1449.7623	2151.7614
HSM	benchmark	MSE	3.2967	0.9065	2.1016
		(β, η, δ)	(1.50, 0.10, 0.10)	(1.50, 0.10, 0.10)	(1.50, 0.10, 0.10)
		(w, γ)	(0.75, 1.00)	(0.75, 1.00)	(0.75, 1.00)
	optimized 1	AIC	2273.3741	1545.1924	2019.4464
		BIC	2273.3741	1545.1924	2019.4464
		MSE	2.3001	0.7329	1.5402
optimized 2	$(\hat{\beta}, \hat{\eta}, \hat{\delta})$	(1.44, 1.00, 0.00)	(0.01, 0.96, 0.34)	(0.02, 0.97, 0.00)	
	(w, γ)	(0.75, 1.00)	(0.75, 1.00)	(0.75, 1.00)	
	AIC	2076.3458	1431.2754	1850.1692	
HSM	optimized 1	BIC	2089.3510	1444.2806	1863.1743
		MSE	1.9225	0.7088	1.5051
		$(\hat{\beta}, \hat{\eta}, \hat{\delta})$	(1.46, 1.00, 0.15)	(0.48, 0.85, 0.00)	(0.37, 0.87, 0.00)
	optimized 2	$(\hat{w}, \hat{\gamma})$	(0.98, -2.00)	(0.83, 0.75)	(0.75, 0.75)
		AIC	1979.2104	1416.4752	1841.1650
		BIC	2000.8857	1438.1504	1862.8403

Table 2 Fit of 9 different models over 47 time periods of the experiment. The fit is performed over all 6 experimental groups of the negative feedback treatment (2nd column), and over 6 out of 7 experimental groups in the positive feedback treatment (3rd column), see footnote 11. The last column displays the joint fit for 12 experimental groups. The entries show the values of three fitness criteria MSE, AIC and BIC, and the values of corresponding parameters.

worse and the fundamental models perform very bad. Consider now a forecasting method of averaging the predictions of the two heuristics (line 'mixed' in Table 2). Performance of this model with heterogeneous expectations but fixed, equal weights is better than the average performance of the heuristics in a given environment. However, in every environment and also for all 12 groups, this model still performs considerably worse than the best of the two simple heuristics. Finally, the heuristic switching model performs better than the other models in the positive feedback environment and is worse than the adaptive heuristic in the negative feedback environment. It is also the best performing among these 5 models when the MSEs are averaged over 12 experimental groups. As opposed to the model with constant equal weights, the HSM makes the best out of two homogeneous heuristics.

A comparison of the models on the basis of the mean squared errors is not completely satisfactory, since the HSM has more additional parameters. To address this issue one can relate minimization of the MSE with maximization of a certain likelihood function and then evaluate the fitness of the models using the two criteria commonly used to penalize models for additional parameters, the Akaike Information Criterion and the Bayesian Information Criterion. See Appendix A for details. According to the Akaike Information Criterion better models should have *lower* value of $AIC = 2n + kG \ln(\text{MSE}_{kG}) - 2C$, where $k = 47$ is the number of observations in one group, G is the number of groups in the sample, n is the number of parameters of the model and C is a constant which depends on the sample size but does not

	0%	+10%	-10%	+20%	-20%	+30%	-30%	+40%	-40%
β	1.5051	1.5052	1.5052	1.5053	1.5055	1.5054	1.5061	1.5056	1.5072
η	1.5051	1.5109	1.5158	1.5156	1.5654	1.5163	1.6037	1.5166	1.6350
δ	1.5051	1.5051	1.5051	1.5051	1.5051	1.5051	1.5051	1.5051	1.5051
β, η	1.5051	1.5114	1.5188	1.5160	1.5745	1.5168	1.6190	1.5171	1.6585
β, δ	1.5051	1.5052	1.5052	1.5053	1.5055	1.5054	1.5061	1.5056	1.5072
η, δ	1.5051	1.5109	1.5158	1.5156	1.5654	1.5163	1.6037	1.5166	1.6350
β, η, δ	1.5051	1.5114	1.5188	1.5160	1.5745	1.5168	1.6190	1.5171	1.6585
MSE	1.5051	1.6556		1.8061		1.9566		2.1072	

