



Munich Personal RePEc Archive

Heuristic Switching Model and Exploration-Exploitation Algorithm to describe long-run expectations in LtFEs: a comparison

Colasante, Annarita and Alfarano, Simone and
Camacho-Cuena, Eva

Universitat Jaume I

2019

Online at <https://mpra.ub.uni-muenchen.de/92391/>
MPRA Paper No. 92391, posted 26 Feb 2019 13:09 UTC

Heuristic Switching Model and Exploration-Exploitation Algorithm to describe long-run expectations in LtFEs: a comparison

Annarita Colasante, Simone Alfarano, Eva Camacho-Cuena

Universitat Jaume I

Abstract

We compare the performance of two learning algorithms in replicating individual short and long-run expectations: the Exploration-Exploitation Algorithm (EEA) and the Heuristic Switching Model (HSM). Individual expectations are elicited in a series of Learning-to-Forecast Experiments (LtFEs) with different feedback mechanisms between expectations and market price: positive and negative feedback markets. We implement the EEA proposed by Colasante et al. (2018c). Moreover, we modify the existing version of the HSM in order to incorporate the long-run predictions. Although the two algorithms provide a fairly good description of market prices in the short-run, the EEA outperforms the HSM in replicating the main characteristics of individual expectation in the long-run, both in terms of coordination of individual expectations and convergence of expectations to the fundamental value.

JEL: D03,G12,C91.

Keywords: Expectations, Experiment, Evolutionary Learning.

1 Introduction

The origin of heterogeneity across individual expectations and the role that it plays in shaping aggregate outcomes is an important topic in theoretical as well as empirical research in macroeconomics. Expectations, however, are not directly observable, differently from prices of stocks, volumes of sold books, interest rates of bonds, downloads or number of likes, which can be precisely measured and recorded in the new world where almost “everything” is now in electronic format. This means that there is a significant limitation when it comes to fully understanding the precise role played by expectations in driving macroeconomic aggregates. One way to circumvent this problem and rigorously model the expectations of individuals is to assume consistent expectations, i.e. rational expectations, following the seminal idea of Muth (1961) which has been further developed by Lucas Jr and Prescott (1971). From a formal point of view, the main advantage of rational expectations is that agents can be rational in one way only. The argument put forward by Friedman (1966) on the irrelevance of “irrational” individuals in the long-run gives to the rationality assumption a further intuitive appeal.

As an alternative to the rational expectations paradigm, it certainly plays a central role the bounded rationality assumption of economic agents, introduced by Simon (1957). In the world of bounded rational agents, we typically lose the uniqueness of the behavior, since agents can be “non-rational” in many different ways. In order to shed some light on the degree of bounded rationality of individuals, laboratory experiments have been largely demonstrated to be an essential methodology. Countless experiments have repetitively shown that in complex environments, subjects follow simple adaptive rules, called heuristics, in order to form expectations, changing their mind as a function of the evolution of the environment and adapting to the new circumstances. The principle of anchor and adjustment, introduced by Kahneman and Tversky (1973), is a sufficiently general and flexible framework that can be certainly cast into the bounded rationality paradigm and able to realistically describe the way individuals form and adapt their expectations.

Laboratory experiments is one of the methodologies that allows us to directly elicit individual expectations using performance-based incentives. In particular, within the experimental literature, Learning-to-Forecat Experiments, in the following LtFEs, (see Marimon et al. (1993)), are designed to study the formation of individual expectations within different expectations feedback systems in a market where the price depends on subjects’ (short-run) predictions. Such experimental framework allows to study under which conditions individual predictions converge to the rational expectations equilibrium. Moreover, it allows to efficiently test alternative formulations of expectation formation models. A large number of LtFEs have been conducted to analyse the way individuals form and adapt their

short-run expectations in different economic environments: in financial markets (Hommes et al. (2005)), real estate markets (Bao and Ding (2016)), commodity markets (Bao et al. (2013)) and in simple macroeconomic frameworks (Assenza et al. (2011); Anufriev et al. (2013); Cornand and M'baye (2016)). The vast majority of those LtFEs focuses on eliciting individual *short-run expectations*, i.e. providing incentives to forecast the next period market price. The novel contribution to this experimental literature of Colasante et al. (2018a,b) is to focus on the entire spectrum of expectations, since they elicit contemporaneously short and *long-run expectations* about the evolution of the market price under positive and negative feedback systems. In particular, they explicitly give incentives to the subjects to submit their predictions on the market price at the beginning of every period, giving the possibility to revise the predictions as new information becomes available. This experimental design allows to study how expectations form and co-evolve with the price at different forecast horizons. Their results concerning short-run predictions are in line with the literature (see Heemeijer et al. (2009)): fast convergence and slow coordination in the markets with negative feedback, whereas slow convergence and fast coordination in the markets with positive feedback. Regarding the long-run predictions, Colasante et al. (2018a,b) observe that in markets with positive feedback treatments the market price plays a pivotal role in the expectations formation process, whereas in markets with negative feedback it turns out that subjects use the fundamental value as their main reference point.

We can find in the literature only few computational learning algorithms to describe individual expectations in LtFEs (see Heemeijer et al. (2009); Assenza et al. (2011); Bao et al. (2013); Hommes and Lux (2013)). The most commonly used is the so-called Heuristic Switching Model, HSM hereafter (see Brock and Hommes (1998)). Using experimental data on long-run expectations, Colasante et al. (2018c) have introduced an alternative adaptive learning model of bounded rationality, the Exploration-Exploitation Algorithm (hereafter EEA) that, contrary to the HSM, accounts contemporaneously for subjects' short- and long-run expectations.

The aim of this paper is to evaluate the performance of the EEA and the HSM to reproduce the long-run expectations in markets with positive and negative feedback. Since the original version of the HSM can account for short-run expectations only, we introduce a modified version of the HSM in order to capture the short- as well as long-run expectations. We take the position that comparing the capability of different algorithms in describing the dynamical properties of the formation mechanism of individual expectations is a valuable contribution. It helps, in fact, by confronting the basic constituents of the algorithms, to understand which are the crucial determinants in designing reliable models of expectations formation. Furthermore, it allows to devise new models by possibly combining the crucial

characteristics of the analysed algorithms in an more efficient architecture, to describe the data and model expectations.

The paper is organized as follows: in the Section 2, we illustrate the experimental setting of the LtFE and the experimental results. In Section 3, we describe the details of the two learning algorithms, namely the HSM and its modified version and the EEA. In Sections 4 and 5, we describe the simulation results and propose a modified version of the EEA to include a different anchor, respectively. Finally, in Section 6 we present the main conclusions.

2 The Learning to Forecast Experiment

2.1 Experimental Design

This paper builds upon the Learning to Forecast Experiments of Colasante et al. (2018a,b). In this section we briefly describe their novel experimental design to elicit subjects' expectations at different time horizons. In those LtFEs the authors elicit subjects' short and long-run expectations in markets characterized by either a positive or a negative expectations feedback system. They conduct a total of 15 sessions: 7 with positive and 8 with negative feedback. In each session, 6 subjects play the role of professional forecasters for 20 periods. More precisely, at the beginning of period t , subject i submits her short-run prediction for the price at the end of period t , denoted as $i p_{t,t}^e$, as well as her long-run predictions for the price at the end of each one of the $20 - t$ remaining periods, denoted as $i p_{t,t+k}^e$, with $1 \leq k \leq 20 - t$. To compute the market price, it is implemented the pricing equation proposed by Heemeijer et al. (2009). In the positive feedback treatment, the law of motion of the price is given by:

$$p_t = p_f + \frac{1}{1+r} (\bar{p}_{t,t}^e - p_f) + \epsilon_t, \quad (1)$$

while in the negative feedback treatment the realized market price is computed as follows:

$$p_t = p_f - \frac{1}{1+r} (\bar{p}_{t,t}^e - p_f) + \epsilon_t, \quad (2)$$

where $r=0.05$ in all sessions and p_f is the constant fundamental value in a given session.¹ The term $\bar{p}_{t,t}^e$ in the equations is the average of the six one-step-ahead

¹To avoid the effects of communication among subjects between sessions, two different values for p_f are implemented, so that there are some markets with a fundamental value equal to 65 and others with a fundamental value equal to 70. In the positive feedback treatment, the fundamental value is computed as $p_f = \frac{d}{r}$, where the average dividend d is equal to 3.5 or 3.25, depending on the session.

predictions submitted at the beginning of period t , $\bar{p}_{t,t}^e = \frac{1}{6} \sum_{i=1}^6 i p_{t,t}^e$. Finally, the term $\epsilon_t \sim N(0, 0.25)$ is an iid normal shock.

The main difference between eqs. (1) and (2) is how expectations affect the price: eq. (1) describes a positive feedback system where subjects predictions are self-fulfilling, i.e. the higher (lower) is the average forecast, the higher (lower) will be the price. Eq. (2) describes, instead, a system in which there is a negative feedback between expectations and price, i.e. the higher (lower) is the average forecast, the lower (higher) will be the price. Even if market price is solely determined by short-run predictions, all predictions are rewarded. Individual earnings at the end of each period depend on forecasting errors and are computed as ${}_i\pi_t = {}_i\pi_t^s + {}_i\pi_t^l$, where ${}_i\pi_t^s$ denotes the pay-off for the short-run predictions:

$${}_i\pi_t^s = \frac{250}{1 + \beta_{i,t}} \quad \text{with} \quad \beta_{i,t} = \left(\frac{{}_i p_{t,t}^e - p_t}{2} \right)^2. \quad (3)$$

and ${}_i\pi_t^l$ denotes the subject's pay-off that depends on her long-run predictions error, being ${}_i\pi_t^l = \sum_{j=1}^{t-1} {}_i\pi_{t-j,t}^l$, where ${}_i\pi_{t-j,t}^l$ represents the individual profit associated with the prediction submitted by subject i at the beginning of period $t-j$ for the price in period t , where $1 \leq j \leq t-1$. The long-run predictions are rewarded according to the following scheme:

$${}_i\pi_{t-j,t}^l = \begin{cases} 25 & \text{if } 0 \leq {}_i\delta_{t-j,t} \leq 5 \\ 12 & \text{if } 5 < {}_i\delta_{t-j,t} \leq 10 \\ 5 & \text{if } 10 < {}_i\delta_{t-j,t} \leq 15 \\ 0 & \text{otherwise} \end{cases} \quad (4)$$

where ${}_i\delta_{t-j,t} = |{}_i p_{t-j,t}^e - p_t|$. Note that subjects receive an immediate feedback on their short-run forecasting accuracy, while they experience a delay in evaluating the accuracy of their long-run predictions. The final payment of each subject is the sum of pay-offs across all periods.² In the positive feedback treatment, subjects are informed about the value of the constant interest rate (r), average dividend (d), the asset prices until period $t-1$ and all their own (short and long-run) past predictions. In the negative feedback, subjects receive only information about their own past predictions and the past market prices. As it is typically done in the literature of LtFEs, the subjects receive some qualitative information on the feedback system affecting the price mechanism and the profit functions. See Colasante et al. (2018a,b) for additional details about the experimental design.

²The parameters of the pay-off functions are calibrated such that approximately $\max \sum_{t=1}^{20} {}_i\pi_t^s = \max \sum_{t=1}^{20} {}_i\pi_t^l$, in order to give to the subjects the same incentive to provide accurate predictions in the short as well as in the long-run.

According to eqs. (1) and (2), the REE predicts that the price p_t converges to the fundamental value with fairly small fluctuations proportional to the idiosyncratic shock term ϵ_t . What would be the REE for long-run expectations? When eliciting their long-run expectations, subjects submit at the beginning of period t their predictions for the end of period $t+k$, for all $k > 0$. The price at the end of period $t+k$ is a function of the subjects' short-run predictions submitted at the beginning of period $t+k$. Each subject has, indeed, to guesstimate k -periods in advance, the short-run predictions of the other subjects. If we assume that all subjects follow rational expectations, then their predictions in each period t and for each forecast horizon k are $i p_{t,t+k}^e \approx p_f$, independently of the expectations feedback system.

2.2 Experimental results

In the following we summarize the experimental results in Colasante et al. (2018a). Figure 1 displays the dynamics of the market price in all the 15 markets. As an illustrative example, Figures 2 and 3 show the individual short-run predictions and the market price dynamics in two representative group for each treatment. From a visual inspection of those figures, we observe that market prices follow qualitatively different patterns, depending on the feedback mechanism implemented in the particular treatment. In the positive feedback treatment, short-run predictions coordinate after few periods, although not necessary to the fundamental value. The market price exhibits an oscillatory pattern without converging to the fundamental value. In the negative feedback treatment, instead, we observe that market prices quickly converge to the fundamental value after few periods of uneven fluctuations, while individual short-run predictions need more periods to coordinate on the fundamental value. Figures 4 and 5 show individual long-run predictions in two representative groups, one for the positive and one for the negative feedback treatment, respectively.