Table 3 Robustness test of the heuristics switching model. For one, two or three parameters we report the value of MSE (over 47 time periods in 12 experimental groups) of the model resulting in a given percentage change of every of these parameters. The last row shows the corresponding percentage increase of the benchmark value of MSE, 1.5051.

depend on the model. According to the Bayesian Information Criterion better models should have *lower* value of $BIC = n \ln(kG) + kG \ln(kMSE_k) - 2C$.

Table 2 shows the MSEs, AICs and BICs for the adaptive and trend homogeneous models and the heuristic switching model. In each case we do not only show the statistics for the model with the benchmark values of the parameters but also optimize the model by searching for the parameter values which minimize the MSE.¹² The corresponding parameters values are also given. The best (i.e., smallest) statistics are shown in bold. Under the negative feedback the optimized model with trend heuristic is better according to AIC and BIC. Notice, however, that the maximum for the HSM model is obtained when constrained on γ is binding by artificially chosen value of -2 . More importantly, the best model results in a negative value of γ , which is difficult to interpret and would lead to very bad fit under the positive feedback. Under the positive feedback according to all criteria, the best performing model is the HSM when it is optimized *both* with respect to the three learning parameters, β , η and δ and to the two parameters of the heuristics, w and γ . Finally, when all 12 groups are taken into account, all three specifications of the HSM outperform the other considered models according to each fitness measure. Even if penalized for 5 parameters, the fully optimized HSM is better in one-period ahead prediction than other considered specifications.

4.3 Robustness

Finally, we briefly discuss the issue of robustness. A model is said to be *robust* if a change in parameter values does not immediately result in a change of the model's characteristics. A good prediction model should yield good results even if parameters are slightly changed. If, for example, the parameter value for β in the HSM is increased by, say, ten percent, this should not lead to a dramatic change in price dynamics. In the experiments it is unlikely that the agents held stringent parameters from which they do not deviate. The general results of the model should, therefore, hold even if the parameters of the model change slightly.

In order to test the model on robustness we compare the MSE of the best-fitted model with the MSEs of $4 \times 2 \times 8 = 64$ other models. Each of these alternative models is obtained by changing one, two or three learning parameters (β , η and δ) by $\pm 10\%$, $\pm 20\%$, $\pm 30\%$ or $\pm 40\%$ with respect to their estimated values reported in the bottom right cell of Table 2. We do not vary the values of two other parameters $w = 0.75$ and $\gamma = 0.75$. Table 3 shows the results of this exercise obtained for the last 47 time periods in all the experimental groups (except for one group from the positive feedback experiment, see footnote 11). For a comparison, the last row shows the corresponding percentage increase of the MSE of the best-fitted model. We observe that changes of parameters of 10% and 20% do not lead to a significant change in the model performance. Even after a 40% change of the optimal values of the parameters, the resulting models perform better than homogeneous models. We conclude, that the heuristic switching model is fairly robust with respect to the changes of the parameter values.

¹² The search is made by `fmincon` routine in MATLAB, and we controlled the results using the grid search. We implemented the following constraints: $w \in [0, 1]$, $\gamma \in [-2, 2]$, $\beta \geq 0$, $\eta \in [0, 1]$, and $\delta \in [0, 1]$.

5 Conclusion

This paper analyzed the experimental results of markets with negative and positive feedbacks in Heemeijer, Hommes, Sonnemans, and Tuinstra (2009). Estimations of individual forecasting rules in that paper showed that participants relied on simple first-order forecasting heuristics, anchoring their expectations to past prices and past predictions and extrapolating past trends. The estimated individual rules tended to attach positive weight to the trend term in the positive feedback treatments but did not attach weight to this term in the negative feedback treatments. The aggregate dynamics was also very different between the two treatments. The heuristic switching model, simulated in this paper, provides an explanation why this difference in individual forecasting and aggregate price behavior took place.

Several models have been compared with the aggregate price dynamics from the experiment. None of the single heuristic models could provide qualitative features observed in both negative and positive feedback experimental environments. In contrast, the switching model was able to reproduce experimental “stylized facts”. In fact, a parsimonious model where agents switch between the adaptive and the trend heuristics does well in explaining the most important characteristics of the price dynamics observed in the experiment. Consistently with the experiment, in the negative feedback market the simulations showed strong oscillations in the first periods followed by quick convergence towards the equilibrium price. In the positive feedback market the model exhibited persistent deviations from the equilibrium price and slow oscillations. The price outcomes of the model were also dependent on the initial price with more persistent oscillations when the initial state deviated more from the fundamental price, similarly to the outcome of the experiments.

APPENDIX

A Likelihood Method

We use likelihood technique to compare the predictive power of different models. Recall that for the price in the experimental group X is denoted as $p_t^{\text{Gr X}}$ and that $p_t^{\text{Mod}}(\theta)$ stays for a price prediction made by a model Mod with parameters θ . We confine ourselves to the one-step ahead predictions using data from an experimental group. In this case the information set consists of all prices in the experimental group X up to and included period $t - 1$ and the prediction error at time t is $p_t^{\text{Mod}}(\theta) - p_t^{\text{Gr X}}$. Assuming that these errors are independent and normally distributed with mean 0 and variance σ^2 , we can define the likelihood function

$$L(\theta, \sigma^2; \text{exp. data}) = \prod_{\text{Groups}} \prod_{t=3}^{49} \frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{(p_t^{\text{Mod}}(\theta) - p_t^{\text{Gr X}})^2}{2\sigma^2}\right), \quad (11)$$

where for each group we include the last $k = 47$ periods in the sample (see footnote 10). The first product in (11) is over all the groups whose data are used in the estimation. In Table 2 three cases are shown: when the likelihood is computed over 6 groups with negative feedback, over 6 groups with positive feedback, and over all 12 groups.

Consistently with the notation in Section 4.2 the mean squared error in one group is given by

$$\text{MSE}_k = \frac{1}{k} \sum_{t=50-k}^{49} (p_t^{\text{Mod}}(\theta) - p_t^{\text{Gr X}})^2.$$

When we compute (11) over data from G groups, we denote the mean squared error over all groups as MSE_{kG} , which is a mean of squared errors over kG data points.

Maximization of the likelihood function is equivalent to maximization of

$$\begin{aligned} \ln L(\theta, \sigma^2; \text{exp. data}) &= \sum_{\text{Groups}} \sum_{t=t_0}^T \left(-\frac{(p_t^{\text{Mod}}(\theta) - p_t^{\text{Gr X}})^2}{2\sigma^2} - \ln(\sigma\sqrt{2\pi}) \right) = \\ &= -\frac{kG}{2\sigma^2} \text{MSE}_{kG}(\theta) - kG \ln(\sigma\sqrt{2\pi}). \end{aligned} \quad (12)$$