To quantify the convergence of the expectations, the Mean Absolute Deviation (MAD) between individual predictions and the fundamental value is computed for both, short and long-run predictions as:

$$MAD_{t,t+k}^{p_f} = \left\langle \frac{\sum_{i=1}^6 |i p_{t,t+k}^e - p_f|}{6} \right\rangle_g, \quad (5)$$

with $k=0,1,2,4,6,9$. The notation $\langle \dots \rangle_g$ denotes the average across groups. To quantify the degree of coordination, the MAD between individual predictions and the (within-group) average predictions is computed for each period t and for a given forecast horizon k , with $k = 0, 1, 2, 4, 6, 9$:

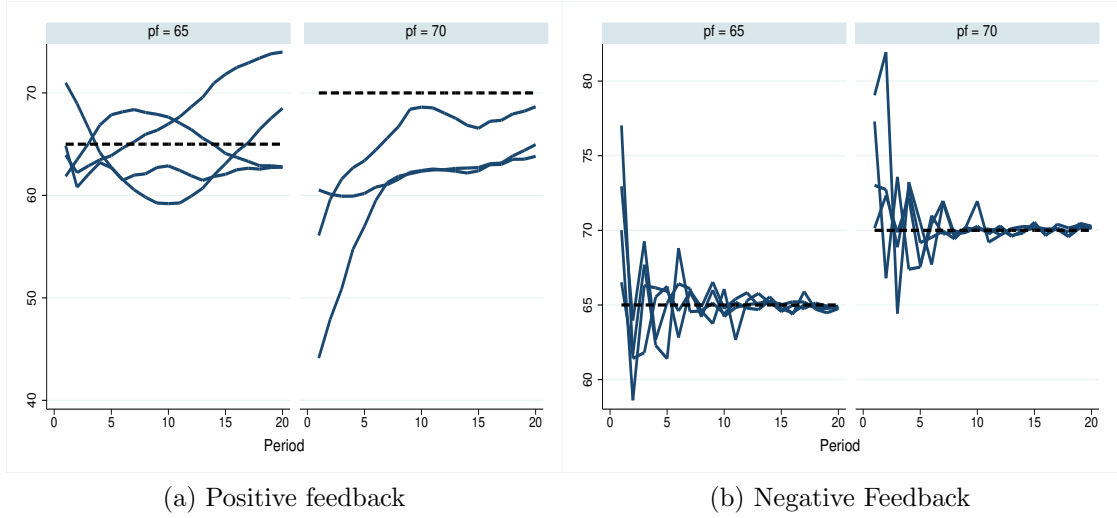


Figure 1: Realized price for all groups in the two treatments. The blue solid lines are market prices, the dashed lines represent the fundamental value.




Figure 2: Realized price and individual short-run predictions of two representative groups in the positive feedback system. The black solid line is the market price, the grey lines are the individual one-step-ahead predictions and the dashed line represents the fundamental value.

$$MAD_{t,t+k}^C = \left\langle \frac{\sum_{i=1}^6 |iP_{t,t+k}^e - \bar{P}_{t,t+k}^e|}{6} \right\rangle_g, \quad (6)$$

Figure 6 summarizes the main results in terms of both convergence and coordination. A comparison between treatments in terms of convergence leads to a well-known conclusion in line with the LtFE literature (see Heemeijer et al. (2009)): the predictions for different forecasting horizons converge to the fundamental value




Figure 3: Realized price and individual short-run predictions of two representative groups in the negative feedback system. The black solid line is the market price, the grey lines are the individual one-step-ahead predictions and the dashed line represents the fundamental value.

only in the negative feedback treatment, while, in the positive feedback treatment, predictions systematically deviate from the fundamental value. Focusing on coordination of expectations, in the positive feedback treatment, subjects' one step-ahead predictions coordinate faster than in the negative feedback treatment, whereas in the long-run predictions, the forecasts disagreement increases with the horizon, i.e. the longer is the forecast horizon, the higher is the dispersion of predictions. This is strongly connected to the absence of a long-term anchor for subjects' expectations. In fact, the long-run predictions are characterized by some sort of cone-shape trajectory, compatible with a linear trend extrapolation rule of subjects with heterogeneous slopes. In the negative feedback treatment, where subjects learn the fundamental value, coordination of short and long-run predictions occur simultaneously, driven by the converge of the market price to the fundamental value. Subjects are able to learn the REE and, as a consequence, both price and expectations converge to it over time. It is just in the negative feedback system that the REE constitutes a good benchmark to describe the market price dynamics and the evolution of subjects' expectations.

3 Learning Algorithms: HSM and EEA

In this section we describe the details of the two learning algorithms that we employ to reproduce the experimental data.

The HSM is a well-established learning algorithm implemented to explain subjects' behavior in many LtFEs, see Assenza et al. (2014) and Hommes (2018). According to this model, subjects forecast future prices by selecting a prediction

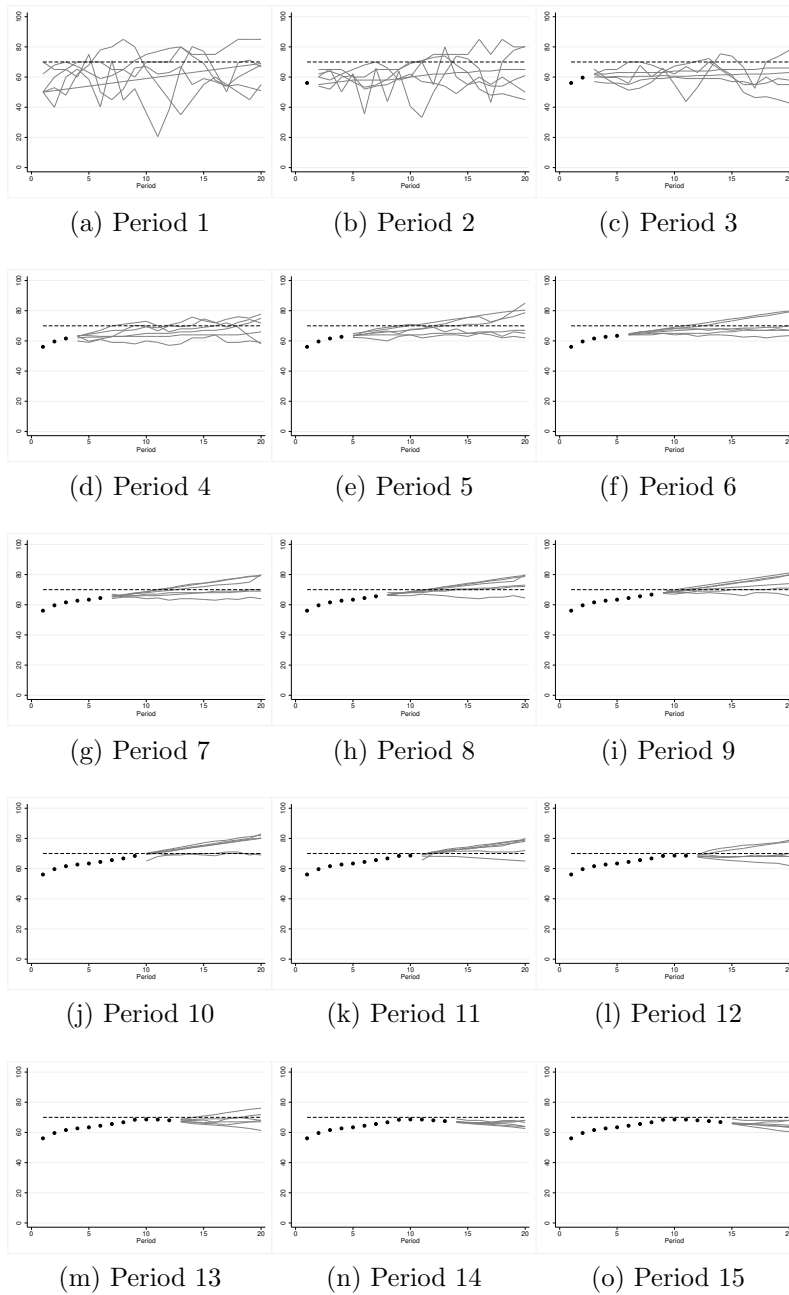


Figure 4: Realized prices (dots), the asset fundamental value (broken line) and all individual predictions (solid lines), submitted from period 1 to 15 of Group 1 in the positive feedback treatment, as a function of the time horizon.

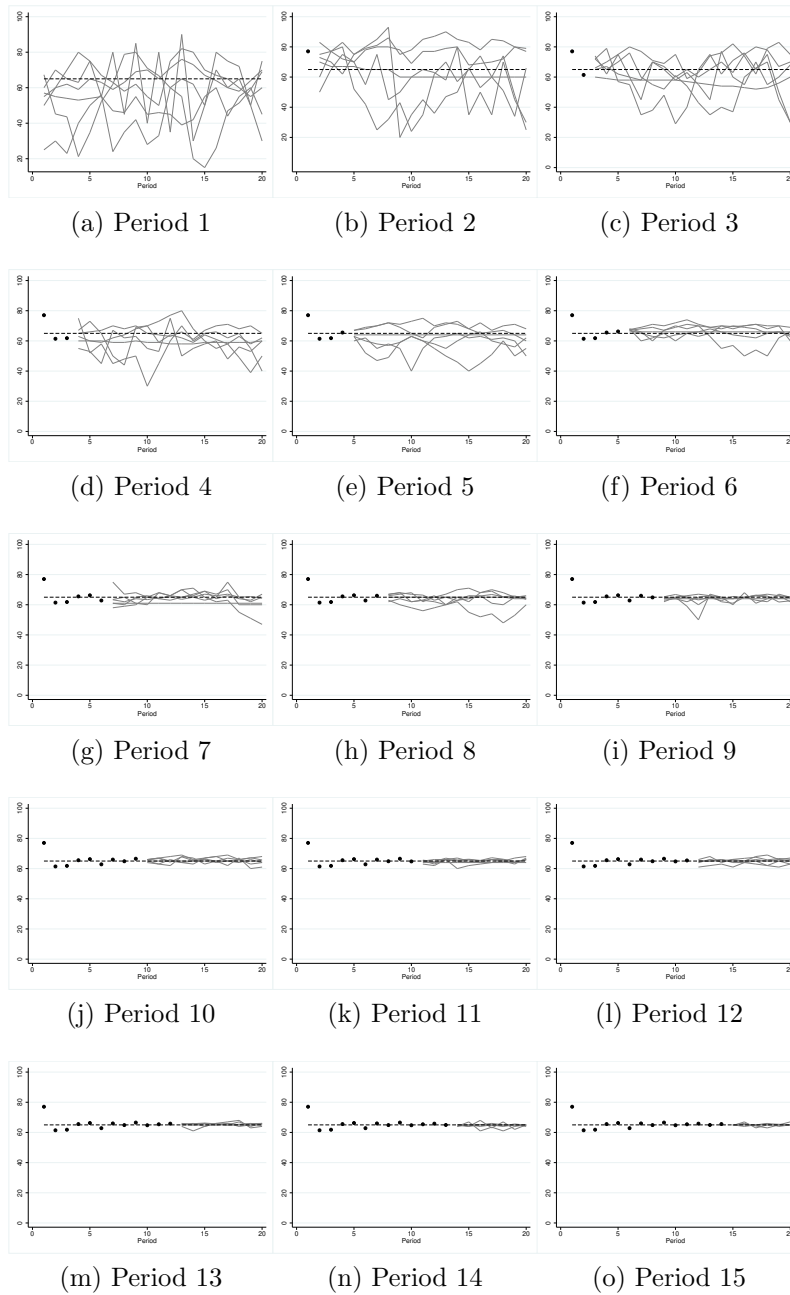

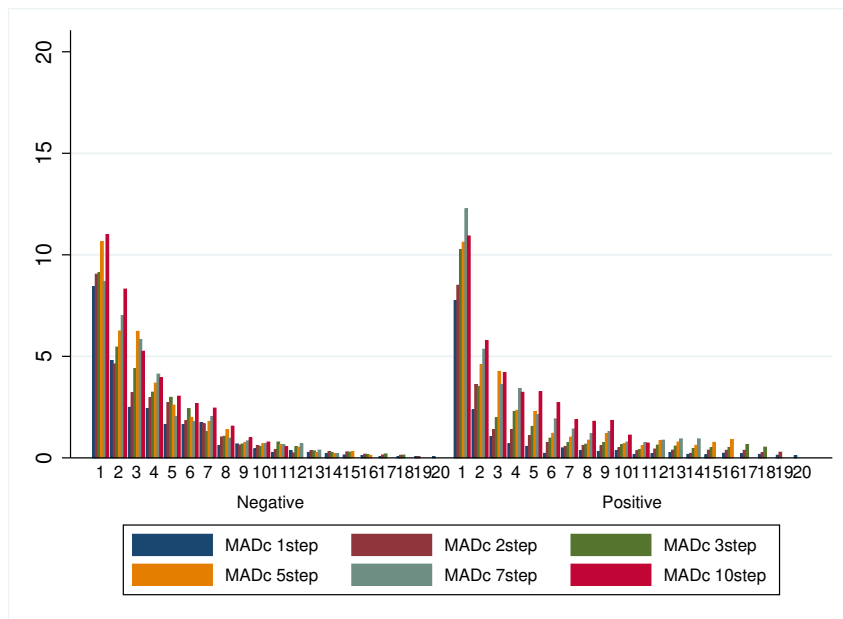


Figure 5: Realized prices (dots), the asset fundamental value (broken line) and all individual predictions (solid lines), submitted from period 1 to 15 of Group 1 in the negative feedback treatment, as a function of the time horizon.