Denote by $\hat{\theta}$, $\hat{\sigma}$ the parameter values under which the log-likelihood is maximized. They satisfy the set of first order conditions (f.o.c.). The f.o.c. with respect to θ do not depend on σ and are identical to the f.o.c. for minimization of the MSE. This justifies that the parameters of the learning model found by minimization of the MSE will also maximize the likelihood. The f.o.c. with respect to σ is given by

$$\frac{kG}{\hat{\sigma}^3} \text{MSE}_{kG}(\hat{\theta}) - \frac{kG}{\hat{\sigma}} = 0,$$

and, therefore, $\hat{\sigma}^2 = \text{MSE}_{kG}(\hat{\theta})$. Plugging this result in (12) we find that

$$\begin{aligned} \max_{\hat{\theta}, \hat{\sigma}^2} [\ln L(\hat{\theta}, \hat{\sigma}^2; \text{exp. data})] &= \ln L(\hat{\theta}, \hat{\sigma}^2; \text{exp. data}) = \\ &= -\frac{kG}{2} - \frac{kG}{2} (\ln \text{MSE}_{kG}(\hat{\theta}) + \ln(2\pi)) = \\ &= -\frac{kG}{2} \ln \text{MSE}_{kG}(\hat{\theta}) + C, \end{aligned} \quad (13)$$

where the constant $C = -kG(1 + \ln(2\pi))/2$ depends on the number of observations but not on the model.

For a model with n parameters the Akaike Information Criterion is defined as

$$\text{AIC} = 2n - 2 \ln L(\hat{\theta}, \hat{\sigma}^2; \text{exp. data}),$$

while the Bayesian Information Criterion is defined as

$$\text{BIC} = n \ln(kG) - 2 \ln L(\hat{\theta}, \hat{\sigma}^2; \text{exp. data}).$$

Plugging the log-likelihood found in (13) into the last two formulas, we find that

$$\text{AIC} = 2n + kG \ln \text{MSE}_{kG}(\hat{\theta}) - 2C, \quad (14)$$

and

$$\text{BIC} = n \ln(kG) + kG \ln \text{MSE}_{kG}(\hat{\theta}) - 2C. \quad (15)$$

Acknowledgements We are grateful to the participants of the workshop “Evolution and market behavior in economics and finance” in Pisa, Italy (October 2009) for fruitful discussion, and to the three anonymous referees for useful suggestions. This work was supported by the EU 7th framework collaborative project “*Monetary, Fiscal and Structural Policies with Heterogeneous Agents (POLHIA)*”, grant no. 225408. The usual caveats apply.

References

- ADAM, K. (2007): “Experimental Evidence on the Persistence of Output and Inflation,” *The Economic Journal*, 117, 520.
- AKERLOF, G., AND R. SHILLER (2009): *Animal spirits: how human psychology drives the economy, and why it matters for global capitalism*. Princeton Univ Press.
- ANUFRIEV, M., J. ARIFOVIC, J. LEDYARD, AND V. PANCHENKO (2011): “Efficiency of Continuous Double Auctions under Individual Evolutionary Learning with Full and Limited Information,” *Journal of Evolutionary Economics*, forthcoming.
- ANUFRIEV, M., T. BAO, AND J. TUINSTRAN (2011): “Switching behavior in the lab,” mimeo, University of Amsterdam.
- ANUFRIEV, M., AND C. HOMMES (2009): “Evolutionary Selection of Individual Expectations and Aggregate Outcomes,” CeNDEF Working Paper 2009-09, University of Amsterdam.
- (2012): “Evolution of market heuristics,” *Knowledge Engineering Review*, forthcoming.
- ARIFOVIC, J. (1996): “The behavior of the exchange rate in the genetic algorithm and experimental economics,” *Journal of Political Economy*, pp. 510–541.
- ARIFOVIC, J., AND J. LEDYARD (2007): “Call market book information and efficiency,” *Journal of Economic Dynamics and Control*, 31(6), 1971–2000.
- BAO, T., C. HOMMES, J. SONNEMANS, AND J. TUINSTRAN (2010): “Individual expectation, limited rationality and aggregate outcome,” CeNDEF Working Paper 2010-07, University of Amsterdam.
- BRANCH, W. (2004): “The Theory of Rationally Heterogeneous Expectations: Evidence from Survey Data on Inflation Expectations,” *The Economic Journal*, 114(497), 592–621.
- BROCK, W. A., AND C. H. HOMMES (1997): “A Rational Route to Randomness,” *Econometrica*, 65(5), 1059–1095.
- COLUCCI, D., AND V. VALORI (2006): “Ways of learning in a simple economic setting: a comparison,” *Chaos, Solitons & Fractals*, 29(3), 653–670.
- CONLISK, J. (1996): “Bounded rationality and market fluctuations,” *Journal of Economic Behavior and Organization*, 29(2), 233–250.
- DAVIDOVITS, P. (2008): *Physics in biology and medicine*. Academic press.
- DAWKINS, R. (1976): *The selfish gene*. Oxford University Press.
- 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. (2006): “Agent-Based Models and Human Subject Experiments,” 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).
- EVANS, G. W., AND S. HONKAPOHJA (2001): *Learning and Expectations in Macroeconomics*. Princeton University Press.
- FEHR, E., AND J. TYRAN (2008): “Limited rationality and strategic interaction: the impact of the strategic environment on nominal inertia,” *Econometrica*, 76(2), 353–394.
- HALTIWANGER, J., AND M. WALDMAN (1989): “Limited rationality and strategic complements: the implications for macroeconomics,” *The Quarterly Journal of Economics*, 104(3), 463–483.