(a) Convergence



(b) Coordination

Figure 6: Convergence and coordination of short and long-run expectations in the positive and negative feedback treatments.

rule that best performed in the recent past among a given set of simple heuristics. The process of learning is based on the feedback they receive from their profits,

which constitutes a fitness measure to rank those rules. In every period, each subject selects the rules with a probability proportional to the fitness measure, so that the best performing rule has higher chances to be selected. Similarly to the HSM, the EEA is based on the very basic principle of anchor-and-adjustment, (Kahneman and Tversky (1973)). Differently from the HSM which can be thought as a “parametric” learning algorithm, the EEA allows a higher degree of flexibility in the range of possible prediction rules, see Colasante et al. (2018c). Its key learning mechanism is the identification of the range of feasible actions around the anchor, represented by the market price in the previous period. As long as more information on the price dynamics is available, each subject adjusts her own range of possible predictions and, then, selects her specific prediction proportionally to the fitness measure. The process of selection consists of two main phases: the exploration phase, in which subjects have few information and try to form their expectations about the future evolution of the market price; the exploitation phase, in which subjects refine their predictions in order to locally “optimize” their performance. Note that in the EEA we are not imposing any precise parametrization of the subjects’ predictions, which are free to vary in a range whose evolution adapts to the past market conditions.

3.0.1 The Heuristic Switching Model

In the following, we list the four heuristics of the HSM as introduced by Bao et al. (2012). We label the four rules according to the index $h = 1, \dots, 4$, indicating the corresponding forecasting price as ${}_h p_{t,t}^e$. Note that the left sub-index h denotes now the heuristic instead of the subject. The heuristics are³:

- Adaptive rule (ADA):

$${}_1 p_{t,t}^e = \alpha p_{t-1} + (1 - \alpha) p_{t-1,t-1}^e \quad 0 < \alpha < 1 . \quad (7)$$

- Trend following rule (TFR):

$${}_2 p_{t,t}^e = p_{t-1} + w(p_{t-1} - p_{t-2}) \quad w > 0 . \quad (8)$$

- Contrarian rule (CR):

$${}_3 p_{t,t}^e = p_{t-1} + s(p_{t-1} - p_{t-2}) \quad s < 0 . \quad (9)$$

- Learning and adjustment rule (LAA):

$${}_4 p_{t,t}^e = 0.5(p_{t-1}^{av} + p_{t-1}) + (p_{t-1} - p_{t-2})$$

³For a detailed description of the rules, their meaning and performances we refer the reader to the pertinent literature.

The learning mechanism is based on the possibility for the subjects to switch among the four given rules, choosing the one providing the relatively highest profitability in the recent past. The following equation is employed to compute the performance of each heuristic:

$$U_{h,t-1} = - \left(p_{t-1} - {}_h p_{t-1,t-1}^e \right)^2 + \eta U_{h,t-2}, \quad h = 1, \dots, 4,$$

where the parameter $0 \leq \eta \leq 1$ represents the “memory” of agents, meaning the weight assigned to past errors. We set $\eta = 0.7$, following Hommes (2013).

It is assumed that only a fraction of agents changes their rule every period. The share of agents using a specific rule is computed by using the discrete choice model with asynchronous updating as in Diks and Van Der Weide (2005):

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

$$Z_{t-1} = \sum_{h=1}^4 \exp(\beta U_{h,t-1}), \quad (11)$$

where $0 < \delta \leq 1$ denotes the share of agents that update their choice; the parameter $\beta \geq 0$ represents the intensity of choice and it determines the switching speed to the most successful rule. Z_{t-1} is a normalization factor. As in Hommes (2013), we consider $\delta = 0.9$ and $\beta = 0.4$. We compute the expected price as a weighted average across the different expectations given by the four rules:

$$\bar{p}_{t,t}^e = \sum_{h=1}^4 n_{h,t-1} {}_h p_{t,t}^e.$$

Plugging the value of $\bar{p}_{t,t}^e$ into eqs. (1) and (2), we compute the market price, in the positive and negative feedback treatment, respectively.

We generalize the HSM in order to account for the long-run expectations. We simply assume that the subjects linearly extrapolate their short-run prediction rules to determine their long-run expectations as follows:

- Long-run adaptive rule (L-ADA):

$${}_1 p_{t,t+k}^e = p_{t-1} + (k + 1) \alpha (p_t - p_{t-1,t-1}^e) \quad k = 1, 2, 3. \quad (12)$$

- Long-run trend following rule (L-TFR):

$${}_2 p_{t,t}^e = p_{t-1} + (k + 1) w (p_{t-1} - p_{t-2}) \quad k = 1, 2, 3. \quad (13)$$

- Long-run contrarian rule (L-CR):

$${}_3p_{t,t}^e = p_{t-1} + (k + 1) s(p_{t-1} - p_{t-2}) \quad k = 1, 2, 3 .$$

- Long-run learning and adjustment rule (L-LAA):

$${}_4p_{t,t}^e = 0.5(p_{t-1}^{av} + p_{t-1}) + (k + 1)(p_{t-1} - p_{t-2}) \quad k = 1, 2, 3 .$$

We assume that, once an agent selects one of the short-run prediction rules, she extrapolates linearly that rule in order to predict prices in horizons up to 4-steps ahead. In other words, we assume that the selection of the rule at time t is based on the feedback from the profits associated to the short-run predictions only. Our simple modification of the HSM to account for long-run predictions is the first step in this direction present in the literature on LtFEs. It is essentially based a simple intuition and in its easy implementation.

3.0.2 The EEA: the Exploration-Exploitation Algorithm

The Exploration-Exploitation Algorithm, as outlined in Colasante et al. (2018c), is a non parametric learning algorithm implemented to characterize the price dynamics of both, short- and long-run expectation in LtFEs. We assume that each subject, represented by an artificial agent, has a set of available actions in every period. To compute short-run predictions, the range of action of the artificial agents evolves adaptively as a function of the past price and expectations. In particular, the range of the set of actions is centred in the last realized price and its range, which represents the exploration space, is proportional to the standard deviation of the past market prices. The probability to take a particular action is proportional to a given distribution, with mean computed as a convex combination between the last observed price and its past prediction, and whose range of variability is proportional to the standard deviation of the market price. During the experiment, the range of variability of subjects' predictions diminishes over time, from a very high level of dispersion in the first few periods, to a very narrow band at the end of the experiment, see Figure 6b. Within the EEA, the behavior of the subjects can be cast into two distinct phases: (i) the exploration phase, i.e. the tendency to explore the range of actions to acquire knowledge on their environment and (ii) the exploitation phase, i.e. learning to coordinate their actions using the acquired information on the price generating process in order to gain higher profits. In the second phase the range of actions significantly reduces around the last market price.

Let us formalize the EEA algorithm. All agents are characterized by a set of $n = 101$ feasible actions, $A_t = \{a_{1t}, a_{2t}, \dots, a_{nt}\}$, where $a_{1t} < a_{2t} < \dots < a_{nt}$

and a_{jt} denotes the single element in the set. The set \mathbf{A}_t changes every period depending on the last realized price and the magnitude of the fluctuations of past prices. More precisely, the range of the set of actions lies in the interval $(p_{t-1} - 5\Lambda_{t-1}, p_{t-1} + 5\Lambda_{t-1})$, and it is common to all agents, where Λ_{t-1} is the standard deviation of the last three market prices. In each period t , agent i selects an action ${}_i\tilde{a}_t \equiv {}_i p_{t,t}^e$ from \mathbf{A}_t , that corresponds to its prediction of the price for the end of period t , i.e. its one-step-ahead prediction.

Agent i chooses three additional actions ${}_i\tilde{a}_t^k \equiv {}_i p_{t,t+k}^e$ from the corresponding set of actions $\mathbf{A}_t^k = \{a_{1t}^k, a_{2t}^k, \dots, a_{nt}^k\}$, where $k \in \{1, 2, 3\}$. They represent the spectrum of agent i 's long-run expectations up to four-steps-ahead. The interval of variability of the range of \mathbf{A}_t^k is centred in p_{t-1} , similarly to the case of short-run predictions, with a constant width independent of the evolution of the price. The maximum range for the agents' long-run predictions is 15, similar to the maximum deviation of a long-run prediction to be rewarded with a positive profit in the experiment, see eq. (4). The range of the actions that belongs to \mathbf{A}_t^k is therefore $(p_{t-1} - 15, p_{t-1} + 15)$.⁴ Once all agents choose their actions, the price is computed according to eqs. (1) and (2).

Agents, then, evaluate the performance of all feasible actions using a fitness function that accounts for the last realized price as an anchor and their last predictions. We introduce two different measures: \mathbf{V}_t to evaluate the individual actions in the set \mathbf{A}_t (short-run predictions); \mathbf{V}_t^k to evaluate the individual actions in the sets \mathbf{A}_t^k (long-run predictions). The value of the fitness measures are:

$${}_i\mathbf{V}_{j,t} = (p_{t-1} - {}_i a_{jt})^2 + \phi_s ({}_i p_{t-1,t-1}^e - {}_i a_{jt})^2, \quad (14)$$

$${}_i\mathbf{V}_{j,t}^k = |p_{t-1} - {}_i a_{jt}^k| + \phi_l |{}_i p_{t-k-1,t-1}^e - {}_i a_{jt}^k|. \quad (15)$$

We can interpret these fitness measures as a sort of adaptive adjustment of the expectations. The parameters ϕ_s and ϕ_l constitute the relative weights assigned to the past prediction with respect to the last realized price. Note that, in the fitness functions, we reproduce the structure of the payoff function of eqs. (3) and (4), a quadratic term to evaluate short-run predictions, and a term proportional to the absolute distance to evaluate the subject's long-run predictions. We then introduce a probability distribution over the range of the possible actions of each agent i . Essentially, all agents have the same set of actions, however the fitness measures and the associated probability distributions are different among agents, depending on their individual past performance. Let ${}_i P_{jt}$ be the probability that agent i selects action ${}_i a_{jt}$ (i.e. a short-run prediction) from the set \mathbf{A}_t , such that $0 \leq {}_i P_{jt} \leq 1$ and $\sum_{j=1}^{101} {}_i P_{jt} = 1$. The probability to select action ${}_i a_{jt}$ is computed

⁴Note that, if we exclude the peculiarity of the first period, the empirical data on long-run predictions lie always in that interval.

as:

$${}_iP_{jt} = \frac{\exp(-\gamma \cdot {}_iV_{j,t})}{\sum_{l=1}^{101} \exp(-\gamma \cdot {}_iV_{l,t})}, \quad (16)$$

where $\gamma \in [0, \infty)$ represents the intensity of choice. It determines the way an agent ranks the relative performance of its actions. We introduce the following formula for the long-run predictions:

$${}_iP_{jt}^k = \frac{\exp(-\gamma \cdot {}_iV_{j,t}^k)}{\sum_{j=l}^{101} \exp(-\gamma \cdot {}_iV_{l,t}^k)}. \quad (17)$$

According to the probability distributions in eqs. (16) and (17), each agent randomly chooses four actions ${}_i\tilde{a}_t$ and ${}_i\tilde{a}_t^k$, where $k \in \{1, 2, 3\}$ (one short-run prediction and three long-run predictions) in each period t .

3.1 Calibration of the learning algorithms

3.1.1 HSM calibration

In order to calibrate the extrapolative trend parameters of the heuristics, we run the regressions given by eqs. (7) and (8) on the time series of predictions of each individual subject.⁵ For the adaptive rule, we obtain similar results in the two treatments (see Tables 1 and 2). Interestingly, in the negative feedback treatment the majority of subject adopts a contrarian behavior, i.e. subjects form their predictions by assigning a negative coefficient to the observed trend. The extrapolative parameter of the TFR is equal to the average of all the estimated significant coefficients for the positive feedback treatment. We follow the same procedure for the negative feedback treatment to compute the parameter of the CR. In order to simulate the HSM, we use the following common specification: $\alpha = 0.63$, $w = 0.44$ and $s = -0.44$. Note that we decide not to implement exactly the heuristics described in the original paper by Anufriev and Hommes (2012).

Table 1: Estimated parameter of the ADA rule and the TFR/CR rule for the positive feedback treatment. The mean is the average across significant coefficients.

Equation	Parameters			
	Mean	Min	Max	SD
ADA	0.65	0.21	0.9	0.19
TFR/CR	0.48	-0.34	1.3	0.66

⁵Note that we eliminate the restriction on the parameter w in eq. (8), so that it can take any real value.

Table 2: Estimated parameter of the ADA rule and the TFR/CR rule for the negative feedback treatment. The mean is the average across significant coefficients.

Equation	Parameters			
	Mean	Min	Max	SD
ADA	0.61	0.3	0.9	0.16
TFR/CR	-0.4	-0.84	0.43	0.31

3.1.2 EEA calibration

We can estimate the main parameters of the EEA algorithm, namely ϕ_s , γ and ϕ_l , for each subject using the maximum likelihood procedure. The two parameters related to the short-run predictions can be expressed in a close-form, while for ϕ_l we rely on a numerical optimization. The probability distribution for one-step ahead predictions of eq. (16) can be approximated by a Gaussian distribution, whose parameters can be expressed as:

$${}_i\mu_{t-1} = \frac{p_{t-1} + {}_i\phi_s {}_iP_{t-1,t-1}^e}{(1 + {}_i\phi_s)}, \quad (18)$$

$$\sigma_i^2 = \frac{1}{2 \gamma_i (1 + {}_i\phi_s)}. \quad (19)$$

The mean ${}_i\mu_{t-1}$ is determined the past prediction and past market price with a given weight (ϕ_s). The variance is invariant over time and depends on the given subject. By rearranging eq. (18), it is possible to express the mean of the distribution as a convex combination between past expectations and past market price:

$${}_i\mu_{t-1} = (1 - \alpha_i)p_{t-1} + \alpha_i {}_iP_{t-1,t-1}^e, \quad (20)$$

where the weight is given by $\alpha_i = \frac{{}_i\phi_s}{1 + {}_i\phi_s}$ (note that ${}_i\phi_s > -1$). Estimating ${}_i\phi_s$ gives us information on how subjects adjust their short-run expectations over time. According to the values for ${}_i\phi_s$, subjects combine in different ways past predictions and prices. In particular, we identify the following behaviors: (i) $0 < {}_i\phi_s < 1$ translates in $0 < \alpha_i < 1/2$ so that subjects try to adjust their future predictions towards the market price; (ii) a value of ${}_i\phi_s > 1$ implies $\alpha_i > \frac{1}{2}$ and, as a consequence, subjects forecast values closed to previous predictions; (iii) for negative values of ${}_i\phi_s$ subjects “overcorrect” their past expectations by submitting predictions opposite to either market price ($-\frac{1}{2} < {}_i\phi_s < 0$) or past predictions ($-\frac{1}{2} < {}_i\phi_s < -\frac{1}{3}$). Figure 7 shows the distribution of the estimated values of ${}_i\hat{\phi}_s$

(panel a), $\hat{\gamma}_i$ (panel b) and ${}_i\hat{\phi}_l$ (panel c) for both positive and negative feedback treatments.

In Table 3, we classify the subjects⁶ according to the above mentioned behaviors. As can be seen, future predictions are computed differently depending on whether subjects participate in a positive or negative feedback treatment. In the positive feedback treatment, almost half of the subjects “overcorrect” their forecasts (18 out of 40) either towards p_{t-1} or ${}_i p_{t-1,t-1}^e$. In the negative feedback treatment, the large majority of subjects (30 out of 44) adjust their predictions to be close to the market price. Interestingly, only a negligible minority (6 out of 44) take into account past expectations.

Table 3: Classification of subjects across the different expectation formation rules.

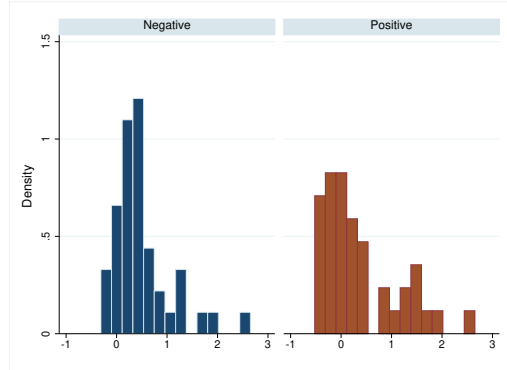
${}_i\phi_s$	α_i	Expectation formation rule	positive feedback	negat.
$0 < {}_i\phi_s < 1$	$0 < \alpha_i < \frac{1}{2}$	smooth correction towards p_{t-1}	14	
${}_i\phi_s > 1$	$\alpha_i > \frac{1}{2}$	smooth correction towards ${}_i p_{t-1,t-1}^e$	8	
$-\frac{1}{3} < {}_i\phi_s < 0$	$-\frac{1}{2} < \alpha_i < 0$	overcorrection towards p_{t-1}	12	
$-\frac{1}{2} < {}_i\phi_s < -\frac{1}{3}$	$-1 < \alpha_i < -\frac{1}{2}$	overcorrection towards ${}_i p_{t-1,t-1}^e$	6	

For the long-run expectations we estimate the individual values of ${}_i\phi_l$, common for two, three and four-steps-ahead predictions of a given subject, and use the value of γ_i obtained from short-run expectations. In order to simulate the EEA we use the following set of parameters $\phi_s = 0.4$, $\phi_l = 1.4$ and $\gamma = 0.4$ for the positive feedback treatment; $\phi_s = 0.15$, $\phi_l = 1.4$ and $\gamma = 0.4$ for the negative feedback treatment. Note that those values are the median values of the estimated coefficients.

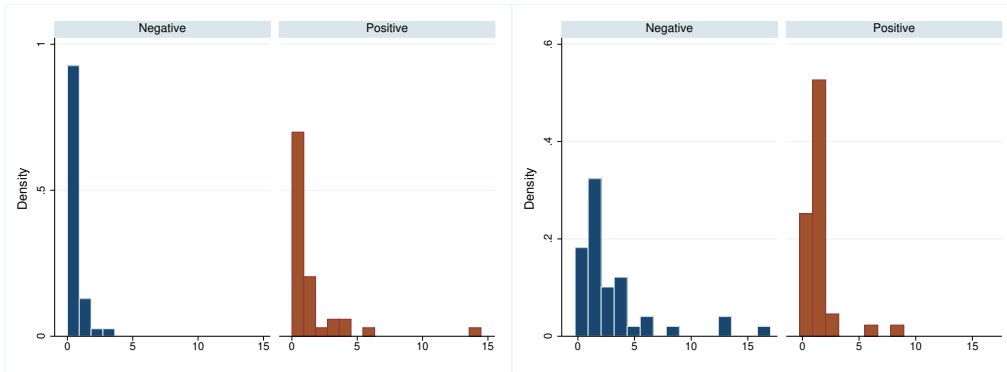
4 Simulation results

After a comprehensive description of the two algorithms, in this section we compare their performance in replicating the experimental results in markets with positive and negative feedback and for short- as well as long-run expectations.

⁶We have eliminated the estimates of six subjects because of the estimated values of ${}_i\phi_s$ are out of range due to an error in typing their predictions, two in the positive feedback and four in the negative treatment.



(a) Estimated values of $i\phi_s$



(b) Estimated values of γ_i

(c) Estimated values of $i\phi_l$

Figure 7: Distribution of the estimated individual values for $i\phi_s$, γ_i and $i\phi_l$. Blue bars refer to Negative feedback and red striped bars refers to Positive feedback.


4.1 Comparing the HSM and EEA to describe short-run expectations

In order to simulate the HSM, in the first three periods we use the experimental data to initiate the algorithm, assigning the same weight to each rule, i.e. $n_{h,1} = 0.25, \forall h$. Starting from period three, we compute the fitness measure and the weights $n_{h,3}$ associated to each heuristic. For the subsequent periods we iterate the algorithm detailed in the previous section.

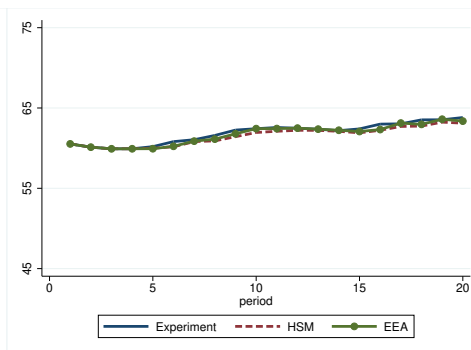
In order to initiate the EEA algorithm, we use the experimental individual predictions and the first three realized prices, since the range of the action sets depends on the three past realized prices. Individual predictions are independent realizations from different distributions, so that we have six (the number of subjects in the group) different distributions and six predictions in every period and for every group. Once the short-term predictions are determined, we compute the market price by using either eq. (1) or eq. (2) according to the treatment the group belongs to.

Figures 8 and 9 show the simulated market prices compared to the experimental market prices for all groups in positive and negative feedback treatment, respectively. From a preliminary inspection, both algorithms perform well in replicating experimental prices in the two treatments. Table 4 displays the values of the Mean Squared Error (MSE) for the two algorithms. The ability to foresight experimental prices is similar for both algorithms. Interestingly, we observe a systematic higher value of the MSE in the negative feedback treatment compared to the positive treatment. Thus, it seems that both algorithms are able to better capture the price time series in the positive than in the negative feedback treatment. This is in line with the literature on the application of the HSM in reproducing the experimental data.

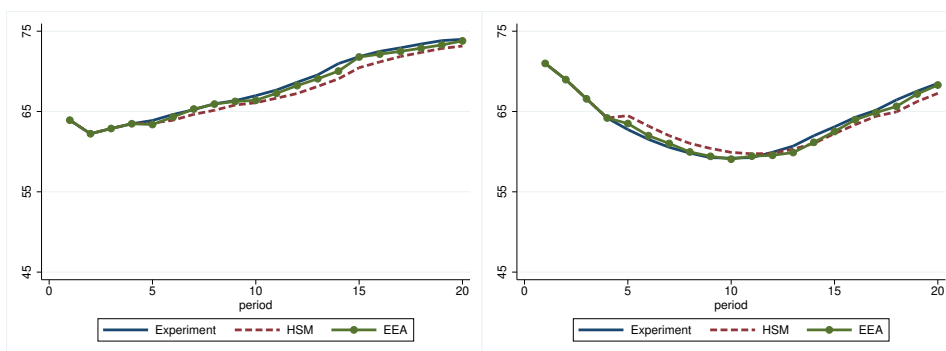
The two algorithms can fairly well reproduce the stylized facts regarding the mutual coordination of expectations and convergence to the fundamental value, described in section 2.2. Figures 17 and 16 compare the experimental data with the simulated data from the two algorithms. Both algorithms capture the faster coordination of short-term predictions in the positive feedback markets as compared to the negative feedback markets. At the same time, the HSM and EEA can reproduce the convergence to the REE of the short-term expectations in the negative feedback treatment. In the positive feedback treatment, both prices and predictions are not converging to the fundamental value, which is well captured by the two algorithms.




(a) Group 1




(b) Group 2




(c) Group 3




(d) Group 4



(e) Group 5



(f) Group 6



(g) Group 7

Figure 8: Simulation results based on the HSM and the EEA in the positive feedback treatment. The continuous black line is the experimental market price, the blue line is the market price generated using the HSM and the dashed grey line is the market price generated using the EEA.