- HEEMELJER, P., C. HOMMES, J. SONNEMANS, AND J. TUINSTRRA (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.
- HEY, J. D. (1994): "Expectation formation: Rational or adaptive or ...?," *Journal of Economic Behavior and Organization*, 25, 329–349.
- HOMMES, C. (2006): "Heterogeneous Agent Models in Economics and Finance," in *Handbook of Computational Economics Vol. 2: Agent-Based Computational Economics*, ed. by K. Judd, and L. Tesfatsion. Elsevier/North-Holland (Handbooks in Economics Series).
- HOMMES, C. (2011): "The Heterogeneous Expectations Hypothesis: Some Evidence from the Lab," *Journal of Economic Dynamics and Control*, 35(1), 1–24.
- 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 (2011): "Individual Expectations and Aggregate Behavior in Learning to Forecast Experiments," *Macroeconomic Dynamics*, forthcoming.
- HOMMES, C., J. SONNEMANS, J. TUINSTRRA, AND H. VAN DE VELDEN (2005): "Coordination of Expectations in Asset Pricing Experiments," *Review of Financial Studies*, 18(3), 955–980.
- HOMMES, C., J. SONNEMANS, J. TUINSTRRA, AND H. VAN DE VELDEN (2007): "Learning in Cobweb Experiments," *Macroeconomic Dynamics*, 11(S1), 8–33.
- HOMMES, C., J. SONNEMANS, J. TUINSTRRA, AND H. VAN DE VELDEN (2008): "Expectations and bubbles in asset pricing experiments," *Journal of Economic Behavior and Organization*, 67(1), 116–133.
- KAHNEMAN, D. (2003): "Maps of Bounded Rationality: Psychology for Behavioral Economics," *American Economic Review*, 93, 1449–1475.
- KIRMAN, A. (2011): *Complex Economics: Individual and Collective Rationality*. Routledge, Abingdon.
- MAILATH, G. (1998): "Do people play Nash equilibrium? Lessons from evolutionary game theory," *Journal of Economic Literature*, pp. 1347–1374.
- MARCET, A., AND T. SARGENT (1989): "Convergence of least squares learning mechanisms in self-referential linear stochastic models," *Journal of Economic Theory*, 48(2), 337–368.
- 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.
- SONNEMANS, J., AND J. TUINSTRRA (2010): "Positive expectations feedback experiments and number guessing games as models of financial markets," *Journal of Economic Psychology*, 31, 964–984.
- SUTAN, A., AND M. WILLINGER (2009): "Guessing with negative feedback: An experiment," *Journal of Economic Dynamics and Control*, 33(5), 1123 – 1133.