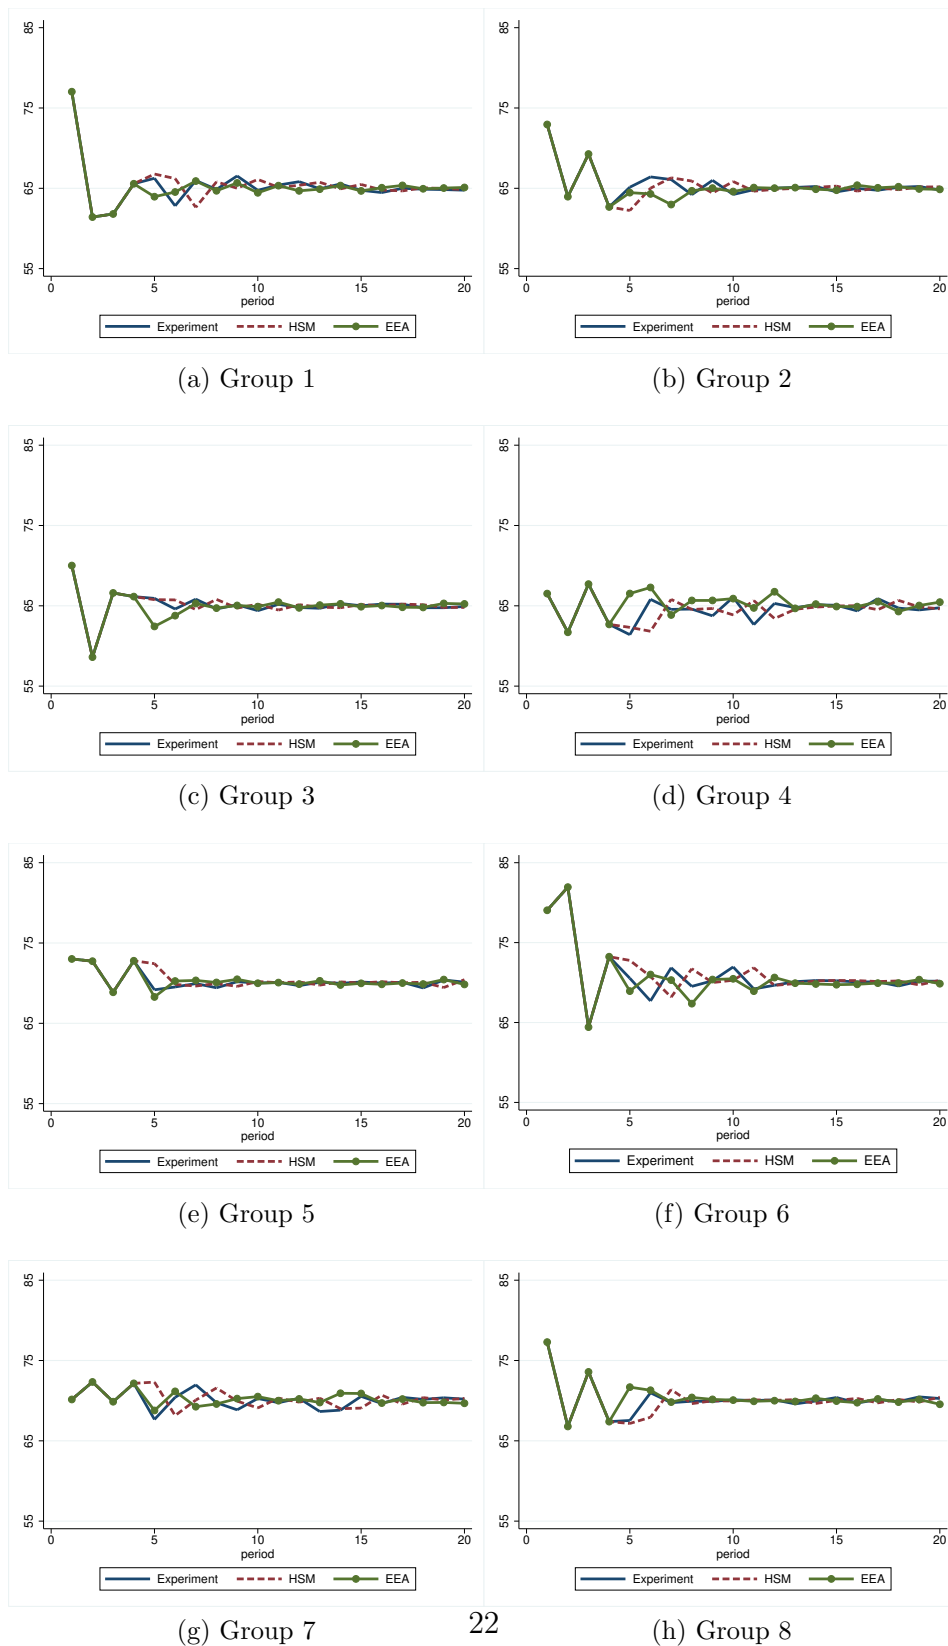


Figure 9: Simulation results based on the HSM and the EEA in the negative feedback treatment. The continuous black line is the experimental market price, the blue line is the market price generated using the HSM and the dashed grey line is the market price generated using the EEA.

Table 4: Mean Squared Error of the HSM and EEA in describing the time series of experimental market prices.

	Positive		Negative	
	HSM	EEA	HSM	EEA
Group 1	0.46	0.29	1.50	0.46
Group 2	0.22	0.06	0.97	3.21
Group 3	0.94	0.13	0.32	0.75
Group 4	0.99	0.13	2.00	1.84
Group 5	0.17	0.38	0.67	0.61
Group 6	0.48	0.16	2.16	1.77
Group 7	0.98	0.77	2.18	1.03
Group 8	-	-	0.69	2.07

4.2 Comparing the HSM and EEA to describe long-run expectations

One interesting contribution of the paper is the extension of the HSM in order to describe long-term expectations, with a simple linear extrapolative modification of the existing heuristics. The comparison with the EEA and its extension can give a rough idea of the goodness of the modified HSM in reproducing the long term-expectations. This is a relevant step to have a reliable framework to model long-term expectations in a realistic environment.

To the best of our knowledge, this is the first attempt in the literature on LtFEs to reproduce individual long-run expectations using the HSM. Despite the fact that in the experiment subjects submit in each period t their expectations for the remaining $20 - t$ periods, we replicate the individual expectations up to four steps ahead. Our choice represents a good compromise between considering the whole time-span and having a sufficient statistics to analyse the properties of the two algorithms as a function of the time horizon and comparing them to the experimental data.

We study the performance of the two algorithms in replicating the main statistical properties of the experimental data, namely (i) the time series of the individual long-term expectations, (ii) the coordination of long-term predictions as a function of the time horizon and (iii) the convergence of long-run expectations to the fundamental value. It is worth mentioning that we do not have an aggregate variable as the market price for the long-run predictions, but just individual long-run predictions. When necessary, therefore, we rely on the average of the long-run expectations.

Figures 10 and 11 illustrate the evolution of the average long-run expectations

(across subjects) for 2, 3 and 4 steps ahead in the case of positive and negative feedback treatments. Once again, both algorithms seems to replicate the experimental data with a reasonable accuracy. Figures from 12 to 15 show the evolution of individual long run expectations in a representative group in the two treatments and for the two algorithms. A first glance, those figures shows that the individual long-run expectations generated by the two algorithms resemble the experimental elicited expectations. In particular, we can observe in the positive feedback treatment, the cone-shape form of the predictions submitted in a given period as a function on the increasing time horizon. In the negative feedback treatment, we can observe the dynamical process of convergence to the REE of short- as well as long-run predictions.

In order to have a more quantitative comparison between the two algorithms, Tables 5 and 6 show the mean squared error of the simulated data (averaged over 100 Monte Carlo simulations) in describing the average long-run expectations. The EEA describes significantly better the long-run predictions than the modified HSM with the linear extrapolation heuristics.

In Figures 16 and 17, we compare the dynamical properties of the mutual coordination of long-run expectations and their convergence to the fundamental value. At a first inspection, it seems that the two algorithms are able to fairly well replicate the behavior of the experimental data, since the degree of coordination resulting from the simulated data is fairly close to the degree of coordination of experimental data. Additionally, the EEA and the HSM are capable of reproducing the more persistent heterogeneity observed for the long-run predictions as compared to the degree of coordination of the one-step-ahead predictions. A closer look at the behavior of the dispersion (measured with the MAD) of the long-run expectations as a function of the time horizon (panel (c) of Figure 16) shows that it increases linearly with the time horizon. This tendency is systematic in both, the negative and the positive feedback treatments. In the positive treatment, such linear behavior is similar to the empirical data. In the negative feedback treatment, instead, such increase of the dispersion over the horizons is much less evident, with some periods showing an absence of such a systematic increasing tendency.

Note that the HSM with the linear extrapolative trend for the long-run predictions has built-in such characteristics. Any (linear) measure of dispersion of the predictions as a function of the time horizon will, therefore, exhibit a linear increase. This property is counter factual if we consider the negative feedback treatment. It is intuitive, in fact, that a simple linear extrapolation of prices cannot predict convergent prices to the fundamental value. Our numerical exercise, thus, shows that a the HSM should be modified with more complex heuristic rules to account for long-run expectations.

Figure 16 displays the standard deviation of individual predictions for the price

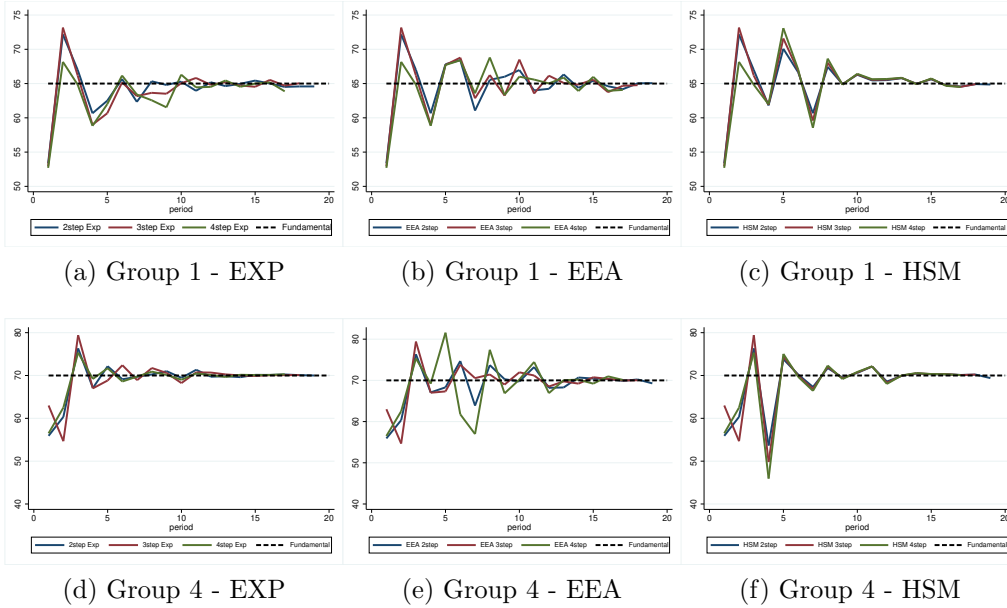


Figure 10: Average long-run expectations of two representative groups in the negative feedback treatment for both experimental and simulated data. Blu, red and green lines represent the average 2, 3 and 4 step-ahead predictions, respectively.

2, 3 and 4 periods ahead. It shows that the degree of coordination resulting from the simulated data is fairly close to the degree of coordination of experimental data. Additionally, the EEA is able to reproduce the more persistent heterogeneity observed for the long-run predictions as compared to the degree of coordination of the one-step-ahead predictions.

5 EEA with the fundamental value as anchor

From our analysis on how subjects form their expectations we observe that subjects follow an anchor and adjustment mechanism where the anchor depends on the feedback mechanism that drives the formation of market prices: in markets with positive feedback subjects' expectations are driven by the past market price dynamics whereas, in markets with a negative feedback markets, subjects learn the fundamental value and use it as anchor to form their short-run expectations on the price. Table 7 illustrates our conclusion. We estimate two equations to explain individual short-run expectations for the positive and negative feedback treatments implemented in the laboratory experiments: in Model (1) we consider as explanatory variables the lagged value of individual short-run expectations and

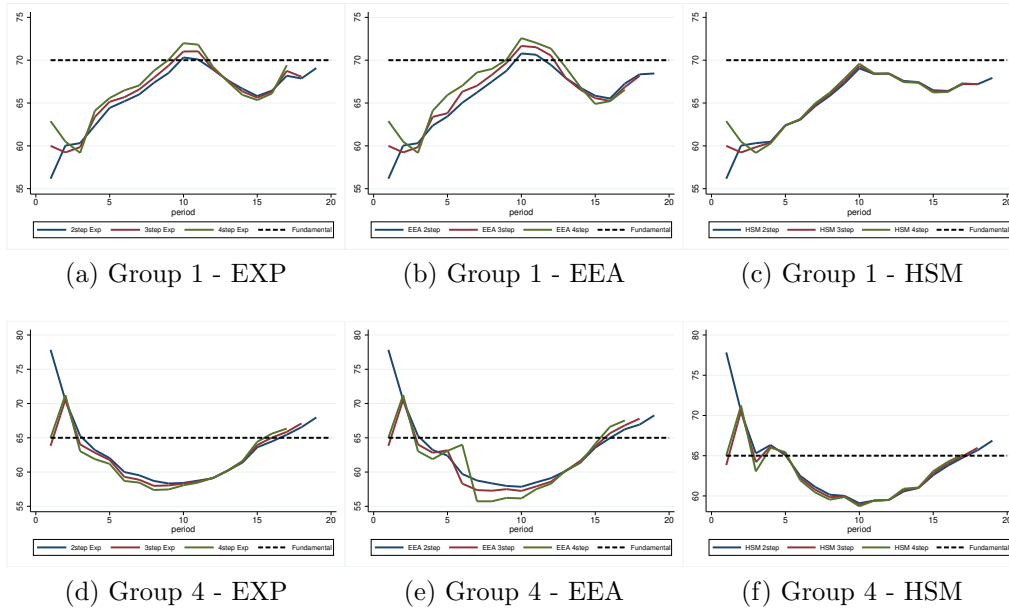


Figure 11: Average long run expectations of two representative groups in the positive feedback treatment for both experimental and simulated data. Blu, red and green lines represent the average 2, 3 and 4 step-ahead predictions, respectively.

Table 5: Mean squared error of simulated long-run predictions in both positive and negative feedback treatments. We compute the average across groups quadratic distance between the average long-run experimental and EEA simulated predictions. Predictions are computed by using last observed price as an anchor.

	Negative			Positive		
	2step	3step	4step	2step	3step	4step
Group 1	2.46	5.16	4.97	0.23	0.59	1.21
Group 2	2.3	1.47	2.81	0.18	0.8	0.94
Group 3	0.3	2.88	3.66	0.26	0.5	1.07
Group 4	1.87	1.17	3.88	0.13	0.55	3.02
Group 5	0.89	0.77	0.93	0.38	0.52	1.18
Group 6	5.41	1.67	22.94	0.2	0.5	1.15
Group 7	2.06	1.31	2.06	1.03	2.37	5.23
Group 8	2.2	0.42	1.91	-	-	-

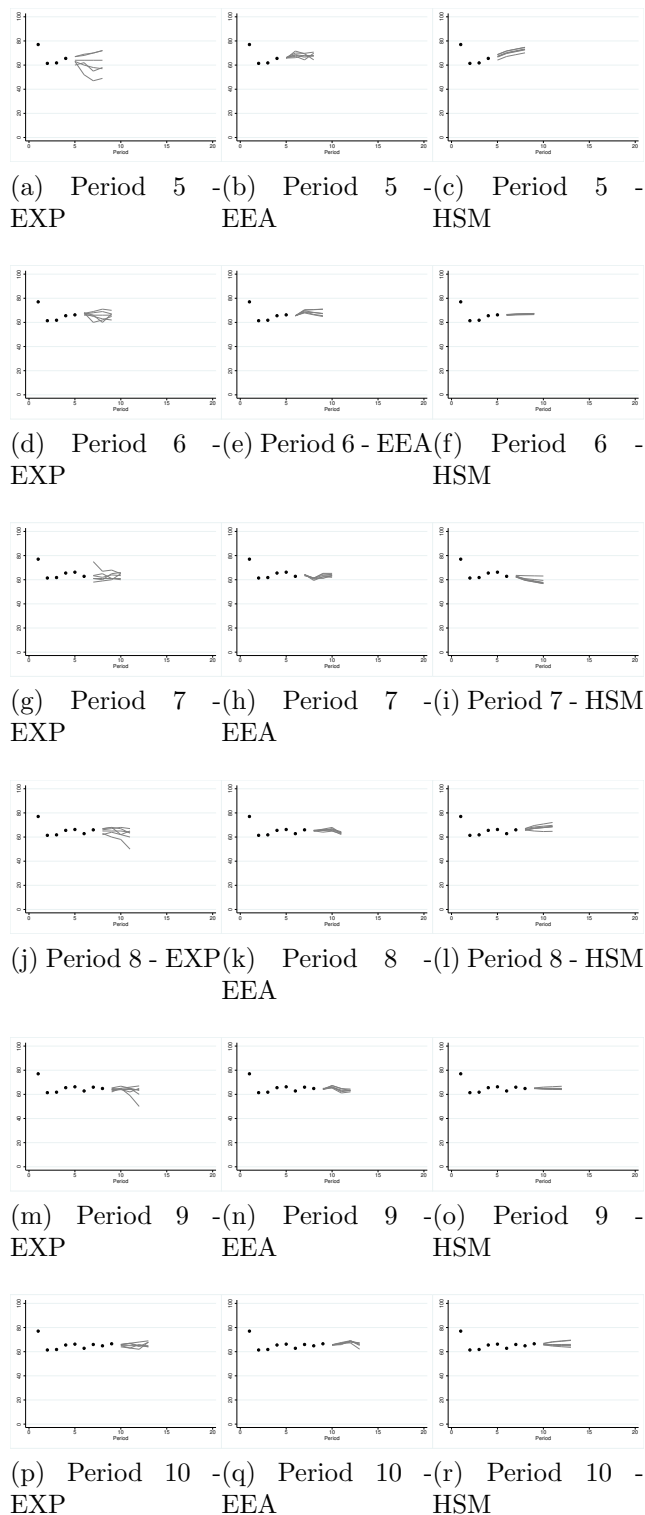


Figure 12: Individual long-run predictions both real and simulated of Group 1 of the negative feedback treatment from period 5 to period 10. The black dots are the market price, the grey lines are the individual forecasts and the dashed line represents the fundamental value.

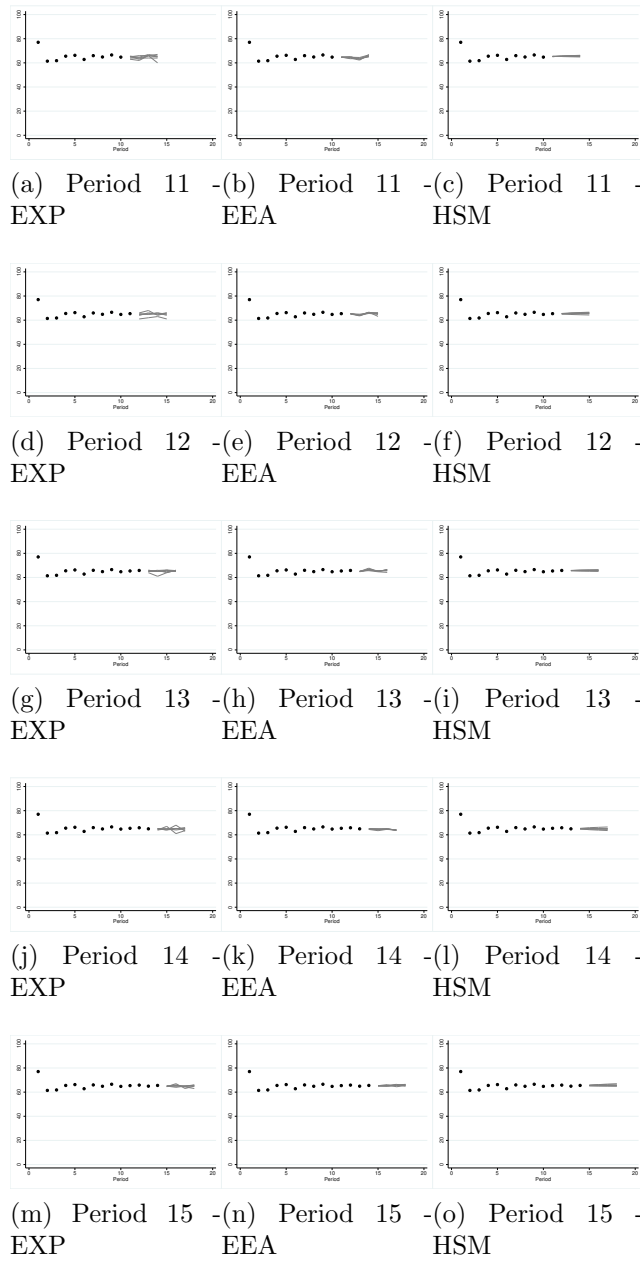


Figure 13: Individual long-run predictions both real and simulated of Group 1 of the negative feedback treatment from period 11 to period 15. The black dots are the market price, the grey lines are the individual forecasts and the dashed line represents the fundamental value.

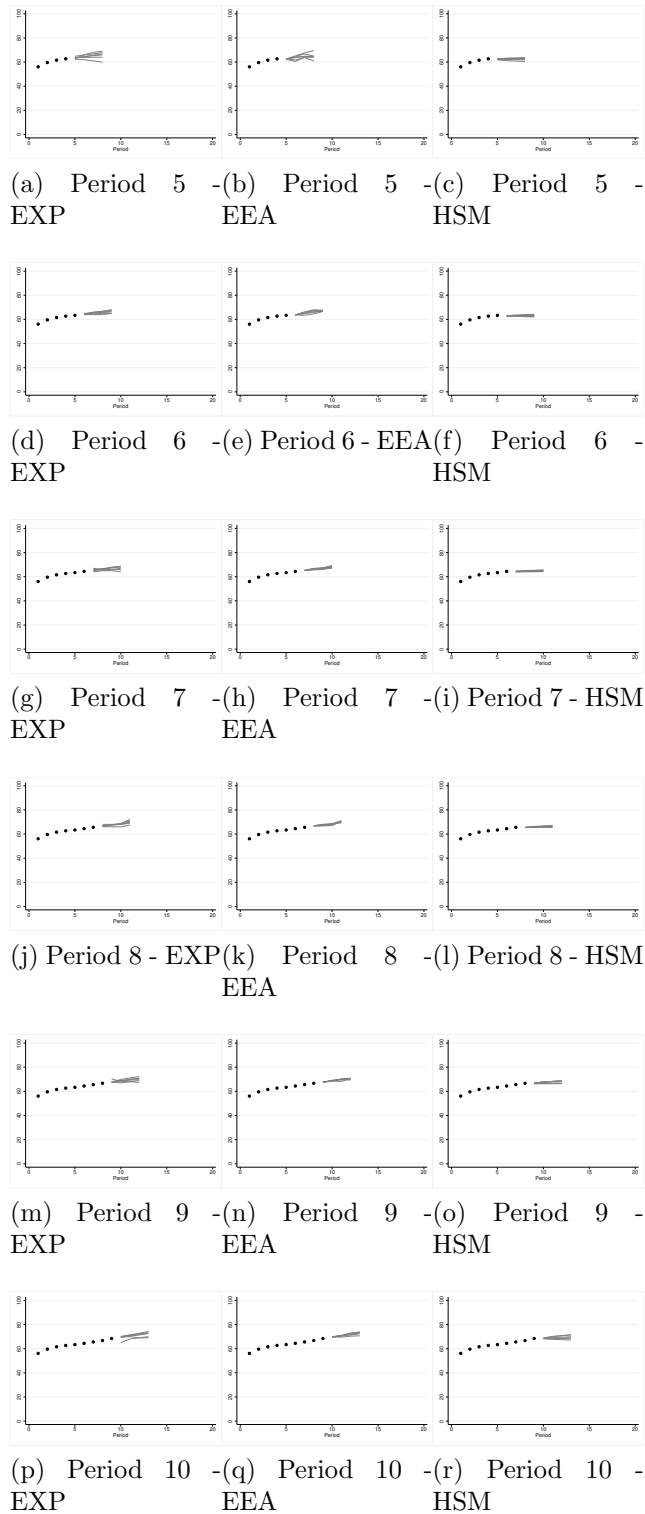


Figure 14: Individual long-run predictions both real and simulated of Group 1 of the positive feedback treatment from period 5 to period 10. The black dots are the market price, the grey lines are the individual forecasts and the dashed line represents the fundamental value.

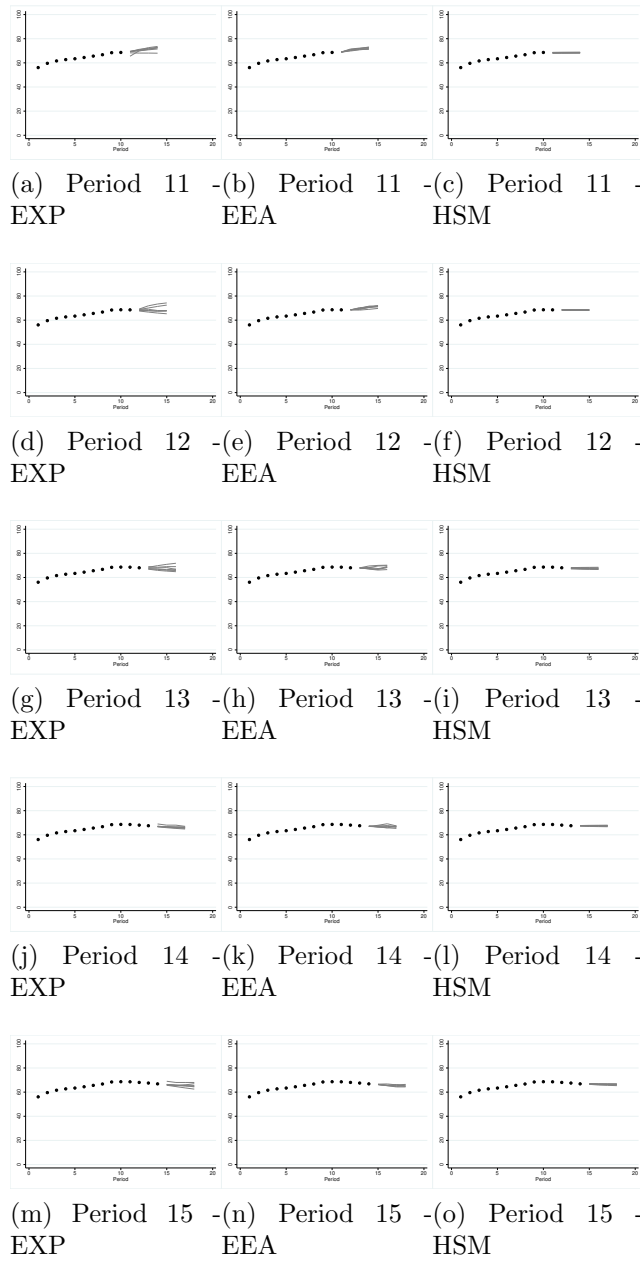
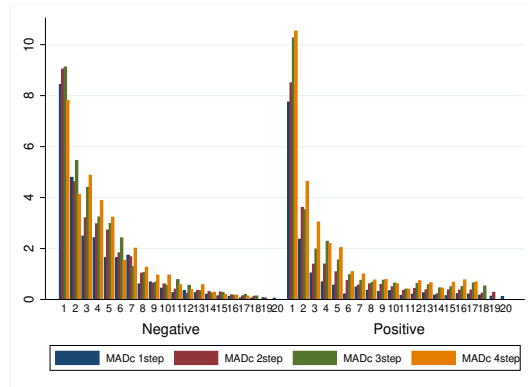
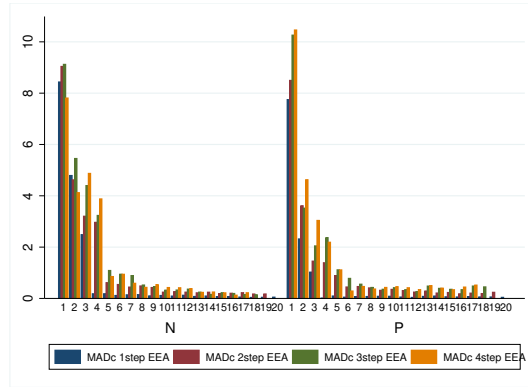


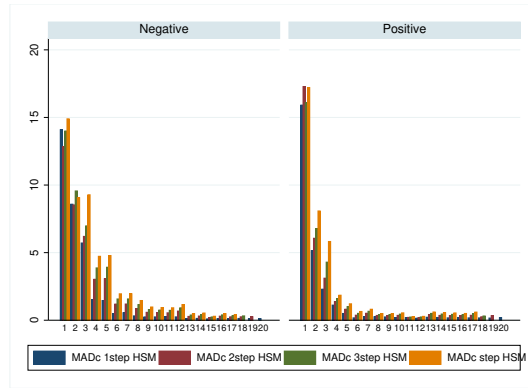
Figure 15: Individual long-run predictions both real and simulated of Group 1 of the positive feedback treatment from period 11 to period 15. The black dots are the market price, the grey lines are the individual forecasts and the dashed line represents the fundamental value.



(a) Experimental

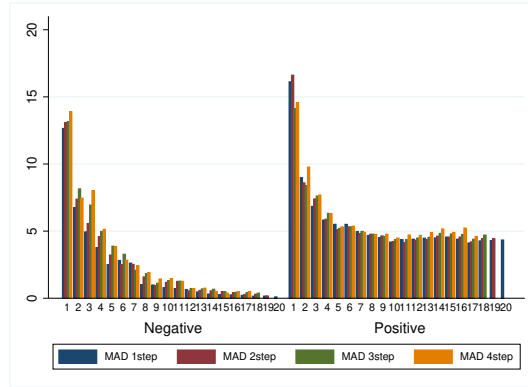


(b) Simulated




(c) Simulated

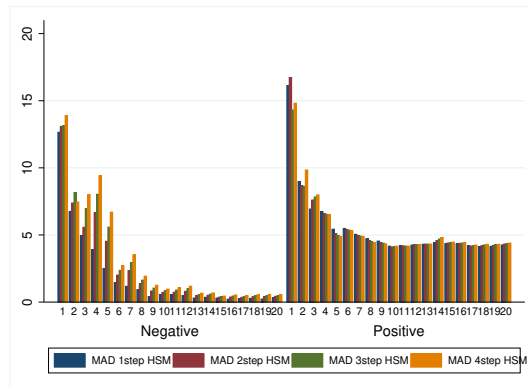
Figure 16: Coordination: comparison between experimental and simulated data of the average across groups of the mean absolute deviation of one, two, three and four-step-ahead predictions per period for both positive and negative treatment.



(a) Experimental



(b) Simulated EEA



(c) Simulated HSM

Figure 17: Convergence to the fundamental value: comparison between experimental and simulated data of the average across groups of the mean absolute deviation of one, two, three and four-step-ahead predictions per period for both positive and negative treatment.

Table 6: Mean squared error of simulated long-run predictions in both positive and negative feedback treatments. We compute the average (across groups) quadratic distance between the average long-run experimental and HSM simulated predictions.

	Negative			Positive		
	2step	3step	4step	2step	3step	4step
Group 1	3.78	9.39	12.31	1.40	2.68	4.51
Group 2	2.42	10.09	13.35	0.05	0.59	0.25
Group 3	2.92	4.87	9.01	0.51	1.01	1.96
Group 4	4.69	7.90	9.08	1.20	2.71	3.95
Group 5	0.59	0.56	0.35	0.29	0.42	0.79
Group 6	10.48	19.64	34.03	1.14	2.37	3.66
Group 7	0.81	1.37	1.76	1.94	2.26	2.19
Group 8	3.05	6.41	6.97	-	-	-

market price. Instead, in Model (2) we use as explanatory variables past individual short-run expectations and the fundamental value. Our results show that, whereas Model (1) performs well in the two feedback mechanism showing the relevance of past prices in the formation of expectations, Model (2) clearly highlights the pivotal role of the fundamental value in determining short-run expectations in negative feedback markets with a coefficient close to 1, compared to the role of the fundamental price in the formation of short-run expectations in markets with positive feedback (with a coefficient of 0.09).

Table 7: Dynamic panel regression of individual short-run expectations. Dependent variable: individual short-run expectations. Standard errors in parentheses.

	Positive		Negative	
	Model(1)	Model(2)	Model(1)	Model(2)
p_{t-1}	0.77*** (0.05)		0.70*** (0.02)	
$iP_{t-1,t-1}^e$	0.23*** (0.05)		0.30*** (0.02)	
pf		0.09* (0.05)		0.99*** (0.04)
$iP_{t-1,t-1}^e$		0.91*** (0.06)		0.01 (0.04)
N	798	798	912	912
Wald χ^2	134952.46***	32720.83***	15010.52***	15494.63***

Focusing now on individual expectations in the long-run, recall that the the EEA was developed on the basis of the anchor-and-adjustment mechanism and the

choice of the right anchor is therefore crucial. Colasante et al. (2018c) show that, using the market price as anchor in markets with positive feedback, they obtain a fairly good replication of the subjects' expectations elicited in the laboratory experiment. In this work, they also compare the performance of the EEA with an alternative model like the so-called *noisy rational expectations*, that is based on the idea that subjects have homogeneous rational expectations.⁷ They conclude that this model provides a good approximation of the expectations elicited in the negative feedback treatment. However, even though (experimental) predictions in markets with negative feedback quickly converge to the fundamental value, they observe a persistent heterogeneity of expectations in the experimental data that is not found in the simulated expectations.

In order to reproduce the main features typically observed in (experimental) subjects' expectations concerning heterogeneity and convergence of long-run expectations, we implement an alternative version of the EEA, where we use the fundamental value as anchor.⁸ We compute then agents' individual predictions for different forecast horizons k , where $k = 1, 2, 3$, using the fundamental value (instead of past market prices) as anchor. In other words, the fitness measure is computed using eqs. (14) and (15) considering instead the fundamental value as a reference point as follows:

$${}_i\tilde{V}_{j,t} = (pf - {}_i a_{jt})^2 + \phi_s ({}_i p_{t-1,t-1}^e - {}_i a_{jt})^2, \quad (21)$$

$${}_i\tilde{V}_{j,t}^k = |pf - {}_i a_{jt}^k| + \phi_l |{}_i p_{t-k-1,t-1}^e - {}_i a_{jt}^k|. \quad (22)$$

Setting $\phi \neq 0$ we are able to obtain heterogeneous predictions as a result of past individual expectations. In order to illustrate our results, Figures 19 and 18 show as an example the average of individual predictions two, three and four steps ahead in two representative groups in market with positive and negative feedback. Table 8 displays the MSE comparing the experimental data and the simulated expectations in markets with positive and negative feedback. The results obtained show that the EEA with the fundamental value as anchor performs better to replicate, on average, the experimental data in the markets with negative feedback compare to its performance in the markets with positive feedback.

With these results we provide further evidence that subjects form their expectations following different rules depending on the feedback system. In markets

⁷In the EEA homogeneous rational expectations are implemented by setting the mean of the probability distribution of each agent (i.e. i_μ) equal to the fundamental value in the centre of the range of actions. They then set $\phi = 0$ and replace p_{t-1} with the fundamental value

⁸Note that we do not perform the same exercise for short-run predictions since the range of the distribution is given by the standard deviation of the market price and we cannot compute a variable range if we include the fundamental value as anchor.

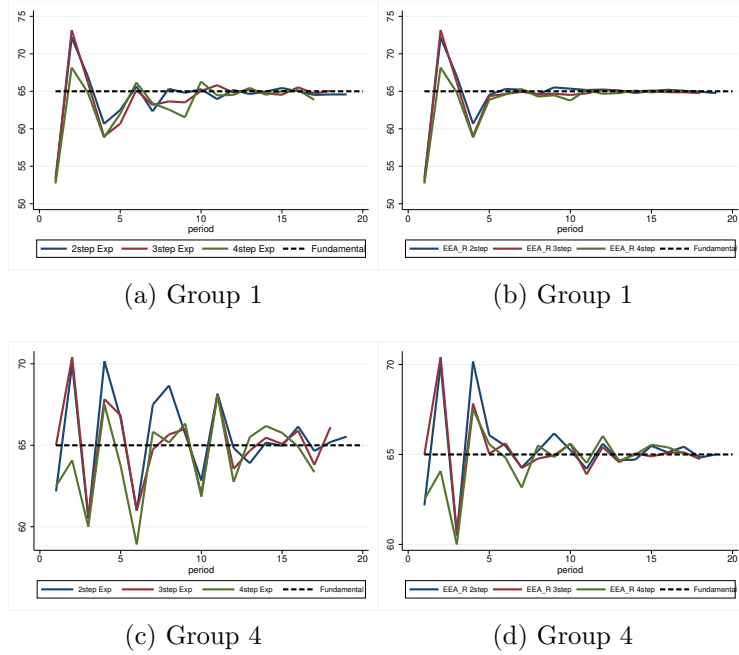


Figure 18: Average long run expectations in the negative feedback treatment.

with negative feedback, subjects' expectations are mainly driven by the fundamental value, while in the positive feedback system the reference point is the market price.

6 Conclusion

In this paper we present the results of the application of two adaptive learning algorithms to describe the experimental data of a LtFEs, in which we elicit short and long-run expectations. In particular, we ask to the subjects to submit their predictions in two distinct environments: positive and negative feedback systems. The main difference between those expectations' feedback systems relies in the sign of the relation between expectations and market price. The descriptive analysis shows that the dynamical properties of the predictions are markedly different in terms of both coordination of expectations and convergence to the fundamental value. In order to understand the process of expectations' formation in both feedback systems, we consider two evolutionary learning algorithms: the Heuristic Switching Model and the Exploration-Exploitation Algorithm. The main difference between these algorithms is that the HSM can be defined as "parametric", meaning that it is based on few predetermined heuristics, and the EEA is instead

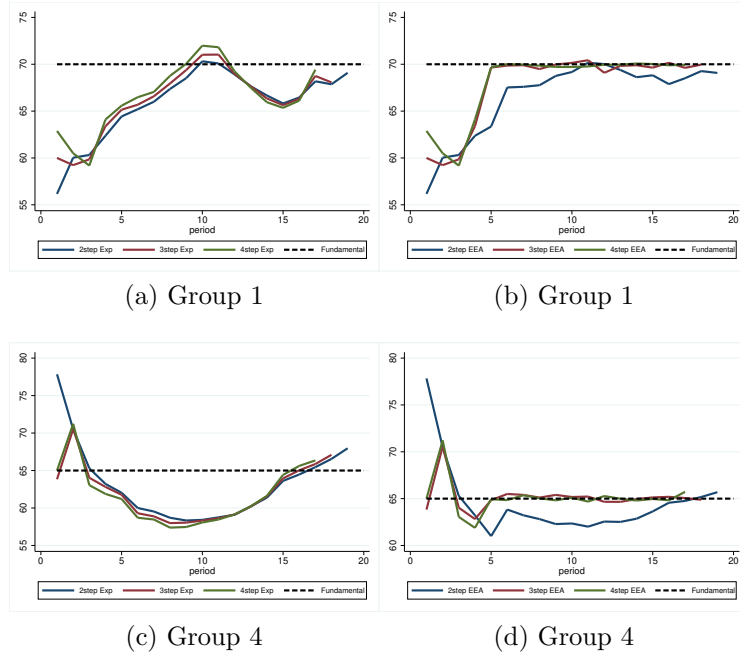


Figure 19: Average long run expectations in the positive feedback treatment.

Table 8: Mean squared error of simulated long run predictions in both positive and negative feedback treatments. We compute the quadratic distance between the average experimental and EEA simulated predictions. Predictions are computed by using fundamental value as an anchor.

	Negative			Positive		
	2step	3step	4step	2step	3step	4step
Group 1	0.84	1.21	1.8	6.86	7.50	8.34
Group 2	0.48	0.61	0.7	29.30	74.54	69.84
Group 3	0.23	0.6	1.62	11.11	28.08	30.28
Group 4	3.59	3.61	5.33	8.71	25.33	28.59
Group 5	0.17	0.23	1.15	4.08	6.62	6.74
Group 6	0.65	0.97	0.73	3.07	8.62	11.65
Group 7	0.8	1.33	1.12	18.85	17.35	17.10
Group 8	0.23	0.34	0.37	-	-	-

“non-parametric”, so that the predictions are chosen according to a specific probability distribution. The two algorithms are based on the common principle of anchor and adjustment rule. Regarding short run predictions, we observe that both algorithms perform well in replicating individual predictions. In order to simulate long-run predictions, we have introduced a straightforward extension of the HSM: a linear extrapolation of short-run predictions across different horizons. A considerable difference emerges between the two algorithms: the linear extrapolation provide a good approximation of the main stylized facts in the positive feedback treatment. In the negative feedback system, the coordination and convergence properties are better explained by the EEA. Moreover, we have performed an exercise to test whether a change in the anchor can lead to better results in replicating the experimental data. We have considered as an alternative anchor the fundamental value instead of the market price. This modification leads to better simulation results in the negative feedback system, while in the positive feedback market we obtain a worse performance. The comparison between such different algorithms helps us to understand that the subjects form their expectations using different anchors. The EEA is a good and flexible tool to replicate short and long-run predictions in both positive and negative feedback systems. The linear extrapolation we have implemented in this paper as an extension of the HSM to compute long-run expectations is not sufficiently flexible to capture the behavior in the different feedback environments: a more complex structure is needed to replicate the properties of long-run predictions, which is the focus of future research.

Acknowledgement

The authors are grateful for funding the Universitat Jaume I under the project UJI-B2018-77, the Spanish Ministry of Science and Technology under the project ECO2015-68469-R and the Generalitat Valenciana for the financial support under the project AICO/2018/036.

References

- Mikhail Anufriev and Cars Hommes. Evolutionary selection of individual expectations and aggregate outcomes in asset pricing experiments. *American Economic Journal: Microeconomics*, 4(4):35–64, 2012.
- Mikhail Anufriev, Cars Assenza, Tiziana and, and Domenico Massaro. Interest rate rules and macroeconomic stability under heterogeneous expectations. *Macroeconomic Dynamics*, 17(08):1574–1604, 2013.
- Tiziana Assenza, Peter Heemeijer, Cars Hommes, and Domenica Massaro. Individual Expectations and Aggregate Macro Behavior. Dnb working papers, Netherlands Central Bank, Research Department, May 2011.
- Tiziana Assenza, Te Bao, Cars Hommes, and Domenico Massaro. Experiments on expectations in macroeconomics and finance. In *Experiments in macroeconomics*, pages 11–70. Emerald Group Publishing Limited, 2014.
- Te Bao and Li Ding. Nonrecourse mortgage and housing price boom, bust, and rebound. *Real Estate Economics*, 44(3):584–605, 2016.
- Te Bao, Cars Hommes, Joep Sonnemans, and Jan Tuinstra. Individual expectations, limited rationality and aggregate outcomes. *Journal of Economic Dynamics and Control*, 36(8):1101–1120, 2012.
- Te Bao, John Duffy, and Cars Hommes. Learning, forecasting and optimizing: An experimental study. *European Economic Review*, 61:186–204, 2013.
- William A Brock and Cars H Hommes. Heterogeneous beliefs and routes to chaos in a simple asset pricing model. *Journal of Economic dynamics and Control*, 22(8-9):1235–1274, 1998.
- Annarita Colasante, Simone Alfarano, Eva Camacho, and Mauro Gallegati. Long-run expectations in a learning-to-forecast experiment. *Applied Economics Letters*, 25(10):681–687, 2018a.
- Annarita Colasante, Simone Alfarano, and Eva Camacho-Cuena. The term structure of cross-sectional dispersion of expectations in a Learning-to-Forecast Experiment. Technical report, 2018b.
- Annarita Colasante, Simone Alfarano, Eva Camacho-Cuena, and Mauro Gallegati. Long-run expectations in a Learning-to-Forecast Experiment: A Simulation Approach. *Journal of Evolutionary Economics*, 2018c. doi: 10.1007/s00191-018-0585-1.

- Camille Cornand and Cheick Kader M'baye. Does inflation targeting matter? an experimental investigation. *Macroeconomic Dynamics*, pages 1–40, 2016.
- Cees Diks and Roy Van Der Weide. Herding, a-synchronous updating and heterogeneity in memory in a cbs. *Journal of Economic Dynamics and Control*, 29(4): 741–763, 2005.
- Milton Friedman. *Essays in Positive Economics*. University of Chicago Press, 1966.
- Peter Heemeijer, Cars Hommes, Joep Sonnemans, and Jan Tuinstra. Price stability and volatility in markets with positive and negative expectations feedback: An experimental investigation. *Journal of Economic Dynamics and Control*, 33(5): 1052–1072, 2009.
- Cars Hommes. *Behavioral rationality and heterogeneous expectations in complex economic systems*. Cambridge University Press, 2013.
- Cars Hommes. Behavioral & experimental macroeconomics and policy analysis: a complex systems approach. Working paper series, European Central Bank, November 2018.
- Cars Hommes and Thomas Lux. Individual expectations and aggregate behavior in learning-to-forecast experiments. *Macroeconomic Dynamics*, 17(02):373–401, 2013.
- Cars Hommes, Joep Sonnemans, Jan Tuinstra, and Henk van de Velden. A strategy experiment in dynamic asset pricing. *Journal of Economic Dynamics and Control*, 29(4):823–843, 2005.
- Daniel Kahneman and Amos Tversky. On the psychology of prediction. *Psychological Review*, 80(4):237, 1973.
- Robert E Lucas Jr and Edward C Prescott. Investment under uncertainty. *Econometrica: Journal of the Econometric Society*, pages 659–681, 1971.
- Ramon Marimon, Stephen E Spear, and Shyam Sunder. Expectationally driven market volatility: an experimental study. *Journal of Economic Theory*, 61(1): 74–103, 1993.
- John F Muth. Rational expectations and the theory of price movements. *Econometrica: Journal of the Econometric Society*, pages 315–335, 1961.
- Herbert Alexander Simon. *Models of man: social and rational; mathematical essays on rational human behavior in society setting*. Wiley, 1957.

The instructions and the scree-shot are not intended to be published. We include this information for the convenience of the referees.

A Instructions and Screenshot

[General instructions]

Welcome to the Laboratory of Experimental Economics! You are participating in an experiment in which you will take decisions in a financial market. The instructions are very simple but, please, read them carefully.

During the whole experiment you will play with experimental currency unit (ECU) and, at the end of the experiment, your final profit, summed to the 3? for the show-up fee, will be converted in Euro according to the following exchange rate: **1 Euro=500 ECU**. The total amount will be paid at the end of the experiment by cash.

[Only in the positive feedback treatment]

You are a financial advisor to a pension fund that wants to invest an amount of money to buy an asset. The pension fund will allocate its money between a bank account which pays fix interest rate and a risky asset that pays dividends. The allocation depends on your forecast on the evolution of the asset price. When making your predictions remember that the asset price each period is affected: positively by the dividend, negatively by the interest rate and positively by the investors' expectations on the asset price in that period.

Your task is to predict the price for 20 periods. In each period (t) you will predict the price for all the remaining $20 - t$ periods, that is, in period 1 you will submit 20 predictions starting from the prediction about the price at the end of period 1, in period 2 you will submit 19 predictions and so on. Your predictions must be between 0 and 100.

In period 1 you will submit predictions having information only about the interest rate and the average dividend. From period 2 on, you will have more information: besides the interest rate and the average dividend, you will see a graph with the time series of your past predictions and the time series of the market prices. The green dots represent the time series of your one-step-ahead predictions, while the blue dots represent the asset price in each period. Additionally, you will see the values of those time series and the time series of all your past predictions. Remember that in any period you will see the information about the asset price

in the previous periods.

$i_{t-1}^{\frac{1}{2}}$

The interest rate will be equal to 5% and the mean dividend will be equal to 3.25 (or 3.5 depending on the session).

[Only in the negative feedback treatment]

You are an advisor to a firm that wants to buy a certain amount of a good. In each period, the manager of the firm decides how many units of that good she wants to buy with the aim to sell it in the next period. To take an optimal decision, the manager needs a good prediction of the market price in the next period. The evolution of the market price will be as follows: if the demand for the good is higher than the supply, the price will rise. Conversely, if the supply will be higher than the demand, the price will decrease. The manager will take her decision based on your predictions about the market price in given period: the higher (lower) is the prediction, the higher (lower) will be the demand and, as a consequence, the market price will fall (rise).

Your task is to predict the price for 20 periods. In each period (t) you will predict the price for all the remaining $20-t$ periods, that is, in period 1 you will submit 20 predictions starting from the prediction about the price at the end of period 1, in period 2 you will submit 19 predictions and so on. Your predictions must be between 0 and 100.

In period 1 you will submit predictions without any information about past prices. From period 2 on, you will see a graph with the time series of your past predictions and the time series of the prices. The green dots represent the time series of your one-step-ahead predictions, while the blue dots represent the market price in each period. Additionally, you will see the values of those time series and the time series of all your past predictions. Remember that in any period you will see the information about the market price of the previous periods.

[General instructions]

Once each subject has submitted her prediction for each period, the price will be computed and it will be shown at the beginning of period 2. The same mechanism will be used for subsequent periods. Once you and the other subjects submit the predictions, the price as well as the profits will be computed according to the forecasting accuracy.

Remember that your profit depends on your forecasting accuracy. The lower your forecasting error (the distance between your predictions and the price in a given period), the higher will be your profit. Your profit will be computed at the end of each period. Besides the profit for your one-step-ahead prediction each period, you will receive an extra profit for your past predictions about the price each period. This extra profit will be computed according to the following table:

Difference between price of period t and your prediction for period t	ECU
± 5	25
± 10	12
± 15	5

At the beginning of each period you will see the profit for all the predictions and the cumulative gains.

A.1 Screenshot

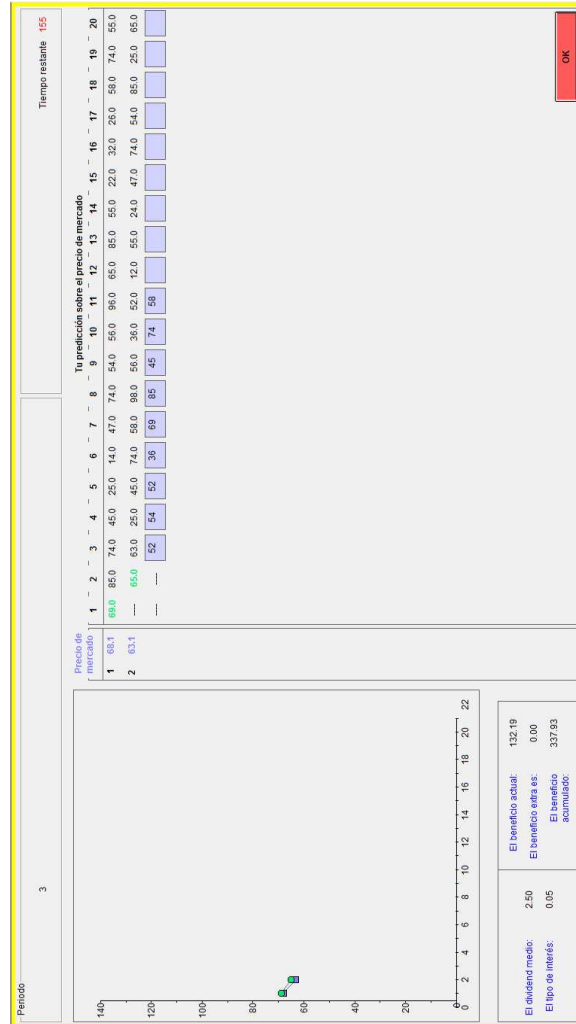


Figure 20: Screen-shot of the experiment.